

The ATLAS Experiment at the CERN Large Hadron Collider

ATLAS Collaboration

ABSTRACT: The ATLAS detector as installed in its experimental cavern at point 1 at CERN is described in this paper. A brief overview of the expected performance of the detector when the Large Hadron Collider begins operation is also presented.

5 KEYWORDS: ATLAS, LHC, CERN.

Contents

	1. Overview of the ATLAS detector	30
	1.1 Introduction	30
10	1.2 Physics requirements and detector overview	30
	1.3 Tracking	34
	1.4 Calorimetry	35
	1.4.1 LAr electromagnetic calorimeter	36
	1.4.2 Hadronic calorimeters	38
15	1.5 Muon system	39
	1.5.1 The toroid magnets	40
	1.5.2 Muon chamber types	41
	1.5.3 Muon chamber alignment and B-field reconstruction	41
	1.6 Forward detectors	42
20	1.7 Trigger, readout, data acquisition, and control systems	42
	1.7.1 Trigger system	43
	1.7.2 Readout architecture and data acquisition	43
	1.8 Radiation, shielding, and interface to the LHC machine	44
	1.8.1 Radiation levels	44
25	1.8.2 Shielding	45
	1.8.3 Beam-pipe	46
	1.8.4 LHC machine interface	46
	1.9 Outline of the paper	46
	2. Magnet system and magnetic field	47
30	2.1 Magnet system	47
	2.1.1 Central solenoid	47
	2.1.2 Barrel toroid	49
	2.1.3 End-cap toroids	51
	2.1.4 Magnet services	52
35	2.1.4.1 Vacuum system	52
	2.1.4.2 Cryogenics	52
	2.1.4.3 Electrical circuits	54
	2.1.4.4 Magnet controls	55
	2.2 Magnetic field determination	56
40	2.2.1 Performance specifications and measurement concepts	56
	2.2.2 B-field modelling	57
	2.2.3 Magnetic field instrumentation and reconstruction	59
	2.2.3.1 B-field sensors	59
	2.2.3.2 B-field reconstruction	60
45	2.2.4 Solenoid-mapping measurements	60

	2.2.4.1	Mapping campaign	60
	2.2.4.2	Mapper geometry, survey and auto-calibration	61
	2.2.4.3	Map fitting	61
	2.2.5	Experimental validation of the field map in the muon spectrometer	62
50	2.2.6	Towards an overall field map for ATLAS data-taking	64
	3.	Background radiation and shielding	65
	3.1	Introduction	65
	3.2	Description of the shielding	66
	3.3	Calculation of particle fluences and absorbed doses	69
55	3.3.1	The inner detector and calorimeter regions	70
	3.3.2	The muon spectrometer region	73
	3.4	Background monitors	74
	3.4.1	Monitors in the inner detector	75
	3.4.2	Monitors in the muon spectrometer	76
60	3.4.3	Network of detectors for radiation measurements	76
	3.5	Activation	77
	4.	Inner detector	80
	4.1	Introduction	80
	4.2	Inner-detector sensors	83
65	4.2.1	Pixel and SCT detector sensors	84
	4.2.2	TRT straw tubes	85
	4.3	Inner-detector modules	87
	4.3.1	Pixel modules and staves	87
	4.3.2	SCT modules	90
70	4.3.3	TRT modules	95
	4.4	Readout of the inner detector	98
	4.4.1	Front-end electronics	98
	4.4.1.1	Pixel front-end electronics	98
	4.4.1.2	SCT front-end electronics	102
75	4.4.1.3	TRT front-end electronics	102
	4.4.2	Data transmission and power-supply services and routing	103
	4.4.2.1	Pixel and SCT readout	105
	4.4.2.2	TRT readout	107
	4.4.2.3	Pixel, SCT and TRT readout drivers	107
80	4.5	Electronics and detector power supplies and services	107
	4.6	Grounding and shielding of the inner detector	110
	4.7	Structure and mechanical integration of the inner detector	111
	4.7.1	Pixel structure and integration	112
	4.7.2	SCT structure and integration	115
85	4.7.3	TRT structure and integration	118
	4.7.4	Integration and installation of the inner-detector components	120

	4.8	Inner detector environment and cooling services	121
	4.8.1	Beam-pipe interface and operational aspects	121
	4.8.2	Inner-detector environmental gas	123
90	4.8.3	Inner-detector cooling	125
	4.8.4	Inner detector controls, safety and interlocks	127
	4.9	Performance status of the integrated inner detector	128
	4.9.1	Electrical performance of the integrated detector	128
	4.9.2	SCT and pixel cooling performance after integration	131
95	4.10	Material distribution of the inner detector	132
	5.	Calorimetry	136
	5.1	Cryostats and associated feed-throughs	137
	5.1.1	Cryostat description	137
	5.1.2	Signal feed-throughs	139
100	5.1.3	High-voltage feed-throughs	139
	5.2	Electromagnetic calorimetry	139
	5.2.1	Accordion geometry	139
	5.2.2	Barrel geometry	140
	5.2.3	End-cap geometry	142
105	5.2.4	Shape and placement of the electromagnetic calorimeters	144
	5.2.5	High-voltage distribution	144
	5.2.6	Electronic boards and cables inside the cryostats	146
	5.2.7	Quality-assurance tests	148
	5.3	Hadronic calorimeters	148
110	5.3.1	Tile calorimeter	148
	5.3.1.1	Overview	148
	5.3.1.2	Mechanical structure	149
	5.3.1.3	Optical components	150
	5.3.1.4	Instrumentation with scintillator and fibres	151
115	5.3.2	Hadronic end-cap calorimeters (HEC)	153
	5.3.2.1	Overview	153
	5.3.2.2	Wheel assembly and installation	155
	5.3.3	Forward calorimeters	156
	5.3.3.1	Overview	156
120	5.3.3.2	Module description	157
	5.3.3.3	Integration into the cryostat	160
	5.3.3.4	Placement of FCal	160
	5.3.3.5	Tests during installation/integration	160
	5.4	Instrumentation in gap between cryostats	161
125	5.5	Calorimeter readout electronics, calibration and services	163
	5.5.1	Readout electronics	163
	5.5.1.1	Overview	163
	5.5.1.2	LAr calorimeter front-end electronics	165

	5.5.1.3	Tile calorimeter front-end electronics	167
130	5.5.1.4	Back-end electronics	169
	5.5.2	Calorimeter calibration systems	171
	5.5.2.1	Calibration of LAr calorimeters	171
	5.5.2.2	Calibration of tile calorimeter	173
	5.5.3	Calorimeter power supplies and services	174
135	5.5.3.1	Low-voltage power supplies	174
	5.5.3.2	High-voltage power supplies	175
	5.5.3.3	Cryogenics and cooling	175
	5.5.3.4	Detector control systems	176
	5.6	Test-beam measurements and results	178
140	5.6.1	Electromagnetic module performance	178
	5.6.2	Hadronic end-cap performance	181
	5.6.3	FCal performance	183
	5.6.4	Tile-calorimeter performance	184
	5.6.4.1	Stand-alone performance	184
145	5.6.4.2	Combined LAr and tile calorimeter test-beam measurements	186
	6.	Muon spectrometer	188
	6.1	Overview	188
	6.2	Precision-tracking chambers	192
	6.3	Monitored drift tube chambers	193
150	6.3.1	Structure and function of the drift tube	193
	6.3.2	Mechanical structure	195
	6.3.3	Signal path and readout electronics	199
	6.3.4	Alignment system	201
	6.3.4.1	Alignment strategies	201
155	6.3.4.2	Optical alignment sensors	203
	6.3.4.3	Layout of the alignment system	204
	6.3.5	Performance of the MDT chambers	206
	6.4	Cathode-strip chambers	207
	6.4.1	Layout of the CSC system	207
160	6.4.2	Spatial and time resolution	208
	6.4.3	Mechanical design	210
	6.4.4	Readout electronics	212
	6.4.5	Performance of the CSC	214
	6.5	Trigger chambers	214
165	6.6	Resistive plate chambers	217
	6.6.1	Principle of operation	218
	6.6.2	Mechanical structure	218
	6.6.3	Signal path and readout electronics	220
	6.7	Thin gap chambers	222
170	6.7.1	Introduction	222

	6.7.2	Principle of operation	223
	6.7.3	Mechanical structure	223
	6.7.4	Signal path, readout, and detector controls	225
	6.8	Commonalities in the muon system	226
175	6.8.1	The gas supplies in the muon system	226
	6.8.2	Electronics services and power consumption	226
	7.	Forward detectors	230
	7.1	The LUCID detector	230
	7.1.1	Detector design	232
180	7.1.2	Experimental results from test-beam and radiation tests	232
	7.1.3	LUCID installation in ATLAS	234
	7.2	The ALFA detector	234
	7.2.1	Experimental results from test-beam	237
	7.2.2	ALFA installation in ATLAS	238
185	7.3	The zero-degree calorimeters	238
	7.3.1	ZDC module description	238
	7.3.2	Calibration and monitoring	239
	7.3.2.1	Flashers	239
	7.3.2.2	Particles	239
190	7.3.3	ZDC installation in ATLAS	239
	8.	Trigger, data acquisition, and controls	242
	8.1	Introduction to event selection and data acquisition	242
	8.2	The L1 trigger	243
	8.2.1	Calorimeter trigger	245
195	8.2.1.1	The analogue front-end	246
	8.2.1.2	The pre-processor	247
	8.2.1.3	The cluster and jet/energy-sum processors	248
	8.2.1.4	The cluster processor module	249
	8.2.1.5	The jet/energy module	250
200	8.2.1.6	The common merger module	251
	8.2.1.7	The processor back-plane	252
	8.2.1.8	The readout driver	252
	8.2.2	Muon trigger	253
	8.2.2.1	Muon barrel trigger	254
205	8.2.2.2	Muon end-cap trigger	256
	8.2.2.3	Muon to central trigger processor interface	259
	8.2.3	Central trigger processor	260
	8.2.3.1	Functional overview	260
	8.2.3.2	System implementation and latency	262
210	8.3	Data acquisition system and high-level trigger	264
	8.3.1	Overview	264

	8.3.2	Control	266
	8.3.3	Configuration	267
	8.3.4	Monitoring and information distribution	268
215	8.3.5	Readout system	269
	8.3.6	L2 trigger	271
	8.3.7	Event-building	273
	8.3.8	Event filter	273
	8.3.9	Event output	274
220	8.4	Implementation and capabilities of the DAQ/HLT	274
	8.5	Detector control system	277
	9.	Integration and installation	280
	9.1	Introduction	280
	9.2	Organisational issues	280
225	9.3	Mechanical integration	281
	9.3.1	Envelopes (individual, global, dynamic)	282
	9.3.2	Survey and placement strategy	282
	9.3.2.1	Survey reference grid in the cavern	282
	9.3.2.2	Stability measurements of the floor and the bed-plates	283
230	9.3.2.3	Placement of ATLAS sub-systems	283
	9.3.2.4	Monitoring of the placement of the detector	285
	9.4	Infrastructure and detector services	288
	9.4.1	Civil engineering	288
	9.4.2	Electrical power distribution	289
235	9.4.3	Air-conditioning and cooling systems	290
	9.4.4	Gas distribution	291
	9.4.5	Cryogenic systems	292
	9.4.6	Racks and cables	294
	9.4.7	Drag-chains and mobile services	296
240	9.4.8	Grounding and electromagnetic compatibility	297
	9.5	Support and access structures	298
	9.5.1	Feet and rail system	298
	9.5.2	Trucks	300
	9.5.3	Surrounding structures (HS and HO)	300
245	9.5.4	Muon barrel access structures	300
	9.5.5	Big wheels	301
	9.6	Detector Installation	302
	9.6.1	Phase 1: infrastructure in the main cavern, feet and rails	302
	9.6.2	Phase 2: barrel calorimetry and barrel toroid	302
250	9.6.3	Phase 3: end-cap calorimeters and muon barrel chambers	303
	9.6.4	Phase 4: muon big wheels, inner detector and completion of muon barrel	303
	9.6.5	Phase 5: end-cap toroid magnets and muon small wheels	305
	9.6.6	Phase 6: beam-pipe and forward shielding	305

	9.7	Access and detector opening	306
255	9.7.1	Access scenarios	306
	9.7.2	Movement system	306
	9.8	Beam-pipe	307
	9.9	Safety in ATLAS	308
	9.9.1	Organisation of safety	309
260	9.9.2	Access control	310
	9.9.2.1	General aspects	310
	9.9.2.2	INB regulations	310
	9.9.3	Safety systems	311
	9.9.3.1	Hazard-detection systems	311
265	9.9.3.2	Foam extinguishing system	312
	9.9.3.3	Finding people inside ATLAS	312
	9.9.4	Detector safety system	312
	9.9.5	Safety during operation	312
	9.10	Interface to the LHC machine	313
270	10.	Expected performance of the ATLAS detector	316
	10.1	Introduction	316
	10.1.1	Realistic data challenge	316
	10.1.2	Combined test-beam	318
	10.2	Reconstruction and identification of charged particles in the inner detector	320
275	10.2.1	Track reconstruction	320
	10.2.2	Alignment of the inner detector	322
	10.2.2.1	Alignment in the combined test-beam	324
	10.2.2.2	Misaligned inner detector in ATLAS simulation	324
	10.2.3	Tracking performance for single particles and particles in jets	326
280	10.2.4	Vertexing performance	330
	10.2.5	Particle identification, reconstruction of electrons and photon conversions	331
	10.3	Muon reconstruction and identification	335
	10.3.1	Introduction	335
	10.3.2	Calibration and alignment	336
285	10.3.2.1	Performance of optical alignment system in test-beam	337
	10.3.2.2	Alignment of the muon spectrometer with tracks	338
	10.3.2.3	Overall calibration and alignment strategy	338
	10.3.3	Reconstruction strategies	338
	10.3.4	Muon reconstruction performance for single muons	340
290	10.3.5	Reconstruction of $Z \rightarrow \mu\mu$ and $H \rightarrow \mu\mu\mu\mu$ decays	342
	10.4	Electrons and photons	343
	10.4.1	Calibration and performance of the electromagnetic calorimeter	344
	10.4.2	Electron and photon reconstruction and identification	348
	10.4.2.1	Electrons	349
295	10.4.2.2	Photons	352

	10.4.2.3	Reconstruction of $H \rightarrow eeee$ and $H \rightarrow \gamma\gamma$ final states	353
	10.4.3	Assessment of performance in situ with initial data	355
	10.5	Jet reconstruction	355
	10.5.1	Jet clustering algorithms	356
300	10.5.2	Input to jet reconstruction	356
	10.5.3	Jet calibration	358
	10.5.4	Jet signal characteristics	359
	10.5.4.1	Jet signal linearity and energy resolution	360
	10.5.4.2	Jet direction measurement	361
305	10.5.4.3	Jet signal uniformity	362
	10.5.5	Jet reconstruction performance	362
	10.5.6	Validation of jet calibration with in-situ measurements	364
	10.6	Missing transverse energy	365
	10.6.1	Reconstruction and calibration of E_T^{miss}	365
310	10.6.2	Evaluation of E_T^{miss} performance	367
	10.6.3	Measurement of E_T^{miss} direction	368
	10.6.4	Use of E_T^{miss} for mass reconstruction	368
	10.6.5	Fake E_T^{miss}	368
	10.7	Hadronic τ -decays	370
315	10.7.1	Track reconstruction in hadronic τ -decays	371
	10.7.2	Electromagnetic clusters in single-prong decays	371
	10.7.3	Identification of hadronic τ -decays and rejection of QCD jets	373
	10.8	Flavour tagging	375
	10.8.1	Ingredients of b -tagging algorithms	375
320	10.8.1.1	Track selection and track impact parameters	376
	10.8.1.2	Secondary vertices	376
	10.8.2	Likelihood-ratio tagging algorithms	377
	10.8.3	Jet activity and jet purification	377
	10.8.4	Expected b -tagging performance	379
325	10.8.5	Soft-lepton tagging	379
	10.9	Trigger performance	381
	10.9.1	Overview	381
	10.9.2	Selection strategy	381
	10.9.3	Trigger menus	382
330	10.9.4	Examples of trigger performance	384
	10.9.4.1	Electrons and photons	384
	10.9.4.2	τ -leptons and E_T^{miss}	385
	10.9.4.3	Jets	386
	10.9.4.4	Muons	387
335	10.9.4.5	B -physics	388
	10.9.5	Trigger commissioning	389
	10.9.6	Evolution to higher luminosities	390

	10.9.7 Measurements of trigger efficiency from data	392
	11. Conclusions and outlook	394
340	11.1 Detector installation and hardware status	394
	11.2 Outlook on commissioning with data	397
	11.3 Future changes to the ATLAS detector system	399
	A. Appendix: ATLAS acronyms	401
	B. Acknowledgements	405
345	C. Annex	406

The ATLAS Collaboration

G. Aad⁸², E. Abat^{18,*}, J. Abdallah¹⁶², A.A. Abdelalim⁴⁷, A. Abdesselam¹¹⁷, B.A. Abi¹¹²,
M. Abolins⁸⁷, H. Abramowicz¹⁵¹, E. Acerbi⁸⁸, B.S. Acharya¹⁵⁹, R. Achenbach⁵⁶, M. Ackers²¹,
350 D.L. Adams²⁴, F. Adamyan¹⁶⁹, T.N. Addy⁵⁴, M. Aderholz⁹⁹, C. Adorisio³⁶, P. Adragna⁷³,
M. Aharrouche⁷⁹, S.P. Ahlen²², F. Ahles⁴⁶, A. Ahmad¹⁴⁷, H. Ahmed², G. Aielli¹³⁴,
P.F. Åkesson²⁹, T.P.A. Åkesson⁷⁷, A.V. Akimov⁹⁴, S.M. Alam¹, J. Albert¹⁶⁴, S. Albrand⁵³,
M. Aleksa²⁹, I.N. Aleksandrov⁶³, M. Aleppo⁸⁸, F. Alessandria⁸⁸, C. Alexa²⁵, G. Alexander¹⁵¹,
T. Alexopoulos⁹, G. Alimonti⁸⁸, M. Aliyev¹⁰, P.P. Allport⁷¹, S.E. Allwood-Spiers⁵¹, A. Aloisio¹⁰²,
355 J. Alonso¹⁴, R. Alves¹²³, M.G. Alviggi¹⁰², K. Amako⁶⁴, P. Amaral²⁹, S.P. Amaral²⁹,
G. Ambrosini¹⁶, G. Ambrosio⁸⁸, C. Amelung²⁹, V.V. Ammosov¹²⁷, A. Amorim¹²³, N. Amram¹⁵¹,
C. Anastopoulos¹⁵², B. Anderson⁷⁵, K.J. Anderson³⁰, E.C. Anderssen¹⁴, A. Andreazza⁸⁸,
V. Andrei⁵⁶, L. Andricek⁹⁹, M-L. Andrieux⁵³, X.S. Anduaga⁶⁸, F. Anghinolfi²⁹, A. Antonaki⁸,
M. Antonelli⁴⁵, S. Antonelli²⁰, R. Apsimon¹²⁸, G. Arabidze⁸, I. Aracena¹⁴³, Y. Arai⁶⁴,
360 J.P. Archambault²⁸, J-F. Arguin¹⁴, E. Arik^{18,*}, M. Arik¹⁸, K.E. Arms¹⁰⁹, S.R. Armstrong²⁴,
M. Arnaud¹³⁶, C. Arnault¹¹⁴, A. Artamonov⁹⁵, S. Asai¹⁵³, S. Ask⁸⁰, B. Åsman¹⁴⁵, D. Asner²⁸,
L. Asquith⁷⁵, K. Assamagan²⁴, A. Astbury¹⁶⁴, B. Athar¹, T. Atkinson⁸⁵, B. Aubert⁴,
B. Auerbach¹⁶⁸, E. Auge¹¹⁴, K. Augsten¹²⁶, V.M. Aulchenko¹⁰⁷, N. Austin⁷¹, G. Avolio²⁹,
R. Avramidou⁹, A. Axen¹⁶³, C. Ay⁵², G. Azuelos^{92,a}, G. Baccaglioni⁸⁸, C. Bacci¹³⁵,
365 H. Bachacou¹³⁶, K. Bachas¹⁵², G. Bachy²⁹, E. Badescu²⁵, P. Bagnaia¹³³, D.C. Bailey¹⁵⁵,
J.T. Baines¹²⁸, O.K. Baker¹⁶⁸, F. Ballester¹⁶², F. Baltasar Dos Santos Pedrosa²⁹, J. Ban³⁴,
E. Banas³⁸, D. Banfi⁸⁸, A. Bangert⁹⁹, V. Bansal¹²², S.P. Baranov⁹⁴, S. Baranov⁵, A. Barashkou⁶³,
E.L. Barberio⁸⁵, D. Barberis⁴⁸, G. Barbier⁴⁷, P. Barclay¹²⁸, D.Y. Bardin⁶³, P. Bargassa¹¹⁷,
T. Barillari⁹⁹, M. Barisonzi⁴⁰, B.M. Barnett¹²⁸, R.M. Barnett¹⁴, S. Baron²⁹, A. Baroncelli¹³⁵,
370 M. Barone⁴⁵, A.J. Barr¹¹⁷, F. Barreiro⁷⁸, P. Barrillon¹¹⁴, A. Barriuso Poy²⁹, N. Barros¹²³,
V. Bartheld⁹⁹, H. Bartko⁹⁹, R. Bartoldus¹⁴³, S. Basiladze⁹⁷, J. Bastos¹²³, L.E. Batchelor¹²⁸,
R.L. Bates⁵¹, J.R. Batley²⁷, S. Batraneanu²⁹, M. Battistin²⁹, G. Battistoni⁸⁸, V. Batusov⁶³,
F. Bauer¹³⁶, B. Bauss⁷⁹, D.E. Baynham¹²⁸, M. Bazalova¹²⁴, A. Bazan⁴, P.H. Beauchemin⁹²,
B. Beaugiraud⁴, R.B. Beccherle⁴⁸, G.A. Beck⁷³, H.P. Beck¹⁶, K.H. Becks¹⁶⁷, I. Bedajane¹²⁶,
375 A.J. Beddall^{18,b}, A. Beddall^{18,b}, P. Bednár¹⁴⁴, V.A. Bednyakov⁶³, C. Bee⁸², S. Behar Harpaz¹⁵⁰,
G.A.N. Belanger²⁸, C. Belanger- Champagne¹⁶⁰, B. Belhorma⁵³, P.J. Bell⁸⁰, W.H. Bell⁵¹,
G. Bella¹⁵¹, F. Bellachia⁴, L. Bellagamba²⁰, F. Bellina¹⁶⁷, G. Bellomo⁸⁸, M. Bellomo¹¹⁸,
O. Beltramello²⁹, A. Belymam⁷³, S. Ben Ami¹⁵⁰, M. Ben Moshe¹⁵¹, O. Benary¹⁵¹,
D. Benchevkroun⁹³, C. Benchouk⁸², M. Bendel⁷⁹, B.H. Benedict⁶², N. Benekos¹⁶¹, J. Benes¹²⁶,
380 Y. Benhammou¹⁵¹, G.P. Benincasa¹²³, D.P. Benjamin⁴³, J.R. Bensinger²³, K. Benslama¹²⁹,
S. Bentvelsen¹⁰⁶, M. Beretta⁴⁵, D. Berge²⁹, E. Bergeaas¹⁴⁵, N. Berger⁴, F. Berghaus¹⁶⁴,
S. Berglund¹⁴⁵, F. Bergsma²⁹, J. Beringer¹⁴, J. Bernabéu¹⁶², K. Bernardet⁸², C. Berriaud¹³⁶,
T. Berry⁷⁴, H. Bertelsen³⁵, A. Bertin²⁰, F. Bertinelli²⁹, S. Bertolucci⁴⁵, N. Besson¹³⁶,
A. Beteille⁵³, S. Bethke⁹⁹, W. Bialas³⁷, R.M. Bianchi⁴⁶, M. Bianco⁷⁰, O. Biebel⁹⁸, M. Bieri¹⁴²,
385 M. Biglietti¹⁰², H. Bilokon⁴⁵, M. Binder⁹⁸, S. Binet¹⁴, N. Bingefors¹⁶⁰, A. Bingul^{18,b}, C. Bini¹³³,
C. Biscarat¹¹, R. Bischof⁶⁰, M. Bischofberger⁸⁵, A. Bitadze²⁹, J.P. Bizzell¹²⁸, K. Black⁵⁵,
R.E. Blair⁵, J.J. Blaising⁴, O. Blanch¹¹, G. Blanchot²⁹, C. Blocker²³, J. Blocki³⁸, A. Blondel⁴⁷,
W. Blum⁷⁹, U. Blumenschein¹¹, C. Boaretto¹³³, G.J. Bobbink¹⁰⁶, A. Bocci⁴³, D. Bocian³⁸,

R. Bock²⁹, M. Boehm⁹⁸, J. Boek¹⁶⁷, J.A. Bogaerts²⁹, A. Bogouch⁸⁹, C. Bohm¹⁴⁵, J. Bohm¹²⁴,
 390 V. Boisvert²⁹, T. Bold⁶², V. Boldea²⁵, V.G. Bondarenko⁹⁶, R. Bonino⁴⁷, J. Bonis¹¹⁴,
 W. Bonivento⁸⁸, P. Bonneau²⁹, M. Boonekamp¹³⁶, G. Boorman⁷⁴, M. Boosten²⁹, C.N. Booth¹³⁹,
 P.S.L. Booth^{71,*}, P. Booth¹³⁹, J.R.A. Booth¹⁷, K. Borer¹⁶, A. Borisov¹²⁷, I. Borjanovic⁷⁰,
 K. Bos¹⁰⁶, D. Boscherini²⁰, F. Bosi¹²¹, M. Bosman¹¹, M. Bosteels²⁹, B. Botchev¹⁴⁷,
 H. Boterenbrood¹⁰⁶, D. Botterill¹²⁸, J. Boudreau¹²², E.V. Bouhova-Thacker⁶⁹,
 395 C. Boulahouache¹²², C. Bourdarios¹¹⁴, M. Boutemeur⁹⁸, K. Bouzakis¹⁵², G.R. Boyd¹¹¹,
 J. Boyd²⁹, B.H. Boyer⁵³, I.R. Boyko⁶³, N.I. Bozhko¹²⁷, S. Braccini⁴⁵, A. Braem²⁹,
 P. Branchini¹³⁵, G.W. Brandenburg⁵⁵, A. Brandt⁹⁸, O. Brandt¹¹⁷, U. Bratzler¹⁵⁴, H.M. Braun¹⁶⁷,
 S. Bravo¹¹, I.P. Brawn¹²⁸, B. Brelier^{92,c}, J. Bremer²⁹, R. Brenner¹⁶⁰, S. Bressler¹⁵⁰, D. Breton¹¹⁴,
 N.D. Brett¹¹⁷, P. Breugnon⁸², P.G. Bright-Thomas¹⁷, F.M. Brochu²⁷, I. Brock²¹, R. Brock⁸⁷,
 400 T.J. Brodbeck⁶⁹, E. Brodet¹⁵¹, F. Broggi⁸⁸, Z. Broklova¹²⁵, C. Bromberg⁸⁷, G. Brooijmans³⁴,
 G. Brouwer¹⁰⁶, J. Broz¹²⁵, E. Brubaker³⁰, P.A. Bruckman de Renstrom^{117,d}, D. Bruncko¹⁴⁴,
 A. Bruni²⁰, G. Bruni²⁰, M. Bruschi²⁰, T. Buanes¹³, N.J. Buchanan², P. Buchholz¹⁴¹,
 I.A. Budagov⁶³, V. Büscher²¹, L. Bugge¹¹⁶, D. Buiria-Clark¹¹⁷, E.J. Buis¹⁰⁶, F. Bujor²⁹,
 T. Buran¹¹⁶, H. Burckhart²⁹, D. Burckhart-Chromek²⁹, S. Burdin⁷¹, R. Burns²⁴, E. Busato³³,
 405 J.J.F. Buskop¹⁰⁶, K.P. Buszello⁴⁶, F. Butin²⁹, J.M. Butler²², C.M. Buttar⁵¹, J. Butterworth¹²⁸,
 J.M. Butterworth⁷⁵, T. Byatt⁷⁵, S. Cabrera Urbán¹⁶², E. Cabruja Casas¹⁶², M. Caccia⁸⁸,
 D. Caforio²⁰, O. Cakir³, P. Calafiura¹⁴, G. Calderini⁷⁶, D. Calderón Terol¹⁶², J. Callahan⁵⁹,
 L.P. Caloba¹³⁰, R. Caloi¹³³, D. Calvet³³, A. Camard⁷⁶, F. Camarena¹⁶², P. Camarri¹³⁴,
 M. Cambiaghi¹¹⁸, D. Cameron¹¹⁶, J. Cammin²¹, F. Campabadal Segura¹⁶², S. Campana²⁹,
 410 V. Canale¹⁰², J. Cantero⁷⁸, M.D.M. Capeans Garrido²⁹, I. Caprini²⁵, M. Caprini²⁵, M. Caprio¹⁰²,
 D. Caracinha¹²³, C. Caramarcu²⁵, Y. Carcagno⁵³, R. Cardarelli¹³⁴, C. Cardeira¹²³,
 L. Cardiel Sas²⁹, A. Cardini¹³³, T. Carli²⁹, G. Carlino¹⁰², L. Carminati⁸⁸, B. Caron^{2,a}, S. Caron⁴⁶,
 C. Carpentieri⁴⁶, F.S. Carr¹²⁸, A.A. Carter⁷³, J.R. Carter²⁷, J. Carvalho¹²³, D. Casadei¹⁰⁸,
 M.P. Casado¹¹, M. Cascella¹²¹, C. Caso^{48,*}, J. Castelo¹⁶², V. Castillo Gimenez¹⁶², N. Castro¹²³,
 415 F. Castrovillari³⁶, G. Cataldi⁷⁰, F. Cataneo⁸⁸, A. Catinaccio²⁹, J.R. Catmore⁶⁹, A. Cattai²⁹,
 S. Caughron³⁴, D. Cauz¹⁵⁹, A. Cavallari¹³³, P. Cavalleri⁷⁶, D. Cavalli⁸⁸, M. Cavalli-Sforza¹¹,
 V. Cavasinni¹²¹, F. Ceradini¹³⁵, C. Cerna⁸², C. Cernoch²⁹, A.S. Cerqueira¹³⁰, A. Cerri²⁹,
 F. Cerutti⁴⁵, M. Cervetto⁴⁸, S.A. Cetin^{18,e}, F. Cevenini¹⁰², M. Chalifour¹³⁶, M. Chamizo Ilatas⁴⁷,
 A. Chan²⁴, J.W. Chapman⁸⁶, D.G. Charlton¹⁷, S. Charron⁹², S.V. Chekulaev¹²⁷, G.A. Chelkov⁶³,
 420 H. Chen²⁴, L. Chen², T. Chen³², X. Chen¹⁶⁶, S. Cheng³², T.L. Cheng⁷⁴, A. Cheplakov⁵¹,
 V.F. Chepurnov⁶³, R. Cherkaoui El Moursli⁹³, D. Chesneau²⁵, E. Cheu⁶, L. Chevalier¹³⁶,
 J.L. Chevalley²⁹, F. Chevallier⁵³, V. Chiarella⁴⁵, G. Chiefari¹⁰², L. Chikovani⁴⁹, A. Chilingarov⁶⁹,
 G. Chiodini⁷⁰, S. Chouridou¹³⁷, D. Chren¹²⁶, T. Christiansen⁹⁸, I.A. Christidi¹⁵², A. Christov⁴⁶,
 M.L. Chu¹⁴⁹, J. Chudoba¹²⁴, A.G. Chuguev¹²⁷, G. Ciapetti¹³³, E. Cicalini¹²¹, A.K. Ciftci³,
 425 V. Cindro⁷², M.D. Ciobotaru⁶², A. Ciochio¹⁴, M. Cirilli⁸⁶, M. Citterio⁸⁸, M. Ciubancan²⁵,
 J.V. Civera¹⁶², A. Clark⁴⁷, W. Cleland¹²², J.C. Clemens⁸², C. Clément¹⁴⁵, D. Clements⁵¹,
 R.W. Clift¹²⁸, M. Cobal¹⁵⁹, A. Coccaro⁴⁸, J. Cochran⁶¹, R. Coco⁹¹, P. Coe¹¹⁷, S. Coelli⁸⁸,
 E. Cogneras³³, C.D. Cojocaru²⁸, J. Colas⁴, A.P. Colijn¹⁰⁶, C. Collard¹¹⁴, C. Collins-Tooth⁵¹,
 J. Collot⁵³, R. Coluccia⁷⁰, G. Comune⁸⁷, P. Conde Muiño¹²³, E. Coniavitis¹⁶⁰, M. Consonni⁴,
 430 S. Constantinescu²⁵, C. Conta¹¹⁸, F.A. Conventi¹⁰², J. Cook²⁹, M. Cooke³⁴,
 N.J. Cooper-Smith⁷⁴, T. Cornelissen²⁹, M. Corradi²⁰, S. Correard⁸², A. Corso-Radu⁶², J. Coss¹⁴⁴,

G. Costa⁸⁸, M.J. Costa¹⁶², D. Costanzo¹³⁹, T. Costin³⁰, R. Coura Torres¹³⁰, L. Courneyea¹⁶⁴,
C. Couyoumtzelis⁴⁷, G. Cowan⁷⁴, B.E. Cox⁸⁰, J. Cox¹¹⁷, D.A. Cragg¹²⁸, K. Cranmer¹⁰⁸,
J. Cranshaw⁵, M. Cristinziani²¹, G. Crosetti³⁶, C. Cuenca Almenar¹⁶², S. Cuneo⁴⁸, A. Cunha²⁴,
435 M. Curatolo⁴⁵, C.J. Curtis¹⁷, P. Cwetanski⁵⁹, Z. Czcyczula³⁵, S. D'Auria⁵¹, M. D'Onofrio⁴⁷,
A. Da Rocha Gesualdi Mello¹³⁰, P.V.M. Da Silva¹³⁰, R. Da Silva²⁹, W. Dabrowski³⁷, A. Dael¹³⁶,
A. Dahlhoff⁴⁶, T. Dai⁸⁶, C. Dallapiccola⁸³, S.J. Dallison¹²⁸, J. Dalmau⁷³, C.H. Daly¹³⁸,
M. Dam³⁵, D. Damazio²⁴, M. Dameri⁴⁸, K.M. Danielsen¹¹⁶, H.O. Danielsson²⁹, R. Dankers¹⁰⁶,
D. Dannheim²⁹, G. Darbo⁴⁸, P. Dargent⁸², C. Daum¹⁰⁶, J.P. Dauvergne²⁹, M. David¹²³,
440 T. Davidek¹²⁵, N. Davidson⁸⁵, R. Davidson⁶⁹, I. Dawson¹³⁹, J.W. Dawson⁵, R.K. Daya³⁹, K. De⁷,
R. de Asmundis¹⁰², R. de Boer¹⁰⁶, S. De Castro²⁰, N. De Groot¹⁰⁵, P. de Jong¹⁰⁶,
X. de La Broise¹³⁶, E. De La Cruz-Burelo⁸⁶, C. De La Taille¹¹⁴, B. De Lotto¹⁵⁹,
M. De Oliveira Branco²⁹, D. De Pedis¹³³, P. de Saintignon⁵³, A. De Salvo¹³³, U. De Sanctis⁸⁸,
A. De Santo⁷⁴, J.B. De Vivie De Regie¹¹⁴, G. De Zorzi¹³³, S. Dean⁷⁵, G. Dedes⁹⁹,
445 D.V. Dedovich⁶³, P.O. Defay³³, R. Degele⁷⁹, M. Dehchar¹¹⁷, M. Deile⁹⁸, C. Del Papa¹⁵⁹,
J. Del Peso⁷⁸, T. Del Prete¹²¹, E. Delagnes¹³⁶, P. Delebecque⁴, A. Dell'Acqua²⁹,
M. Della Pietra¹⁰², D. della Volpe¹⁰², M. Delmastro²⁹, P. Delpierre⁸², N. Delruelle²⁹,
P.A. Delsart⁴, C. Deluca Silberberg¹¹, S. Demers¹⁴³, M. Demichev⁶³, P. Demierre⁴⁷,
B. Demirköz²⁹, W. Deng²⁴, S.P. Denisov¹²⁷, C. Dennis¹¹⁷, C.J. Densham¹²⁸, M. Dentan¹³⁶,
450 J.E. Derkaoui⁹³, F. Derue⁷⁶, P. Dervan⁷¹, K.K. Desch²¹, A. Dewhurst⁶⁹, A. Di Ciaccio¹³⁴,
L. Di Ciaccio⁴, A. Di Domenico¹³³, A. Di Girolamo¹³³, B. Di Girolamo²⁹, S. Di Luise¹³⁵,
A. Di Mattia⁸⁷, A. Di Simone²⁹, M.M. Diaz Gomez⁴⁷, E.B. Diehl⁸⁶, H. Dietl⁹⁹, W. Dietsche²¹,
S. Diglio¹³⁵, M. Dima¹⁶⁷, K. Dindar³⁹, B. Dinkespiler³⁹, C. Dionisi¹³³, R. Dipanjan⁷, P. Dita²⁵,
S. Dita²⁵, F. Dittus²⁹, S.D. Dixon¹³⁹, F. Djama⁸², R. Djilkibaev¹⁰⁸, T. Djobava⁴⁹,
455 M.A.B. do Vale¹³⁰, M. Dobbs⁸⁴, R. Dobinson^{29,*}, D. Dobos²⁹, E. Dobson¹¹⁷, M. Dobson²⁹,
J. Dodd³⁴, O.B. Dogan^{18,*}, T. Doherty⁵¹, Y. Doi⁶⁴, J. Dolejsi¹²⁵, I. Dolenc⁷², Z. Dolezal¹²⁵,
B.A. Dolgoshein⁹⁶, E. Domingo¹¹, M. Donega²⁹, J. Dopke¹⁶⁷, D.E. Dorfan¹³⁷, O. Dorholt¹¹⁶,
A. Doria¹⁰², A. Dos Anjos¹⁶⁶, M. Dosil¹¹, A. Dotti¹²¹, M.T. Dova⁶⁸, J.D. Dowell¹⁷, A.T. Doyle⁵¹,
G. Drake⁵, D. Drakoulakos²⁹, Z. Drasal¹²⁵, J. Drees¹⁶⁷, N. Dressnandt¹¹⁹, H. Drevermann²⁹,
460 C. Driouichi³⁵, M. Dris⁹, J.G. Drohan⁷⁵, J. Dubbert⁹⁹, T. Dubbs¹³⁷, E. Duchovni¹⁶⁵,
G. Duckeck⁹⁸, A. Dudarev²⁹, M. Dührssen⁴⁶, H. Dür⁶⁰, I.P. Duerdoth⁸⁰, S. Duffin²⁴, L. Dufflot¹¹⁴,
M-A. Dufour⁸⁴, N. Dumont Dayot⁴, H. Duran Yildiz^{3,f}, D. Durand¹³⁶, A. Dushkin²³,
R. Duxfield¹³⁹, M. Dwuznik³⁷, F. Dydak²⁹, D. Dzahini⁵³, S. Díez Cornell¹⁶², M. Düren⁵⁰,
W.L. Ebenstein⁴³, S. Eckert⁴⁶, S. Eckweiler⁷⁹, P. Eerola⁷⁷, I. Efthymiopoulos⁴⁷, U. Egede⁷⁷,
465 K. Egorov⁵⁹, W. Ehrenfeld^{40,g}, T. Eifert⁴⁷, G. Eigen¹³, K. Einsweiler¹⁴, E. Eisenhandler⁷³,
T. Ekelof¹⁶⁰, L.M. Eklund²⁹, M. El Kacimi⁹³, M. Ellert¹⁶⁰, S. Elles⁴, N. Ellis²⁹, J. Elmsheuser⁹⁸,
M. Elsing²⁹, R. Ely¹⁴, D. Emel'yanov¹²⁸, R. Engelmann¹⁴⁷, M. Engström¹⁴⁵, P. Ennes¹⁰⁶,
B. Epp⁶⁰, A. Eppig⁸⁶, V.S. Epshteyn⁹⁵, A. Ereditato¹⁶, V. Eremin⁹⁷, D. Eriksson¹⁴⁵, I. Ermoline⁸⁷,
J. Ernwein¹³⁶, D. Errede¹⁶¹, S. Errede¹⁶¹, M. Escalier⁷⁶, C. Escobar¹⁶², X. Espinal Curull¹¹,
470 B. Esposito⁴⁵, F. Esteves¹²³, F. Etienne⁸², A.I. Etienvre¹³⁶, E. Etzion¹⁵¹, H. Evans⁵⁹,
V.N. Evdokimov¹²⁷, P. Evtoukhovitch⁶³, A. Eyring²¹, L. Fabbri²⁰, C.W. Fabjan²⁹, C. Fabre²⁹,
P. Faccioli²⁰, K. Facius³⁵, V. Fadeyev¹³⁷, R.M. Fakhruddinov¹²⁷, S. Falciano¹³³, I. Falleau¹¹⁴,
A.C. Falou¹¹⁴, Y. Fang¹⁶⁶, M. Fanti⁸⁸, A. Farbin⁷, A. Farilla¹³⁵, J. Farrell²⁴, P. Farthouat²⁹,
D. Fasching¹⁶⁶, F. Fassi¹⁶², P. Fassnacht²⁹, D. Fassouliotis⁸, F. Fawzi⁸⁵, L. Fayard¹¹⁴, F. Fayette⁷⁶,

475 R. Febbraro³³, O.L. Fedin¹²⁰, I. Fedorko²⁹, L. Feld⁴⁶, G. Feldman⁵⁵, L. Feligioni⁸², C. Feng³²,
 E.J. Feng³⁰, J. Fent⁹⁹, A.B. Fenyuk¹²⁷, J. Ferencei¹⁴⁴, D. Ferguson¹⁶⁶, J. Ferland⁹²,
 W. Fernando¹⁰⁹, S. Ferrag⁵¹, A. Ferrari⁵³, P. Ferrari²⁹, R. Ferrari¹¹⁸, A. Ferrer¹⁶², M.L. Ferrer⁴⁵,
 D. Ferrere⁴⁷, C. Ferretti⁸⁶, F. Ferro⁴⁸, M. Fiascarini¹¹⁷, S. Fichet⁷⁶, F. Fiedler⁹⁸, V. Filimonov¹²⁰,
 A. Filipčič⁷², A. Filippas⁹, F. Filthaut¹⁰⁵, M. Fincke-Keeler¹⁶⁴, G. Finocchiaro¹⁴⁷, L. Fiorini¹¹,
 480 A. Firan³⁹, G. Fischer¹⁵, P. Fischer²¹, M.J. Fisher¹⁰⁹, S.M. Fisher¹²⁸, V. Flaminio¹²¹,
 J. Flammer²⁹, M. Flechl¹⁶⁰, I. Fleck¹⁴¹, W. Flegel²⁹, P. Fleischmann¹³⁶, S. Fleischmann²¹,
 C.M. Fleta Corral¹⁶², F. Fleuret⁷⁶, T. Flick¹⁶⁷, J. Flix¹¹, L.R. Flores Castillo¹⁶⁶, M.J. Flowerdew⁷¹,
 F. Föhlisch⁵⁶, M. Fokitis⁹, T.M. Fonseca Martin²⁹, J. Fopma¹¹⁷, D.A. Forbush¹³⁸, A. Formica¹³⁶,
 J.M. Foster⁸⁰, D. Fournier¹¹⁴, A. Foussat²⁹, A.J. Fowler⁴³, H. Fox⁴⁶, P. Francavilla¹²¹,
 485 D. Francis²⁹, S. Franz²⁹, J.T. Fraser⁷⁵, M. Fraternali¹¹⁸, S. Fratianni²⁹, J. Freestone⁸⁰,
 R.S. French¹³⁹, K. Fritsch⁹⁹, D. Froidevaux²⁹, J.A. Frost²⁷, C. Fukunaga¹⁵⁴, J. Fulachier⁵³,
 E. Fullana Torregrosa⁵, J. Fuster¹⁶², C. Gabaldon⁷⁸, S. Gadomski^{47,d}, G. Gagliardi⁴⁸, P. Gagnon⁵⁹,
 E.J. Gallas¹¹⁷, M.V. Gallas²⁹, B.J. Gallop¹²⁸, K.K. Gan¹⁰⁹, F.C. Gannaway⁷³, Y.S. Gao^{143,h},
 V.A. Gapienko¹²⁷, A. Gaponenko¹⁴, M. Garcia-Sciveres¹⁴, C. García¹⁶², J.E. García Navarro¹⁶²,
 490 V. Garde³³, R.W. Gardner³⁰, N. Garelli⁴⁸, H. Garitaonandia¹¹, V.G. Garonne²⁹, J. Garvey¹⁷,
 C. Gatti⁴⁵, G. Gaudio¹³⁸, O. Gaumer⁴⁷, V. Gautard¹³⁶, P. Gauzzi¹³³, I.L. Gavrilenko⁹⁴, C. Gay¹⁶³,
 J-C. Gayde²⁹, E.N. Gazis⁹, E. Gazo¹⁴⁴, C.N.P. Gee¹²⁸, C. Geich-Gimbel²¹, K. Gellerstedt¹⁴⁵,
 C. Gemme⁴⁸, M.H. Genest⁹⁸, S. Gentile¹³³, M.A. George⁷¹, S. George⁷⁴, P. Gerlach¹⁶⁷,
 Y. Gernizky¹⁵⁰, C. Geweniger⁵⁶, H. Ghazlane⁹³, V.M. Ghete⁶⁰, P. Ghez⁴, N. Ghodbane⁹⁹,
 495 B. Giacobbe²⁰, S. Giagu¹³³, V. Giakoumopoulou⁸, V. Giangiobbe¹²¹, F. Gianotti²⁹, B. Gibbard²⁴,
 A. Gibson¹⁵⁵, M.D. Gibson¹²⁸, S.M. Gibson¹¹⁷, G.F. Gieraltowski⁵, I. Gil Botella¹⁶²,
 L.M. Gilbert¹¹⁷, M. Gilchriese¹⁴, O. Gildemeister²⁹, V. Gilewsky⁹⁰, A.R. Gillman¹²⁸,
 D.M. Gingrich^{2,a}, J. Ginzburg¹⁵¹, N. Giokaris⁸, M.P. Giordani¹⁵⁹, C.G. Girard⁴, P.F. Giraud¹³⁶,
 P. Girtler⁶⁰, D. Giugni⁸⁸, P. Giusti²⁰, B.K. Gjølsten¹⁶, C. Glasman⁷⁸, A. Glazov⁴⁰, K.W. Glitza¹⁶⁷,
 500 G.L. Glonti⁶³, K.G. Gnanvo⁷³, J. Godlewski²⁹, T. Göpfert⁴², C. Gössling⁴¹, T. Göttfert⁹⁹,
 S. Goldfarb⁸⁶, D. Goldin³⁹, N. Goldschmidt¹⁶⁶, T. Golling¹⁴, N.P. Gollub²⁹, P.J. Golonka³⁷,
 S.N. Golovnia¹²⁷, A. Gomes¹²³, J. Gomes¹²³, R. Gonçalo⁷⁴, A. Gongadze⁶³, A. Gonidec²⁹,
 S. Gonzalez¹⁶⁶, S. González de la Hoz¹⁶², V. González Millán¹⁶², M.L. Gonzalez Silva²⁶,
 B. Gonzalez-Pineiro⁸⁷, S. González-Sevilla¹⁶², M.J. Goodrick²⁷, J.J. Goodson¹⁴⁷, L. Goossens²⁹,
 505 P.A. Gorbounov¹⁵⁵, A. Gordeev²⁴, H. Gordon²⁴, I. Gorelov¹⁰⁴, G. Gorfine¹⁶⁷, B. Gorini²⁹,
 E. Gorini⁷⁰, A. Gorišek⁷², E. Gornicki³⁸, S.A. Gorokhov¹²⁷, B.T. Gorski²⁹, S.V. Goryachev¹²⁷,
 V.N. Goryachev¹²⁷, M. Gosselink¹⁰⁶, M.I. Gostkin⁶³, M. Gouanère⁴, I. Gough Eschrich⁶²,
 D. Goujdami⁹³, M. Goulette²⁹, I. Gousakov⁶³, J. Gouveia¹²³, S. Gowdy¹⁴³, C. Goy⁴,
 I. Grabowska-Bold²⁹, V. Grabski¹⁶⁹, P. Grafström²⁹, C. Grah¹⁶⁷, K-J. Grahn¹⁴⁶, F. Grancagnolo⁷⁰,
 510 S. Grancagnolo⁷⁰, H. Grassmann¹⁵⁹, V. Gratchev¹²⁰, H.M. Gray³⁴, E. Graziani¹³⁵, B. Green⁷⁴,
 A. Greenall⁷¹, D. Greenfield¹²⁸, D. Greenwood^{24,i}, I.M. Gregor¹⁶⁷, A. Grewal¹¹⁷,
 E. Griesmayer⁴⁴, N. Grigalashvili⁶³, C. Grigson¹³⁹, A.A. Grillo¹³⁷, F. Grimaldi²⁰, K. Grimm¹⁴⁷,
 P.L.Y. Gris³³, Y. Grishkevich⁹⁷, H. Groenstege¹⁰⁶, L.S. Groer¹⁵⁵, J. Grognoz²⁹, M. Groh⁹⁹,
 E. Gross¹⁶⁵, J. Grosse-Knetter²¹, M.E.M. Grothe²⁹, J. Grudzinski⁵, C. Gruse²⁵, M. Gruwe²⁹,
 515 K. Grybel¹⁴¹, P. Grybos³⁷, E.M. Gschwendtner²⁹, V.J. Guarino⁵, C.J. Guicheney³³,
 G. Guilhem¹¹⁴, T. Guillemin⁴, J. Guimarães da Costa⁵⁵, J. Gunther¹²⁴, B. Guo¹⁵⁵, A. Gupta³⁰,
 L. Gurriana¹²³, N. Guseynov¹⁰, V.N. Gushchin¹²⁷, P. Gutierrez¹¹¹, L. Guy⁸⁵, C. Guyot¹³⁶,

C. Gwenlan⁷⁵, C.B. Gwilliam⁷¹, A. Haas³⁴, S. Haas²⁹, C. Haber¹⁴, G. Haboubi¹²²,
 R. Hackenburg²⁴, E. Hadash¹⁵⁰, H.K. Hadavand³⁹, C. Haeberli¹⁶, R. Härtel⁹⁹, R. Haggerty⁵⁵,
 520 F. Hahn²⁹, S. Haider²⁹, Z. Hajduk³⁸, M. Hakimi⁹³, H. Hakobyan⁴⁷, H. Hakobyan¹⁶⁹, J. Haller^{40.g},
 G.D. Hallewell⁸², B. Hallgren²⁹, K. Hamacher¹⁶⁷, A. Hamilton⁴⁷, H. Han³², L. Han³²,
 K. Hanagaki¹¹⁵, M. Hance¹¹⁹, P. Hanke⁵⁶, C.J. Hansen¹⁶⁰, F.H. Hansen³⁵, J.R. Hansen³⁵,
 J.B. Hansen³⁵, J.D. Hansen³⁵, P.H. Hansen³⁵, T. Hansl-Kozanecka¹³⁷, G. Hanson^{59.j},
 P. Hansson¹⁴⁶, K. Hara¹⁵⁷, S. Harder⁵⁵, A. Harel¹⁵⁰, T. Harenberg¹⁶⁷, R. Harper¹³⁹, J.C. Hart¹²⁸,
 525 R.G.G. Hart¹⁰⁶, F. Hartjes¹⁰⁶, N. Hartman¹⁴, T. Haruyama⁶⁴, A. Harvey⁵⁴, Y. Hasegawa¹⁴⁰,
 K. Hashemi²³, S. Hassani¹³⁶, M. Hatch²⁹, R.W. Hatley¹²⁸, T.G. Haubold⁹⁹, D. Hauff⁹⁹, F. Haug²⁹,
 S. Haug¹⁶, M. Hauschild²⁹, R. Hauser⁸⁷, C. Hauviller²⁹, M. Havranek¹²⁴, B.M. Hawes¹¹⁷,
 R.J. Hawkings²⁹, D. Hawkins⁶², T. Hayler¹²⁸, H.S. Hayward⁷¹, S.J. Haywood¹²⁸, E. Hazen²²,
 M. He³², Y.P. He³⁹, S.J. Head⁸⁰, V. Hedberg⁷⁷, L. Heelan²⁸, F.E.W. Heinemann¹¹⁷,
 530 M. Heldmann⁴⁶, S. Hellman¹⁴⁵, C. Helsens¹³⁶, R.C.W. Henderson⁶⁹, P.J. Hendriks¹⁰⁶,
 A.M. Henriques Correia²⁹, S. Henrot-Versille¹¹⁴, F. Henry-Couannier⁸², T. Henß¹⁶⁷, G. Herten⁴⁶,
 R. Hertenberger⁹⁸, L. Hervas²⁹, M. Hess¹⁶, N.P. Hessey¹⁰⁶, A. Hicheur¹²⁸, A. Hidvegi¹⁴⁵,
 E. Higón-Rodríguez¹⁶², D. Hill^{5,*}, J. Hill¹¹⁷, J.C. Hill²⁷, N. Hill⁵, S.J. Hillier¹⁷, I. Hinchliffe¹⁴,
 D. Hindson¹¹⁷, C. Hinkelbein⁸¹, T.A. Hodges¹⁵⁶, M.C. Hodgkinson¹³⁹, P. Hodgson¹³⁹,
 535 A. Hoecker²⁹, M.R. Hoefkamp¹⁰⁴, J. Hoffman³⁹, A.E. Hoffmann²⁴, D. Hoffmann⁸²,
 H.F. Hoffmann²⁹, M. Holder¹⁴¹, T.I. Hollins¹⁷, A. Holloway¹⁴, G. Hollyman⁷⁴, A. Holmes¹¹⁷,
 S.O. Holmgren¹⁴⁵, R. Holt¹²⁸, E. Holtom¹²⁸, T. Holy¹²⁶, R.J. Homer¹⁷, Y. Homma⁶⁵,
 P. Homola¹²⁶, W. Honerbach²¹, A. Honma^{164.a}, I. Hooton²⁹, T. Horazdovsky¹²⁶, C. Horn¹⁴³,
 S. Horvat⁹⁹, J.-Y. Hostachy⁵³, T. Hott⁹⁹, S. Hou¹⁴⁹, M.A. Houlden⁷¹, A. Houmada⁹³, J. Hover²⁴,
 540 D.F. Howell¹¹⁷, J. Hrivnac¹¹⁴, I. Hruska¹²⁴, T. Hryn'ova²⁹, G.S. Huang¹¹¹, Z. Hubacek¹²⁶,
 F. Hubaut⁸², F. Huegging²¹, B.T. Huffman¹¹⁷, E. Hughes³⁴, G. Hughes⁶⁹, R.E. Hughes-Jones⁸⁰,
 W. Hulsbergen²⁹, P. Hurst⁵⁵, M. Hurwitz³⁰, T. Huse¹¹⁶, J. Huston⁸⁷, J. Huth⁵⁵, G. Iacobucci²⁰,
 M. Ibbotson⁸⁰, I. Ibragimov¹⁴¹, R. Ichimiya⁶⁵, L. Iconomidou-Fayard¹¹⁴, J. Idarraga⁹²,
 M. Idzik³⁷, P. Iengo²⁹, M.C. Iglesias Escudero¹¹, O. Igonkina¹¹³, Y. Ikegami⁶⁴, M. Ikeno⁶⁴,
 545 Y. Ilchenko³⁹, Y. Ilyushenka⁶³, D. Imbault⁷⁶, P. Imbert¹¹⁴, M. Imhaeuser¹⁶⁷, M. Imori¹⁵³,
 T. Ince¹⁶⁴, J. Inigo-Golfin²⁹, K. Inoue¹⁵⁷, P. Ioannou⁸, M. Iodice¹³⁵, G. Ionescu⁴, K. Ishii⁶⁴,
 M. Ishino¹⁵³, Y. Ishizawa¹⁵⁶, R. Ishmukhametov³⁹, C. Issever¹¹⁷, H. Ito²⁴, A.V. Ivashin¹²⁷,
 W. Iwanski³⁸, H. Iwasaki⁶⁴, J. Izen^{24.k}, V. Izzo¹⁰², J. Jackson¹⁰⁹, J.N. Jackson⁷¹, M. Jaekel²⁹,
 S. Jagielski³⁷, M. Jahoda¹²⁴, V. Jain⁵⁹, K. Jakobs⁴⁶, J. Jakubek¹²⁶, E. Jansen¹⁰⁵,
 550 P.P.M. Jansweijer¹⁰⁶, R.C. Jared¹⁶⁶, G. Jarlskog⁷⁷, S. Jarp²⁹, P. Jarron²⁹, K. Jelen³⁷,
 I. Jen-La Plante³⁰, P. Jenni²⁹, A. Jeremie⁴, P. Jez¹²⁴, S. Jézéquel⁴, Y. Jiang³², G. Jin³², S. Jin³²,
 O. Jinnouchi⁶⁴, D. Joffe³⁹, L.G. Johansen¹³, M. Johansen¹⁴⁵, K.E. Johansson¹⁴⁵, P. Johansson¹³⁹,
 K.A. Johns⁶, K. Jon-And¹⁴⁵, M. Jones¹¹⁷, R. Jones²⁹, R.W.L. Jones⁶⁹, T.W. Jones⁷⁵, T.J. Jones⁷¹,
 A. Jones¹²⁸, O. Jonsson²⁹, K.K. Joo^{155.l}, D. Joos⁴⁶, M. Joos²⁹, C. Joram²⁹, S. Jorgensen¹¹,
 555 J. Joseph¹⁶⁶, P. Jovanovic¹⁷, S.S. Junnarkar²⁴, V. Juranek¹²⁴, P. Jussel⁶⁰, V.V. Kabachenko¹²⁷,
 S. Kabana¹⁶, M. Kaci¹⁶², A. Kaczmarek³⁸, M. Kado¹¹⁴, H. Kagan¹⁰⁹, S. Kagawa⁶⁴, S. Kaiser⁹⁹,
 E. Kajomovitz¹⁵⁰, S. Kakurin⁶³, L.V. Kalinovskaya⁶³, S. Kama⁴⁰, H. Kambara⁴⁷, N. Kanaya¹⁵³,
 A. Kandasamy²⁴, S. Kandasamy²⁴, M. Kaneda¹⁵³, V.A. Kantserov⁹⁶, J. Kanzaki⁶⁴, B. Kaplan¹⁶⁸,
 A. Kapliy³⁰, J. Kaplon²⁹, M. Karagounis²¹, M. Karagoz Unel¹¹⁷, K. Karr⁵, P. Karst⁸²,
 560 V. Kartvelishvili⁶⁹, A.N. Karyukhin¹²⁷, L. Kashif⁵⁵, A. Kasmi³⁹, R.D. Kass¹⁰⁹, A. Kastanas¹³,

M. Kataoka²⁹, Y. Kataoka¹⁵³, S. Katounine¹²⁰, E. Katsoufis⁹, K. Kawagoe⁶⁵, M. Kawai⁶⁴,
T. Kawamoto¹⁵³, F. Kayumov⁹⁴, V.A. Kazanin¹⁰⁷, M.Y. Kazarinov⁶³, A. Kazarov¹²⁰, S.I. Kazi⁸⁵,
J.R. Keates⁸⁰, R. Keeler¹⁶⁴, P.T. Keener¹¹⁹, R. Kehoe³⁹, M. Keil⁴⁷, G.D. Kekelidze⁶³, M. Kelly⁸⁰,
J. Kennedy⁹⁸, M. Kenyon⁵¹, O. Kepka¹²⁴, N. Kerschen¹³⁹, B.P. Kerševan⁷², S. Kersten¹⁶⁷,
565 C. Ketterer⁴⁶, M. Khakzad²⁸, H. Khandanyan¹⁶¹, A. Khanov¹¹², D. Kharchenko⁶³,
A. Khodinov¹⁴⁷, A.G. Kholodenko¹²⁷, A. Khomich⁸¹, V.P. Khomutnikov¹²⁰, G. Khoraiuli¹²⁹,
N. Khovanskiy⁶³, V. Khovanskiy⁹⁵, E. Khramov⁶³, J. Khubua⁴⁹, G. Kieft¹⁰⁶, J.A. Kierstead²⁴,
G. Kilvington⁷⁴, H. Kim⁷, H. Kim⁷, S.H. Kim¹⁵⁷, P. Kind¹⁶⁷, B.T. King⁷¹, J. Kirk¹²⁸,
G.P. Kirsch¹¹⁷, L.E. Kirsch²³, A.E. Kiryunin⁹⁹, D. Kisielewska³⁷, B. Kisielewski³⁸,
570 T. Kittelmann¹²², A.M. Kiver¹²⁷, H. Kiyamura⁶⁵, E. Kladiva¹⁴⁴, J. Klaiber-Lodewigs⁴¹,
K. Kleinknecht⁷⁹, A. Klier¹⁶⁵, A. Klimentov²⁴, C.R. Kline⁵⁹, R. Klingenberg⁴¹, E.B. Klinkby³⁵,
T. Klioutchnikova²⁹, P.F. Klok¹⁰⁵, S. Klous¹⁰⁶, E.-E. Kluge⁵⁶, P. Kluit¹⁰⁶, M. Klute²¹, S. Kluth⁹⁹,
N.K. Knecht¹⁵⁵, E. Kneringer⁶⁰, E. Knezo²⁹, J. Knobloch²⁹, B.R. Ko⁴³, T. Kobayashi¹⁵³,
M. Kobel⁴², P. Kodys¹²⁵, A.C. König¹⁰⁵, S. König⁴⁶, L. Köpke⁷⁹, F. Koetsveld¹⁰⁵, T. Koffas²⁹,
575 E. Koffeman¹⁰⁶, Z. Kohout¹²⁶, T. Kohriki⁶⁴, T. Kokott²¹, G.M. Kolachev¹⁰⁷, H. Kolanoski¹⁵,
V. Kolesnikov⁶³, I. Koletsou¹¹⁴, M. Kollfrath⁴⁶, S. Kolos⁶², S.D. Kolya⁸⁰, A.A. Komar⁹⁴,
J.R. Komaragiri¹⁴², T. Kondo⁶⁴, Y. Kondo⁶⁴, N.V. Kondratyeva⁹⁶, T. Kono²⁹, A.I. Kononov⁴⁶,
R. Konoplich¹⁰⁸, S.P. Konovalov⁹⁴, N. Konstantinidis⁷⁵, A. Kootz¹⁶⁷, S. Koperny³⁷,
S.V. Kopikov¹²⁷, K. Korcyl³⁸, K. Kordas¹⁶, V. Koreshev¹²⁷, A. Korn¹⁴, I. Korolkov¹¹,
580 V.A. Korotkov¹²⁷, H. Korsmo⁷⁷, O. Kortner⁹⁹, M.E. Kostrikov¹²⁷, V.V. Kostyukhin⁴⁸,
M.J. Kotamäki²⁹, D. Kotchetkov²³, S. Kotov⁹⁹, V.M. Kotov⁶³, K.Y. Kotov¹⁰⁷, C. Kourkoumelis⁸,
A. Koutsman¹⁰⁶, S. Kovalenko¹²⁰, R. Kowalewski¹⁶⁴, H. Kowalski⁴⁰, T.Z. Kowalski³⁷,
W. Kozanecki¹³⁶, A.S. Kozhin¹²⁷, V. Kral¹²⁶, V. Kramarenko⁹⁷, G. Kramberger⁷², A. Kramer⁴⁶,
O. Krasel⁴¹, M.W. Krasny⁷⁶, A. Krasznahorkay²⁹, A. Krepouri¹⁵², P. Krieger¹⁵⁵, P. Krivkova¹²⁵,
585 G. Krobath⁹⁸, H. Kroha⁹⁹, J. Krstic¹², U. Kruchonak⁶³, H. Krüger²¹, K. Kruger²⁹,
Z.V. Krumshteyn⁶³, P. Kubik¹²⁵, W. Kubischta²⁹, T. Kubota¹⁵³, L.G. Kudin¹²⁰, J. Kudlaty⁴¹,
A. Kugel⁸¹, T. Kuhl⁷⁹, D. Kuhn⁶⁰, V. Kukhtin⁶³, Y. Kulchitsky⁸⁹, N. Kundu¹¹⁷, A. Kupco¹²⁴,
M. Kupper¹⁶⁵, H. Kurashige⁶⁵, L.L. Kurchaninov¹⁵⁶, Y.A. Kurochkin⁸⁹, V. Kus¹²⁴,
W. Kuykendall¹³⁸, P. Kuzhir⁹⁰, E.K. Kuznetsova¹³³, R. Kwee¹⁵, D. La Marra⁴⁷, M. La Rosa⁸⁵,
590 L. La Rotonda³⁶, L. Labarga⁷⁸, J.A. Labbe⁵³, C. Lacasta¹⁶², F. Lacava¹³³, H. Lacker⁴²,
D. Lacour⁷⁶, V.R. Lacuesta¹⁶², E. Ladygin⁶³, R. Lafaye⁴, B. Laforge⁷⁶, T. Lagouri⁷⁸, S. Lai⁴⁶,
E. Lamanna³⁶, M. Lambacher⁹⁸, F. Lambert⁵³, W. Lampl⁶, E. Lancon¹³⁶, U. Landgraf⁴⁶,
M.P.J. Landon⁷³, H. Landsman¹⁵⁰, R.R. Langstaff¹⁵⁶, A.J. Lankford⁶², F. Lanni²⁴, K. Lantzschi¹⁶⁷,
A. Lanza¹¹⁸, V.V. Lapin^{127,*}, S. Laplace⁴, J.F. Laporte¹³⁶, V. Lara¹⁶², T. Lari⁸⁸, A.V. Larionov¹²⁷,
595 C. Lasseur²⁹, W. Lau¹¹⁷, P. Laurelli⁴⁵, A. Lavorato¹¹⁷, W. Lavrijsen¹⁴, A.B. Lazarev⁶³,
A-C. Le Bihan²⁹, O. Le Dortz⁷⁶, C. Le Maner¹⁵⁵, M. Le Vine²⁴, L. Leahu²⁹, M. Leahu²⁹,
C. Lebel⁹², M. Lechowski¹¹⁴, T. LeCompte⁵, F. Ledroit-Guillon⁵³, H. Lee²⁹, J.S.H. Lee¹⁴⁸,
S.C. Lee¹⁴⁹, M. Lefebvre¹⁶⁴, R.P. Lefevre⁴⁷, M. Legendre¹³⁶, A. Leger⁴⁷, B.C. LeGeyt¹¹⁹,
C. Leggett¹⁴, M. Lehmacher²¹, G. Lehmann Miotto²⁹, M. Lehto¹³⁹, R. Leitner¹²⁵, D. Lelas¹⁶⁴,
600 D. Lellouch¹⁶⁵, M. Leltchouk³⁴, V. Lendermann⁵⁶, K.J.C. Leney⁷¹, T. Lenz¹⁶⁷, G. Lenzen¹⁶⁷,
J. Lepidis¹⁶⁷, C. Leroy⁹², J-R. Lessard¹⁶⁴, J. Lesser¹⁴⁵, C.G. Lester²⁷, M. Letheren²⁹,
A. Leung Fook Cheong¹⁶⁶, J. Levêque⁸², D. Levin⁸⁶, L.J. Levinson¹⁶⁵, M.S. Levitski¹²⁷,
M. Leyton¹⁴, J. Li⁷, W. Li¹²⁸, M. Liabline⁶³, Z. Liang³⁹, Z. Liang¹⁴⁹, B. Liberti¹³⁴, P. Lichard²⁹,

W. Liebig¹⁰⁶, R. Lifshitz¹⁵⁰, D. Liko²⁹, H. Lim⁵, M. Limper¹⁰⁶, S.C. Lin¹⁴⁹, A. Lindahl³⁵,
 605 F. Linde¹⁰⁶, L. Lindquist¹⁶⁰, S.W. Lindsay⁷¹, V. Linhart¹²⁶, A.J. Lintern¹²⁸, A. Liolios¹⁵²,
 A. Lipniacka¹³, T.M. Liss¹⁶¹, A. Lissauer²⁴, J. List¹⁶⁷, A.M. Litke¹³⁷, S. Liu², T. Liu³⁹, Y. Liu⁴⁷,
 M. Livan¹¹⁸, A. Lleres⁵³, G. Llosá Llácer¹⁶², S.L. Lloyd⁷³, F. Lobkowicz^{132,*}, P. Loch⁶,
 W.S. Lockman¹³⁷, T. Loddenkoetter²¹, F.K. Loebinger⁸⁰, A. Loginov¹⁶⁸, C.W. Loh¹⁶³, T. Lohse¹⁵,
 K. Lohwasser¹¹⁷, M. Lokajicek¹²⁴, J. Loken¹¹⁷, S. Lokwitz¹⁶⁸, M.C. Long⁵⁴, L. Lopes¹²³,
 610 D. Lopez Mateos³⁴, M.J. Losty¹⁵⁶, X. Lou^{24,k}, K.F. Loureiro¹⁰⁹, L. Lovas¹⁴⁴, J. Love²²,
 A. Lowe⁷⁴, M. Lozano Fantoba¹⁶², F. Lu³², J. Lu², L. Lu³⁹, H.J. Lubatti¹³⁸, S. Lucas²⁹,
 C. Luci¹³³, A. Lucotte⁵³, A. Ludwig⁴², I. Ludwig⁴⁶, J. Ludwig⁴⁶, F. Luehring⁵⁹, D. Lüke⁴¹,
 G. Luijckx¹⁰⁶, L. Luisa¹⁵⁹, D. Lumb⁴⁶, L. Luminari¹³³, E. Lund¹¹⁶, B. Lund-Jensen¹⁴⁶,
 B. Lundberg⁷⁷, J. Lundquist³⁵, A. Lupi¹²¹, N. Lupu¹⁵⁰, G. Lutz⁹⁹, D. Lynn²⁴, J. Lynn¹¹⁷, J. Lys¹⁴,
 615 V. Lysan⁶³, E. Lytken²⁹, J.M. López-Amengual¹⁶², H. Ma²⁴, L.L. Ma¹⁵⁵, M. Maaßen⁴⁶,
 G. Maccarrone⁴⁵, G.G.R. Mace¹¹⁴, D. Macina⁴⁷, R. Mackeprang³⁵, A. Macpherson^{2,m},
 D. MacQueen², C. Macwaters¹²⁸, R.J. Madaras¹⁴, W.F. Mader⁴², R. Maenner⁸¹, T. Maeno²⁴,
 P. Mättig¹⁶⁷, S. Mättig⁴⁰, C.A. Magrath¹⁰⁵, Y. Mahalalel¹⁵¹, K. Mahboubi⁴⁶, G. Mahout¹⁷,
 C. Maidantchik¹³⁰, A. Maio¹²³, G.M. Mair⁶⁰, K. Mair²⁹, Y. Makida⁶⁴, D. Makowiecki²⁴,
 620 P. Malecki³⁸, V.P. Maleev¹²⁰, F. Malek⁵³, D. Malon⁵, S. Maltezos⁹, V. Malychhev¹⁰⁷,
 S. Maljukov⁶³, M. Mambelli³⁰, R. Mameghani⁹⁸, J. Mamuzic⁴⁰, A. Manabe⁶⁴, A. Manara⁵⁹,
 G. Manca⁷¹, L. Mandelli⁸⁸, I. Mandić⁷², M. Mandl²⁹, J. Maneira¹²³, M. Maneira¹²³,
 P.S. Mangeard⁸², M. Mangin-Brinet⁴⁷, I.D. Manjavidze⁶³, W.A. Mann¹⁵⁸, S. Manolopoulos¹³⁹,
 A. Manousakis-Katsikakis⁸, B. Mansoulie¹³⁶, A. Manz⁹⁹, A. Mapelli²⁹, L. Mapelli²⁹,
 625 L. March¹⁶², J.F. Marchand⁴, M. Marchesotti²⁹, M. Marcisovsky¹²⁴, A. Marin^{22,*},
 C.N. Marques¹²³, F. Marroquim¹³⁰, R. Marshall⁸⁰, Z. Marshall³⁴, F.K. Martens¹⁵⁵,
 S. Marti i Garcia¹⁶², A.J. Martin¹⁶⁸, B. Martin²⁹, B. Martin⁸⁷, F.F. Martin¹¹⁹, J.P. Martin⁹²,
 Ph. Martin⁵³, G. Martinez²¹, C. Martínez Lacambra¹⁶², V. Martinez Outschoorn⁵⁵, A. Martini⁴⁵,
 J. Martins¹²³, T. Maruyama¹⁵⁷, F. Marzano¹³³, T. Mashimo¹⁵³, R. Mashinistov⁹⁶, J. Masik⁸⁰,
 630 A.L. Maslennikov¹⁰⁷, M. Maß⁴¹, I. Massa²⁰, G. Massaro¹⁰⁶, N. Massol⁴, M. Mathes²¹,
 J. Matheson¹²⁸, P. Matricon¹¹⁴, H. Matsumoto¹⁵³, H. Matsunaga¹⁵³, J.M. Maugain²⁹,
 S.J. Maxfield⁷¹, E.N. May⁵, J.K. Mayer¹⁵⁵, C. Mayri¹³⁶, R. Mazini¹⁵⁵, M. Mazzanti⁸⁸,
 P. Mazzanti²⁰, E. Mazzoni¹²¹, F. Mazzucato⁴⁷, S.P. Mc Kee⁸⁶, R.L. McCarthy¹⁴⁷,
 C. McCormick⁶², N.A. McCubbin¹²⁸, J. McDonald¹²², K.W. McFarlane⁵⁴, S. McGarvie⁷⁴,
 635 H. McGlone⁵¹, R.A. McLaren²⁹, S.J. McMahan¹²⁸, T.R. McMahan⁷⁴, T.J. McMahan¹⁷,
 R.A. McPherson^{164,n}, M. Mechtel¹⁶⁷, D. Meder-Marouelli¹⁶⁷, M. Medinnis⁴⁰,
 R. Meera-Lebbai¹¹¹, C. Meessen⁸², R. Mehdiyev⁹², A. Mehta⁷¹, K. Meier⁵⁶, H. Meinhard²⁹,
 J. Meinhardt⁴⁶, C. Meirosu²⁹, F. Meisel⁴⁶, A. Melamed-Katz¹⁶⁵, B.R. Mellado Garcia¹⁶⁶,
 P. Mendez⁹⁸, S. Menke⁹⁹, C. Menot²⁹, E. Meoni³⁶, D. Merkl⁹⁸, L. Merola¹⁰², C. Meroni⁸⁸,
 640 F.S. Merritt³⁰, I. Messmer⁴⁶, J. Metcalfe¹⁰⁴, S. Meuser²¹, J-P. Meyer¹³⁶, T.C. Meyer²⁹,
 W.T. Meyer⁶¹, V. Mialkovski⁶³, M. Michelotto²⁹, L. Micu²⁵, R. Middleton¹²⁸, P. Miele²⁹,
 A. Migliaccio¹⁰², L. Mijović⁷², G. Mikenberg¹⁶⁵, M. Mikestikova¹²⁴, B. Mikulec⁴⁷, M. Mikuž⁷²,
 D.W. Miller¹⁴³, R.J. Miller⁸⁷, W. Miller¹⁴, M. Milosavljevic¹², D.A. Milstead¹⁴⁵, S. Mima¹¹⁰,
 A.A. Minaenko¹²⁷, M. Minano¹⁶², I.A. Minashvili⁶³, A.I. Mincer¹⁰⁸, B. Mindur³⁷, M. Mineev⁶³,
 645 L.M. Mir¹¹, G. Mirabelli¹³³, L. Miralles Verge¹¹, S. Misawa²⁴, S. Miscetti⁴⁵, A. Misiejuk⁷⁴,
 A. Mitra¹¹⁷, G.Y. Mitrofanov¹²⁷, V.A. Mitsou¹⁶², P.S. Miyagawa⁸⁰, Y. Miyazaki¹⁴⁰,

J.U. Mjörnmark⁷⁷, S. Mkrtychyan¹⁶⁹, D. Mladenov²⁹, T. Moa¹⁴⁵, M. Moch¹³³, A. Mochizuki¹⁵⁷,
 P. Mockett¹³⁸, P. Modesto¹⁶², S. Moed⁵⁵, K. Mönig⁴⁰, N. Möser²¹, B. Mohn¹³, W. Mohr⁴⁶,
 S. Mohrdieck-Möck⁹⁹, A.M. Moisseev^{127,*}, R.M. Moles Valls¹⁶², J. Molina-Perez²⁹, A. Moll⁴⁰,
 650 G. Moloney⁸⁵, R. Mommsen⁶², L. Moneta⁴⁷, E. Monnier⁸², G. Montarou³³, S. Montesano⁸⁸,
 F. Monticelli⁶⁸, R.W. Moore², T.B. Moore⁸³, G.F. Moorhead⁸⁵, A. Moraes¹⁵⁶, J. Morel⁵³,
 A. Moreno¹⁶², D. Moreno¹⁹, P. Morettini⁴⁸, D. Morgan¹³⁹, M. Morii⁵⁵, J. Morin⁷³,
 A.K. Morley⁸⁵, G. Mornacchi²⁹, M-C. Morone⁴⁷, S.V. Morozov⁹⁶, E.J. Morris⁵⁹, J. Morris⁷³,
 M.C. Morrissey¹²⁸, H.G. Moser⁹⁹, M. Mosidze⁴⁹, A. Moszczynski³⁸, S.V. Mouraviev⁹⁴,
 655 T. Mouthuy⁸², T.H. Moye¹⁷, E.J.W. Moyses⁸³, J. Mueller¹²², M. Müller⁸¹, A. Muijs¹⁰⁶,
 T.R. Muller²⁴, A. Munar¹¹⁹, D.J. Munday²⁷, K. Murakami⁶⁴, R. Murillo Garcia⁶²,
 W.J. Murray¹²⁸, A.G. Myagkov¹²⁷, K. Nagai⁷³, Y. Nagai¹⁵⁷, K. Nagano⁶⁴, Y. Nagasaka⁵⁸,
 A.M. Nairz²⁹, D. Naito¹¹⁰, K. Nakamura¹⁵⁷, Y. Nakamura¹⁵⁷, I. Nakano¹¹⁰, G. Nanava⁶³,
 M. Nassiakou²⁹, I. Nasteva⁸⁰, N.R. Nation²², T. Naumann⁴⁰, F. Nauyock⁸⁰, S.K. Nderitu²¹,
 660 H.A. Neal⁸⁶, E. Nebot⁷⁸, P. Nechaeva⁹⁴, A. Neganov⁶³, A. Negri⁶², S. Negroni³⁴, C. Nelson¹²⁸,
 S. Nemecek¹²⁴, P. Nemethy¹⁰⁸, A.A. Nepomuceno¹³⁰, M. Nessi²⁹, S.Y. Nesterov¹²⁰,
 L. Neukermans⁴, P. Nevski²⁴, F.M. Newcomer¹¹⁹, A. Nichols¹²⁸, C. Nicholson⁵¹,
 R. Nicholson¹³⁹, R.B. Nickerson¹¹⁷, R. Nicolaidou¹³⁶, G. Nicoletti⁴⁵, B. Nicquevert²⁹,
 M. Niculescu²⁵, J. Nielsen¹³⁷, T. Niinikoski²⁹, M.J. Niinimäki¹¹⁶, N. Nikitin⁹⁷, K. Nikolaev⁶³,
 665 I. Nikolic-Audit⁷⁶, K. Nikolopoulos⁸, H. Nilsen⁴⁶, B.S. Nilsson³⁵, P. Nilsson⁷, A. Nisati¹³³,
 R. Nisius⁹⁹, L.J. Nodulman⁵, M. Nomachi¹¹⁵, H. Nomoto¹⁵³, J-M. Noppe¹¹⁴, M. Nordberg²⁹,
 O. Norniella Francisco¹¹, P.R. Norton¹²⁸, J. Novakova¹²⁵, M. Nowak¹⁴, M. Nozaki⁶⁴, R. Nunes²⁹,
 G. Nunes Hanninger²¹, T. Nunnemann⁹⁸, T. Nyman²⁹, P. O'Connor¹⁴, S.W. O'Neale^{17,*},
 D.C. O'Neil¹⁴², M. O'Neill²⁸, V. O'Shea⁵¹, F.G. Oakham^{28,a}, H. Oberlack⁹⁹, M. Obermaier⁹⁸,
 670 P. Oberson¹³³, A. Ochi⁶⁵, W. Ockenfels²¹, S. Odaka⁶⁴, I. Odenthal²¹, G.A. Odino⁴⁸, H. Ogren⁵⁹,
 S.H. Oh⁴³, T. Ohshima¹⁰¹, H. Ohshita¹⁴⁰, H. Okawa¹⁵³, M. Olcese⁴⁸, A.G. Olchevski⁶³,
 C. Oliver⁷⁸, J. Oliver⁵⁵, M. Olivo Gomez⁹⁹, A. Olszewski³⁸, J. Olszowska³⁸, C. Omachi⁶⁵,
 A. Onea²⁹, A. Onofre¹²³, C.J. Oram¹⁵⁶, G. Ordonez¹⁰⁵, M.J. Oreglia³⁰, F. Orellana⁴⁷, Y. Oren¹⁵¹,
 D. Orestano¹³⁵, I.A. Orlov¹⁰⁷, R.S. Orr¹⁵⁵, F. Orsini⁷⁶, L.S. Osborne⁹¹, B. Osculati⁴⁸, C. Osuna¹¹,
 675 R. Otec¹²⁶, R. Othegraven⁷⁹, B. Ottewell¹¹⁷, F. Ould-Saada¹¹⁶, A. Ouraou¹³⁶, Q. Ouyang³²,
 O.K. Øye¹³, V.E. Ozcan⁷⁵, K. Ozone⁶⁴, N. Ozturk⁷, A. Pacheco Pages¹¹, S. Padhi¹⁶⁶,
 C. Padilla Aranda¹¹, E. Paganis¹³⁹, F. Paige²⁴, P.M. Pailler²⁹, K. Pajchel¹¹⁶, S. Palestini²⁹,
 J. Palla²⁹, D. Pallin³³, M.J. Palmer²⁷, Y.B. Pan¹⁶⁶, N. Panikashvili¹⁵⁰, V.N. Panin¹⁰⁷, S. Panitkin²⁴,
 D. Pantea²⁵, M. Panuskova¹²⁴, V. Paolone¹²², A. Paoloni¹³⁴, I. Papadopoulos²⁹,
 680 T. Papadopoulou⁹, I. Park¹¹, W. Park^{24,o}, M.A. Parker²⁷, S. Parker¹⁴, C. Parkman²⁹, F. Parodi⁴⁸,
 J.A. Parsons³⁴, U. Parzefall⁴⁶, E. Pasqualucci¹³³, G. Passardi²⁹, A. Passeri¹³⁵, M.S. Passmore²⁹,
 F. Pastore¹³⁵, Fr. Pastore¹³³, S. Patariaia⁹⁹, D. Pate²⁴, J.R. Pater⁸⁰, S. Patricelli¹⁰², T. Pauly²⁹,
 E. Pauna¹⁶⁷, L.S. Peak¹⁴⁸, S.J.M. Peeters¹⁰⁶, M. Peetz⁷⁸, E. Pei³², S.V. Peleganchuk¹⁰⁷,
 G. Pellegrini¹⁶², R. Pengo²⁹, J. Pequeno¹³⁷, M. Perantoni¹³⁰, A. Perazzo¹⁴³, A. Pereira¹²³,
 685 E. Perepelkin⁶³, V.J.O. Perera¹²⁸, E. Perez Codina¹¹, V. Perez Reale²⁹, I. Peric²¹, L. Perini⁸⁸,
 H. Pernegger²⁹, E. Perrin⁴⁷, R. Perrino⁷⁰, P. Perrodo⁴, G. Perrot⁴, P. Perus¹¹⁴,
 V.D. Peshekhonov⁶³, E. Petereit⁵, J. Petersen²⁹, T.C. Petersen²⁹, P.J.F. Petit⁵³, C. Petridou¹⁵²,
 E. Petrolo¹³³, F. Petrucci¹³⁵, R. Petti^{24,o}, M. Pezzetti²⁹, B. Pfeifer⁴⁶, A. Phan⁸⁵, A.W. Phillips²⁷,
 P.W. Phillips¹²⁸, G. Piacquadio⁴⁶, M. Piccinini²⁰, A. Pickford⁵¹, R. Piegaia²⁶, S. Pier⁶²,

690 J.E. Pilcher³⁰, A.D. Pilkington⁸⁰, M.A. Pimenta Dos Santos²⁹, J. Pina¹²³, J.L. Pinfeld², J. Ping³²,
J. Pinhão¹²³, B. Pinto¹²³, O. Pirotte²⁹, R. Placakyte⁴⁰, A. Placci²⁹, M. Plamondon¹¹⁴,
W.G. Plano⁸⁰, M-A. Pleier²¹, A.V. Pleskach¹²⁷, S. Podkladkin⁷⁰, F. Podlyski³³, P. Poffenberger¹⁶⁴,
L. Poggioli¹¹⁴, M. Pohl⁴⁷, I. Polak¹²⁴, G. Polesello¹¹⁸, A. Policicchio³⁶, A. Polini²⁰,
V. Polychronakos²⁴, D.M. Pomarede¹³⁶, K. Pommès²⁹, P. Ponsot¹³⁶, L. Pontecorvo¹³³,
695 B.G. Pope⁸⁷, R. Popescu²⁴, D.S. Popovic¹², A. Poppleton²⁹, J. Popule¹²⁴, X. Portell Bueso⁴⁶,
C. Posch²², G.E. Pospelov⁹⁹, P. Pospichal²⁹, S. Pospisil¹²⁶, M. Postranecky⁷⁵, I.N. Potrap⁹⁹,
C.J. Potter⁷⁴, G. Poulard²⁹, A. Pousada¹²³, J. Poveda¹⁶², R. Prabhu²¹, P. Pralavorio⁸², S. Prasad⁵⁵,
J. Prast⁴, S. Prat¹¹⁴, M. Prata¹¹⁸, R. Pravahan⁷, T. Preda²⁵, K. Pretzl¹⁶, L. Pribyl¹²⁴, D. Price⁶⁹,
L.E. Price⁵, M.J. Price²⁹, P.M. Prichard⁷¹, D. Prieur¹²⁸, M. Primavera⁷⁰, D. Primor²⁹,
700 K. Prokofiev¹³⁹, E. Proso¹⁶⁵, J. Proudfoot⁵, H. Przysieznik⁴, C. Puigdemengoles¹¹, J. Purdham⁸⁶,
M. Purohit^{24.o}, P. Puzo¹¹⁴, A.N. Pylaev¹²⁷, Y. Pylypchenko¹¹⁶, M. Qi³², J. Qian⁸⁶, W. Qian¹²⁸,
Z. Qian⁸², D. Qing¹⁴⁹, A. Quadt⁵², D.R. Quarrie¹⁴, W.B. Quayle¹⁶⁶, J.J. Rabbers²⁹, V. Radeka²⁴,
J.M. Rafi¹⁶², F. Ragusa⁸⁸, A.M. Rahimi¹⁰⁹, D. Rahm²⁴, C. Raine⁵¹, B. Raith²¹, S. Rajagopalan²⁴,
S. Rajek⁴¹, H. Rammer²⁹, M. Ramstedt¹⁴⁵, S. Rangod²⁹, P.N. Ratoff⁶⁹, T. Rauber⁴⁶, F. Rauscher⁹⁸,
705 E. Rauter⁹⁹, M. Raymond²⁹, A.L. Read¹¹⁶, D. Rebutzi¹¹⁸, G.R. Redlinger²⁴, K. Reeves¹⁶⁷,
M. Rehak²⁴, A. Reichold¹⁰⁶, E. Reinherz-Aronis¹⁵¹, I. Reisinger⁴¹, D. Reljic¹², C. Rembser²⁹,
Z. Ren¹⁴⁹, P. Renkel³⁹, B. Rensch³⁵, S. Rescia²⁴, M. Rescigno¹³³, S. Resconi⁸⁸, B. Resende⁸²,
P. Rewiersma¹⁰⁶, J. Rey¹³⁶, M. Rey-Campagnolle⁵³, E. Rezaie¹⁴², P. Reznicek¹²⁵,
R.A. Richards⁸⁷, J-P. Richer¹¹⁴, R.H. Richter⁹⁹, R. Richter⁹⁹, E. Richter-Was^{38.p}, M. Ridel⁷⁶,
710 W. Riegler⁵⁵, S. Rieke⁷⁹, M. Rijpstra¹⁰⁶, M. Rijssenbeek¹⁴⁷, A. Rimoldi¹¹⁸, R.R. Rios³⁹,
I. Riu Dachs¹¹, M. Rivline¹⁶⁵, G. Rivoltella⁸⁸, F.R. Rizatdinova¹¹², S.H. Robertson^{84.n},
A. Robichaud-Veronneau⁴⁷, S. Robins¹³³, D. Robinson²⁷, A. Robson⁵¹, J.H. Rochford¹²⁸,
C. Roda¹²¹, S. Rodier⁷⁸, S. Roe²⁹, O. Røhne¹¹⁶, F. Rohrbach²⁹, J. Roldán¹⁶², S. Rolli¹⁵⁸,
J.B. Romance¹⁶², A. Romaniouk⁹⁶, V.M. Romanov⁶³, G. Romeo²⁶, L. Roos⁷⁶, E. Ros¹⁶²,
715 S. Rosati¹³³, F. Rosenbaum¹³⁷, G.A. Rosenbaum¹⁵⁵, E.I. Rosenberg⁶¹, L. Rosselet⁴⁷, L.P. Rossi⁴⁸,
L. Rossi⁸⁸, M. Rotaru²⁵, J. Rothberg¹³⁸, I. Rottländer²¹, D. Rousseau¹¹⁴, A. Rozanov⁸²,
Y. Rozen¹⁵⁰, R. Ruber^{160.m}, B. Ruckert⁹⁸, G. Rudolph⁶⁰, F. Rühr⁵⁶, F. Ruggieri¹³⁵, G. Ruggiero²⁹,
H. Ruiz¹¹, A. Ruiz-Martinez¹⁶², E. Rulikowska-Zarebska³⁷, V. Rumiantsev^{90.*}, L. Rumyantsev⁶³,
K. Runge⁴⁶, O. Runolfsson²¹, N.A. Rusakovich⁶³, D.R. Rust⁵⁹, J.P. Rutherford⁶, C. Ruwiedel²¹,
720 Y.F. Ryabov¹²⁰, V. Ryadovikov¹²⁷, P. Ryan⁸⁷, G. Rybkine¹¹⁴, J. Sá da Costa¹²³, A.F. Saavedra¹⁴⁸,
S. Saboumazrag⁵³, H.F-W. Sadrozinski¹³⁷, R. Sadykov⁶³, H. Sakamoto¹⁵³, P. Sala⁸⁸,
A. Salamon¹³⁴, M. Saleem¹¹¹, D. Salihagic⁹⁹, J. Salt¹⁶², O. Saltó Bauza¹¹, B.M. Salvachúa
Ferrando¹⁶², D. Salvatore³⁶, A. Salzburger⁶⁰, D. Sampsonidis¹⁵², B.H. Samset¹¹⁶, C.A. Sánchez
Sánchez¹¹, M.A. Sanchis Lozano¹⁶², E. Sanchis Peris¹⁶², H. Sandaker¹³, H.G. Sander⁷⁹,
725 M. Sandhoff¹⁶⁷, S. Sandvoss¹⁶⁷, D.P.C. Sankey¹²⁸, B. Sanny¹⁶⁷, S. Sansone⁵⁵, A. Sansoni⁴⁵,
C. Santamarina Rios⁸⁴, J. Santander¹⁶², L. Santi¹⁵⁹, C. Santoni³³, R. Santonico¹³⁴, J. Santos¹²³,
M. Sapinski³⁸, J.G. Saraiva¹²³, F. Sarri¹²¹, O. Sasaki⁶⁴, T. Sasaki⁶⁴, N. Sasao⁶⁶,
I. Satsounkevitch⁸⁹, D. Sauvage^{82.*}, G. Sauvage⁴, P. Savard^{155.a}, A.Y. Savine⁶, V. Savinov¹²²,
A. Savoy-Navarro⁷⁶, P. Savva⁹, D.H. Saxon⁵¹, L.P. Says³³, C. Sbarra²⁰, E. Sbrissa²⁹, A. Sbrizzi²⁰,
730 D.A. Scannicchio¹¹⁸, J. Schaarschmidt⁴², P. Schacht⁹⁹, U. Schäfer⁷⁹, A.C. Schaffer¹¹⁴,
D. Schaile⁹⁸, O. Schaile⁹⁸, M. Schaller²⁹, A.G. Schamov¹⁰⁷, V.A. Schegelsky¹²⁰, D. Scheirich¹²⁵,
M. Schernau⁶², M.I. Scherzer¹⁴, C. Schiavi⁴⁸, H. Schick⁸⁶, J. Schieck⁹⁹, P. Schieferdecker⁹⁸,

M. Schioppa³⁶, G. Schlager²⁹, S. Schlenker²⁹, J.L. Schlereth⁵, P. Schmid²⁹, M.P. Schmidt^{168,*},
 C. Schmitt²⁹, K. Schmitt⁵⁶, M. Schmitz²¹, H. Schmücker⁹⁹, T. Schoerner²⁹, R.C. Scholte¹⁰⁶,
 735 M. Schott⁹⁸, D. Schouten¹⁴², A. Schricker²³, D. Schroff⁴⁶, S. Schuh²⁹, H.W. Schuijlenburg¹⁰⁶,
 G. Schuler²⁹, J. Schultes¹⁶⁷, H-C. Schultz-Coulon⁵⁶, J. Schumacher⁴², M. Schumacher¹⁴¹,
 Ph. Schune¹³⁶, A. Schwartzman¹⁴³, D. Schweiger²⁹, Ph. Schwemling⁷⁶, C. Schwick²⁹,
 R. Schwienhorst⁸⁷, R. Schwierz⁴², J. Schwindling¹³⁶, W.G. Scott¹²⁸, H. Secker⁷⁹, E. Sedykh¹²⁰,
 N. Seguin-Moreau¹¹⁴, E. Segura¹¹, S.C. Seidel¹⁰⁴, A. Seiden¹³⁷, J.M. Seixas¹³⁰,
 740 G. Sekhniaidze¹⁰², D.M. Seliverstov¹²⁰, B. Selldén¹⁴⁵, M. Seman¹⁴⁴, N. Semprini-Cesari²⁰,
 C. Serfon⁹⁸, L. Serin¹¹⁴, R. Seuster¹⁶⁴, H. Severini¹¹¹, M.E. Sevier⁸⁵, K.A. Sexton²⁴, A. Sfyrla⁴⁷,
 T.P. Shah¹²⁸, L. Shan³², J.T. Shank²², M. Shapiro¹⁴, P.B. Shatalov⁹⁵, L. Shaver⁶, C. Shaw⁵¹,
 T.G. Shears⁷¹, P. Sherwood⁷⁵, A. Shibata⁷³, P. Shield¹¹⁷, S. Shilov⁶³, M. Shimojima¹⁰⁰, T. Shin⁵⁴,
 M. Shiyakova⁶³, A. Shmeleva⁹⁴, M. Shoa¹⁶⁵, M.J. Shochet³⁰, M.A. Shupe⁶, P. Sicho¹²⁴,
 745 A. Sidoti¹³³, A. Siebel¹⁶⁷, M. Siebel²⁹, J. Siegrist¹⁴, D. Sijacki¹², J. Silva¹²³, S.B. Silverstein¹⁴⁵,
 V. Simak¹²⁶, Lj. Simic¹², S. Simion¹¹⁴, B. Simmons⁷⁵, M. Simonyan¹⁶⁹, P. Sinervo¹⁵⁵,
 V. Sipica¹⁴¹, G. Siragusa⁷⁰, A.N. Sisakyan⁶³, S. Sivoklov⁹⁷, J. Sjölin¹⁴⁵, P. Skubic¹¹¹,
 N. Skvorodnev²³, P. Slattery¹³², T. Slavicek¹²⁶, K. Sliwa¹⁵⁸, T.J. Sloan⁶⁹, J. Sloper²⁹,
 V. Smakhtin¹⁶⁵, A. Small⁶⁹, S.Yu. Smirnov⁹⁶, Y. Smirnov³⁰, L. Smirnova⁹⁷, O. Smirnova⁷⁷,
 750 N.A. Smith⁷¹, B.C. Smith⁵⁵, D.S. Smith¹⁰⁹, J. Smith²⁴, K.M. Smith⁵¹, B. Smith¹²⁸,
 M. Smizanska⁶⁹, K. Smolek¹²⁶, A.A. Snesarev⁹⁴, S.W. Snow⁸⁰, J. Snow¹¹¹, J. Snuverink¹⁰⁶,
 S. Snyder²⁴, M. Soares⁷⁸, S. Soares¹²³, R. Sobie^{164,n}, J. Sodomka¹²⁶, M. Söderberg⁷⁷,
 A. Soffer¹⁵¹, C.A. Solans¹⁶², M. Solar¹²⁶, D. Sole¹²⁸, E. Solfaroli Camillocci¹³³,
 A.A. Solodkov¹²⁷, O.V. Solov'yanov¹²⁷, I. Soloviev¹²⁰, R. Soluk², J. Sondericker²⁴, V. Sopko¹²⁶,
 755 B. Sopko¹²⁶, M. Sorbi⁸⁸, J. Soret Medel¹⁶², M. Sosebee⁷, V.V. Sosnovtsev⁹⁶,
 L. Sospedra Suay¹⁶², A. Soukharev¹⁰⁷, J. Soukup², S. Spagnolo⁷⁰, F. Spano³⁴, P. Speckmayer²⁹,
 M. Spegel²⁹, E. Spencer¹³⁷, R. Spighi²⁰, G. Spigo²⁹, F. Spila¹³³, E. Spiriti¹³⁵, R. Spiwoks²⁹,
 L. Spogli¹³⁵, M. Spousta¹²⁵, G. Sprachmann²⁹, B. Spurlock⁷, R.D. St. Denis⁵¹, T. Stahl¹⁴¹,
 R.J. Staley¹⁷, R. Stamen⁵⁶, S.N. Stancu⁶², E. Stanecka³⁸, R.W. Stanek⁵, C. Stanescu¹³⁵,
 760 S. Stapnes¹¹⁶, E.A. Starchenko¹²⁷, P. Staroba¹²⁴, J. Stastny¹²⁴, A. Staude⁹⁸, P. Stavina¹⁴⁴,
 M. Stavrianakou²⁹, G. Stavropoulos¹⁴, E. Stefanidis⁷⁵, J.L. Steffens¹⁴⁷, I. Stekl¹²⁶, H.J. Stelzer²⁹,
 H. Stenzel⁵⁰, G. Stewart⁵¹, T.D. Stewart¹⁴², W. Stiller⁹⁹, T. Stockmanns²¹, M. Stodulski³⁸,
 S. Stonjek⁹⁹, A. Stradling¹⁶⁶, A. Straessner⁴⁷, A. Strandlie¹¹⁶, M. Strauss¹¹¹, V. Strickland^{28,a},
 D. Striegel⁹⁹, P. Strizenec¹⁴⁴, R. Ströhmer⁹⁸, D.M. Strom¹¹³, J.A. Strong^{74,*}, R. Stroynowski³⁹,
 765 B. Stugu¹³, I. Stumer²⁴, D. Su¹⁴³, S. Subramania⁵⁹, S.I. Suchkov⁹⁶, Y. Sugaya¹¹⁵, T. Sugimoto¹⁰¹,
 M. Suk¹²⁵, V.V. Sulin⁹⁴, S. Sultanov^{3,q}, Z. Sun¹³⁶, B. Sundal¹¹⁶, S. Sushkov¹¹, G. Susinno³⁶,
 P. Sutcliffe⁷¹, M.R. Sutton⁷⁵, Yu.M. Sviridov¹²⁷, I. Sykora¹⁴⁴, R.R. Szczygiel³⁸, B. Szeless²⁹,
 T. Szymocha³⁸, J. Sánchez¹⁶², D. Ta²¹, S. Taboada Gameiro²⁹, M. Tadel⁷², R. Tafirout¹⁵⁶,
 A. Taga¹¹⁶, H. Takai²⁴, R. Takashima⁶⁷, H. Takeda⁶⁵, T. Takeshita¹⁴⁰, M. Talby⁸², A. Talyshev¹⁰⁷,
 770 M.C. Tamsett⁷⁴, J. Tanaka¹⁵³, K. Tanaka⁶⁴, R. Tanaka¹¹⁰, S. Tanaka¹³¹, S. Tanaka⁶⁴, Y. Tanaka¹⁰⁰,
 G.P. Tappern²⁹, S. Tapprogge⁷⁹, S. Tarem¹⁵⁰, F. Tarrade²⁴, J. Tarrant¹²⁸, G. Tartarelli⁸⁸, P. Tas¹²⁵,
 M. Tasevsky¹²⁴, Y. Tayalati¹⁴⁶, F.E. Taylor⁹¹, G. Taylor¹³⁷, G.N. Taylor⁸⁵, R.P. Taylor¹⁶⁴,
 V. Tcherniatine²⁴, F. Tegenfeldt⁶¹, P. Teixeira-Dias⁷⁴, H. Ten Kate²⁹, P.K. Teng¹⁴⁹,
 R. Ter-Antonyan¹⁰⁹, S. Terada⁶⁴, J. Terron⁷⁸, M. Terwort^{40,g}, R.J. Teuscher^{155,n}, C.M. Tevlin⁸⁰,
 775 J. Thadome¹⁶⁷, J. Thion⁴, M. Thioye¹⁴⁷, A. Thomas⁴³, J.P. Thomas¹⁷, T.L. Thomas¹⁰⁴,

E. Thomas¹⁶, R.J. Thompson⁸⁰, A.S. Thompson⁵¹, R.P. Thun⁸⁶, T. Tic¹²⁴, V.O. Tikhomirov⁹⁴,
 Y.A. Tikhonov¹⁰⁷, S. Timm¹, C.J.W.P. Timmermans¹⁰⁵, P. Tipton¹⁶⁸, F.J. Tique Aires Viegas²⁹,
 S. Tisserant⁸², M. Titov⁴⁶, J. Tobias⁴⁶, V.M. Tocut¹¹⁴, B. Toczek³⁷, S. Todorova-Nova¹⁵⁸,
 J. Tojo⁶⁴, S. Tokár¹⁴⁴, K. Tokushuku⁶⁴, L. Tomasek¹²⁴, M. Tomasek¹²⁴, F. Tomasz¹⁴⁴,
 780 M. Tomoto¹⁰¹, D. Tompkins⁶, L. Tompkins¹⁴, K. Toms⁹⁷, A. Tonazzo¹³⁵, G. Tong³²,
 A. Tonoyan¹³, C. Topfel¹⁶, N.D. Topilin⁶³, E. Torrence¹¹³, J.G. Torres Pais¹⁶², J. Toth^{82,r},
 F. Touchard⁸², D.R. Tovey¹³⁹, S.N. Tovey⁸⁵, E.F. Towndrow¹²⁸, T. Trefzger⁷⁹, M. Treichel²⁹,
 J. Treis²¹, L. Tremblet²⁹, W. Tribanek⁹⁹, A. Tricoli¹²⁸, I.M. Trigger¹⁵⁶, G. Trilling¹⁴,
 S. Trincaz-Duvoid⁷⁶, M.F. Tripiana⁶⁸, W. Trischuk¹⁵⁵, Z. Trka¹²⁵, B. Trocme⁵³, C. Troncon⁸⁸,
 785 J.C-L. Tseng¹¹⁷, I. Tsiafis¹⁵², P.V. Tsireshka⁸⁹, G. Tsipolitis⁹, E.G. Tskhadadze⁴⁹,
 I.I. Tsukerman⁹⁵, V. Tsulaia¹²², S. Tsuno⁶⁴, M. Turala³⁸, I. Turk Cakir^{3,s}, E. Turlay¹¹⁴,
 P.M. Tuts³⁴, M.S. Twomey¹³⁸, M. Tyndel¹²⁸, D. Typaldos¹⁷, H. Tyrvaianen²⁹, E. Tzamarioudaki⁹,
 G. Tzanakos⁸, I. Ueda¹⁵³, M. Uhrmacher⁵², F. Ukegawa¹⁵⁷, M. Ullán Comes¹⁶², G. Unal²⁹,
 D.G. Underwood⁵, A. Undrus²⁴, G. Unel⁶², Y. Unno⁶⁴, E. Urkovsky¹⁵¹, G. Usai³⁰, Y. Usov⁶³,
 790 L. Vacavant⁸², V. Vacek¹²⁶, B. Vachon⁸⁴, S. Vahsen¹⁴, C. Valderanis⁹⁹, J. Valenta¹²⁴, P. Valente¹³³,
 A. Valero¹⁶², S. Valkar¹²⁵, J.A. Valls Ferrer¹⁶², H. Van der Bij²⁹, H. van der Graaf¹⁰⁶,
 E. van der Kraaij¹⁰⁶, B. Van Eijk¹⁰⁶, N. van Eldik⁸³, P. van Gemmeren⁵, Z. van Kesteren¹⁰⁶,
 I. van Vulpen¹⁰⁶, R. VanBerg¹¹⁹, W. Vandelli²⁹, G. Vandoni²⁹, A. Vaniachine⁵, F. Vannucci⁷⁶,
 M. Varanda¹²³, F. Varela Rodriguez²⁹, R. Vari¹³³, E.W. Varnes⁶, D. Varouchas¹¹⁴, A. Vartapetian⁷,
 795 K.E. Varvell¹⁴⁸, V.I. Vassilakopoulos⁵⁴, L. Vassilieva⁹⁴, E. Vataga¹⁰⁴, L. Vaz¹²³, F. Vazeille³³,
 P. Vedrine¹³⁶, G. Vegni⁸⁸, J.J. Veillet¹¹⁴, C. Vellidis⁸, F. Veloso¹²³, R. Veness²⁹, S. Veneziano¹³³,
 A. Ventura⁷⁰, S. Ventura⁴⁵, V. Vercesi¹¹⁸, M. Verducci²⁹, W. Verkerke¹⁰⁶, J.C. Vermeulen¹⁰⁶,
 L. Vertogardov¹¹⁷, M.C. Vetterli^{142,a}, I. Vichou¹⁶¹, T. Vickey¹⁶⁶, G.H.A. Viehhauser¹¹⁷,
 E. Vigeolas⁸², M. Villa²⁰, E.G. Villani¹²⁸, J. Villate¹²³, I. Villella¹⁰², E. Vilucchi⁴⁵, P. Vincent⁷⁶,
 800 H. Vincke²⁹, M.G. Vincter²⁸, V.B. Vinogradov⁶³, M. Virchaux^{136,*}, S. Viret⁵³, J. Virzi¹⁴,
 A. Vitale²⁰, I. Vivarelli¹²¹, R. Vives¹⁶², F. Vives Vagues¹¹, S. Vlachos⁹, H. Vogt⁴⁰, P. Vokac¹²⁶,
 C.F. Vollmer⁹⁸, M. Volpi¹¹, G. Volpini⁸⁸, R. von Boehn-Buchholz²⁹, H. von der Schmitt⁹⁹,
 E. von Toerne²¹, V. Vorobel¹²⁵, A.P. Vorobiev¹²⁷, A.S. Vorozhtsov⁶³, S.B. Vorozhtsov⁶³,
 M. Vos¹⁶², K.C. Voss¹⁶⁴, R. Voss²⁹, J.H. Vosseveld⁷¹, A.S. Vovenko¹²⁷, N. Vranjes¹², V. Vrba¹²⁴,
 805 M. Vreeswijk¹⁰⁶, T. Vu Anh⁴⁷, B. Vuaridel⁴⁷, M. Vudragovic¹², V. Vuillemin²⁹, R. Vuillermet²⁹,
 A. Wänanen³⁵, H. Wahlen¹⁶⁷, J. Walbersloh⁴¹, R. Walker¹⁵⁶, W. Walkowiak¹⁴¹, R. Wall¹⁶⁸,
 R.S. Wallny²⁹, S. Walsh¹³⁹, C. Wang⁴³, J.C. Wang¹³⁸, F. Wappler¹, A. Warburton⁸⁴, C.P. Ward²⁷,
 G.P. Warner¹²⁸, M. Warren⁷⁵, M. Warsinsky⁴², R. Wastie¹¹⁷, P.M. Watkins¹⁷, A.T. Watson¹⁷,
 G. Watts¹³⁸, A.T. Waugh¹⁴⁸, B.M. Waugh⁷⁵, C. Weaverdyck⁸⁶, M. Webel⁴⁶, G. Weber⁷⁹,
 810 J. Weber⁴¹, M. Weber¹²⁸, P. Weber⁵⁶, A.R. Weidberg¹¹⁷, P.M. Weilhammer²⁹, J. Weingarten⁴¹,
 C. Weiser⁴⁶, H. Wellenstein²³, H.P. Wellisch¹⁵⁶, P.S. Wells²⁹, A. Wemans¹²³, M. Wen⁴⁵,
 T. Wenaus²⁴, S. Wendler¹²², T. Wengler⁸⁰, S. Wenig²⁹, N. Wermes²¹, P. Werneke¹⁰⁶, P. Werner²⁹,
 U. Werthenbach¹⁴¹, S.J. Wheeler-Ellis⁶², S.P. Whitaker²², A. White⁷, M.J. White²⁷, S. White²⁴,
 D. Whittington⁵⁹, F. Wicek¹¹⁴, D. Wicke¹⁶⁷, F.J. Wickens¹²⁸, W. Wiedenmann¹⁶⁶, M. Wielers¹²⁸,
 815 P. Wienemann²¹, M. Wiesmann²⁹, M. Wiesmann⁹⁹, T. Wijnen¹⁰⁵, A. Wildauer²⁹, I. Wilhelm¹²⁵,
 H.G. Wilkens²⁹, H.H. Williams¹¹⁹, W. Willis³⁴, S. Willocq⁸³, I. Wilmut¹²⁸, J.A. Wilson¹⁷,
 A. Wilson⁸⁶, I. Wingerter-Seez⁴, L. Winton⁸⁵, W. Witzeling²⁹, T. Wlodek²⁴, E. Woehrling¹⁷,
 M.W. Wolter³⁸, H. Wolters¹²³, B. Wosiek³⁸, J. Wotschack²⁹, M.J. Woudstra⁸³, C. Wright⁵¹,

S.L. Wu¹⁶⁶, X. Wu⁴⁷, J. Wuestenfeld⁴¹, R. Wunstorf⁴¹, S. Xella-Hansen³⁵, A. Xiang³⁹, S. Xie⁴⁶,
 820 Y. Xie³², G. Xu³², N. Xu¹⁶⁶, A. Yamamoto⁶⁴, S. Yamamoto¹⁵³, H. Yamaoka⁶⁴, Y. Yamazaki⁶⁵,
 Z. Yan²², H. Yang⁸⁶, J.C. Yang³⁹, S. Yang¹¹⁷, U.K. Yang⁸⁰, Y. Yang³², Z. Yang²⁸, W-M. Yao¹⁴,
 Y. Yao², K. Yarradoddi^{24,i}, Y. Yasu⁶⁴, J. Ye³⁹, M. Yilmaz^{3,t}, R. Yoosooomiya¹²², K. Yorita³⁰,
 H. Yoshida¹⁰³, R. Yoshida⁵, C. Young¹⁴³, S.P. Youssef²², D. Yu²⁴, J. Yu⁷, M. Yu⁸¹, X. Yu³²,
 J. Yuan⁹⁹, A. Yurkewicz¹⁴⁷, V.G. Zaets¹²⁷, R. Zaidan⁸², A.M. Zaitsev¹²⁷, J. Zajac³⁷,
 825 Z. Zajacova²⁹, A.Yu. Zalite¹²⁰, Yo.K. Zalite¹²⁰, L. Zanello¹³³, P. Zarzhitsky³⁹, A. Zaytsev¹⁰⁷,
 M. Zdrzil¹⁴, C. Zeitnitz¹⁶⁷, M. Zeller¹⁶⁸, P.F. Zema²⁹, C. Zender²¹, A.V. Zenin¹²⁷, T. Zenis¹⁴⁴,
 Z. Zenonos¹²¹, S. Zenz¹⁴, D. Zerwas¹¹⁴, H. Zhang^{32,u}, J. Zhang⁵, W. Zheng¹²², X. Zhang³²,
 L. Zhao¹⁰⁸, T. Zhao¹³⁸, X. Zhao²⁴, Z. Zhao⁸⁶, A. Zhelezko⁹⁶, A. Zhemchugov⁶³, S. Zheng³²,
 L. Zhichao¹, B. Zhou⁸⁶, N. Zhou³⁴, S. Zhou¹⁴⁹, Y. Zhou¹⁴⁹, C.G. Zhu³², H.Z. Zhu¹³⁹,
 830 X.A. Zhuang⁹⁹, V. Zhuravlov⁶³, B. Zilka¹⁴⁴, N.I. Zimin⁶³, S. Zimmermann⁴⁶, M. Ziolkowski¹⁴¹,
 R. Zitoun⁴, L. Zivkovic¹⁶⁵, V.V. Zmouchko^{127,*}, G. Zobernig¹⁶⁶, A. Zoccoli²⁰, M.M. Zoeller¹⁰⁹,
 Y. Zolnierowski⁴, A. Zsenei²⁹, M. zur Nedden¹⁵, V. Zychacek¹²⁶.

¹ University at Albany, 1400 Washington Ave, Albany, NY 12222, United States of America

² Centre for Particle Physics, Department of Physics, University of Alberta, Edmonton, AB T6G
 835 2G7, Canada

³ Faculty of Sciences, Department of Physics, TR-061000 Tandogan, Ankara, Turkey

⁴ Laboratoire d'Annecy-le-Vieux de Physique des Particules LAPP, IN2P3/CNRS, Université de Savoie, France

⁵ Argonne National Laboratory, High Energy Physics Division, Building 362, 9700 S. Cass
 840 Avenue, Argonne - IL 60439, United States of America

⁶ University of Arizona, Department of Physics, Tucson, AZ 85721, United States of America

⁷ The University of Texas at Arlington, Department of Physics, Box 19059, Arlington, TX 76019, United States of America

⁸ Nuclear & Particle Physics Department of Physics, University of Athens, Panepistimiopouli,
 845 Zografou, GR- 15771 Athens, Greece

⁹ National Technical University of Athens, Physics Dept, 9-Iroon Polytechniou, GR - 15780 Zografou, Greece

¹⁰ Institute of Physics, Azerbaijan Academy of Sciences, H. Javid Avenue 33, AZ - 143 Baku, Azerbaijan

¹¹ Institut de Física d'Altes Energies, Universitat Autònoma de Barcelona, Edifici Cn, ES - 08193
 850 Bellaterra (Barcelona), Spain

¹² Institute of Physics, University of Belgrade, P.O. Box 57, 11001 Belgrade, Serbia and Montenegro

¹³ University of Bergen, Department for Physics and Technology, Allegaten 55, NO - 5007
 855 Bergen, Norway

¹⁴ Lawrence Berkeley National Laboratory and University of California, Physics Division, MS50B-6227, 1 Cyclotron Road, Berkeley, CA 94720, United States of America

¹⁵ Institute of Physics, Humboldt University, Berlin, Newtonstr. 15, D-12489 Berlin, Germany

¹⁶ Laboratory for High Energy Physics, University of Bern, Sidlerstrasse 5, CH - 3012 Bern,
 860 Switzerland

- ¹⁷ School of Physics and Astronomy, University of Birmingham, Edgbaston, Birmingham B15 2TT, United Kingdom
- ¹⁸ Faculty of Sciences, Department of Physics, Bogazici University, TR - 80815 Bebek-Istanbul, Turkey
- 865 ¹⁹ Universidad Antonio Narino, Centro de Investigaciones, Cra 3 Este No.47A-15, Bogota, Colombia
- ²⁰ INFN Bologna and Università di Bologna, Dipartimento di Fisica , viale C. Berti Pichat, 6/2, IT - 40127 Bologna, Italy
- ²¹ Physikalisches Institut der Universitaet Bonn, Nussallee 12, D - 53115 Bonn, Germany
- 870 ²² Boston University, Physics Department, , 590 Commonwealth Avenue, Boston, MA 02215, United States of America
- ²³ Brandeis University, Department of Physics, MS057, 415 South Street, Waltham, MA 02454, United States of America
- ²⁴ Brookhaven National Laboratory, Physics Department, Bldg. 510A, Upton, NY 11973, United States of America
- 875 ²⁵ National Institute of Physics and Nuclear Engineering, Bucharest, P.O. Box MG-6, R-077125, Romania
- ²⁶ Universidad de Buenos Aires, FCEyN, Dto. Fisica, Pab I - C. Universitaria, 1428 Buenos Aires, Argentina
- 880 ²⁷ Cavendish Laboratory, University of Cambridge, J J Thomson Avenue, Cambridge CB3 0HE, United Kingdom
- ²⁸ Carleton University, Department of Physics, 1125 Colonel By Drive, Ottawa ON K1S 5B6, Canada
- ²⁹ CERN, CH - 1211 Geneva 23, Switzerland
- 885 ³⁰ University of Chicago, Enrico Fermi Institute, 5640 S. Ellis Avenue, Chicago, IL 60637, United States of America
- ³¹ Pontificia Universidad Católica de Chile, Facultad de Fisica, Departamento de Fisica, Avda. Vicuna Mackenna 4860, San Joaquin, Santiago; Universidad Tecnica Federico Santa Maria, Departamento de Fisica, Avda. Espana 1680, Casilla 110-V, Valparaiso, Chile
- 890 ³² Institute of HEP, Chinese Academy of Sciences, P.O. Box 918, CN-100049 Beijing; USTC, Department of Modern Physics, Hefei, CN-230026 Anhui; Nanjing University, Department of Physics, CN-210093 Nanjing; Shandong University, HEP Group, CN-250100 Shadong, China
- ³³ Laboratoire de Physique Corpusculaire, IN2P3-CNRS, Universite Blaise Pascal, FR - 63177 Aubiere Cedex, France
- 895 ³⁴ Columbia University, Nevis Laboratory, 136 So. Broadway, Irvington, NY 10533, United States of America
- ³⁵ Niels Bohr Institute, University of Copenhagen, Blegdamsvej 17, DK - 2100 Kobenhavn 0, Denmark
- ³⁶ INFN Gruppo Collegato di Cosenza and Università della Calabria, Dipartimento di Fisica, IT-87036 Arcavacata di Rende, Italy
- 900 ³⁷ Faculty of Physics and Applied Computer Science of the AGH-University of Science and Technology, (FPACS, AGH-UST), al. Mickiewicza 30, PL-30059 Cracow, Poland
- ³⁸ The Henryk Niewodniczanski Institute of Nuclear Physics, Polish Academy of Sciences, ul.

- Radzikowskiego 152, PL - 31342 Krakow, Poland
- 905 ³⁹ Southern Methodist University, Physics Department, 106 Fondren Science Building, Dallas,
TX 75275-0175, United States of America
- ⁴⁰ DESY, Hamburg and Zeuthen, Notkestr. 85, D-22603 Hamburg, Germany
- ⁴¹ Universitaet Dortmund, Experimentelle Physik IV, DE - 44221 Dortmund, Germany
- ⁴² Institut fuer Kern- und Teilchenphysik, Technical University Dresden, Zellescher Weg 19,
910 D-01069 Dresden, Germany
- ⁴³ Duke University, Department of Physics, Durham, NC 27708, United States of America
- ⁴⁴ Fachhochschule Wiener Neustadt; Johannes Gutenbergstrasse 3 AT - 2700 Wiener Neustadt,
Austria
- ⁴⁵ INFN Laboratori Nazionali di Frascati, via Enrico Fermi 40, IT-00044 Frascati, Italy
- 915 ⁴⁶ Physikalisches Institut, Universitaet Freiburg, Hermann-Herder Str. 3, D - 79104 Freiburg i.Br.,
Germany
- ⁴⁷ Université de Genève, Section de Physique, 24 rue Ernest Ansermet, CH - 1211 Geneve 4,
Switzerland
- ⁴⁸ INFN Genova and Università di Genova, Dipartimento di Fisica, via Dodecaneso 33, IT -
920 16146 Genova, Italy
- ⁴⁹ Institute of Physics of the Georgian Academy of Sciences, 6 Tamarashvili St., GE - 380077
Tbilisi; Tbilisi State University, HEP Institute, University St. 9, GE - 380086 Tbilisi, Georgia
- ⁵⁰ II Physikalisches Institut, Justus-Liebig-Universitaet Giessen, Heinrich-Buff Ring 16, D-35392
Giessen, Germany
- 925 ⁵¹ University of Glasgow, Department of Physics and Astronomy, Glasgow G12 8QQ, United
Kingdom
- ⁵² Georg-August-Universität, II. Physikalisches Institut, Friedrich-Hund Platz 1, D-37077
Goettingen, Germany
- ⁵³ Laboratoire de Physique Subatomique et de Cosmologie, CNRS-IN2P3, Universite Joseph
930 Fourier, INPG, 53 avenue des Martyrs, FR - 38026 Grenoble Cedex, France
- ⁵⁴ Hampton University, Department of Physics, Hampton, VA 23668, United States of America
- ⁵⁵ Harvard University, Laboratory for Particle Physics and Cosmology, 18 Hammond Street,
Cambridge, MA 02138, United States of America
- ⁵⁶ Kirchhoff-Institut fuer Physik, Ruprecht-Karls-Universitaet Heidelberg, Im Neuenheimer Feld
935 227, DE - 69120 Heidelberg, Germany
- ⁵⁷ Hiroshima University, Faculty of Science, 1-3-1 Kagamiyama, Higashihiroshima-shi, JP -
Hiroshima 739-8526, Japan
- ⁵⁸ Hiroshima Institute of Technology, Faculty of Applied Information Science, 2-1-1 Miyake
Saeki-ku, Hiroshima-shi, JP - Hiroshima 731-5193, Japan
- 940 ⁵⁹ Indiana University, Department of Physics, Swain Hall West, Room 117, 727 East Third St.,
Bloomington, IN 47405-7105, United States of America
- ⁶⁰ Institut fuer Astro- und Teilchenphysik, Technikerstrasse 25, A - 6020 Innsbruck, Austria
- ⁶¹ Iowa State University, Dept of Physics and Astronomy, Ames High Energy Physics Group,
Ames, IA 50011-3160, United States of America
- 945 ⁶² University of California, Irvine, Department of Physics & Astronomy, Irvine ,CA 92697-4575,
United States of America

- ⁶³ Joint Institute for Nuclear Research, JINR Dubna, RU - 141 980 Moscow Region, Russia
- ⁶⁴ KEK, High Energy Accelerator Research Organization, 1-1 Oho, Tsukuba-shi, Ibaraki-ken 305-0801, Japan
- ⁹⁵⁰ ⁶⁵ Kobe University, Graduate School of Science, 1-1 Rokkodai-cho, Nada-ku, JP Kobe 657-8501, Japan
- ⁶⁶ Kyoto University, Faculty of Science, Oiwake-cho, Kitashirakawa, Sakyou-ku, Kyoto-shi, JP - Kyoto 606-8502, Japan
- ⁶⁷ Kyoto University of Education, 1 Fukakusa, Fujimori, fushimi-ku, Kyoto-shi, JP - Kyoto 612-8522, Japan
- ⁹⁵⁵ ⁶⁸ Universidad Nacional de La Plata, FCE, Departamento de Física, IFLP (CONICET-UNLP), C.C. 67, 1900 La Plata, Argentina
- ⁶⁹ Physics Department, Lancaster University, Lancaster LA1 4YB, United Kingdom
- ⁷⁰ INFN Lecce and Università del Salento, Dipartimento di Fisica, Via Arnesano IT - 73100 Lecce, Italy
- ⁹⁶⁰ ⁷¹ Oliver Lodge Laboratory, University of Liverpool, P.O. Box 147, Oxford Street, Liverpool L69 3BX, United Kingdom
- ⁷² Jožef Stefan Institute and Department of Physics, University of Ljubljana, SI-1000 Ljubljana, Slovenia
- ⁹⁶⁵ ⁷³ Department of Physics, Queen Mary, University of London, Mile End Road, London E1 4NS, United Kingdom
- ⁷⁴ Royal Holloway, University of London, Department of Physics, Egham Hill, Egham, Surrey TW20 0EX, United Kingdom
- ⁷⁵ University College London, Department of Physics and Astronomy, Gower Street, London WC1E 6BT, United Kingdom
- ⁹⁷⁰ ⁷⁶ Laboratoire de Physique Nucléaire et de Hautes Energies, Université Pierre et Marie Curie (Paris 6), Université Denis Diderot (Paris-7), IN2P3-CNRS, Tour 33, 4 place Jussieu, FR - 75252 Paris Cedex 05, France
- ⁷⁷ Lunds Universitet, Fysiska Institutionen, Box 118, SE - 221 00 Lund, Sweden
- ⁹⁷⁵ ⁷⁸ Departamento de Física Teórica, Facultad de Ciencias, Universidad Autónoma de Madrid, ES - 28049 Madrid, Spain
- ⁷⁹ Universitaet Mainz, Institut fuer Physik, Staudinger Weg 7, DE - 55099 Mainz, Germany
- ⁸⁰ School of Physics and Astronomy, University of Manchester, Manchester M13 9PL, United Kingdom
- ⁹⁸⁰ ⁸¹ Universitaet Mannheim, Lehrstuhl fuer Informatik V, B6, 23-29, DE - 68131 Mannheim, Germany
- ⁸² CPPM, Aix-Marseille Université, CNRS/IN2P3, Marseille, France
- ⁸³ Department of Physics, University of Massachusetts, 710 North Pleasant Street, Amherst, MA 01003, United States of America
- ⁹⁸⁵ ⁸⁴ McGill University, High Energy Physics Group, 3600 University Street, Montreal, Quebec H3A 2T8, Canada
- ⁸⁵ School of Physics, University of Melbourne, AU - Parkvill, Victoria 3010, Australia
- ⁸⁶ The University of Michigan, Department of Physics, 2477 Randall Laboratory, 500 East University, Ann Arbor, MI 48109-1120, United States of America

- 990 ⁸⁷ Michigan State University, High Energy Physics Group, Department of Physics and
Astronomy, East Lansing, MI 48824-2320, United States of America
- ⁸⁸ INFN Milano and Università di Milano, Dipartimento di Fisica, via Celoria 16, IT - 20133
Milano, Italy
- ⁸⁹ B.I. Stepanov Institute of Physics, National Academy of Sciences of Belarus, Independence
995 Avenue 68, Minsk 220072, Republic of Belarus
- ⁹⁰ National Scientific & Educational Centre of Particle & High Energy Physics, NC PHEP BSU,
M. Bogdanovich St. 153, Minsk 220040, Republic of Belarus
- ⁹¹ Massachusetts Institute of Technology, Department of Physics, Room 24-516, Cambridge, MA
02139, United States of America
- 1000 ⁹² University of Montreal, Group of Particle Physics, C.P. 6128, Succursale Centre-Ville,
Montreal, Quebec, H3C 3J7 , Canada
- ⁹³ Universite Hassan II, Faculte des Sciences Ain Chock, B.P. 5366, MA - Casablanca; Universite
Mohammed V, Faculte des Sciences, BP 1014, MO - Rabat, Morocco
- ⁹⁴ P.N. Lebedev Institute of Physics, Academy of Sciences, Leninsky pr. 53, RU - 117 924
1005 Moscow, Russia
- ⁹⁵ Institute for Theoretical and Experimental Physics (ITEP), B. Cheremushkinskaya ul. 25, RU
117 259 Moscow, Russia
- ⁹⁶ Moscow Engineering & Physics Institute (MEPhI), Kashirskoe Shosse 31, RU - 115409
Moscow, Russia
- 1010 ⁹⁷ Lomonosov Moscow State University, Skobeltsyn Institute of Nuclear Physics, RU - 119 991
GSP-1 Moscow Lenskie gory 1-2, Russia
- ⁹⁸ Fakultät für Physik der Ludwig-Maximilians-Universität München, Am Coulombwall 1, DE -
85748 Garching, Germany
- ⁹⁹ Max-Planck-Institut für Physik, (Werner-Heisenberg-Institut), Föhringer Ring 6, 80805
1015 München, Germany
- ¹⁰⁰ Nagasaki Institute of Applied Science, 536 Aba-machi, JP Nagasaki 851-0193, Japan
- ¹⁰¹ Nagoya University, Graduate School of Science, Furo-Cho, Chikusa-ku, Nagoya, 464-8602,
Japan
- ¹⁰² INFN Napoli and Università di Napoli, Dipartimento di Scienze Fisiche, Complesso
1020 Universitario di Monte Sant'Angelo, via Cinthia, IT - 80126 Napoli, Italy
- ¹⁰³ Naruto University of Education, Takashima, Naruto-cho, Naruto-shi, JP - Tokushima 772,
Japan
- ¹⁰⁴ Department of Physics and Astronomy, University of New Mexico, Albuquerque, NM 87131,
United States of America
- 1025 ¹⁰⁵ Radboud University Nijmegen/NIKHEF, Dept. of Exp. High Energy Physics, Toernooiveld 1,
NL - 6525 ED Nijmegen , Netherlands
- ¹⁰⁶ Nikhef National Institute for Subatomic Physics, Kruislaan 409, P.O. Box 41882, NL - 1009
DB Amsterdam, Netherlands
- ¹⁰⁷ Budker Institute of Nuclear Physics (BINP), RU - Novosibirsk 630 090, Russia
- 1030 ¹⁰⁸ Department of Physics, New York University, 4 Washington Place, New York NY 10003,
USA, United States of America
- ¹⁰⁹ Ohio State University, 191 West Woodruff Ave, Columbus, OH 43210-1117, United States of

America

- 110 Okayama University, Faculty of Science, Tsushimanaka 3-1-1, Okayama 700-8530, Japan
- 1035 111 University of Oklahoma, Homer L. Dodge Department of Physics, 440 West Brooks, Room 100, Norman, OK 73019-0225, United States of America
- 112 Department of Physics, Oklahoma State University, 145 Physical Sciences Building, Stillwater, OK 74078-3072, United States of America
- 113 1274 University of Oregon, Eugene, OR 97403-1274, United States of America
- 1040 114 LAL, Univ. Paris-Sud, IN2P3/CNRS, Orsay, France
- 115 Graduate School of Science, Osaka University, Machikaneyama-machi 1-1, Toyonaka, Osaka 560-0043, Japan
- 116 Department of Physics, University of Oslo, Blindern, NO - 0316 Oslo 3, Norway
- 117 Department of Physics, Oxford University, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, United Kingdom
- 1045 118 INFN Pavia and Università di Pavia, Dipartimento di Fisica Nucleare e Teorica, Via Bassi 6, IT-27100 Pavia, Italy
- 119 University of Pennsylvania, High Energy Physics Group, Dept. of Physics, 209 S. 33rd Street, Philadelphia, PA 19104, United States of America
- 1050 120 Petersburg Nuclear Physics Institute, RU - 188 300 Gatchina, Russia
- 121 INFN Pisa and Università di Pisa, Dipartimento di Fisica E. Fermi, Largo B.Pontecorvo 3, IT - 56127 Pisa, Italy
- 122 University of Pittsburgh, Department of Physics and Astronomy, 3941 O'Hara Street, Pittsburgh, PA 15260, United States of America
- 1055 123 Laboratorio de Instrumentacao e Fisica Experimental de Particulas - LIP (and Univ. de Coimbra, Univ. de Lisboa and Univ. Nova de Lisboa), Avenida Elias Garcia 14-1, PT - 1000-149 Lisboa, Portugal
- 124 Institute of Physics, Academy of Sciences of the Czech Republic, Na Slovance 2, CZ - 18221 Praha 8, Czech Republic
- 1060 125 Charles University in Prague, Faculty of Mathematics and Physics, Institute of Particle and Nuclear Physics, V Holesovickach 2, CZ - 18000 Praha 8, Czech Republic
- 126 Czech Technical University in Prague, Zikova 4, CZ - 166 35 Praha 6, Czech Republic
- 127 Institute for High Energy Physics (IHEP), Federal Agency of Atom. Energy, Moscow Region, RU - 142 284 Protvino, Russia
- 1065 128 Rutherford Appleton Laboratory, Science and Technology Facilities Council, Harwell Science and Innovation Campus, Didcot OX11 0QX, United Kingdom
- 129 Physics Department, University of Regina, Canada
- 130 Universidade Federal do Rio De Janeiro, Instituto de Fisica, Caixa Postal 68528, Ilha do Fundao, BR - 21945-970 Rio de Janeiro, Brazil
- 1070 131 Ritsumeikan University, Noji Higashi 1 chome 1-1, JP - Kusatsu, Shiga 525-8577, Japan
- 132 University of Rochester, Department of Physics & Astronomy, Rochester, NY 14627, United States of America
- 133 INFN Roma I and Università La Sapienza, Dipartimento di Fisica, Piazzale A. Moro 2, IT-00185 Roma, Italy
- 1075 134 INFN Roma Tor Vergata and Università di Roma Tor Vergata, Dipartimento di Fisica, via della

- Ricerca Scientifica, IT-00133 Roma, Italy
- 135 INFN Roma Tre and Università Roma Tre, Dipartimento di Fisica, via della Vasca Navale 84, IT-00146 Roma, Italy
- 136 CEA, DSM/DAPNIA, Centre d'Etudes de Saclay, FR - 91191 Gif-sur-Yvette, France
- 1080 137 University of California Santa Cruz, Santa Cruz Institute for Particle Physics (SCIPP), Santa Cruz, CA 95064, United States of America
- 138 Department of Physics, Department of Mechanical Engineering, University of Washington, Box 351560, Seattle, WA 98195, United States of America
- 139 University of Sheffield, Department of Physics & Astronomy, Hounsfield Road, Sheffield S3 7RH, United Kingdom
- 1085 140 Shinshu University, Dept. of Physics, Faculty of Science, 3-1-1 Asahi, Matsumoto-shi, JP - Nagano 390-8621, Japan
- 141 Universitaet Siegen, Fachbereich Physik, DE - 57068 Siegen, Germany
- 142 Department of Physics, Simon Fraser University, 8888 University Drive, CA - Burnaby, BC V5A 1S6, Canada
- 1090 143 Stanford Linear Accelerator Center, Stanford, California 94309, United States of America
- 144 Comenius University, Faculty of Mathematics, Physics & Informatics, Mlynska dolina F2, SK - 84248 Bratislava; Institute of Experimental Physics of the Slovak Academy of Sciences, Dept. of Subnuclear Physics, Watsonova 47, SK - 04353 Kosice, Slovak Republic
- 1095 145 Stockholm University, Department of Physics, AlbaNova, SE - 106 91 Stockholm, Sweden
- 146 Royal Institute of Technology (KTH), Physics Department, SE - 106 91 Stockholm, Sweden
- 147 Stony Brook University, Dept. of Physics and Astronomy, Nicolls Road, Stony Brook, NY 11794-3800, United States of America
- 148 School of Physics, University of Sydney, AU - Sydney NSW 2006, Australia
- 1100 149 Insitute of Physics, Academia Sinica, TW - Taipei 11529, Taiwan
- 150 Technion, Israel Inst. of Technology, Department of Physics, Technion City, IL - 32000 Haifa, Israel
- 151 Raymond and Beverly Sackler School of Physics and Astronomy, Tel Aviv University, Ramat Aviv, IL - Tel Aviv 69978, Israel
- 1105 152 Aristotle University of Thessaloniki, Department of Physics, Division of Nuclear & Particle Physics, GR - Thessaloniki 54124, Greece
- 153 The University of Tokyo, International Center for Elementary Particle Physics and Department of Physics, 7-3-1 Hongo, Bunkyo-ku, JP - Tokyo 113-0033, Japan
- 154 Graduate School of Science and Technology, Tokyo Metropolitan University, 1-1 Minami-Osawa, Hachioji, Tokyo 192-0397, Japan
- 1110 155 Department of Physics, University of Toronto, 60 Saint George Street, Toronto M5S 1A7, Ontario, Canada
- 156 TRIUMF, 4004 Wesbrook Mall, Vancouver, B.C. V6T 2A3, Canada
- 157 University of Tsukuba, Institute of Pure and Applied Sciences, 1-1-1 Tennoudai, Tsukuba-shi, JP - Ibaraki 305-8571, Japan
- 1115 158 Tufts University, Science & Technology Center, 4 Colby Street, Medford, MA 02155, United States of America
- 159 INFN Gruppo Collegato di Udine and Università di Udine, Dipartimento di Fisica, via delle

- Scienze 208, IT - 33100 Udine; INFN Gruppo Collegato di Udine and ICTP, Strada Costiera 11,
 1120 IT - 34014 Trieste, Italy
- ¹⁶⁰ University of Uppsala , Dept of Nuclear and Particle Physics, P.O. Box 535, SE- 75121
 Uppsala, Sweden
- ¹⁶¹ University of Illinois , Department of Physics, 1110 West Green Street, Urbana, Illinois 61801,
 United States of America
- 1125 ¹⁶² Instituto de Física Corpuscular (IFIC), Centro Mixto UVEG-CSIC, Apdo. 22085, ES-46071
 Valencia; Dept. Física At., Mol. y Nuclear, Univ. of Valencia and Instituto de Microelectrónica de
 Barcelona (IMB-CNM-CSIC), 08193 Bellaterra, Barcelona, Spain
- ¹⁶³ University of British Columbia, Dept of Physics, 6224 Agricultural Road, CA - Vancouver,
 B.C. V6T 1Z1, Canada
- 1130 ¹⁶⁴ University of Victoria, Dept., of Physics and Astronomy P.O. Box 3055 STN CSC Victoria
 B.C. V8W 3P6, Canada
- ¹⁶⁵ The Weizmann Institute of Science, Department of Particle Physics, P.O. Box 26, IL - 76100
 Rehovot, Israel
- ¹⁶⁶ University of Wisconsin, Department of Physics, Madison, WI 53706, United States of
 1135 America
- ¹⁶⁷ Bergische Universitaet, Fachbereich C, Physik, Postfach 100127, Gauss-Strasse 20, DE-
 42097 Wuppertal, Germany
- ¹⁶⁸ Yale University, Department of Physics, PO Box 208121, New Haven CT, 06520-8121 ,
 United States of America
- 1140 ¹⁶⁹ Yerevan Physics Institute, Alikhanian Brothers Street 2, AM - 375036 Yerevan, Armenia
- ^a Also at TRIUMF
- ^b Also at Gaziantep University, Turkey
- ^c Also at LPSC
- ^d Also at H. Niewodniczanski Institute of Nuclear Physics PAN, Cracow, Poland
- 1145 ^e Currently at Dogus University, Istanbul
- ^f Currently at Dumlupinar University, Kutahya
- ^g Also at Hamburg University
- ^h At Department of Physics, California State University, Fresno, 2345 E. San Ramon Avenue,
 Fresno, CA 93740-8031, United States of America
- 1150 ⁱ Louisiana Tech University
- ^j Currently at UC Riverside
- ^k UT Dallas
- ^l Now at Chonnam National University, Chonnam, Korea 500-757
- ^m Also at CERN
- 1155 ⁿ Also at Institute of Particle Physics (IPP), Canada
- ^o University of South Carolina
- ^p Also at Institute of Physics, Jagiellonian University, Cracow, Poland
- ^q Currently at TOBB University, Ankara
- ^r Also at KFKI Research Institute for Particle and Nuclear Physics, Budapest
- 1160 ^s Currently at TAEA, Ankara
- ^t Currently at Gazi University, Ankara

^u Also at CPPM, Aix-Marseille université, CNRS/IN2P3, Marseille, France

* Deceased

1. Overview of the ATLAS detector

1.1 Introduction

The Large Hadron Collider (LHC) at CERN will extend the frontiers of particle physics with its unprecedented high energy and luminosity. Inside the LHC, bunches of up to 10^{11} protons (p) will collide 40 million times per second to provide 14 TeV proton-proton collisions at a design luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$. The LHC will also collide heavy ions (A), in particular lead nuclei, at 5.5 TeV per nucleon pair, at a design luminosity of $10^{27} \text{ cm}^{-2}\text{s}^{-1}$.

The high interaction rates, radiation doses, particle multiplicities and energies, as well as the requirements for precision measurements have set new standards for the design of particle detectors. Two general purpose detectors, ATLAS (A Toroidal LHC ApparatuS) and CMS (Compact Muon Solenoid) have been built for probing p - p and A - A collisions.

This paper presents a comprehensive overview of the ATLAS detector prior to the first LHC collisions, written as the installation of the ATLAS detector is nearing completion. This detector represents the work of a large collaboration of several thousand physicists, engineers, technicians, and students over a period of fifteen years of dedicated design, development, fabrication, and installation.

1.2 Physics requirements and detector overview

The coordinate system and nomenclature used for describing the ATLAS detector and the particles emerging from the p - p collisions are briefly summarised here as they are used repeatedly throughout this paper. The beam direction defines the z -axis and the x - y plane is transverse to the beam direction. The positive x -axis is defined as pointing from the interaction point to the centre of the LHC ring and the positive y -axis is defined as pointing upwards. The side-A of the detector is defined as that with positive z and side-C is that with negative z . The azimuthal angle ϕ is measured around the beam axis, and the polar angle θ is the angle from the beam axis. The pseudorapidity is defined as $\eta \equiv -\ln \tan(\theta/2)$. The transverse momentum p_T , the transverse energy E_T , and the missing transverse energy E_T^{miss} are defined in the x - y plane unless stated otherwise. The distance ΔR in the pseudorapidity-azimuthal angle space is defined as $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}$.

The LHC is extensively reviewed in another article of this volume. It will provide a rich physics potential, ranging from more precise measurements of Standard Model parameters to the search for new physics phenomena. Furthermore, nucleus-nucleus collisions at the LHC provide an unprecedented opportunity to study the properties of strongly interacting matter at extreme energy density, including the possible phase transition to a colour-deconfined state: the quark-gluon plasma. Requirements for the ATLAS detector system [1] have been defined using a set of processes covering much of the new phenomena which one can hope to observe at the TeV scale.

The high luminosity and increased cross-sections at the LHC enable further high precision tests of QCD, electroweak interactions, and flavour physics. The top quark will be produced at the LHC at a rate of a few tens of Hz, providing the opportunity to test its couplings and spin.

The search for the Standard Model Higgs boson has been used as a benchmark to establish the performance of important sub-systems of ATLAS. It is a particularly important process since there is a range of production and decay mechanisms, depending on the mass of the Higgs boson, H . At

low masses ($m_H < 2m_Z$), the natural width would only be a few MeV, and so the observed width
1205 would be defined by the instrumental resolution. The predominant decay mode into hadrons would
be difficult to detect due to QCD backgrounds, and the two-photon decay channel would be an
important one. Other promising channels could be, for example, associated production of H such
as $t\bar{t}H$, WH , and ZH , with $H \rightarrow b\bar{b}$, using a lepton from the decay of one of the top quarks or of
the vector boson for triggering and background rejection. For masses above 130 GeV, Higgs-boson
1210 decays, $H \rightarrow ZZ^{(*)}$, where each Z decays to a pair of oppositely charged leptons, would provide
the experimentally cleanest channel to study the properties of the Higgs boson. For masses above
approximately 600 GeV, WW and ZZ decays into jets or involving neutrinos would be needed to
extract a signal. The tagging of forward jets from the WW or ZZ fusion production mechanism has
also been shown to be important for the discovery of the Higgs boson. Searches for the Higgs boson
1215 beyond the Standard Model, for such particles as the A and H^\pm of the minimal supersymmetric
extension of the Standard Model, require sensitivity to processes involving τ -leptons and good b -
tagging performance. Should the Higgs boson be discovered, it would need to be studied in several
modes, regardless of its mass, in order to fully disentangle its properties and establish its credentials
as belonging to the Standard Model or an extension thereof.

1220 New heavy gauge bosons W' and Z' could be accessible for masses up to ~ 6 TeV. To study
their leptonic decays, high-resolution lepton measurements and charge identification are needed in
the p_T -range of a few TeV. Another class of signatures of new physics may be provided by very
high- p_T jet measurements. As a benchmark process, quark compositeness has been used, where
the signature would be a deviation in the jet cross-sections from the QCD expectations. Searches
1225 for flavour-changing neutral currents and lepton flavour violation through $\tau \rightarrow 3\mu$ or $\tau \rightarrow \mu\gamma$, as
well as measurements of $B_s^0 \rightarrow \mu\mu$ and triple and quartic-gauge couplings may also open a window
onto new physics.

The decays of supersymmetric particles, such as squarks and gluinos, would involve cascades
which, if R-parity is conserved, always contain a lightest stable supersymmetric particle (LSP). As
1230 the LSP would interact very weakly with the detector, the experiment would measure a significant
missing transverse energy, E_T^{miss} , in the final state. The rest of the cascade would result in a number
of leptons and jets. In schemes where the LSP decays into a photon and a gravitino, an increased
number of hard isolated photons is expected.

Several new models propose the existence of extra dimensions leading to a characteristic en-
ergy scale of quantum gravity in the TeV region. In terms of experimental signatures, this could
1235 lead to the emission of gravitons which escape into extra dimensions and therefore generate E_T^{miss} ,
or of Kaluza-Klein excitations which manifest themselves as Z -like resonances with \sim TeV sep-
arations in mass. Other experimental signatures could be anomalous high-mass di-jet production,
and miniature black-hole production with spectacular decays involving democratic production of
1240 fundamental final states such as jets, leptons, photons, neutrinos, W 's, and Z 's.

The formidable LHC luminosity and resulting interaction rate are needed because of the small
cross-sections expected for many of the processes mentioned above. However, with an inelastic-
proton-proton cross-section of 80 mb, the LHC will produce a total rate of 10^9 inelastic events/s
at design luminosity. This presents a serious experimental difficulty as it implies that every candi-
1245 date event for new physics will on the average be accompanied by 23 inelastic events per bunch-
crossing.

The nature of proton-proton collisions imposes another difficulty. QCD jet production cross-sections dominate over the rare processes mentioned above, requiring the identification of experimental signatures characteristic of the physics processes in question, such as E_T^{miss} or secondary vertices. Identifying such final states for these rare processes imposes further demands on the integrated luminosity needed, and on the particle-identification capabilities of the detector.

Viewed in this context, these benchmark physics goals can be turned into a set of general requirements for the LHC detectors.

- Due to the experimental conditions at the LHC, the detectors require fast, radiation-hard electronics and sensor elements. In addition, high detector granularity is needed to handle the particle fluxes and to reduce the influence of overlapping events.
- Large acceptance in pseudorapidity with almost full azimuthal angle coverage is required.
- Good charged-particle momentum resolution and reconstruction efficiency in the inner tracker are essential. For offline tagging of τ -leptons and b -jets, vertex detectors close to the interaction region are required to observe secondary vertices.
- Very good electromagnetic (EM) calorimetry for electron and photon identification and measurements, complemented by full-coverage hadronic calorimetry for accurate jet and missing transverse energy measurements, are important requirements, as these measurements form the basis of many of the studies mentioned above.
- Good muon identification and momentum resolution over a wide range of momenta and the ability to determine unambiguously the charge of high p_T muons are fundamental requirements.
- Highly efficient triggering on low transverse-momentum objects with sufficient background rejection, is a prerequisite to achieve an acceptable trigger rate for most physics processes of interest.

The overall ATLAS detector layout is shown in Fig. 1 and its main performance goals are listed in Table 1. It is important to note that, for high- p_T muons, the muon-spectrometer performance as given in Table 1 is independent of the inner-detector system.

The magnet configuration comprises a thin superconducting solenoid surrounding the inner-detector cavity, and three large superconducting toroids (one barrel and two end-caps) arranged with an eight-fold azimuthal symmetry around the calorimeters. This fundamental choice drove the design of the rest of the detector.

The inner detector is immersed in a 2 T solenoidal field. Pattern recognition, momentum and vertex measurements, and electron identification are achieved with a combination of discrete, high-resolution semiconductor pixel and strip detectors in the inner part of the tracking volume, and straw-tube tracking detectors with the capability to generate and detect transition radiation in its outer part.

High granularity liquid-argon (LAr) electromagnetic sampling calorimeters, with excellent performance in terms of energy and position resolution, cover the pseudorapidity range $|\eta| < 3.2$. The hadronic calorimetry in the range $|\eta| < 1.7$ is provided by a scintillator-tile calorimeter, which




Figure 1. Cut-away view of the ATLAS detector. The dimensions of the detector are 25 m in height and 44 m in length. The overall weight of the detector is approximately 7000 tonnes.

is separated into a large barrel and two smaller extended barrel cylinders, one on either side of the central barrel. In the end-caps ($|\eta| > 1.5$), LAr technology is also used for the hadronic calorimeters, matching the outer $|\eta|$ limits of end-cap electromagnetic calorimeters. The LAr forward calorimeters provide both electromagnetic and hadronic energy measurements, and extend the pseudorapidity coverage to $|\eta| = 4.9$.

The calorimeter is surrounded by the muon spectrometer. The air-core toroid system, with a long barrel and two inserted end-cap magnets, generates strong bending power in a large volume within a light and open structure. Multiple-scattering effects are thereby minimised, and excellent muon momentum resolution is achieved with three layers of high precision tracking chambers. The muon instrumentation includes, as a key component, trigger chambers with timing resolution of the order of 1.5-4 ns. The muon spectrometer defines the overall dimensions of the ATLAS detector.

The proton-proton interaction rate at the design luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$ is approximately 1 GHz, while the event data recording, based on technology and resource limitations, is limited to about 200 Hz. This requires an overall rejection factor of 5×10^6 against minimum-bias processes while maintaining maximum efficiency for the new physics. The Level-1 (L1) trigger system uses a subset of the total detector information to make a decision on whether or not to continue processing an event, reducing the data rate to approximately 75 kHz (limited by the bandwidth of the readout system, which is upgradeable to 100 kHz). The subsequent two levels, collectively known as the high-level trigger, are the Level-2 (L2) trigger and the event filter. They provide the reduction to a final data-taking rate of approximately 200 Hz.

Due to budgetary constraints, some detector systems had to be staged. They will be completed and installed as soon as technically and financially feasible. These include, in particular, a sig-

Detector component	Required resolution	η coverage	
		Measurement	Trigger
Tracking	$\sigma_{p_T}/p_T = 0.05\% p_T \oplus 1\%$	± 2.5	
EM calorimetry	$\sigma_E/E = 10\%/\sqrt{E} \oplus 0.7\%$	± 3.2	± 2.5
Hadronic calorimetry (jets)			
barrel and end-cap	$\sigma_E/E = 50\%/\sqrt{E} \oplus 3\%$	± 3.2	± 3.2
forward	$\sigma_E/E = 100\%/\sqrt{E} \oplus 10\%$	$3.1 < \eta < 4.9$	$3.1 < \eta < 4.9$
Muon spectrometer	$\sigma_{p_T}/p_T = 10\%$ at $p_T = 1$ TeV	± 2.7	± 2.4

Table 1. General performance goals of the ATLAS detector. Note that, for high- p_T muons, the muon-spectrometer performance is independent of the inner-detector system. The units for E and p_T are in GeV.

nificant part of the high-level trigger processing farm. The initial input capacity will be limited to a L1 trigger rate of about 40 kHz. This capacity will be increased as needed to deal with the LHC luminosity profile during the first years. The ultimate goal is to be able to handle 100 kHz if needed. Some parts of the muon spectrometer are staged, most noticeably part of the precision chambers in the transition region between the barrel and the end-caps. In addition, some of the forward shielding elements will be completed later, as the LHC approaches design luminosity.

1.3 Tracking

Approximately 1000 particles will emerge from the collision point every 25 ns within $|\eta| < 2.5$, creating a very large track density in the detector. To achieve the momentum and vertex resolution requirements imposed by the benchmark physics processes, high-precision measurements must be made with fine detector granularity. Pixel and silicon microstrip (SCT) trackers, used in conjunction with the straw tubes of the Transition Radiation Tracker (TRT), offer these features.

The layout of the Inner Detector (ID) is illustrated in Fig. 2 and detailed in Chapter 4. Its basic parameters are summarised in Table 2 (also see intrinsic accuracies in Table 10). The ID is immersed in a 2 T magnetic field generated by the central solenoid, which extends over a length of 5.3 m with a diameter of 2.5 m. The precision tracking detectors (pixels and SCT) cover the region $|\eta| < 2.5$. In the barrel region, they are arranged on concentric cylinders around the beam axis while in the end-cap regions they are located on disks perpendicular to the beam axis. The highest granularity is achieved around the vertex region using silicon pixel detectors. All pixel sensors are identical and have a minimum pixel size of $50 \times 400 \mu\text{m}^2$. The pixel layers are segmented in $R - \phi$ and z with typically three pixel layers crossed by each track. The intrinsic accuracies in the barrel are $10 \mu\text{m}$ ($R - \phi$) and $115 \mu\text{m}$ (z) and in the disks are $10 \mu\text{m}$ ($R - \phi$) and $115 \mu\text{m}$ (R). The pixel detector has approximately 80.4 million readout channels. For the SCT, eight strip layers (four space points) are crossed by each track. In the barrel region, this detector uses small-angle (40 mrad) stereo strips to measure both coordinates, with one set of strips in each layer parallel to the beam direction, measuring $R - \phi$. They consist of two 6.4 cm long daisy-chained sensors with a strip pitch of $80 \mu\text{m}$. In the end-cap region, the detectors have a set of strips running radially and a set of stereo strips at an angle of 40 mrad. The mean pitch of the strips is also approximately $80 \mu\text{m}$. The intrinsic accuracies per module in the barrel are $17 \mu\text{m}$ ($R - \phi$) and $580 \mu\text{m}$ (z) and in




Figure 2. Cut-away view of the ATLAS inner detector.

the disks are $17 \mu\text{m}$ ($R - \phi$) and $580 \mu\text{m}$ (R). The total number of readout channels in the SCT is approximately 6.3 million.

1340 A large number of hits (typically 36 per track) is provided by the 4 mm diameter straw tubes of the TRT, which enables track-following up to $|\eta| = 2.0$. The TRT only provides $R - \phi$ information, for which it has an intrinsic accuracy of $130 \mu\text{m}$ per straw. In the barrel region, the straws are parallel to the beam axis and are 144 cm long, with their wires divided into two halves, approximately at $\eta = 0$. In the end-cap region, the 37 cm long straws are arranged radially in wheels. The total number of TRT readout channels is approximately 351,000.

1345 The combination of precision trackers at small radii with the TRT at a larger radius gives very robust pattern recognition and high precision in both $R - \phi$ and z coordinates. The straw hits at the outer radius contribute significantly to the momentum measurement, since the lower precision per point compared to the silicon is compensated by the large number of measurements and longer measured track length.

1350 The inner detector system provides tracking measurements in a range matched by the precision measurements of the electromagnetic calorimeter. The electron identification capabilities are enhanced by the detection of transition-radiation photons in the xenon-based gas mixture of the straw tubes. The semiconductor trackers also allow impact parameter measurements and vertexing for heavy-flavour and τ -lepton tagging. The secondary vertex measurement performance is enhanced
1355 by the innermost layer of pixels, at a radius of about 5 cm.

1.4 Calorimetry

A view of the sampling calorimeters is presented in Fig. 3, and the pseudorapidity coverage, granularity, and longitudinal segmentation of the calorimeters are summarised in Table 3 (see also

Item		Radial extension (mm)	Length (mm)
Overall ID envelope		$0 < R < 1150$	$0 < z < 3512$
Beam-pipe		$29 < R < 36$	
Pixel	Overall envelope	$45.5 < R < 242$	$0 < z < 3092$
3 cylindrical layers	Sensitive barrel	$50.5 < R < 122.5$	$0 < z < 400.5$
2×3 disks	Sensitive end-cap	$88.8 < R < 149.6$	$495 < z < 650$
SCT	Overall envelope	$255 < R < 549$ (barrel)	$0 < z < 805$
		$251 < R < 610$ (end-cap)	$810 < z < 2797$
4 cylindrical layers	Sensitive barrel	$299 < R < 514$	$0 < z < 749$
2×9 disks	Sensitive end-cap	$275 < R < 560$	$839 < z < 2735$
TRT	Overall envelope	$554 < R < 1082$ (barrel)	$0 < z < 780$
		$617 < R < 1106$ (end-cap)	$827 < z < 2744$
73 straw planes	Sensitive barrel	$563 < R < 1066$	$0 < z < 712$
160 straw planes	Sensitive end-cap	$644 < R < 1004$	$848 < z < 2710$

Table 2. Main parameters of the inner-detector system.

Chapter 5). These calorimeters cover the range $|\eta| < 4.9$, using different techniques suited to the
1360 widely varying requirements of the physics processes of interest and of the radiation environment
over this large η -range. Over the η region matched to the inner detector, the fine granularity of
the EM calorimeter is ideally suited for precision measurements of electrons and photons. The
coarser granularity of the rest of the calorimeter is sufficient to satisfy the physics requirements for
jet reconstruction and E_T^{miss} measurements.

1365 Calorimeters must provide good containment for electromagnetic and hadronic showers, as
well as limit punch-through into the muon system. Hence, calorimeter depth is an important con-
sideration. The total thickness of the EM calorimeter is > 22 radiation lengths (X_0) in the barrel
and $> 24 X_0$ in the end-caps. The approximate 9.7 interaction lengths (λ) of active calorimeter in
the barrel (10λ in the end-caps) are adequate to provide good resolution for high-energy jets (see
1370 Table 1). The total thickness, including 1.3λ from the outer support, is 11λ at $\eta = 0$ and has been
shown both by measurements and simulations to be sufficient to reduce punch-through well below
the irreducible level of prompt or decay muons. Together with the large η -coverage, this thickness
will also ensure a good E_T^{miss} measurement, which is important for many physics signatures and in
particular for SUSY particle searches.

1375 1.4.1 LAr electromagnetic calorimeter

The EM calorimeter is divided into a barrel part ($|\eta| < 1.475$) and two end-cap components
($1.375 < |\eta| < 3.2$), each housed in their own cryostat. The position of the central solenoid in
front of the EM calorimeter demands optimisation of the material in order to achieve the de-

	Barrel		End-cap	
EM calorimeter				
Number of layers and $ \eta $ coverage				
Presampler	1	$ \eta < 1.52$	1	$1.5 < \eta < 1.8$
Calorimeter	3	$ \eta < 1.35$	2	$1.375 < \eta < 1.5$
	2	$1.35 < \eta < 1.475$	3	$1.5 < \eta < 2.5$
			2	$2.5 < \eta < 3.2$
Granularity $\Delta\eta \times \Delta\phi$ versus $ \eta $				
Presampler	0.025×0.1	$ \eta < 1.52$	0.025×0.1	$1.5 < \eta < 1.8$
Calorimeter 1st layer	$0.025/8 \times 0.1$	$ \eta < 1.40$	0.050×0.1	$1.375 < \eta < 1.425$
	0.025×0.025	$1.40 < \eta < 1.475$	0.025×0.1	$1.425 < \eta < 1.5$
			$0.025/8 \times 0.1$	$1.5 < \eta < 1.8$
			$0.025/6 \times 0.1$	$1.8 < \eta < 2.0$
			$0.025/4 \times 0.1$	$2.0 < \eta < 2.4$
			0.025×0.1	$2.4 < \eta < 2.5$
0.1×0.1	$2.5 < \eta < 3.2$			
Calorimeter 2nd layer	0.025×0.025	$ \eta < 1.40$	0.050×0.025	$1.375 < \eta < 1.425$
	0.075×0.025	$1.40 < \eta < 1.475$	0.025×0.025	$1.425 < \eta < 2.5$
			0.1×0.1	$2.5 < \eta < 3.2$
Calorimeter 3rd layer	0.050×0.025	$ \eta < 1.35$	0.050×0.025	$1.5 < \eta < 2.5$
Number of readout channels				
Presampler	7808		1536 (both sides)	
Calorimeter	101760		62208 (both sides)	
LAr hadronic end-cap				
$ \eta $ coverage			$1.5 < \eta < 3.2$	
Number of layers			4	
Granularity $\Delta\eta \times \Delta\phi$			0.1×0.1	$1.5 < \eta < 2.5$
			0.2×0.2	$2.5 < \eta < 3.2$
Readout channels			5632 (both sides)	
LAr forward calorimeter				
$ \eta $ coverage			$3.1 < \eta < 4.9$	
Number of layers			3	
Granularity $\Delta x \times \Delta y$ (cm)			FCal1: 3.0×2.6	$3.15 < \eta < 4.30$
			FCal1: \sim four times finer	$3.10 < \eta < 3.15,$ $4.30 < \eta < 4.83$
			FCal2: 3.3×4.2	$3.24 < \eta < 4.50$
			FCal2: \sim four times finer	$3.20 < \eta < 3.24,$ $4.50 < \eta < 4.81$
			FCal3: 5.4×4.7	$3.32 < \eta < 4.60$
			FCal3: \sim four times finer	$3.29 < \eta < 3.32,$ $4.60 < \eta < 4.75$
Readout channels			3524 (both sides)	
Scintillator tile calorimeter				
	Barrel		Extended barrel	
$ \eta $ coverage	$ \eta < 1.0$		$0.8 < \eta < 1.7$	
Number of layers	3		3	
Granularity $\Delta\eta \times \Delta\phi$	0.1×0.1		0.1×0.1	
	Last layer 0.2×0.1		0.2×0.1	
Readout channels	5760		4092 (both sides)	

Table 3. Main parameters of the calorimeter system.

sired calorimeter performance. As a consequence, the central solenoid and the LAr calorimeter share a common vacuum vessel, thereby eliminating two vacuum walls. The barrel calorimeter consists of two identical half-barrels, separated by a small gap (4 mm) at $z = 0$. Each end-cap calorimeter is mechanically divided into two coaxial wheels: an outer wheel covering the region




Figure 3. Cut-away view of the ATLAS calorimeter system.

1.375 < $|\eta|$ < 2.5, and an inner wheel covering the region $2.5 < |\eta| < 3.2$. The EM calorimeter is a lead-LAr detector with accordion-shaped kapton electrodes and lead absorber plates over its full coverage. The accordion geometry provides complete ϕ symmetry without azimuthal cracks. The lead thickness in the absorber plates has been optimised as a function of η in terms of EM calorimeter performance in energy resolution. Over the region devoted to precision physics ($|\eta| < 2.5$), the EM calorimeter is segmented into three longitudinal sections. For the end-cap inner wheel, the calorimeter is segmented in two longitudinal sections and has a coarser lateral granularity than for the rest of the acceptance.

In the region of $|\eta| < 1.8$, a presampler detector is used to correct for the energy lost by electrons and photons upstream of the calorimeter. The presampler consists of an active LAr layer of thickness 1.1 cm (0.5 cm) in the barrel (end-cap) region.

1.4.2 Hadronic calorimeters

Tile calorimeter The tile calorimeter is placed directly outside the EM calorimeter envelope. Its barrel covers the region $|\eta| < 1.0$, and its two extended barrels the range $0.8 < |\eta| < 1.7$. It is a sampling calorimeter using steel as the absorber and scintillating tiles as the active material. The barrel and extended barrels are divided azimuthally into 64 modules. Radially, the tile calorimeter extends from an inner radius of 2.28 m to an outer radius of 4.25 m. It is longitudinally segmented in three layers approximately 1.5, 4.1 and 1.8 interaction lengths thick for the barrel and 1.5, 2.6, and 3.3 λ for the extended barrel. The total detector thickness at the outer edge of the tile-instrumented region is 9.7 λ at $\eta = 0$. The tiles are 3 mm thick and the total thickness of the steel plates in one period is 14 mm. Two sides of the scintillating tiles are read out by wavelength shifting fibres

into two separate photomultiplier tubes. In η , the readout cells built by grouping fibres into the
1405 photomultipliers are pseudo-projective towards the interaction region.

LAr hadronic end-cap calorimeter The Hadronic End-cap Calorimeter (HEC) consists of two independent wheels per end-cap, located directly behind the end-cap electromagnetic calorimeter and sharing the same LAr cryostats. To reduce the drop in material density at the transition between the end-cap and the forward calorimeter (around $|\eta| = 3.1$), the HEC extends out to
1410 $|\eta| = 3.2$, thereby overlapping with the forward calorimeter. Similarly, the HEC η range also slightly overlaps that of the tile calorimeter ($|\eta| < 1.7$) by extending to $|\eta| = 1.5$. Each wheel is built from 32 identical wedge-shaped modules, assembled with fixtures at the periphery and at the central bore. Each wheel is divided into two longitudinal segments, for a total of four layers per end-cap. The wheels closest to the interaction point are built from 25 mm parallel copper plates,
1415 while those further away use 50 mm copper plates (for all wheels the first plate is half-thickness). The outer radius of the copper plates is 2.03 m, while the inner radius is 0.475 m (except in the overlap region with the forward calorimeter where this radius becomes 0.372 m). The copper plates are interleaved with 8.5 mm LAr gaps, providing the active medium for this sampling calorimeter.

LAr forward calorimeter The Forward Calorimeter (FCal) is integrated into the end-cap
1420 cryostats, as this provides clear benefits in terms of uniformity of the calorimetric coverage as well as reduced radiation background levels in the muon spectrometer. In order to reduce the amount of neutron albedo in the inner detector cavity, the front face of the FCal is recessed by about 1.2 m with respect to the EM calorimeter front face. This severely limits longitudinal space and therefore calls for a high-density design. The FCal is approximately 10 interaction lengths
1425 deep, and consists of three modules in each end-cap: the first, made of copper, is optimised for electromagnetic measurements, while the other two, made of tungsten, measure predominantly the energy of hadronic interactions. Each module consists of a metal matrix, with regularly spaced longitudinal channels filled with the electrode structure consisting of concentric rods and tubes parallel to the beam axis. The LAr in the gap between the rod and the tube is the sensitive medium.
1430 This geometry allows for excellent control of the gaps, which are as small as 0.25 mm in the first section, in order to avoid problems due to ion buildup.

1.5 Muon system

The conceptual layout of the muon spectrometer is shown in Fig. 4 and the main parameters of the muon chambers are listed in Table 4 (see also Chapter 6). It is based on the magnetic deflection of
1435 muon tracks in the large superconducting air-core toroid magnets, instrumented with separate trigger and high-precision tracking chambers. Over the range $|\eta| < 1.4$, magnetic bending is provided by the large barrel toroid. For $1.6 < |\eta| < 2.7$, muon tracks are bent by two smaller end-cap magnets inserted into both ends of the barrel toroid. Over $1.4 < |\eta| < 1.6$, usually referred to as the transition region, magnetic deflection is provided by a combination of barrel and end-cap fields.
1440 This magnet configuration provides a field which is mostly orthogonal to the muon trajectories, while minimising the degradation of resolution due to multiple scattering. The anticipated high level of particle flux has had a major impact on the choice and design of the spectrometer instrumentation, affecting performance parameters such as rate capability, granularity, ageing properties, and radiation hardness.




Figure 4. Cut-away view of the ATLAS muon system.

1445 In the barrel region, tracks are measured in chambers arranged in three cylindrical layers
 around the beam axis; in the transition and end-cap regions, the chambers are installed in planes
 perpendicular to the beam, also in three layers.

1.5.1 The toroid magnets

A system of three large air-core toroids generates the magnetic field for the muon spectrometer.
 1450 The two end-cap toroids are inserted in the barrel toroid at each end and line up with the central
 solenoid. Each of the three toroids consists of eight coils assembled radially and symmetrically
 around the beam axis. The end-cap toroid coil system is rotated by 22.5° with respect to the barrel
 toroid coil system in order to provide radial overlap and to optimise the bending power at the
 interface between the two coil systems.

1455 The barrel toroid coils are housed in eight individual cryostats, with the linking elements
 between them providing the overall mechanical stability. Each end-cap toroid consists of eight
 racetrack-like coils in an aluminium alloy housing. Each coil has two double-pancake type wind-
 ings. They are cold-linked and assembled as a single cold mass, housed in one large cryostat.
 Therefore the internal forces in the end-cap toroids are taken by the cold supporting structure be-
 1460 tween the coils, a different design solution than in the barrel toroid.

The performance in terms of bending power is characterised by the field integral $\int B dl$, where B
 is the field component normal to the muon direction and the integral is computed along an infinite-

momentum muon trajectory, between the innermost and outermost muon-chamber planes. The barrel toroid provides 1.5 to 5.5 Tm of bending power in the pseudorapidity range $0 < |\eta| < 1.4$, and the end-cap toroids approximately 1 to 7.5 Tm in the region $1.6 < |\eta| < 2.7$. The bending power is lower in the transition regions where the two magnets overlap ($1.4 < |\eta| < 1.6$).

1.5.2 Muon chamber types

Over most of the η -range, a precision measurement of the track coordinates in the principal bending direction of the magnetic field is provided by Monitored Drift Tubes (MDT's). The mechanical isolation in the drift tubes of each sense wire from its neighbours guarantees a robust and reliable operation. At large pseudorapidities, Cathode Strip Chambers (CSC's, which are multiwire proportional chambers with cathodes segmented into strips) with higher granularity are used in the innermost plane over $2 < |\eta| < 2.7$, to withstand the demanding rate and background conditions. The stringent requirements on the relative alignment of the muon chamber layers are met by the combination of precision mechanical-assembly techniques and optical alignment systems both within and between muon chambers.

The trigger system covers the pseudorapidity range $|\eta| < 2.4$. Resistive Plate Chambers (RPC's) are used in the barrel and Thin Gap Chambers (TGC's) in the end-cap regions. The trigger chambers for the muon spectrometer serve a threefold purpose: provide bunch-crossing identification, provide well-defined p_T thresholds, and measure the muon coordinate in the direction orthogonal to that determined by the precision-tracking chambers.

1.5.3 Muon chamber alignment and B-field reconstruction

The overall performance over the large areas involved, particularly at the highest momenta, depends on the alignment of the muon chambers with respect to each other and with respect to the overall detector.

The accuracy of the stand-alone muon momentum measurement necessitates a precision of $30 \mu\text{m}$ on the relative alignment of chambers both within each projective tower and between consecutive layers in immediately adjacent towers. The internal deformations and relative positions of the MDT chambers are monitored by approximately 12000 precision-mounted alignment sensors, all based on the optical monitoring of deviations from straight lines. Because of geometrical constraints, the reconstruction and/or monitoring of the chamber positions rely on somewhat different strategies and sensor types in the end-cap and barrel regions, respectively.

The accuracy required for the relative positioning of non-adjacent towers to obtain adequate mass resolution for multi-muon final states, lies in the few millimetre range. This initial positioning accuracy is approximately established during the installation of the chambers. Ultimately, the relative alignment of the barrel and forward regions of the muon spectrometer, of the calorimeters and of the inner detector will rely on high-momentum muon trajectories.

For magnetic field reconstruction, the goal is to determine the bending power along the muon trajectory to a few parts in a thousand. The field is continuously monitored by a total of approximately 1800 Hall sensors distributed throughout the spectrometer volume. Their readings are compared with magnetic-field simulations and used for reconstructing the position of the toroid coils in space, as well as to account for magnetic perturbations induced by the tile calorimeter and other nearby metallic structures.

Monitored drift tubes - Coverage - Number of chambers - Number of channels - Function	MDT $ \eta < 2.7$ (innermost layer: $ \eta < 2.0$) 1088 (1150) 339 000 (354 000) Precision tracking
Cathode strip chambers - Coverage - Number of chambers - Number of channels - Function	CSC $2.0 < \eta < 2.7$ 32 31 000 Precision tracking
Resistive plate chambers - Coverage - Number of chambers - Number of channels - Function	RPC $ \eta < 1.05$ 544 (606) 359 000 (373 000) Triggering, second coordinate
Thin gap chambers - Coverage - Number of chambers - Number of channels - Function	TGC $1.05 < \eta < 2.7$ (2.4 for triggering) 3588 318 000 Triggering, second coordinate

Table 4. Main parameters of the muon spectrometer. Numbers in brackets for the MDT's and the RPC's refer to the final configuration of the detector in 2009.

1.6 Forward detectors

1505 Three smaller detector systems cover the ATLAS forward region (see Chapter 7). The main func-
tion of the first two systems is to determine the luminosity delivered to ATLAS. At ± 17 m from the
interaction point lies LUCID (LUminosity measurement using Cerenkov Integrating Detector). It
detects inelastic p - p scattering in the forward direction, and is the main online relative-luminosity
monitor for ATLAS. The second detector is ALFA (Absolute Luminosity For ATLAS). Located at
1510 ± 240 m, it consists of scintillating fibre trackers located inside Roman pots which are designed to
approach as close as 1 mm to the beam. The third system is the Zero-Degree Calorimeter (ZDC),
which plays a key role in determining the centrality of heavy-ion collisions. It is located at ± 140 m
from the interaction point, just beyond the point where the common straight-section vacuum-pipe
divides back into two independent beam-pipes. The ZDC modules consist of layers of alternating
1515 quartz rods and tungsten plates which will measure neutral particles at pseudorapidities $|\eta| \geq 8.2$.

1.7 Trigger, readout, data acquisition, and control systems

The Trigger and Data Acquisition (collectively TDAQ) systems, the timing- and trigger-control
logic, and the Detector Control System (DCS) are partitioned into sub-systems, typically associated
with sub-detectors, which have the same logical components and building blocks (see Chapter 8).

1520 The trigger system has three distinct levels: L1, L2, and the event filter. Each trigger level
refines the decisions made at the previous level and, where necessary, applies additional selection
criteria. The data acquisition system receives and buffers the event data from the detector-specific
readout electronics, at the L1 trigger accept rate, over 1600 point-to-point readout links. The first
level uses a limited amount of the total detector information to make a decision in less than $2.5 \mu\text{s}$,
1525 reducing the rate to about 75 kHz. The two higher levels access more detector information for a
final rate of up to 200 Hz with an event size of approximately 1.3 Mbyte.

1.7.1 Trigger system

The L1 trigger searches for high transverse-momentum muons, electrons, photons, jets, and τ -
leptons decaying into hadrons, as well as large missing and total transverse energy. Its selection is
1530 based on information from a subset of detectors. High transverse-momentum muons are identified
using trigger chambers in the barrel and end-cap regions of the spectrometer. Calorimeter selections
are based on reduced-granularity information from all the calorimeters. Results from the L1 muon
and calorimeter triggers are processed by the central trigger processor, which implements a trigger
'menu' made up of combinations of trigger selections. Pre-scaling of trigger menu items is also
1535 available, allowing optimal use of the bandwidth as luminosity and background conditions change.
Events passing the L1 trigger selection are transferred to the next stages of the detector-specific
electronics and subsequently to the data acquisition via point-to-point links.

In each event, the L1 trigger also defines one or more Regions-of-Interest (RoI's), i.e. the
geographical coordinates in η and ϕ , of those regions within the detector where its selection process
1540 has identified interesting features. The RoI data include information on the type of feature identified
and the criteria passed, e.g. a threshold. This information is subsequently used by the high-level
trigger.

The L2 selection is seeded by the RoI information provided by the L1 trigger over a dedicated
data path. L2 selections use, at full granularity and precision, all the available detector data within
1545 the RoI's (approximately 2% of the total event data). The L2 menus are designed to reduce the
trigger rate to approximately 3.5 kHz, with an event processing time of about 40 ms, averaged over
all events. The final stage of the event selection is carried out by the event filter, which reduces
the event rate to roughly 200 Hz. Its selections are implemented using offline analysis procedures
within an average event processing time of the order of four seconds.

1550 1.7.2 Readout architecture and data acquisition

The Readout Drivers (ROD's) are detector-specific functional elements of the front-end systems,
which achieve a higher level of data concentration and multiplexing by gathering information from
several front-end data streams. Although each sub-detector uses specific front-end electronics and
ROD's, these components are built from standardised blocks and are subject to common require-
1555 ments. The front-end electronics sub-system includes different functional components:

- the front-end analogue or analogue-to-digital processing;
- the L1 buffer in which the (analogue or digital) information is retained for a time long enough
to accommodate the L1 trigger latency;

- the derandomising buffer in which the data corresponding to a L1 trigger accept are stored before being sent to the following level. This element is necessary to accommodate the maximum instantaneous L1 rate without introducing significant deadtime (maximum 1%);
- the dedicated links or buses which are used to transmit the front-end data stream to the next stage.

After an event is accepted by the L1 trigger, the data from the pipe-lines are transferred off the detector to the ROD's. Digitised signals are formatted as raw data prior to being transferred to the DAQ system. The ROD's follow some general ATLAS rules, including the definition of the data format of the event, the error detection/recovery mechanisms to be implemented, and the physical interface for the data transmission to the DAQ system.

The first stage of the DAQ, the readout system, receives and temporarily stores the data in local buffers. It is subsequently solicited by the L2 trigger for the event data associated to RoI's. Those events selected by the L2 trigger are then transferred to the event-building system and subsequently to the event filter for final selection. Events selected by the event filter are moved to permanent storage at the CERN computer centre. In addition to the movement of data, the data acquisition also provides for the configuration, control and monitoring of the hardware and software components which together provide the data-taking functionality.

The DCS permits the coherent and safe operation of the ATLAS detector hardware, and serves as a homogeneous interface to all sub-detectors and to the technical infrastructure of the experiment. It controls, continuously monitors and archives the operational parameters, signals any abnormal behaviour to the operator, and allows automatic or manual corrective actions to be taken. Typical examples are high- and low-voltage systems for detector and electronics, gas and cooling systems, magnetic field, temperatures, and humidity. The DCS also enables bi-directional communication with the data acquisition system in order to synchronise the state of the detector with data-taking. It also handles the communication between the sub-detectors and other systems which are controlled independently, such as the LHC accelerator, the CERN technical services, the ATLAS magnets, and the detector safety system.

1.8 Radiation, shielding, and interface to the LHC machine

The experimental conditions at the LHC will be challenging. The unprecedented radiation levels are expected to result in activated detector components and severe beam-induced backgrounds. To mitigate their impact, the ATLAS shielding layout and the LHC beam-pipe design were carefully optimised. Good communication between the LHC and ATLAS operators will be needed during both injection and data-taking, especially to prevent beam-related accidents which could damage the detector.

1.8.1 Radiation levels

At the LHC, the primary source of radiation at full luminosity comes from collisions at the interaction point. In the inner detector, charged hadron secondaries from inelastic proton-proton interactions dominate the radiation backgrounds at small radii while further out other sources, such as neutrons, become more important. Table 5 shows projected radiation levels in key areas of the detector (see Chapter 3).

Inner detector					
Location	F_{neq} (10^{14} cm^{-2})		Dose (kGy)		Charged-particle flux above 10 MeV (Hz/cm^2)
Pixel layer 0	13.5		790		40×10^6
SCT layer 1	0.8		38		1.5×10^6
SCT disk 9	0.6		23		10^6
TRT outer radius	0.25		3.5		10^5
Calorimeters					
Location	$ \eta $		Maximum dose (kGy)		
EM barrel	1.475		1.2		
EM end-cap	3.2		150		
Tile	1.2		0.15		
HEC	3.2		30		
FCal	4.9		1000		
Muon spectrometer					
Location	Flux				Single-plane rates (Hz/cm^2)
	(kHz/cm ²)		(Hz/cm ²)		
	n	γ	μ	p	
Barrel chambers	2.6–4.0	1.0–1.5	0.3–4.5	0.4–3.2	6.0–11.0
Inner edge of inner wheel	79	25	21	64	347
Inner edge of outer wheel	2.7	1.5	3	0.9	12

Table 5. The 1 MeV neutron equivalent fluence (F_{neq}) and doses in key areas of the detector after 500 fb^{-1} of data (estimated to be approximately seven years of operation). Also given are the charged-particle fluxes in the inner detector and fluxes and single-plane rates in the muon spectrometer.

In ATLAS, most of the energy from primaries is dumped into two regions: the TAS (Target
1600 Absorber Secondaries) collimators protecting LHC quadrupoles and the forward calorimeters. The
beam vacuum system spans the length of the detector and in the forward region is a major source
of radiation backgrounds. Primary particles from the interaction point strike the beam-pipe at
very shallow angles, such that the projected material depth is large. Studies have shown that the
beam-line material contributes more than half of the radiation backgrounds in the muon system.
1605 The deleterious effects of background radiation fall into a number of general categories: increased
background and occupancies, radiation damage and ageing of detector components and electronics,
single event upsets and single event damage, and creation of radionuclides which will impact access
and maintenance scenarios.

1.8.2 Shielding

1610 In order to limit the effects of radiation on the detector, ATLAS relies on the use of almost
3000 tonnes of shielding (see Chapter 3). The shielding procedure is based on a three-layer con-
cept. The inner layer is designed to stop high energy hadrons and secondaries. It is built from
materials such as iron or copper which pack a large number of interaction lengths into a limited

1615 volume. A second layer, consisting of doped polyethylene rich in hydrogen, is used to moderate
the neutron radiation escaping from the first layer; the low energy neutrons are then captured by
a boron dopant. Photon radiation is created in the neutron capture process and these photons are
stopped in the third shielding layer, which consists of steel or lead.

1.8.3 Beam-pipe

1620 The 38 m long beam-pipe section in the ATLAS experimental area consists of seven parts, bolted
together with flanges to form a fully in-situ bakeable ultra-high vacuum system (see Chapter 9).
The central chamber is centred around the interaction point and is integrated and installed with the
pixel detector. It has a 58 mm inner diameter and is constructed from 0.8 mm thick beryllium. The
remaining six chambers, made of stainless steel, are installed symmetrically on both sides of the
interaction point. They are supported by the end-cap LAr cryostats, the end-cap toroids and the
1625 forward shielding, respectively.

1.8.4 LHC machine interface

The LHC machine and ATLAS must continually exchange information to ensure the safe and op-
timal operation of the machine (see Chapter 9). The LHC machine gives ATLAS such information
as beam and bunch intensities, and other characteristics of the beam such as its position. It also
1630 provides the 40.08 MHz bunch clock of the accelerator, needed for the L1 trigger and detector
sub-systems. ATLAS provides information on total luminosity and luminosity per bunch obtained
from its luminosity detectors, and indications of the quality of the collisions based on informa-
tion from the detector and from the beam conditions monitors (see Chapter 3). The ATLAS beam
interlock system provides information on whether or not to safely inject or dump the beams, or
1635 to move from one mode of operation (e.g. filling) to the next (e.g. ramping). This information is
exchanged through the detector control system information server, and dedicated hardware links
for such critical information as beam permission signals and timing.

1.9 Outline of the paper

This paper is structured as follows. Chapter 2 summarises the key features of the solenoid and
1640 toroid magnet system and the B-field determination methodology. The beam-line shielding and
the expected radiation levels are presented in Chapter 3. This is followed, in Chapters 4, 5, and 6,
by the description of the inner detector, calorimetry, and muon spectrometer, respectively. The
forward detectors, not shown in Fig. 1, are discussed in Chapter 7. Chapter 8 outlines the hardware
aspects of the trigger and data acquisition systems. Chapter 9 then presents the main features of the
1645 infrastructure in the ATLAS cavern, including the overall integration of the experiment, the sub-
detector positioning strategy and the corresponding survey results, the services (including cables,
pipes, gas and cooling systems, cryogenics, back-up power infrastructure, etc.), the beam-pipe, and
the access and maintenance scenarios. Chapter 10 presents an overview of the global performance
expected from the ATLAS detector, as obtained from combined test-beam measurements and from
1650 the latest analysis results based on the large-scale simulations done for the commissioning of the
computing system. Finally, Chapter 11 briefly summarises the current status of installation and
commissioning, and the expectations for the ultimate completion of the detector and its operation.

2. Magnet system and magnetic field

This chapter begins with a brief description of the ATLAS magnet system (Section 2.1), which
1655 consists of one solenoid and three toroids (one barrel and two end-caps). Section 2.2 proceeds with
a description of the current understanding of the magnetic field across the whole apparatus. This
includes mapping of the solenoid field and first measurements of the barrel toroid field with the
Hall-probe system, as well as calculations to determine the detailed field map with the required
accuracy and performance specifications to be used in ATLAS simulation and reconstruction ap-
1660 plications.

2.1 Magnet system

ATLAS features a unique hybrid system of four large superconducting magnets. This magnetic
system is 22 m in diameter and 26 m in length, with a stored energy of 1.6 GJ. After approximately
15 years of design, construction in industry, and system integration at CERN, the system is installed
1665 and operational in the underground cavern. This section presents the properties of the magnets and
their related services. More details can be found in [2] for the solenoid and in [3] for the three
toroids.

Figure 1 shows the general layout, the four main layers of detectors and the four superconduct-
ing magnets which provide the magnetic field over a volume of approximately 12,000 m³ (defined
1670 as the region in which the field exceeds 50 mT). The spatial arrangement of the coil windings is
shown in Fig. 5. The ATLAS magnet system, whose main parameters are listed in Table 6, consists
of:

- a solenoid (Section 2.1.1), which is aligned on the beam axis and provides a 2 T axial mag-
netic field for the inner detector, while minimising the radiative thickness in front of the
1675 barrel electromagnetic calorimeter;
- a barrel toroid (Section 2.1.2) and two end-cap toroids (Section 2.1.3), which produce a
toroidal magnetic field of approximately 0.5 T and 1 T for the muon detectors in the central
and end-cap regions, respectively.

The first conceptual design of the magnet system was sketched in the early 1990's, and the
1680 technical design reports [4, 5, 6, 7] were published in 1997. Regular project overviews and status
reports of design and production were made available [8, 9] throughout the design and manufactur-
ing phases. The cold-mass and cryostat integration work began in 2001. The first barrel toroid coil
was lowered in the cavern in fall 2004, immediately followed by the solenoid (embedded inside the
LAr barrel calorimeter). The remaining seven barrel-toroid coils were installed in 2004 and 2005,
1685 and the end-cap toroids in the summer of 2007.

2.1.1 Central solenoid

The central solenoid [2] is displayed in Fig. 6, and its main parameters are listed in Table 6. It
is designed to provide a 2 T axial field (1.998 T at the magnet's centre at the nominal 7.730 kA
operational current). To achieve the desired calorimeter performance, the layout was carefully op-
1690 timised to keep the material thickness in front of the calorimeter as low as possible, resulting in




Figure 5. Geometry of magnet windings and tile calorimeter steel. The eight barrel toroid coils, with the end-cap coils interleaved are visible. The solenoid winding lies inside the calorimeter volume. The tile calorimeter is modelled (Section 2.2.2) by four layers with different magnetic properties, plus an outside return yoke. For the sake of clarity the forward shielding disk (Section 3.2) is not displayed.



Figure 6. Bare central solenoid in the factory after completion of the coil winding.

the solenoid assembly contributing a total of ~ 0.66 radiation lengths [10] at normal incidence. This required, in particular, that the solenoid windings and LAr calorimeter share a common vacuum vessel, thereby eliminating two vacuum walls. An additional heat shield consisting of 2 mm thick aluminium panels is installed between the solenoid and the inner wall of the cryostat. The single-layer coil is wound with a high-strength Al-stabilised NbTi conductor, specially developed to achieve a high field while optimising thickness, inside a 12 mm thick Al 5083 support cylinder. The inner and outer diameters of the solenoid are 2.46 m and 2.56 m and its axial length is 5.8 m. The coil mass is 5.4 tonnes and the stored energy is 40 MJ. The stored-energy-to-mass ratio of only 7.4 kJ/kg at nominal field [2] clearly demonstrates successful compliance with the design requirement of an extremely light-weight structure. The flux is returned by the steel of the ATLAS hadronic calorimeter and its girder structure (see Fig. 5). The solenoid is charged and discharged in about 30 minutes. In the case of a quench, the stored energy is absorbed by the enthalpy of the cold mass which raises the cold mass temperature to a safe value of 120 K maximum. Re-cooling to 4.5 K is achieved within one day.

The electromagnetic forces are counteracted by the combination of the coil and warm-to-cold mechanical support, which maintains the concentricity of the windings. All solenoid services pass through an S-shaped chimney at the top of the cryostat, routing the service lines to the corresponding control dewar (Section 2.1.4.2).

The coil was manufactured and pre-tested in the factory [11], came to CERN for integration in the LAr cryostat, underwent an on-surface acceptance test in its semi-final configuration [12], and was installed in its final central position in ATLAS in October 2005. The one week cool-down and a commissioning test up to nominal field were successfully completed in the summer of 2006 [13]. The solenoid is now ready for detector operation.

Property	Feature	Unit	Solenoid	Barrel toroid	End-cap toroids
Size	Inner diameter	m	2.46	9.4	1.65
	Outer diameter	m	2.56	20.1	10.7
	Axial length	m	5.8	25.3	5.0
Mass	Number of coils		1	8	2 × 8
	Conductor	t	3.8	118	2 × 20.5
	Cold mass	t	5.4	370	2 × 140
Coils	Total assembly	t	5.7	830	2 x 239
	Turns per coil		1154	120	116
	Nominal current	kA	7.73	20.5	20.5
	Magnet stored energy	GJ	0.04	1.08	2 x 0.25
	Peak field in the windings	T	2.6	3.9	4.1
	Field range in the bore	T	0.9–2.0	0.2–2.5	0.2–3.5
Conductor	Overall size	mm ²	30 x 4.25	57 x 12	41 x 12
	Ratio Al:Cu:NbTi		15.6:0.9:1	28:1.3:1	19:1.3:1
	Number of strands (NbTi)		12	38–40	40
	Strand diameter (NbTi)	mm	1.22	1.3	1.3
	Critical current (at 5 T and 4.2 K)	kA	20.4	58	60
	Operating/critical-current ratio at 4.5 K	%	20	30	30
	Residual resistivity ratio (RRR) for Al		> 500	> 800	> 800
	Temperature margin	K	2.7	1.9	1.9
	Number of units × length	m	4 × 2290	8 × 4 × 1730	2 × 8 × 2 × 800
Heat load	Total length (produced)	km	10	56	2 x 13
	At 4.5 K	W	130	990	330
	At 60–80 K	kW	0.5	7.4	1.7
	Liquid helium mass flow	g/s	7	410	280

Table 6. Main parameters of the ATLAS magnet system.

2.1.2 Barrel toroid

1715 The main parameters of the magnet are listed in Table 6. The cylindrical volume surrounding
the calorimeters and both end-cap toroids (see Fig. 1) is filled by the magnetic field of the barrel
toroid, which consists of eight coils encased in individual racetrack-shaped, stainless-steel vacuum
vessels (see Fig. 7). The coil assembly is supported by eight inner and eight outer rings of struts.
The overall size of the barrel toroid system as installed is 25.3 m in length, with inner and outer
1720 diameters of 9.4 m and 20.1 m, respectively.

The conductor and coil-winding technology is essentially the same in the barrel and end-cap
toroids; it is based on winding a pure Al-stabilised Nb/Ti/Cu conductor [14] into pancake-shaped
coils, followed by vacuum impregnation.

1725 The cold-mass integration [15] and the cryostat integration [16] were performed at CERN
over a period of approximately three years, and were completed in summer 2005. In parallel, all
coils successfully underwent on-surface acceptance test procedures [17]. Cool down and testing
of the barrel toroid in the cavern took place in 2006. The cool down of the 360-tonne cold mass
to 4.6 K takes five weeks. The test programme included normal ramps, up to nominal current (in
2 hours) followed by either a slow dump (in 2 hours) or a fast dump (in 2 minutes) in the case
1730 of a provoked quench. The ultimate test sequence that proved the system's health is shown in




Figure 7. Barrel toroid as installed in the underground cavern; note the symmetry of the supporting structure. The temporary scaffolding and green platforms were removed once the installation was complete. The scale is indicated by the person standing in between the two bottom coils. Also visible are the stainless-steel rails carrying the barrel calorimeter with its embedded solenoid, which await translation towards their final position in the centre of the detector.

Fig. 8. The magnet current is raised in steps up to its nominal value of 20.5 kA and then finally up to 21.0 kA, demonstrating the ability of the system to withstand at least an additional 500 A. The current is then allowed to decay back to its design value; the magnet is finally turned off by a deliberate fast dump. After re-cooling the cycle was repeated, demonstrating that no degradation had occurred up to the nominal operating current. During a fast dump, triggered either manually or by the quench detection system, the stored energy of 1.1 GJ is absorbed by the enthalpy of the cold mass following the activation of four quench heaters per coil and in all eight coils, which forces the entire magnet into the normal conducting state within less than two seconds. This leads to a very safe global cold mass temperature of about 58 K and a hot-spot temperature in the windings of about 85 K maximum. The uniform quench heating system also ensures that the internal voltage in the toroid is kept at a low value of about 70 V. After a fast dump the magnet cooling system needs about 50 hours to re-cool the toroid to 4.6 K whereafter normal operation can re-start. The details of the coil testing are published elsewhere, in [18] for the first coil, in [19] for an overall summary, and in [20] and [21] for quench behaviour and quench losses, respectively.

The net Lorentz forces of approximately 1400 tonnes per coil directed inwards and the self-weight of the toroids are counteracted by the warm structure of Al-alloy struts mounted in between the eight coils. However, the barrel toroid structure still deflects significantly under its own weight. After release of the temporary support structure and systematic loading of the toroid with its own weight of 830 tonnes and the additional 400 tonnes of weight of the muon chambers, the final shape of the toroid bore was designed to be cylindrical. The toroid coils were installed in calculated positions on an oval, longer by 30 mm in the vertical direction, to allow for structure deflection during




Figure 8. Time history of the barrel toroid current during an excitation test up to 102% of the nominal value. The current drops back to zero within two minutes of the deliberately-provoked quench.



Figure 9. End-cap toroid cold mass inserted into the cryostat. The eight flat, square coil units and eight keystone wedges (with the circular holes) are visible.

load transfer from the temporary support structure. Since the release and removal of the installation supports, the upper edge of the toroid moved down by about 26 mm, which demonstrates that the design values had been well established and that the installation was precise to within a few millimetres.

The installation of the barrel toroid in the ATLAS cavern commenced in October 2004. It took about 11 months to install the complete toroid, as depicted in Fig. 7. This is discussed in more detail in Section 9.6 within the context of the overall ATLAS installation, for which this toroid installation phase was one of the most demanding ones. The overall structure design and installation experience are reported in [22].

2.1.3 End-cap toroids

The main parameters of the two end-cap toroids are listed in Table 6. These toroids generate the magnetic field required for optimising the bending power in the end-cap regions of the muon spectrometer system. They are supported off and can slide along the central rails, which facilitates the opening of the detector for access and maintenance. Each end-cap toroid consists of a single cold mass built up from eight flat, square coil units and eight keystone wedges, bolted and glued together into a rigid structure to withstand the Lorentz forces (see Fig. 9). Design details are given elsewhere [23], and the production in industry of the coil modules and vacuum vessels is described in [24].

The cold masses were assembled and inserted into their cryostats at CERN. Figure 9 shows the first end-cap toroid interior just prior to the closing of the vacuum vessel. A crucial step in the integration process is the adjustment of the cold mass supports [25]. The weights of cold mass and vacuum vessel are 140 and 80 tonnes respectively. With the exception of windings, coil supports,

and bore tube, the entire structure is made of Al alloy. With a weight of 240 tonnes, the end-cap
1775 toroids were some of the heaviest objects to be lowered into the cavern.

The end-cap-toroid cold masses will each be subject to a Lorentz force of 240 tonnes, pushing
them against the stops mounted on the eight barrel toroid coils. Achieving the correct sharing of the
forces in the axial tie-rods has therefore been a critical design goal. Prior to their installation in the
cavern in summer 2007, both end-cap toroids passed tests at 80 K to check the magnet mechanics
1780 and electrical insulation after thermal shrinkage. Once the end-cap toroids are powered in series
with the barrel toroid, the peak stress in the barrel-toroid windings, in the areas where the magnetic
fields overlap, will increase by about 30%. before they were powered in combination with the
barrel toroid. After a four-week cooldown, both end-cap toroids were successfully tested at half
current, albeit one at a time and in stand-alone mode. The final tests at full field will take place in
1785 the spring of 2008, after the installation of the shielding disks and with the end-cap calorimeters at
their nominal position.

2.1.4 Magnet services

2.1.4.1 Vacuum system The insulating vacuum is achieved with diffusion pumps directly at-
tached to the barrel and the end-cap toroids, two per coil for all toroids, each with a capacity
1790 of 3000 m³/h. In addition, two roughing and three backing pumps are used in the low stray-field
area at the cavern wall. Under normal conditions, with a leak rate less than 10⁻⁴ mbar · l/s, a single
pump would be sufficient. However, for redundancy and in order to minimise detector down-time,
extra pumping units were installed. Since the solenoid is installed inside the cryostat of the LAr bar-
rel calorimeter, the insulation vacuum is controlled by the LAr cryogenic system (Section 9.4.5)
1795 rather than by the magnet control system (see Section 2.1.4.4).

2.1.4.2 Cryogenics The overall cryogenic systems in ATLAS are described in Section 9.4. Here,
details are provided on the system specific to the magnets.

The overall magnet cryogenic system is divided into external, proximity, and internal cryogen-
ics, which are connected via transfer lines. The lines serving the solenoid and barrel toroid remain
1800 fixed, whereas those of the end-cap toroids are partially flexible, as these toroids have to be moved
to access the calorimeters and inner detector for maintenance and repairs (Section 9.7).

The layout of the various cryogenic systems is shown in Fig. 10. The external cryogenics
consist of two refrigerators (the main refrigerator and the shield refrigerator), a distribution transfer
line, and a distribution valve box. The main refrigerator cold box has a refrigeration capacity of
1805 6 kW at 4.5 K equivalent, while the shield refrigerator cold box has a refrigeration capacity of
20 kW at 40–80 K.

The gas buffers are located on the surface with the refrigerator compressors, while the refrig-
erator cold boxes are installed in the USA15 side cavern. The common distribution transfer line
makes the link to the distribution valve box in the main cavern. All proximity cryogenics equip-
1810 ment, including the storage dewar, cold pumps, cryostat phase separator, and distribution valve
box (except for the valve unit of the solenoid) are positioned near the wall of the main cavern, as
schematically shown in Fig. 11 (left).

The distribution valve box channels the fluids to two independent proximity cryogenic sys-
tems, one for the toroids (barrel cryo-ring and two end-caps) and one for the solenoid. For the




Figure 10. Layout of the magnet cryogenics system in the surface hall (compressors) and service cavern (shield refrigerator and helium liquefier). They deliver cold gas and liquid to the distribution valve box in the experimental cavern, from which the solenoid and the toroid proximity cryogenics are fed (see Fig. 11).

1815 toroids, there is a storage dewar with a capacity of 11,000 litres of liquid helium. There also exist a
 distribution valve box, a phase separator dewar with two centrifugal pumps and a storage capacity
 of 600 litres of liquid helium. The solenoid has a control dewar with a storage capacity of 250 litres
 of liquid helium, positioned at the top of the detector.

1820 The proximity cryogenic equipment supplies coolant to the magnet internal cryogenics, which
 consist mainly of cooling pipes attached to the cold mass and the thermal shield. The aluminium
 cooling tubes are either welded to the outer surface of the Al-alloy support cylinder (solenoid) or
 embedded and glued inside and on top of the Al-alloy coil casings enclosing the pancake coils
 (toroids).

1825 The toroids are cooled with a forced flow of boiling helium, which enters the magnets from
 the top. In the case of the barrel toroid (see Fig. 11), helium is supplied from the current lead
 cryostat positioned on the top sector, runs down to the distribution valve box at floor level with
 eight control valves regulating the flow in the eight coils, then goes up and enters the eight coils




Figure 11. Left: Layout of underground service connections to the solenoid and toroid systems. The two large helium dewars can be seen on the side of the main cavern. Also shown are the fixed cryogenic lines supplying the solenoid and the cryo-ring for the barrel toroid coils at the top. The cryogenics lines in the flexible chains supply the two end-cap toroids and follow them whenever they move for detector access and maintenance. Right: schematic of the liquid-helium supply in the barrel toroid. The cryo-ring contains six standard sectors; a bottom sector with a valve box where the input flow per coil is measured and controlled; and the top sector where all lines come together and which is connected to the current lead cryostat.

separately, while the return line returns to the top. A total of 1200 g/s of slightly sub-cooled liquid helium is circulated by means of centrifugal pumps, which take the liquid from the phase separator dewar. The system is equipped with two pumps for redundancy. The second pump is called into operation if the first one fails. The liquid helium in the storage dewar will be used in the event of a failure with the main refrigerator to provide the required cooling capacity to safely ramp down the toroids over a two-hour period.

The solenoid, with a cold mass of approximately five tonnes, is cooled by a direct Joule-Thompson flow from the main refrigerator and is slightly sub-cooled via a heat exchanger in the 250 litre helium control dewar.

The flow in the solenoid and the ten toroid cold masses is controlled individually to cope with variations in flow resistance and to guarantee helium quality in all coils. Given that the end-cap toroids and solenoid each have a single cold mass, there is a single flow control and the branches of cooling pipes (two for the solenoid and sixteen for each end-cap toroid) are arranged in parallel.

2.1.4.3 Electrical circuits The three toroids are connected in series to the 20.5 kA/16 V power supply shown schematically in Fig. 12 (left). They are however individually voltage-protected by the two diode/resistor ramp-down units. The electrical circuit of the central solenoid is similar and shown in Fig. 12 (right). It has a 8 kA/8 V power supply. The power supply, switches, and diode/resistor units are located in the side cavern and approximately 200 m of aluminium bus-bars provide the connections to the magnets in the cavern. Ramping up is accomplished at a rate




Figure 12. Electrical circuit showing the barrel (BT) and end-cap (ECT) toroids connected in series, fed by a 20.5 kA power converter and protected by a voltage-limiting diode/resistor ramp-down unit (left). Electrical circuit of the central solenoid (CS), fed by a 8 kA power converter (right).

of 3 A/s, leading to a maximum ramp-up time of two hours. In the case of a slow dump, the magnets are de-energised across the diode/resistor units in about 2.5 hours. Quench detection is by classical bridge connections across the entire barrel toroid, across the end-cap toroids and across the solenoid, as well as across individual coils, using differential voltage measurements with inductive voltage compensation.

There is a six-fold redundancy in the toroid quench detection grouped in two physically-separated units and cable routings. Quench protection is arranged by firing heaters in all toroid coils so that a uniform distribution of the cold-mass heating is achieved. Given the normal-zone propagation of 10–15 m/s, a toroid coil is switched back to the normal state within 1–2 seconds. As for the quench detection, the quench-protection heater circuits including power supply, cabling, and heaters embody a two-fold redundancy. A similar system is used for the solenoid. An overview of the magnet services can be found in [26].

2.1.4.4 Magnet controls A magnet control system steers and executes automatically the various running modes of the magnet system. Its implementation is realised as part of the overall ATLAS detector control system, as described in Section 8.5. The hardware designs rely on a three-layer model, using distributed input/output connected via field-networks or directly by wiring to a process-control layer, the last layer being the supervisor.

The main control functions are:

- performing automatic operational sequences on a given magnet (sub-system tests);
- providing a communication interface with the power converter;
- regulating the helium flow in the magnet current leads as a function of the magnet current;

- enabling information exchange between the control system and other sub-systems such as vacuum or cryogenics;
- 1870 • monitoring of all critical parameters in the coil (temperatures, strain and displacement gauges);
- performing calculations of non-linear sensor corrections (temperature sensors, vacuum gauges).

The supervision system displays a synopsis of the main process parameters, communicates with the power supply, collects both continuous and transient data, allows visualisation of any collected data on trend charts and archives collected data. For long-term storage and for correlation
 1875 of data between different systems, a central data-logging system will regularly receive a pre-defined number of data items from each magnet system. A subset of the main control parameters is sent to the ATLAS detector safety system and also to the LHC machine (see Section 9.10).

2.2 Magnetic field determination

The specifications on the determination of the magnetic field (Section 2.2.1) are rather different
 1880 in the inner detector (ID) and the muon spectrometer. In the ID cavity, the driving consideration is the absolute accuracy of the momentum scale. In the muon spectrometer, the field is highly non-uniform: residual bending-power uncertainties, if large enough, would translate primarily into degraded muon momentum resolution. Detailed magnetic modelling (Section 2.2.2) and novel instrumentation (Section 2.2.3) have allowed a high-precision mapping of the solenoid field (Sec-
 1885 tion 2.2.4) as well as a preliminary experimental validation of the field measurement and reconstruction strategy in the muon spectrometer (Section 2.2.5). Studies are in progress to combine magnetic models with field measurements into an overall field map for ATLAS data-taking (Section 2.2.6).

2.2.1 Performance specifications and measurement concepts

1890 In the inner detector, the systematic error affecting the momentum measurement of charged tracks is dominated by the relative alignment of detector components and by bending-power uncertainties, the former being the more demanding. A high-precision measurement of the W -boson mass is clearly the most challenging goal for such measurements: a lepton from W decay carries typically a transverse momentum of 40 GeV, resulting in a sagitta of approximately 1 mm as the lepton
 1895 traverses the ID cavity. The systematic alignment uncertainties in the ID are unlikely to improve beyond the $1\ \mu\text{m}$ level or 0.1% of the sagitta. This suggests setting a target of $\sim 5 \times 10^{-4}$ for the fractional bending power uncertainty, so that it remains negligible in the determination of the absolute momentum scale. Such stringent requirements can only be achieved reliably by in-situ mapping, using dedicated instrumentation inside the ID cavity, with all the relevant magnetic ma-
 1900 terials in place and just before the final installation of the ID itself. Eventual long-term drifts of the absolute scale will be detected to a much higher accuracy using permanently installed NMR probes.

In the muon spectrometer, the expected sagitta is approximately 0.5 mm for a muon with a momentum of 1 TeV. The extraction of the momentum from the Monitored Drift Tube (MDT) chamber measurements requires a precise knowledge of the field integral between consecutive
 1905 chambers along the muon trajectory. Because the field gradient can reach 1 mT/mm, local bending-power uncertainties translate into fluctuations of the momentum scale from one region in space to

Criterion	Bending-power accuracy	MDT drift properties
Performance	$\Delta\sigma_{pT}/\sigma_{pT} < 5\%$ overall	Single-wire resolution degraded by $< 5\%$
Field measurement accuracy	$\Delta B_\phi/B_\phi < 2 - 5 \times 10^{-3}$	$\Delta B_{x,y,z} < 4$ mT (relative over chamber)
Reconstructed position of toroid conductors with respect to MDT tower	$\Delta R \sim 1 - 12$ mm, $\Delta(R - \phi) \sim 1 - 6$ mm, $\Delta z \sim 2 - 30$ mm	-
Muon chamber 2nd-coordinate resolution	1.7-5.5 mm	6 to ~ 100 mm

Table 7. Summary of magnetic-field-related performance specifications in the muon spectrometer. The quoted spread reflects the $\eta - \phi$ variations in field gradient and/or strength.

another, adding in quadrature to the overall momentum resolution. In addition, the interpretation, in terms of spatial coordinates, of the drift time measured in the MDT's is sensitive to the local electric and magnetic fields experienced by the ionisation electrons in each tube. The corresponding functional requirements are extensively discussed in [27] and summarised in Table 7.

For a given muon trajectory, three sources of uncertainty affect the measured curvature: field measurement errors; accuracy on the relative position of muon chambers and magnet coils; and trajectory measurement errors, in particular along the direction of MDT wires. For the purpose of setting specifications, it has been required (somewhat arbitrarily) that the combined effect of these sources degrade the momentum resolution by no more than 5% in relative terms; each source should then contribute no more than $\sim 3\%$ of fractional resolution degradation, anywhere in the spectrometer volume.

In-situ mapping of the spectrometer by conventional techniques would have been impractical because of the rapidly-varying field and very large volume. Instead, the muon system is equipped with a total of approximately 1840 B-field sensors; their readings are compared with magnetic simulations and used for reconstructing the field in space. This strategy was shown [28] to meet the field-map specifications above, provided the B-sensor readings, after correcting for perturbations induced by magnetic materials, are accurate to ~ 1 mT (absolute) and the field direction is measured to within ± 3 mrad.

2.2.2 B-field modelling

The total field in the ID cavity, the calorimeters, and the muon spectrometer is computed as the superposition of the Biot-Savart contributions of all magnet windings (see Fig. 5) with those of the magnetised calorimeter and with the localised perturbations induced by other ferromagnetic structures. In order to reach the required accuracy, the calculation combines numerical integration of the contributions of the solenoid, barrel-toroid and end-cap-toroid windings with finite-element modelling of magnetic structures.

The solenoid conductor model is described in Section 2.2.4. The magnetised steel (tile calorimeter and solenoid flux-return girder), which surrounds the ID cavity, is predicted to modify the field by 4.1% at the geometrical centre of the coil. At nominal current, the total measured field is 1.998 T at the interaction point, and drops steeply from ~ 1.8 T at $z = 1.7$ m to ~ 0.9 T at the end of the ID cavity (see Fig. 13).

The toroid windings are, at this stage, described using their nominal geometry. The mesh density of the stored field map is tailored to the local field gradient to ensure an accurate representation of field variations (as also done for the solenoid). Depending on the radius R and azimuth ϕ ,




Figure 13. Longitudinal and radial dependence of the magnetic field in the inner detector cavity. The symbols denote the measured axial and radial field components and the lines are the result of the fit described in Section 2.2.4.




Figure 14. Predicted field integral as a function of $|\eta|$ from the innermost to the outermost MDT layer in one toroid octant, for infinite-momentum muons. The curves correspond to the azimuthal angles $\phi = 0$ (red) and $\phi = \pi/8$ (black).

1940 the field varies from 0.15 T to 2.5 T, with an average value of 0.5 T, in the barrel region, and from 0.2 to 3.5 T in the end-cap region [29]. The analysing performance of the toroid system can be roughly quantified by the field integral experienced by particles originating from the interaction point and propagating in a straight line (the ultimate criterion is the momentum resolution: a zero field integral does not necessarily imply infinite resolution). This available bending power is shown
 1945 in Fig. 14 as a function of $|\eta|$. It shows good magnetic field coverage up to $|\eta| \sim 2.6$. The regions with low field integral, between $|\eta| = 1.4$ and $|\eta| = 1.6$, correspond to trajectories in the plane of an end-cap coil or of a barrel coil, where the fringe field of one magnet largely cancels the bending power of the other.

A number of large magnetisable components, shown schematically in Fig. 15, distort the Biot-Savart field at different levels. Although amenable to experimental spot-checks (Section 2.2.5),
 1950 such perturbations can only be determined using field simulations.

The highly anisotropic structure of the tile calorimeter cannot be satisfactorily modelled using only a scalar permeability and an effective steel-packing factor: a formalism incorporating a magnetic permeability tensor, as well as a more sophisticated treatment of magnetic discontinuities at material boundaries, is called for. The problem is compounded by the superposition of the solenoid and toroid fields in the partially-saturated flux-return girder and in the tile calorimeter itself. A novel approach to magnetic-field modelling in such structures has therefore been developed and implemented in the B-field simulation package ATLM [30]. This package, which incorporates a careful description of the toroid and solenoid conductors as well as a detailed mathematical model
 1955 of the tile calorimeter, is used both to compute the Biot-Savart field by numerical integration (as described above), and to predict, by a finite-element method, the field distortions caused by the tile calorimeter, the flux-return girder and the shielding disk in both the ID cavity and the muon spectrometer. Altogether, these distortions affect the field integral in the muon spectrometer by up to 4%, depending on $|\eta|$ and ϕ ; in addition, they induce, at the level of the inner MDT layers, local field distortions of up to $|\Delta\mathbf{B}| \sim 0.2$ T.
 1960

A few discrete magnetic structures, either inside the muon spectrometer or close to its outer




Figure 15. Sources of magnetic perturbations induced by metallic structures in or near the muon spectrometer.




Figure 16. Schematic representation of the magnetic-sensor layout and coil deformation model, used to reconstruct the magnetic field inside a barrel octant. The MDT nomenclature is defined in Tables 34 and 35 (see Section 6.3).

layers, induce additional, localised magnetic perturbations. Their impact has been evaluated using the 3D finite-element magnetostatics package TOSCA [31]. The largest perturbations are caused by the air pads, jacks and traction cylinders which allow the calorimeters, the shielding disks, and the end-cap toroids to slide along the rails. These affect primarily the field distribution across the BIS and BIL chambers in the three bottom sectors (see Figs. 15 and 121), and in addition impact the field integral at the level of up to 10% over small islands in $\eta - \phi$ space.

The field perturbations caused in the outside MDT layers by the massive steel frame and platforms (HS structure described in Section 9.5), which surround the detector, range from $|\Delta\mathbf{B}| \sim 2$ mT up to ~ 50 mT and rapidly decrease as one moves inwards from the outer to the middle chamber layer. While their impact on B-sensor readings and MDT drift properties does need to be taken into account, they barely affect the bending power, except possibly in a few narrow regions.

The other components in Fig. 15 have much less of an impact because either they lie in a low-field region, they intercept a very small fraction of the end-cap muons, or they are made of stainless steel with a high-field relative permeability very close to 1.

2.2.3 Magnetic field instrumentation and reconstruction

2.2.3.1 B-field sensors

The inner detector is equipped with four NMR probes fixed to the wall

of the inner warm vessel near $z \sim 0$ and equally spaced in azimuth. These probes measure the field strength with an accuracy of around 0.01 mT and will remain in place to monitor the ID field strength throughout the lifetime of ATLAS.

Because NMR probes only measure $|\mathbf{B}|$ and because they cease functioning in a gradient of a few tenths of mT/cm, the solenoid mapper, described in Section 2.2.4, and the muon chambers are equipped instead with 3D Hall cards [32, 33]. These consist of a rigid printed-circuit board carrying a small glass cube, with a Hall probe on each of three orthogonal faces to measure each field component. Every card includes its own readout electronics, as well as a thermistor for local temperature compensation.

All the Hall cards were calibrated in a highly uniform field monitored by a NMR probe. The achieved absolute Hall-card accuracy on $|\mathbf{B}|$ is 0.2 mT up to $|\mathbf{B}| = 1.4$ T and 1 mT up to 2.5 T; and the angular accuracy achieved on the measured field direction is 2 mrad.

2.2.3.2 B-field reconstruction In an air-core magnet, the magnetic field can in principle be calculated by direct application of the Biot-Savart law, once the geometry of all conductors is known and assuming material-induced magnetic perturbations are negligible. In practice however, the conductor position and shape are known only approximately, owing to fabrication tolerances and to deformations of the magnet structure under gravitational and magnetic loads. The exact location of each magnet coil, as well as the relative positions of the end-cap and barrel toroids, will be reproducible, after a power cycle or an access period, to a finite precision only. Therefore, the field must be measured under running conditions, with all detector components in place and under the mutual influence of the different magnets and magnetic structures.

The muon spectrometer is equipped with an array of approximately 1730 Hall cards, which remain mounted permanently and precisely on the MDT chambers and continuously measure all three field components (an additional 64 cards are mounted on the inner and outermost faces of the end-cap toroid cryostats to complement the MDT sensor system in the forward region). Two NMR probes, installed at low-gradient locations in the barrel toroid, complement the system, with the aim of detecting eventual long-term drifts in the response of the Hall cards. The 3-D sensor readings are compared with field calculations which include both the contributions of the magnet windings and those of nearby magnetised structures, and are used for reconstructing the position and the shape of the toroid conductors with respect to the muon chambers (see Fig. 16). Once the geometry of the coils is known, the field can be calculated anywhere in the muon spectrometer. Simulation studies using a simplified coil deformation model have shown that the magnetic field can be reconstructed to a relative accuracy of 0.2% [28].

2.2.4 Solenoid-mapping measurements

2.2.4.1 Mapping campaign The field was mapped [34] in August 2006 by a machine, which scanned a Hall-card array over a volume slightly larger than that now occupied by the inner detector. During this mapping campaign, the barrel and end-cap calorimeters were all in their final positions. Although the shielding disks were not yet installed, their differential contribution is small enough (< 0.2 mT in the ID tracking volume) that it can be reliably accounted for later. The same is true of corrections for the absence of toroid excitation during mapping.

Mapping data were recorded with solenoid currents of 7730, 7850, 7000, and 5000 A, with a final set of data back at the nominal operating current of 7730 A. Each data set contains at least 20,000 points, and is sufficient by itself to fit the field with negligible statistical uncertainty. Each map took about four hours, during which the solenoid current remained stable to within 0.1 A, as confirmed by the NMR probes.

2.2.4.2 Mapper geometry, survey and auto-calibration The mapping machine had four arms mounted on a common axle in a windmill configuration, with twelve Hall cards on each arm, at radii ranging from 0.118 to 1.058 m, which directly measured the field components B_z , B_R and B_ϕ . The machine could be rotated around its axle and translated in z along the ID rails by means of pneumatic motors. Optical encoders allowed control of the mapper movements and readout of its stop positions with an accuracy of 0.1 mm. A number of surveys were necessary to determine the positions of each individual Hall sensor for all possible longitudinal mapper positions and azimuthal settings of the windmill arms. After combining all the information, the estimated overall accuracy on the position of a map point in the cryostat coordinate system is approximately 0.3 mm.

The redundancy and internal consistency of the mapping measurements makes it possible to extract individual probe misalignments from the data themselves to an accuracy of ± 0.1 mrad. The strong constraints from Maxwell's equations on physically realisable fields in the absence of any current sources or magnetic materials, combined with the fact that the field at the origin can be almost completely determined from the measurements of a single Hall probe, allow all three probe alignment angles to be determined and the B_z component to be normalised to a common scale for all probes.

The NMR probes, which were operational throughout the field-mapping campaigns, are used to set the overall scale of the Hall sensors with an accuracy of about 0.4 mT, the limitation coming from the extrapolation uncertainty from the mapper arms out to the position of the NMR probes. The NMR data also show that there is negligible hysteresis in the solenoid system: the field at 7730 A remained constant within ± 0.01 mT from the first excitation cycle onwards, provided that this current was approached from below. A small saturation effect is visible in the NMR data, with the field at 5000 A being 0.34 mT higher than would be expected by simply scaling down from 7730 A.

2.2.4.3 Map fitting Using the measured magnet current and a detailed model of the solenoid geometry, the Biot-Savart law is integrated to produce a field model which should account for most of the measured field. The conductor model is based on engineering drawings, with as many parameters as possible taken from surveys of the as-built solenoid. The coil cross-section is assumed to be perfectly circular. The winding was mechanically assembled from four separate sections, each with a slightly different average pitch, and joined together by welds which are represented electrically by turns having just under twice the average pitch. Also modelled are the welds at the coil ends and the return conductor which runs axially along the outside of the support cylinder. The expected distortion of the solenoid, relative to the room-temperature survey and caused by thermal shrinkage and magnetic pressure, is also taken into account.

The geometrical fit to the mapping data has 11 free parameters. Two overall scale factors allow fine tuning of the conductor model: one common to all longitudinal dimensions, and an indepen-

Fitted parameters	Fit results
Scale factors in conductor model	(R scale, z scale) = 0.9993, 1.0012
Fitted offsets from solenoid centre to centre of cryostat	($\Delta x, \Delta y, \Delta z$) = 0.26, -2.42, 0.51 (mm)
Fitted rotations of solenoid around cryostat x and y axes	(θ_x, θ_y) = -0.08, 0.19 (mrad)
Resulting fit residuals	$\sigma(\Delta B_z, \Delta B_R, \Delta B_\phi)$ = 0.44, 0.35, 0.30 (mT)

Table 8. Typical fit results of solenoid-mapping measurement at 7730 A.

dent one for the radial dimension. Five more free parameters quantify the three offsets and two
2065 rotations of the conductor relative to the mapper coordinate system. The calorimeter-steel contri-
bution is modelled by a Fourier-Bessel series with four terms. These parameters are determined by
minimising a χ^2 function which includes the longitudinal and radial field components at all mapped
points. The RMS residuals of the geometrical fit alone are just over 0.5 mT. This field model is fur-
2070 ther improved by parametrising the difference between the data and the geometrical model with a
general series which can represent any field obeying Maxwell's equations. This brings the residuals
down to about 0.4 mT, as shown in Table 8.

Systematic uncertainties are estimated by fitting to several representative data sets under vary-
ing assumptions, with and without implementing various corrections (such as Hall-card alignment,
 z -dependent carriage tilt, residual perturbations induced by slightly magnetic mapper components,
2075 number of Fourier-Bessel terms etc.). The geometrical scale factors emerge as very close to unity
(Table 8), suggesting that the coil survey data are well understood. The fitted offsets and rotations
with respect to the centre of the reference coordinate system (barrel LAr cryostat) are stable at
the 0.2 mm and 0.1 mrad level respectively, confirming the vertical -2 mm offset of the solenoid
axis indicated by the survey results before and after installation in the main cavern (see Table 53
2080 in Section 9.3.2.3).

The overall fit is excellent, as illustrated in Fig. 13 and confirmed by the resulting RMS resid-
uals of ~ 0.4 mT for all three field components (Table 8). The on-axis fractional steel contribution,
as estimated from the Fourier-Bessel series, is consistent with the magnetic-field simulation to
better than 2 mT, although the latter does not perfectly reproduce the measured z -dependence of
2085 this perturbation. The fit quality is best measured in terms of the fractional sagitta residual, $\delta S/S$,
evaluated along an infinite-momentum trajectory from the interaction point to the point where the
track crosses the outer radial or longitudinal boundary of the inner detector. The total uncertainty,
estimated by combining the overall scale error, the fit residuals and the systematic uncertainties, is
shown as a function of $|\eta|$ in Fig. 17.

2090 **2.2.5 Experimental validation of the field map in the muon spectrometer**

The tests carried out in fall 2006 for the barrel toroid provided the first full-scale test of the B-
sensor system, and an initial validation of the magnetic models and field-reconstruction strategy in
the muon spectrometer. The end-cap toroids were not yet installed at the time and the solenoid was
turned off. Since the muon-chamber installation was still in progress, only 400 MDT Hall cards
2095 were available for readout, thus providing sensitivity for field reconstruction in about one third of
the barrel region.




Figure 17. Fractional sagitta error due to uncertainties in the solenoid field vs. $|\eta|$.




Figure 18. Field reconstruction residual ΔB_ϕ for one middle (green, solid), outer (blue, dashed) and inner (red, dot-dashed) MDT layer.

The sensor signals were extremely clean (~ 0.01 mT of RMS noise at full field), and reproducible to ~ 0.05 mT between magnet cycles separated by up to one week. Non-linear effects remain very small (< 4 mT in the BIS layer, close to the calorimeter steel, over the full current range). The absolute field scale, as determined by an NMR probe located in the azimuthal mid-plane of coil 3, at a point where steel-induced perturbations are negligible and the field gradient below 0.2 mT/cm, agrees with the Biot-Savart prediction to better than 0.2% .

The field reconstruction algorithm outlined in Section 2.2.3 and detailed in [28] has been applied to B-sensor data collected at nominal field in the barrel toroid. Because the muon alignment system was still being commissioned and the MDT survey not yet completed, it is necessary, at this stage, to assume that all muon chambers and B-sensors are in their nominal position. For the three coils bracketed by the available sensors, the reconstructed conductor shape is qualitatively consistent with that measured at room temperature before insertion of the windings into their respective cryostats. Figure 18 displays the difference, at each active sensor in sector 2 (see Fig. 121) of the muon spectrometer, between the azimuthal component of the measured field (corrected for perturbations from magnetic materials) and that of the Biot-Savart contribution predicted by the field-reconstruction fit. A perfect description of the conductor geometry and of magnetic perturbations should yield $\Delta B_\phi = 0$. The agreement is best in the middle chambers (BM), where the gradients are smallest: the distribution is well centred and exhibits a spread $\Delta B_\phi^{RMS} \sim 1.2$ mT. In the outer chamber layer (BOS), the distribution of ΔB_ϕ shows a moderate bias of 2.2 mT and a spread of 2.6 mT. In view of the larger field gradient in these chambers, such a spread is consistent with the current ± 5 mm uncertainty on the as-installed MDT chamber positions. The situation is similar but somewhat worse in the inner chambers (BIS). These preliminary results reflect the cumulative effect of errors in the assumed sensor and chamber geometry, of residual imperfections in the magnetic model of the calorimeter steel, and of the performance of the reconstruction fit.

Validation of the TOSCA simulations, which describe the distortions induced by other support and service structures was carried out using 40 dedicated Hall cards temporarily installed at critical locations in the bottom muon sector and between the outer muon chambers and the HS structure (see Fig. 15). The agreement between measured and predicted perturbations typically ranges from 2 to 5 mT at the location of the Hall cards and should be better within the spectrometer volume. It

is satisfactory at most locations, although discrepancies as large as 50 mT are observed very close to a few localised and well-identified steel supports. A more extensive magnetic characterisation campaign is planned during the next full magnet-system test.

2.2.6 Towards an overall field map for ATLAS data-taking

2130 The default field map in the ID tracking volume will mirror the very accurate fit obtained for the solenoid mapping data and illustrated in Fig. 13. This approach automatically takes into account the magnetised steel surrounding the ID cavity without having to rely on any field calculations. The fit function is required to satisfy Maxwell's equations and will include empirical corrections to match the measured map as closely as possible, as well as small (< 0.2 mT) additional corrections
2135 for the shielding disks (which were absent at the time of mapping) and barrel-toroid contributions.

In the calorimeters, the map will be based on the ATLM simulation, with the magnetic parameters describing the calorimeter steel adjusted to fit the solenoid-only and toroid-only field measurements performed in 2006. This simulated map will be smoothly connected to the fitted solenoid map in the future: the potential discontinuity remains to be characterised, but is estimated
2140 not to exceed 2 mT over a very narrow interface region.

In the muon spectrometer, the map will reflect the superposition of the winding contributions with the predicted distortions associated with the calorimeter steel and other significant magnetic structures inside or near the spectrometer volume. So far, the Biot-Savart calculation presented above has been performed only in a 1/16th slice, which spans 45° in azimuth and is longitudinally
2145 symmetric with respect to the interaction point: this is the minimum angular size required to handle correctly the symmetries of the full toroid system. Extending it to the case of an arbitrary geometry (without any symmetry assumptions) is currently in progress and the final implementation will depend on the extent to which the actual coil geometry, as eventually revealed by the field-reconstruction procedure, deviates from the ideal configuration. Similarly, studies are in progress
2150 to assess the magnetic impact of shape or position imperfections in the tile-calorimeter geometry: their outcome will indicate to which extent such deviations from the ideal configuration must be taken into account when describing the field inside the calorimeter and/or muon spectrometer.

3. Background radiation and shielding

3.1 Introduction

2155 In contrast to previous and existing colliders, the dominant primary source of background radiation
at the LHC, when operating at design luminosity, arises from collisions at the interaction point.
The rates expected from beam-halo particles and beam-gas interactions are negligible in compar-
ison. In the inner detector, charged hadrons from inelastic proton-proton interactions dominate
the radiation backgrounds at small radii, while the effects of other backgrounds, such as neutrons,
2160 become more important further out (see [35] for detailed studies of the various radiation sources,
radiation levels, neutron fluences and activation levels expected in ATLAS throughout the lifetime
of the experiment).

In ATLAS, most of the energy from the primary particles is dumped into two regions: the
Target Absorber Secondaries (TAS) collimators, and the Forward Calorimeters (FCal) depicted
2165 in Fig. 19, which are therefore among the strongest sources of secondary radiation. These two
sources are somewhat self-shielding, and since they are compact, they have been further shielded
with layers of dense material and cladding. The beam-vacuum system, on the other hand, spans
the whole length of the detector. In the forward regions, it is another major source of radiation
background, since the primary particles produced at large values of absolute pseudorapidity $|\eta|$
2170 strike the beam-pipe at very shallow angles. Through this mechanism, the beam-pipe becomes an
extended line source illuminating the interior of the forward cavity. Detailed studies have shown
that the beam-line material is responsible for more than half of the fluences expected in the muon
system [35].

A thorough understanding of the impact of background radiation has been a critical element
2175 in the design phase of most of the components of the detector and a number of deleterious effects
have been considered:

1. Increased detector occupancy may be an issue. In tracking detectors, this can lead to ineffi-
ciencies, degraded resolutions, and increased rates of fake tracks. In calorimeters, the pile-up
fluctuations at high luminosity degrade the energy resolution.
- 2180 2. Hits generated by slow neutrons dominate the occupancy of the muon spectrometer system.
This effect has not been of any concern at previous colliders.
3. Spurious trigger rates will increase if the background radiation consists of penetrating tracks.
Also, increased occupancies can increase the rates of random triggers.
4. Radiation may damage silicon detectors and readout electronics.
- 2185 5. Interactions leading to anomalous deposits of local radiation can change the logical status
of electronic signals (single-event upset) or permanently destroy components (single-event
damage).
- 2190 6. Wire detectors can experience “ageing” (reduced gain and therefore efficiency) due to poly-
merised deposits on the wires caused by radiation interacting with certain components of the
detector gas.

7. The large fluences expected at the LHC design luminosity may lead to a significant radiation hazard from the prompt component of the radiation, when the accelerator is operating.
8. Nuclear interactions in dense materials will lead to the creation of residual radio-nuclides. The resulting dose rates from radio-activation of certain materials will lead to radiological hazards, which impact access and maintenance scenarios.

2195

The largest impact from background radiation is of course to be expected close to the beam-pipe, in particular in the region of the inner detector and the forward calorimeters. Given the lack of available space and the large contribution from primaries, only a limited amount of moderator shielding could be installed to minimise the impact of background radiation, as described in Section 3.2.

2200

Very large reductions in the expected background rates in the muon spectrometer have been achieved by designing a large amount of shielding around the TAS. A total shielding weight of 2825 tonnes (1887 tonnes of metal, 920 tonnes of concrete, and 18 tonnes of plastic) has thus been added to the detector. Since different types of radiation are best stopped with different types of shielding materials, a multi-layered shielding approach has been used. The inner layer's purpose is to stop high-energy hadrons and their secondaries. This layer is made of materials such as iron or copper, which provide a large number of interaction lengths. In the case of iron, studies have shown that a minimum carbon content of a few percent is advantageous since it efficiently moderates the neutron energies down to lower values. A second layer, consisting of boron-doped polyethylene, is used to moderate the neutron radiation escaping from the first layer and the low-energy neutrons are then captured by the dopant. Photon radiation is created in the neutron-capture process and these photons are stopped in the third shielding layer which consists of steel or lead. Lead is more effective in stopping photons, but induces more neutron radiation than steel. Figure 19 shows the locations of the different shielding components in ATLAS.

2205

2210

2215

3.2 Description of the shielding

The moderator shielding (Fig. 20a) on the front face of each of the end-cap and forward LAr calorimeters reduces the neutron fluences in the volume of the inner detector by protecting the inner detector from back-splash of neutrons from the calorimeter. It is made of polyethylene, doped with 5% boron in the form of B_4C . Reactor tests have demonstrated that this choice for the dopant results in a plastic which is more radiation-hard than if other boron dopants had been used. This is important since the shielding in front of the forward calorimeters is exposed to a very large ionising dose over the lifetime of the ATLAS experiment.

2220

There are three brass shielding elements inside each of the end-cap calorimeter cryostats, located directly behind the calorimeters (Fig. 20b). The largest one is attached to the rear end-plate of the cryostats and has a diameter of 387 cm. Closer to the beam-line are two other shielding plugs. One of these is a cylindrically-shaped extension of the forward calorimeters. The main purpose of these shielding elements is to protect the end-cap inner muon stations from the background radiation.

2225

The next protection element is the shielding disk (Fig. 20c), which serves in fact a threefold purpose: it supports the muon chambers in the first end-cap muon station, it shields these chambers from background radiation emerging from the calorimeters, and it provides a well-defined path for

2230




Figure 19. Schematic view of major ATLAS detector systems and of the main shielding components (see text).

the magnetic field flux return from the solenoid magnet. The bulk of this shielding disk consists of a vertical steel disk with a diameter of 872 cm. This disk supports end-cap muon trigger chambers (see Section 6.7). At the centre of the disk and surrounding the beam-pipe is a stainless steel tube containing a set of cylindrical shielding pieces made of leaded red brass (85% Cu, 5% Pb, 5% Sn, 5% Zn). This tube also supports Cathode Strip Chambers (CSC) and Monitored Drift Tubes (MDT). Brass shielding has been added to the disk in order to protect the CSC chambers. There is a polyethylene layer on the outside of this brass shielding, which is doped with B_2O_3 , to moderate the neutrons, while photons created in the neutron absorption process are stopped in a third layer made of lead.

The next protection element is the end-cap toroid shielding (Fig. 20d), which consists of two parts, one located outside the toroid and enclosing the beam-pipe and one inside the cryostat:

- the first one is a cylindrical structure made of ductile cast iron, which surrounds the beam-pipe on the inside of the two end-cap toroid cryostats. The front piece has a large hole in the centre, into which the stainless steel tube of the shielding disk fits. On the outside of the cast iron is a polyethylene layer doped with B_2O_3 (5%). The photons created in the polyethylene layer are stopped by the stainless-steel bore tube, which supports the shielding in the end-cap toroid;
- the second part of the toroid shielding consists of various polyethylene structures, which are located in the vacuum of the end-cap toroid cryostats. The polyethylene is doped with B_4C , which causes fewer out-gassing problems than other dopants. Photons created when the neutrons are absorbed by the boron are stopped by the aluminium of the cryostat itself.




Figure 20. Details of the shielding components as described in the text: a) moderator, b) LAr calorimeter plugs, c) disk, d) toroid, e) forward, and f) nose shielding.

2255 The purpose of the two forward shielding assemblies (Fig. 20e) is to protect the middle and outer end-cap muon stations from background particles created in secondary interactions in the beam-pipe, the calorimeters and the TAS collimators. These shielding elements, which are removable and will be stored in the surface building during maintenance of ATLAS, consist of two parts: a cylindrical core and a set of octagonal pieces in the rear. All pieces are made of cast ductile iron, surrounded by a layer of polyethylene doped with boron in the form of H_3BO_3 and followed by a 3 cm thick steel layer. The core pieces are enclosed in a 5 cm thick polyethylene layer,

2260 while an 8 cm thick layer surrounds the octagonal pieces. These polyethylene layers are made of 10,000 bricks of three different shapes.

The final shielding element, or nose shielding as depicted in Fig. 20f, supports the TAS collimator and protects ATLAS from the radiation created in this collimator, which is designed to prevent the first LHC quadrupole from quenching due to the energy deposited by the particles emerging from the interactions in ATLAS. The nose shielding is permanently installed in ATLAS and, unlike the forward shielding assemblies, cannot be removed during shutdowns. The main component of this shielding is the cylindrical 117 tonne heavy “monobloc”, which has an outer diameter of 295 cm. It is made of cast iron and supported by a tube, which is anchored in a 460 tonne concrete structure. The 200 tonne heavy “washers”, which are located around the support tube, increase the radial thickness of the iron shielding by 112 cm in a region where the monobloc is thin.

3.3 Calculation of particle fluences and absorbed doses

A vast and systematic effort has been made in the design phase to optimise the shielding in ATLAS by using different simulation programs [35] for simulating hundreds of different geometrical options. These studies have required significant computing resources, since the secondary particles in the hadronic showers had to be followed down to very low energies. Different event generators and transport codes have been used in an attempt to assess the systematic uncertainties in the calculations. When optimising the shielding configuration and materials in the limited space available in ATLAS, it was very often necessary to make trade-offs between different background types, e.g. neutrons versus photons. It has therefore been quite important to also understand the detector response to different types of background radiation, typically particles in the MeV range, in order to converge to the optimal solution [36, 37].

The expected particle fluences (integrated over energy) agree to typically better than 20%, as was shown by comparing two of the most commonly used minimum-bias event generators, PHOJET1.12 [38, 39, 40] and PYTHIA6.2 [41]. Larger differences of up to 50% were observed for pions, kaons, and muons with energies above several GeV. However, these particles provide only a small contribution to the total fluence. The program most used for the shielding optimisation in ATLAS has been the GCALOR package [42], which contains the CALOR code [43] with an interface to GEANT3 [44]. FLUKA2001 [45] is another transport code, which is widely used for studies of hadronic and electromagnetic cascades induced by high-energy particles, and which has been extensively used in simulations of background radiation in ATLAS. In order to investigate transport-code differences, GCALOR was compared not only to FLUKA but also to MARS14(2002) [46]. Comparisons for simplified geometries as well as for the most detailed descriptions of the detector have been carried out.

The results of these studies are extensively reported in [35]: the overall conclusion is that the predictions of FLUKA, MARS and GCALOR are in good agreement for energy-integrated neutron, charged hadron, photon and e^+e^- fluences. For most regions in the inner detector, the difference between the FLUKA and GCALOR values is below 40%. In the pixel vertexing layer differences as large as 80% are however observed for charged hadrons. An excellent agreement, typically to within 20%, between the respective photon and neutron fluences in the muon spectrometer is observed when comparing the FLUKA and the GCALOR results. The charged hadron and lepton




Figure 21. The total ionising dose per year calculated by GCALOR (see text) in one quarter of the central part of the detector. The locations of the inner detector sub-systems, of the different calorimeters and of the inner end-cap muon stations are indicated. The scale on the left gives the integrated dose per year corresponding to the various iso-lines.

fluences in the muon spectrometer show much larger discrepancies, but the differences are always within a factor of 2.5. An overall safety factor of five has been used in the design of the ATLAS muon spectrometer.

2305 The absorbed dose is the mean energy deposited per unit mass, taking into account all energy-loss mechanisms (but corrected for rest-mass effects). The dominant energy-loss mechanism is usually ionisation, but non-ionising energy loss is also important for understanding detector and electronic damage effects. The ionising dose is defined in the following as the integrated dE/dx energy loss in the detector material from charged particles, excluding ionisation energy loss from nuclear recoils. It is given in units of Gy/y, where one year corresponds to 8×10^{15} inelastic proton-proton collisions (assuming an inelastic cross-section of 80 mb, a luminosity of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and a data-taking period of 10^7 s). Comparisons of the calculated ionising dose in the inner detector between FLUKA and GCALOR show differences of up to a factor of two.

3.3.1 The inner detector and calorimeter regions

2315 Figure 21 shows a GCALOR calculation of the ionising dose in the region closest to the interaction

Region	R (cm)	Particle rates (kHz/cm ²)						F _{neq} (×10 ⁺¹² cm ⁻²)	Ionisation dose (Gy/y)
		γ	Protons	Neutrons	π [±]	μ [±]	e ⁻		
		> 30 keV	> 10 MeV	> 100 keV	> 10 MeV	> 10 MeV	> 0.5 MeV		
Pixel layer 0	5.05	45800	2030	4140	34100	300	8140	270	158000
Pixel layer 2	12.25	9150	280	1240	4120	190	1730	46	25400
SCT barrel layer 1	29.9	4400	80	690	990	130	690	16	7590
SCT barrel layer 4	51.4	3910	36	490	370	67	320	9	2960
SCT end-cap disk 9	43.9	7580	73	840	550	110	470	14	4510
TRT outer radius	108.0	2430	10	380	61	7	53	5	680

Table 9. Particle rates, fluences and doses in key locations of the inner detector sub-systems (see Fig. 29 for the definitions and positions of the inner detector layers). Here, F_{neq} is the 1 MeV neutron equivalent fluence (see text). The FLUKA program has been used for this calculation and the statistical uncertainties are typically less than 10%. One year corresponds to 8×10^{15} inelastic proton-proton collisions (assuming an inelastic cross-section of 80 mb, a luminosity of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and a data-taking period of 10^7 s).

point. The forward calorimeters will be exposed to up to 160 kGy/y, whereas the corresponding number for the end-cap electromagnetic calorimeters is 30 kGy/y. This will lead to very large integrated doses over the full lifetime of the experiment and is one of the main reasons why only the LAr technology with its intrinsically high resistance to radiation is used in the end-cap and forward regions. The main concern in the design phase has been for the electrode materials, primarily polymers such as polyimide, which had to be chosen with care and thoroughly tested for radiation hardness [47, 48].

The tile calorimeter, with its scintillator samplings read out by wavelength-shifting fibres, is protected by the LAr electromagnetic calorimeter and is exposed to less than 30 Gy/y, i.e. 5,000 times less than the forward calorimeters. The scintillators and fibres were nevertheless also thoroughly studied under irradiation [49, 50, 51, 52] in order to determine their degradation during the lifetime of ATLAS.

In the inner detector, a very large effort had to be devoted over many years to the understanding of the impact of irradiation on silicon sensors, on front-end electronics circuits and on ageing phenomena in the ionising gas used for the straw tubes.

Two main mechanisms lead to the degradation of the performance of silicon devices under irradiation. First, there is the effect of damage to the devices due to ionising energy loss. This can lead to the creation of trapped charges, in particular in the oxide layer of the sensor, which alters its electric properties. The second effect is bulk damage, or displacement damage, which is caused by the displacement of silicon atoms in the lattice. In the study of bulk damage to silicon devices, it is useful to introduce a quantity called the 1 MeV neutron equivalent fluence (F_{neq}). This fluence is obtained by convoluting the various particle energy spectra and fluences with silicon displacement-damage functions, normalised using the non-ionising energy loss (NIEL) cross-sections to the expected damage of 1 MeV neutrons [53].

Table 9 lists the particle rates, F_{neq} values and ionising doses predicted by FLUKA in the inner detector regions shown in Fig. 21. In the pixel detector, the particle rates are dominated by charged




Figure 22. Particle fluxes in the various muon spectrometer stations at high luminosity ($10^{34} \text{ cm}^{-2} \text{ s}^{-1}$) as predicted by GCalOR. The neutron and photon fluxes are in units of kHz/cm^2 and the muon and proton fluxes in Hz/cm^2 .

pions and photons. The latter are produced mostly in neutron capture processes but also directly from the primary collisions and from interactions in the beam-pipe and its related equipment. The predicted ionising dose in the innermost layer of the barrel pixel detector is 160 kGy/y , which is the same as for the forward calorimeter, while the F_{neq} is expected to be $3 \times 10^{14} \text{ cm}^{-2}/\text{y}$. In the SCT detector, the charged hadron and neutron rates are comparable and the F_{neq} and the ionising dose are reduced by about a factor of 20 with respect to the first pixel layer.

While most of the charged hadrons originate from the interaction point, most of the neutrons in the inner detector are the result of albedo (back-splash from the calorimeters). The purpose of the moderator shielding described in Section 3.2 is to moderate the neutrons from the end-cap and forward calorimeters to lower their energies to values for which their contribution to the total F_{neq} is minimised. The polyethylene in the moderator shielding is doped with boron, which has a large cross-section for the capture of thermal neutrons. Nevertheless, the inner detector cavity will be filled during LHC operation by an almost uniform “gas” of thermal neutrons with a flux of 1–2 MHz/cm^2 and the sensitive detectors will be exposed to fluxes of 2–10 MHz/cm^2 of low-energy photons originating from the interactions themselves and from neutron capture. The dominant long-term impact of these particle fluences is not only radiation damage but also activation of the detector components (see Section 3.5).




Figure 23. Average expected single-plane counting rates in Hz/cm² at 10³⁴ cm⁻² s⁻¹ and for various regions in the muon spectrometer.

3.3.2 The muon spectrometer region

2360 The effects of the absorbed ionising dose in the most critical muon spectrometer regions have been studied [54]. The CSC's in the inner end-cap stations will be exposed to the highest dose. Figure 21 shows that in this region the ionising dose will vary between 3 and 20 Gy/y. The chambers closest to the beam-line in the middle end-cap stations are expected to see at most 10 Gy/y. Most of the muon spectrometer will, however, be exposed to less than 1 Gy/y.

2365 Although care had also to be applied to the choices of materials, to the design of the front-end electronics circuits and to the choice of the ionising gases for all the muon chamber technologies, radiation damage due to the ionising dose is not the only concern in the muon spectrometer region. Signals in the detectors from background particles are the main issue, because these background signals may significantly reduce the muon track-finding efficiency and, more importantly, introduce large rates of fake triggers. The expected particle fluxes based on a simulation of the background radiation in the ATLAS muon spectrometer using the GCALOR program are shown in Fig. 22. The energy thresholds used in these simulations are 10 keV, 10⁻¹¹ MeV (well below the thermal neutron energy range), 1 MeV, and 1 MeV for photons, neutrons, muons, and protons, respectively. The highest fluxes are expected in the innermost end-cap muon stations, in particular in the CSC's, which will have to cope with a large background counting rate, although the estimation of this rate is subject to systematic uncertainties between approximately 15% and 25%.

Background hit rates caused by neutrons and photons in the relevant energy range have been




Figure 24. Left: Top view of a BCM module, showing the diamond sensors (left side of picture), the HV supply and signal-transmission lines, the two amplification stages and the signal connector (right side of picture). Right: Close-up view of one BCM station installed at 184 cm from the centre of the pixel detector, which can be seen at the far end of the picture. Each one of the four modules can be seen in position at a radius of 5.5 cm, very close to the beam-pipe.

calculated with GEANT3 for all muon-chamber technologies using detailed geometrical descriptions of the muon chamber setup. Energy-dependent efficiency curves have been estimated for neutrons, photons, and electrons [35]. Calculations have been done for the various chamber types taking into account the angular distributions of the particles at the chamber locations. Fluxes tend to be isotropic in the barrel, while in the end-cap a substantial fraction of the particles originates from the interaction region and from the beam-pipe in the region of the end-cap toroids, which is the main local source of secondary radiation. The fake L1 trigger rate in the presence of these background hits was studied in simulation including large contingency factors to account for the various uncertainties in the predictions.

The single-plane efficiency curves have been compared between existing experimental data and simulation and found to be in good agreement. Average single-plane chamber efficiencies have been obtained by folding efficiency curves with the energy spectra predicted at each chamber location. Uncertainties due to the shape of the energy spectrum, the angular distribution, and the surrounding material have been studied and amount to a factor of 1.5. Predicted counting rates in the barrel stations are of the order of 10–12 Hz/cm² for both the MDT's and the resistive plate chambers. These rates are dominated by the photon contribution (80%), followed by neutrons and protons (10% each). In the inner barrel stations, the contribution from muons rises to about 15% and that from punch-through pions to a few percent. In the end-cap regions, photons contribute less to the counting rate. In the CSC's for example, photons account for about half of the rate, while muons account for 30% and protons for 10%. The predicted single-plane counting rates in the muon spectrometer are summarised in Fig. 23.

3.4 Background monitors

Measurements of particle fluences in ATLAS will provide a precise bench-marking of the particle transport codes used in the calculations and will also directly monitor the absorbed doses in the various detectors. Possible beam losses near the detector have to be monitored with specific de-

tectors designed to provide fast feedback to the accelerator operations team. The motivation for equipping ATLAS with a reliable set of background monitors in various regions of the detector is therefore obvious.

2405

3.4.1 Monitors in the inner detector

The inner detector region of ATLAS contains a set of small detectors, which are sensitive to dose, to the 1 MeV neutron equivalent fluence (F_{neq}) and to thermal neutrons. These detectors consist of:

2410

1. Field-effect transistors (RADFET's), which measure the total ionising dose;
2. PIN-diodes, which measure F_{neq} ;
3. Radiation-hardened transistors, which measure thermal neutron fluences.

2415

These detectors will measure the integrated doses and fluences in the inner detector and will also to some extent provide benchmarking estimates of the different contributions (charged particles, neutrons and photons).

2420

One of the worst-case scenarios during LHC operation arises if several proton bunches hit the collimators in front of the detectors. While the accumulated radiation dose from such unlikely accidents corresponds to that acquired during a few days of normal operation, and as such provides no major contribution to the integrated dose, the enormous instantaneous rate might cause detector damage. The ATLAS Beam Conditions Monitor (BCM) [55] system consists of a set of detectors designed to detect such incidents and trigger an abort in time to prevent serious damage to the detector (see also Section 9.10). These incidents need to be distinguished from the stray protons and beam-gas backgrounds which frequently initiate charged particle showers, which originate well upstream (or downstream) from the ATLAS interaction point. Due to their very fast response time and intrinsically very high resistance to radiation, the BCM detectors will be used throughout the lifetime of the experiment to distinguish these stray beam particles from those originating from proton-proton interactions.

2425

2430

The BCM system, designed to tolerate doses of up to 500 kGy and in excess of 10^{15} charged particles per cm^2 over the lifetime of the experiment, consists of two stations, each with four modules. Each module, as depicted in Fig. 24 (left), includes two radiation-hard diamond sensors [56, 57] read out in parallel by radiation-tolerant electronics with a 1 ns rise-time [58]. Fig. 24 (right) shows a close-up view of one station installed around the beryllium beam-pipe. The stations are located symmetrically around the interaction point at $z = \pm 184$ cm and $R = 5.5$ cm, which corresponds to a $|\eta| = 4.2$. The difference in time-of-flight between the two stations, Δt , distinguishes particles from normal collisions ($\Delta t = 0, 25, 50$ ns, etc.) from those arising from stray protons ($\Delta t = 12.5, 37.5$ ns, etc.). The in-time and out-of-time multi-module coincidences are determined by an FPGA-based back-end, which digitises the signals, monitors the detector perfor-

2435

2440




Figure 25. Picture of one set of background monitors, to be installed in the TGC layer of the middle end-cap muon station. The eight different types of detectors are described in the text.

2445 mance and generates beam-abort signals if warranted. Preliminary analysis of data on one of the
modules in a high-energy pion test-beam shows a signal-to-noise ratio of 11 ± 2 in an operational
geometry, where minimum ionising particles are incident on the BCM sensors at a 45° angle. A
full description of the design, construction and test-beam characterisation of the BCM system can
be found in [55].

2450 **3.4.2 Monitors in the muon spectrometer**

Several sets of detectors have been installed in the end-cap muon stations to monitor the back-
ground fluences and thus to constrain further the particle transport codes used in the calculations
described in Section 3.3. These detectors are installed in the inner, middle, and outer end-cap sta-
tions. Figure 25 shows one set of the detectors which have been installed. They were chosen to
2455 provide a reliable response to neutrons or photons in various energy ranges:

1. Boron-lined proportional tubes operating with Ar/CO₂ gas are used to measure thermal and
slow neutrons (energies below 10^{-5} MeV). Each interaction $n + {}^{10}\text{B} \rightarrow \text{Li} + \alpha$ sends a
slow Li or α -particle into the tube. The large ionisation pulse associated with the Li or α -
particle is used for pulse-height discrimination against Compton electrons and minimum-
2460 ionising particles. These detectors are therefore relatively insensitive to photons and charged
particles.
2. Boron-loaded plastic scintillator (BC-454) is sensitive to the neutron interactions described
above and is also used to study thermal and slow neutrons.
3. Detectors with a plastic disk loaded with LiF and coated with a thin layer of ZnS(Ag) scin-
2465 tillator are sensitive to the tritium and α -particles produced in the neutron capture process in
lithium.
4. Another ZnS(Ag) scintillator embedded in plastic is used to study fast neutrons (with ener-
gies of a few MeV). The plastic is rich in hydrogen, from which incoming neutrons scat-
ter to produce recoil protons. These protons produce large ionisation pulses compared to
2470 minimum-ionising particles or low-energy electrons. Pulse-height discrimination schemes
should therefore provide good rejection against these backgrounds.
5. A liquid scintillator, with pulse-shape discrimination electronics, is used in combination with
plastics to measure fast neutrons.
6. Scintillator detectors with NaI and LSO crystals are used to measure the low-energy photon
2475 spectrum (0 to 10 MeV). The spectrum is dominated by photons, but also contains a neutron
component, which can be separated out using fitting techniques and detailed simulations.
7. Small ionisation chambers measure the total ionising dose.

3.4.3 Network of detectors for radiation measurements

A system of small silicon pixel detectors has been developed for radiation measurements in the
2480 experimental environment [59, 60]. This detector network will form a stand-alone system fully ca-
pable of delivering real-time images of fluxes and spectral composition of different particle species,
including slow and fast neutrons.

These silicon detectors will be operated via active USB cables and USB-ethernet extenders by a PC placed in the underground USA15 counting room, located next to the main cavern. The hybrid silicon pixel device consists of a silicon detector chip, 300 μm thick with 256×256 pixels, bonded to a readout chip. Each of the $55 \mu\text{m} \times 55 \mu\text{m}$ pixels is connected to its respective readout chain integrated on the chip. Settings of the pulse height discriminators determine the input energy window and at the same time provide noise suppression. The pixel counter determines the number of interacting quanta of radiation falling within this window. These devices can be used for position and energy sensitive (from 5 keV up to tens of MeV) spectroscopic detection of radiation. They are also capable of counting particle fluxes at rates in excess of GHz/cm².

This system can be used in both tracking and counting modes, to record tracks or counts caused by x-rays, gamma-radiation, neutrons, electrons, minimum ionising particles and ions. For neutron detection, the silicon detectors are partially covered by neutron converters (⁶LiF and polyethylene for slow and fast neutrons, respectively). The tracking mode is based on electronic visualisation of tracks and traces of individual quanta of radiation in the sensitive silicon volume. In the case of count rates above 5×10^3 events/cm² s, the devices are operated in counting mode, in which charge deposition in the pixels is counted at different threshold settings. Calibration of the devices enables the conversion of the individual tracks observed and/or counts measured into fluxes of respective types of radiation and dose rates. At least 14 of these pixel devices will be placed inside ATLAS: four devices on the LAr calorimeter facing the inner detector, four devices on the tile calorimeter, four devices near the muon chambers in the inner end-cap muon station, and two devices near the forward shielding and close to the outer end-cap muon station.

3.5 Activation

Induced radioactivity will be a major problem at the LHC, and ATLAS is the experiment with the highest levels of induced radiation. This is due to the small radius of the ATLAS beam-pipe, the small bore of the forward calorimeters, and to the shielding elements close to the beam-pipe. A comprehensive study has been made of the expected activation in different regions and for different data-taking and cooling-off scenarios. The methods and assumptions used in the calculation of the induced activity are given in [35]. The main conclusion of these studies is that the beam-pipe will be the major source of radiation in ATLAS.

Three different access scenarios are foreseen for ATLAS during shutdowns, as described in more detail in Section 9.7. They are described below and two of the scenarios are depicted in Figs. 26 and 27.

- (a) In the very short access scenario, all detector components remain in place and the magnetic fields remain on. These accesses are typically on the order of a few hours long.
- (b) In the short access scenario, the beam-pipe remains in place, but then acts as a linear source of photon radiation as can be seen in Fig. 26. Because of the high level of radiation, the area around the beam-pipe, out to a radius of about one metre, has to be fenced off after high-luminosity running. This will ensure that people working in ATLAS during short access will not be exposed to dose rates larger than 0.1 mSv/h (maintenance work in ATLAS will be designed to limit the yearly dose to 6 mSv per person). The only detector which is truly inside




Figure 26. The inner region of the detector during one of the short access scenarios. The predicted dose rates have been calculated for 10 years of operation at $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and for five days of cooling off. The short access scenario (a) has the beam-pipe in place.




Figure 27. The inner region of the detector during one of the main long access scenarios. The predicted dose rates have been calculated for 10 years of operation at $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and for five days of cooling off. The long access scenario (b) has only the inner detector section of the beam-pipe in place. The expected dose rates are greatly reduced in this access scenario.

the barrier is the inner detector. During a short access, maintenance of the inner detector will therefore be severely limited.

- 2525 (c) In the long access scenario, all the beam-pipe sections except the one inside the inner detector
- 2530 volume are removed as well as the small muon wheel (or inner end-cap muon stations) and the end-cap toroids. Two hot spots can clearly be seen in the final configuration, as shown in Fig. 27. One is the end-piece of the inner detector beam-pipe, which is made of aluminium, whereas the rest of the inner detector beam-pipe is made of beryllium. The expected dose rate can reach 0.2 mSv/h at this location. The other hot spot is in front of the forward

calorimeters, where the dose rate is predicted to reach very high values of up to 0.5 mSv/h. These relatively small-size regions will therefore be temporarily shielded with lead blocks during maintenance of the inner detector.

2535 While the beam-pipe section inside the inner detector is mostly made of beryllium, the rest
of the beam-pipe is made of stainless steel and has to be removed in the case of the long access
scenario, since it will become very radioactive with a contact dose rate of 3–5 mSv/h. This could in
certain cases inflict several mSv of integrated dose to personnel performing the intervention. One
way of reducing the dose to personnel would be to make the beam-pipe out of aluminium instead of
stainless steel. This is expected to give a factor 10–50 reduction of the dose levels. If the beam-pipe
2540 material were instead to be changed to beryllium over the whole length of the detector, the dose
rate would decrease by a factor of 100–1000 and would no longer be a problem. This is, however,
very costly and will only be discussed further in the context of the LHC upgrade programme.

4. Inner detector

4.1 Introduction

2545 The ATLAS Inner Detector (ID) is designed to provide hermetic and robust pattern recognition, excellent momentum resolution and both primary and secondary vertex measurements [61, 62] for charged tracks above a given p_T threshold (nominally 0.5 GeV, but as low as 0.1 GeV in some ongoing studies of initial measurements with minimum-bias events) and within the pseudorapidity range $|\eta| < 2.5$. It also provides electron identification over $|\eta| < 2.0$ and a wide range of energies (between 0.5 GeV and 150 GeV). This performance, which is required even at the highest luminosities expected from LHC collisions, is consequently at the limit of existing technology.

2550 The ID layout, as shown in Fig. 28, reflects the performance requirements. The ID is contained within a cylindrical envelope of length ± 3512 mm and of radius 1150 mm, within a solenoidal magnetic field of 2 T (see Section 2.2.4). Figures 29 and 30 show the sensors and structural elements traversed by 10 GeV tracks in respectively the barrel and end-cap regions.

2555




Figure 28. Plan view of a quarter-section of the ATLAS inner detector showing each of the major detector elements with its active dimensions and envelopes. The labels PP1, PPB1 and PPF1 indicate the patch-panels for the ID services.

The ID consists of three independent but complementary sub-detectors. The envelopes of each sub-detector are listed in Table 2 (see Section 1.3) and shown in Fig. 28. At inner radii, high-resolution pattern recognition capabilities are available using discrete space-points from silicon pixel layers and stereo pairs of silicon microstrip (SCT) layers. At larger radii, the transition




Figure 29. Drawing showing the sensors and structural elements traversed by a charged track of 10 GeV p_T in the barrel inner detector ($\eta = 0.3$). The track traverses successively the beryllium beam-pipe, the three cylindrical silicon-pixel layers with individual sensor elements of $50 \times 400 \mu\text{m}^2$, the four cylindrical double layers (one axial and one with a stereo angle of 40 mrad) of barrel silicon-microstrip sensors (SCT) of pitch $80 \mu\text{m}$, and approximately 36 axial straws of 4 mm diameter contained in the barrel transition-radiation tracker modules within their support structure.

2560 radiation tracker (TRT) comprises many layers of gaseous straw tube elements interleaved with
 transition radiation material. With an average of 36 hits per track, it provides continuous tracking to
 enhance the pattern recognition and improve the momentum resolution over $|\eta| < 2.0$ and electron
 identification complementary to that of the calorimeter over a wide range of energies.

2565 The high-radiation environment imposes stringent conditions on the inner-detector sensors,
 on-detector electronics, mechanical structure and services. Over the ten-year design lifetime of the
 experiment, the pixel inner vertexing layer must be replaced after approximately three years of oper-
 ation at design luminosity. The other pixel layers and the pixel disks must withstand a 1 MeV neu-
 tron equivalent fluence F_{neq} , defined in Section 3.3.1 (see also [53]), of up to $\sim 8 \times 10^{14} \text{cm}^{-2}$. The
 innermost parts of the SCT must withstand F_{neq} of up to $2 \times 10^{14} \text{cm}^{-2}$. To maintain an adequate
 2570 noise performance after radiation damage, the silicon sensors must be kept at low temperature (ap-




Figure 30. Drawing showing the sensors and structural elements traversed by two charged tracks of 10 GeV p_T in the end-cap inner detector ($\eta = 1.4$ and 2.2). The end-cap track at $\eta = 1.4$ traverses successively the beryllium beam-pipe, the three cylindrical silicon-pixel layers with individual sensor elements of $50 \times 400 \mu\text{m}^2$, four of the disks with double layers (one radial and one with a stereo angle of 40 mrad) of end-cap silicon-microstrip sensors (SCT) of pitch $\sim 80 \mu\text{m}$, and approximately 40 straws of 4 mm diameter contained in the end-cap transition radiation tracker wheels. In contrast, the end-cap track at $\eta = 2.2$ traverses successively the beryllium beam-pipe, only the first of the cylindrical silicon-pixel layers, two end-cap pixel disks and the last four disks of the end-cap SCT. The coverage of the end-cap TRT does not extend beyond $|\eta| = 2$.

proximately -5 to -10°C) implying coolant temperatures of $\sim -25^\circ\text{C}$. In contrast, the TRT is designed to operate at room temperature.

The above operating specifications imply requirements on the alignment precision which are summarised in Table 10 and which serve as stringent upper limits on the silicon-module build precision, the TRT straw-tube position, and the measured module placement accuracy and stability. This leads to:

- a) a good build accuracy with radiation-tolerant materials having adequate detector stability and well understood position reproducibility following repeated cycling between temperatures of -20°C and $+20^\circ\text{C}$, and a temperature uniformity on the structure and module mechanics which minimises thermal distortions;
- b) an ability to monitor the position of the detector elements using charged tracks and, for the SCT, laser interferometric monitoring [63];
- c) a trade-off between the low material budget needed for optimal performance and the significant material budget resulting from a stable mechanical structure with the services of a highly granular detector.

The inner-detector performance requirements imply the need for a stability between alignment periods which is small compared with the alignment precision. Quantitatively, the track precision

Item	Intrinsic accuracy (μm)	Alignment tolerances (μm)		
		Radial (R)	Axial (z)	Azimuth (R- ϕ)
Pixel				
Layer-0	10 (R- ϕ) 115 (z)	10	20	7
Layer-1 and -2	10 (R- ϕ) 115 (z)	20	20	7
Disks	10 (R- ϕ) 115 (R)	20	100	7
SCT				
Barrel	17 (R- ϕ) 580 (z) ¹	100	50	12
Disks	17 (R- ϕ) 580 (R) ¹	50	200	12
TRT	130			30 ²
<p>1. Arises from the 40 mrad stereo angle between back-to-back sensors on the SCT modules with axial (barrel) or radial (end-cap) alignment of one side of the structure. The result is pitch-dependent for end-cap SCT modules.</p> <p>2. The quoted alignment accuracy is related to the TRT drift-time accuracy.</p>				

Table 10. Intrinsic measurement accuracies and mechanical alignment tolerances for the inner-detector sub-systems, as defined by the performance requirements of the ATLAS experiment. The numbers in the table correspond to the single-module accuracy for the pixels, to the effective single-module accuracy for the SCT and to the drift-time accuracy of a single straw for the TRT.

should not deteriorate by more than 20% between alignment periods.

This chapter describes the construction and early performance of the as-built inner detector. In Section 4.2, the basic detector sensor elements are described. Section 4.3 describes the detector modules. Section 4.4 details the readout electronics of each sub-detector, Section 4.5 describes the detector power and control and Section 4.6 describes the ID grounding and shielding. Section 4.7 discusses the mechanical structure for each sub-detector, as well as the integration of the detectors and their cooling and electrical services. The overall ID environmental conditions and general services are briefly summarised in Section 4.8. Finally, Section 4.9 indicates some initial results on the operational performance and Section 4.10 catalogues the material budget of the ID, which is significantly larger than that of previous large-scale tracking detectors.

4.2 Inner-detector sensors

This section describes the detector sensors of the pixel, SCT and TRT sub-systems - silicon pixel and micro-strip sensors, and straw tubes filled with a Xe/CO₂/O₂ gas mixture. As discussed in Section 3.3, the detector sensors are subject to large integrated radiation doses. They have therefore

been developed and controlled to withstand the expected irradiation, with a safety factor of approximately two.

4.2.1 Pixel and SCT detector sensors

2605 The pixel and SCT sensors [64, 65] are required to maintain adequate signal performance over the detector lifetime at design luminosity (with the exception of the pixel vertexing layer, as discussed above). The integrated radiation dose has important consequences for the sensors of both detectors. In particular the required operating voltage, determined by the effective doping concentration, depends on both the irradiation and the subsequent temperature-sensitive annealing. The
2610 sensor leakage current also increases linearly with the integrated radiation dose. The n-type bulk material effectively becomes p-type after a fluence F_{neq} of $\sim 2 \times 10^{13} \text{ cm}^{-2}$. The effective doping concentration then grows with time in a temperature-dependent way. To contain this annealing and to reduce the leakage current, the sensors will, as noted above, be operated in the temperature range -5°C to -10°C . The sensors must further meet significant geometrical constraints on their
2615 thickness, granularity and charge-collection efficiency.

The pixel sensors required the most leading-edge and novel technology to meet the very stringent specifications on radiation hardness, resolution and occupancy in the innermost layers. The sensors are $250 \mu\text{m}$ thick detectors, using oxygenated n-type wafers with readout pixels on the n^+ -implanted side of the detector. Despite its higher cost and complexity, this novel design involving
2620 double-sided processing was used because:

- a) the n^+ implants allow the detector to operate with good charge-collection efficiency after type inversion, even when operated below the depletion voltage, because the depletion zone grows from the pixel side;
- b) highly oxygenated material has been shown to give increased radiation tolerance to charged
2625 hadrons, with improved charge collection after type inversion and lower depletion voltage.

All of the 1744 pixel sensors (external dimensions $19 \times 63 \text{ mm}^2$) are identical. The minimum pixel size is $50 \times 400 \mu\text{m}^2$, dictated by the readout pitch of the front-end electronics. There are 47232 pixels on each sensor, but for reasons of space there are four ganged pixels in each column of the front-end chip, thus leading to a total of 46080 readout channels. A common bias grid
2630 ensures a DC connection to each pixel implant during wafer testing of sensor tiles, but it is only AC-coupled in production modules due to concerns about the effects of high radiation doses. To guarantee optimal post-irradiation performance and to maintain a small feature size, a p-spray isolation technology has been used [66]. Each pixel of a sensor is bump-bonded through a hole in the passivation layer to an element of the front-end readout integrated circuit as part of the module.

2635 For reasons of cost and reliability, the 15912 sensors of the SCT use a classic single-sided p-in-n technology with AC-coupled readout strips. The sensors will initially operate at $\sim 150 \text{ V}$ bias voltage, but operating voltages of between 250 and 450 V will be required for good charge collection efficiency after ten years of operation, depending on the sensor radius, the integrated luminosity and the length of warm-up periods. The sensor thickness of $285 \pm 15 \mu\text{m}$ is a compromise between the required operating voltage, the primary signal ionisation and the simplicity
2640 of fabrication. The strip pitch was determined by the required digitising precision, granularity,

Module type	Sensor type	Cut length (mm)	Outer width (mm)	Inner width (mm)	Strip pitch (μm)	Inter-strip angle (μrad)
Barrel	Barrel	63.960	63.560	63.560	80.0	0
End-cap inner	W12	61.060	55.488	45.735	56.9–69.2	207.0
End-cap middle	W21	65.085	66.130	55.734	69.9–83.0	207.0
	W22	54.435	74.847	66.152	83.4–94.2	207.0
End-cap outer	W31	65.540	64.635	56.475	70.9–81.1	161.5
	W32	57.515	71.814	64.653	81.5–90.4	161.5

Table 11. External cut dimensions of the SCT barrel and end-cap (EC) sensors. The tolerance on all external dimensions is $\pm 25 \mu\text{m}$; the mask accuracy is at the level of $\pm 1 \mu\text{m}$. The inter-strip angle is that between adjacent strips of the sensor. The sensors are fabricated from 4-inch wafers.

particle occupancy and noise performance. A strip pitch of $80 \mu\text{m}$ with two 6 cm-long sensors daisy-chained was chosen for the rectangular barrel sensors and radial strips of constant azimuth with mean pitch of $\sim 80 \mu\text{m}$ were chosen for the trapezoidal end-cap sensors. There are thus a total of 768 active strips of 12 cm length per sensor, plus two strips at bias potential to define the sensor edge. The detector dimensions are summarised in Table 11 (see also Table 16). Using the binary readout electronics described in Section 4.4, a noise occupancy per channel of $< 5 \times 10^{-4}$ for a threshold of 1 fC is specified for fully irradiated modules. The detectors were required to operate stably at 500 V bias and to have $< 1\%$ of bad readout strips at 350 V bias before and after irradiation to 3×10^{14} 24 GeV protons/cm², equivalent to the damage expected for a fluence F_{neq} of 2×10^{14} cm⁻².

The SCT sensors were fabricated by two suppliers. For one supplier (7.8% of sensors), the p-spray isolation technology causes difficulties of micro-discharges at low humidity. All sensors were carefully tested, and a sub-sample was subjected to extended pixel or strip-by-strip studies [64, 65]. A sub-sample was also used for extensive post-irradiation performance studies [64, 67, 68, 69]. Apart from precautions related to the humidity sensitivity noted above, the rejected strip sensors were at the level of 1%. The pixel sensor rejection rate was somewhat higher because of a requirement to carefully control the profile and doping concentration of the p-spray isolation.

4.2.2 TRT straw tubes

Polyimide drift (straw) tubes of 4 mm diameter are the basic TRT detector elements [70]. The straw tube wall, especially developed to have good electrical and mechanical properties with minimal wall thickness, is made of two $35 \mu\text{m}$ thick multi-layer films bonded back-to-back. The bare material, a $25 \mu\text{m}$ thick polyimide film, is coated on one side with a $0.2 \mu\text{m}$ Al layer which is protected by a $5\text{--}6 \mu\text{m}$ thick graphite-polyimide layer. The other side of the film is coated by a $5 \mu\text{m}$ polyurethane layer used to heat-seal the two films back-to-back. Mechanically, the straws are stabilised using carbon fibres. After fabrication, the straws were cut to length (144 cm for the barrel

and 37 cm for the end-caps) and leak-tested at 1 bar over-pressure. The straw (cathode) resistance was required to be $< 300 \Omega/\text{m}$.

2670 The straw anodes are $31 \mu\text{m}$ diameter tungsten (99.95%) wires plated with $0.5\text{--}0.7 \mu\text{m}$ gold, supported at the straw end by an end-plug. They are directly connected to the front-end electronics (see Section 4.4) and kept at ground potential. The anode resistance is approximately $60 \Omega/\text{m}$ and the assembled straw capacitance is $< 10 \text{ pF}$. The signal attenuation length is $\sim 4 \text{ m}$ and the signal propagation time is $\sim 4 \text{ ns/m}$. The cathodes are operated typically at -1530V to give a gain of 2.5×10^4 for the chosen gas mixture of 70% Xe, 27% CO_2 and 3% O_2 with $5\text{--}10 \text{ mbar}$ over-pressure. A mixture containing 70% Ar and 30% CO_2 has been used during quality-control and cosmic-ray studies, resulting in an increased electron mobility and a degraded electron identification performance. Under normal operating conditions, the maximum electron collection time is $\sim 48 \text{ ns}$ and the operational drift-time accuracy is $\sim 130 \mu\text{m}$. Low-energy transition radiation (TR) photons are absorbed in the Xe-based gas mixture, and yield much larger signal amplitudes than minimum-ionising charged particles. The distinction between TR and tracking signals is obtained on a straw-by-straw basis using separate low and high thresholds in the front-end electronics (see Section 4.4.1.3).

2680 For the barrel straws, the anode wires (active length $\pm 71.2 \text{ cm}$) are read out from each end. Near their centre, the wires are supported mechanically by a plastic insert glued to the inner wall of the straw and split electrically by a fused glass capillary of 6 mm length and 0.254 mm diameter to reduce the occupancy. Each long barrel straw is therefore inefficient near its centre over a length of 2 cm . In the inner nine layers of type-1 barrel modules (see Section 4.3.3), the wires are subdivided into three segments keeping only the 31.2 cm -long end-segments on each side active.

2690 To guarantee stable operation, the wire offset with respect to the straw centre is required to be less than $300 \mu\text{m}$. This is essentially a requirement on straw straightness since the wire sag is $< 15 \mu\text{m}$. To maintain straw straightness in the barrel, alignment planes made of polyimide with a matrix of 4.3 mm diameter holes are positioned each 25 cm along the module. Wires with offsets larger than $400 \mu\text{m}$ (amounting to $\sim 0.1\%$ of all wires) have been disconnected in the final barrel and end-cap acceptance tests.

2695 The stable operation of TRT straws with the Xe-based gas mixture requires a re-circulating gas system with continuous monitoring of the gas quality. To avoid pollution from permeation through the straw walls or through leaks, the straws are operated in an envelope of CO_2 .

2700 At LHC rates, significant heat is generated in the straws by the ionisation current in the gas. The heat dissipation is proportional to the single straw counting rate and is estimated to be 10 mW to 20 mW per straw for the highest-occupancy straws (inner barrel layers) at the LHC design luminosity. To preserve gas-gain uniformity, the temperature gradient along each straw is required to be $< 10^\circ\text{C}$. The heat is evacuated differently for the barrel modules and end-cap wheels (see Section 4.3.3).

2705 At the LHC design luminosity, the straw counting rate will reach 20 MHz in the most critical detector regions, and the ionisation current density will reach $0.15 \mu\text{A}$ per cm of anode wire. The total accumulated charge after ten years of operation will reach $\sim 10 \text{ C/cm}$ in the most exposed straws. Many studies, including direct ageing tests lasting thousands of hours, have demonstrated stable straw-tube operation over their full operational lifetime [71, 72, 73]. Minute levels of pollution cannot be excluded and organo-silicone impurities, for which the relative concentrations must

Barrel	Radius (mm)	Staves	Modules	Pixels
Layer-0	50.5	22	286	13.2×10^6
Layer-1	88.5	38	494	22.8×10^6
Layer-2	122.5	52	676	31.2×10^6
End-cap (one side)	z (mm)	Sectors	Modules	Pixels
Disk 1	495	8	48	2.2×10^6
Disk 2	580	8	48	2.2×10^6
Disk 3	650	8	48	2.2×10^6
Barrel and both end-caps			1744	80.4×10^6

Table 12. Parameters of the pixel detector. The numbers of pixels, modules, staves/sectors are shown for each of the barrel layers and end-cap disks. The pixel envelope is shown in Table 2. The active barrel length is 801 mm and the inner and outer active radii of each disk are respectively 88.8 and 149.6 mm. The quoted barrel radii are average values since the barrel staves are tilted at 20° with respect to a tangent vector at the given radius. The disk z -positions are also average values.

2710 be kept below 10^{-11} , are particularly harmful. A gas filter has been incorporated in the TRT gas system for this reason. This filter is also effective in removing the ozone produced during gas amplification. In the case where ageing would significantly affect the detector performance, the use of a Ar/CO₂/CF₄ gas mixture during a few days of normal LHC operation has been shown to clean Si-based deposits from the anode wire.

2715 4.3 Inner-detector modules

4.3.1 Pixel modules and staves

There are 1744 modules in the pixel detector [64]. As shown in Fig. 28, the pixel modules are arranged in three barrel layers and two end-caps each with three disk layers. The basic detector parameters are listed in Table 12. A total of 112 barrel staves and 48 end-cap sectors (8 sectors per disk) form the barrel and disk layers. Details of the pixel mechanical structure are described in Section 4.7.

A schematic view and photograph of a pixel module are shown in Fig. 31. A pixel module consists of a stack, from the bottom up, of the following components:

- 2725 a) 16 front-end electronics chips thinned to $180 \mu\text{m}$ thickness, each with 2880 electronics channels;
- b) bump bonds (In or PbSn), which connect the electronics channels to pixel sensor elements;
- c) the sensor tile of area $63.4 \times 24.4 \text{ mm}^2$ and approximately $250 \mu\text{m}$ thick;
- d) a flexible polyimide printed-circuit board (flex-hybrid) with a module-control chip glued to the flex-hybrid;




Figure 31. Schematic view of a barrel pixel module (top) illustrating the major pixel hybrid and sensor elements, including the MCC (module-control chip), the front-end (FE) chips, the NTC thermistors, the high-voltage (HV) elements and the Type0 signal connector. Also shown (middle) is a plan view showing the bump-bonding of the silicon pixel sensors to the polyimide electronics substrate. The photograph at the bottom shows a barrel pixel module.

2730 e) a polyimide pig-tail with Cu lines and a connector (barrel modules) or a wire micro-cable (end-cap modules) bonded to the flex-hybrid.

In all cases for the barrel, the top of the module faces the beam-pipe, whereas the end-cap disks have modules facing both ways. This means that the pixel sensors see in general the minimum amount of material in front of them, i.e. basically only the flex hybrid.

2735 The sensors and electronics chips are connected by bump-bonding technology to form bare modules. Both solder (PbSn) and indium (In) bump-bonding technologies have been used to make pixel modules. Including reworked modules, the production statistics of bare modules are summarised in Table 13. Flex hybrids with attached and tested module-control chips are glued to accepted bare modules. In total, 2312 modules were available for final electrical and mechanical




Figure 32. Close up of a bi-stave loaded with modules. The insert shows the U-link cooling connection between staves.




Figure 33. Pixel disk sector during the attachment of modules. There are also three modules on the back of the sector.

Module yield	Number	Percentage
Module starts	2625	
Modules accepted after bump-bonding	2418	
Modules accepted after flex-hybrid glued	2312	100
Modules accepted for staves	2100	90.8
Module sub-set acceptable for layer-0	726	31.4

Table 13. Yield of pixel modules after bump-bonding of the electronics channels to the pixels. A sub-set of the highest quality modules were selected for the layer-0. Two materials for bump-bonding were used: indium (In) and lead-tin solder (PbSn) balls. At all stages of the fabrication, the use of PbSn gave significantly better yields.

2740 characterisation.

After construction, the pixel modules were tested electrically at room temperature and at the approximate operating temperature of -10°C . Thermal cycling was performed on each module prior to completion of electrical testing. A ranking was made to separate the modules into those acceptable for the layer-0 radius (highest quality) and for the outer barrel and disk regions. Modules with the best ranking were loaded on barrel staves and disk sectors, leaving modules with the poorest ranking as spares. A brief summary of the production statistics and ranking after testing is given in Table 13.

In the barrel region, 13 pixel modules are mounted on each stave using robotic tools and then glued (see Fig. 32). The staves are themselves mounted on carbon-fibre structures (see Section 4.7.1). A bare stave consists of machined carbon-carbon (C-C) plates, an aluminium cooling tube and a carbon-fibre composite piece that captures the aluminium tube and is glued to the C-C pieces. The C-C pieces are precisely machined in a step pattern and are one-half the length of a stave. They are joined in the middle during the stave fabrication. A custom extrusion was used to

make the aluminium tube with a flat surface at the interface with the C-C material. A custom fitting
2755 is welded to the end of the aluminium pipe. A thermal compound is used to conduct heat between
the C-C plate and the tube. The tube is held in place by a carbon-fibre piece glued to the C-C
plate. Because of corrosion, the cooling tubes associated with initial loaded staves were reworked
through the insertion of a smaller cooling tube with some deterioration of the thermal performance.
The cooling tubes were replaced on bare staves. The electrical and thermal performance of each
2760 stave was measured after assembly. As illustrated in Fig. 32, the staves are joined to form bi-staves,
which form the cooling unit in the barrel region. A custom-welded aluminium U-link is attached to
each bi-stave. Custom low-mass cables are connected to each stave via a connector on each module
and attached to the back of the stave to reach the ends of the stave.

The end-cap equivalent of the stave is a sector. The two pixel end-caps each have three identi-
2765 cal disks. Each disk is composed of eight sectors. Six pixel modules are directly mounted on each
sector, as shown in Fig. 33. The sectors are composed of thin, C-C faceplates with a rectangular
aluminium cooling tube and vitreous carbon foam between the faceplates. The cooling tube is bent
into a W-like shape to fit within the sector and makes contact with the faceplates with a compliant,
thermally conducting adhesive. Each cooling circuit in the disk region serves two sectors.

2770 Tested end-cap modules are positioned and glued on each sector with a precision of 1–2 μm in
the plane of the module and about 10 μm perpendicular to the module plane. The module locations
on each sector were optically surveyed using fiducial marks in the corners of the pixel sensors and
other information. The survey precision with respect to the mounting bushings is estimated to be
better than 5 μm in the plane of the module and about 15 μm in the direction perpendicular to this
2775 plane.

The spatial resolution of individual pixel modules has been measured in a test beam, for both
non-irradiated and fully irradiated modules [64, 74]. At normal incidence, a spatial resolution
of 12 μm is measured and approximately 80% of the tracks have a single pixel hit. The resolution
is not significantly degraded after irradiation. The optimal resolution of 4.7 μm (before) and 6.0 μm
2780 (after) irradiation is obtained for incident angles of 10–15°. The pixel barrel staves overlap and are
mounted at a tilt angle of -20° , defined by geometrical constraints. The Lorentz angle, however,
for a magnetic field of 2T, varies between 12° (before) and 6° (after) irradiation [64, 74].

4.3.2 SCT modules

As indicated in Fig. 28, the SCT consists of 4088 modules [68, 69] tiling four coaxial cylindrical
2785 layers in the barrel region and two end-caps each containing nine disk layers [61]. The modules
cover a surface of 63 m^2 of silicon and provide almost hermetic coverage with at least four precision
space-point measurements over the fiducial coverage of the inner detector. Tables 14 and 15 show
the SCT detector parameters in detail. Details of the SCT structure are described in Section 4.7.

The 2112 barrel SCT modules [68] use 80 μm pitch micro-strip sensors [65], which are de-
2790 scribed in Section 4.2. The sensors are connected to binary signal readout chips [75], which are
described in Section 4.4.1.2. The barrel module is shown, with its components, in Fig. 34. The
module parameters are shown in Table 16. The four sensors, two each on the top and bottom side,
are rotated with their hybrids by ± 20 mrad around the geometrical centre of the sensors. They
are glued on a 380 μm -thick thermal pyrolytic graphite (TPG) base-board [76], which provides the
2795 thermal and mechanical structure. This extends sideways to include beryllia facings. A polyimide




Figure 34. Photograph (left) and drawing (right) of a barrel module, showing its components. The thermal pyrolytic graphite (TPG) base-board provides a high thermal conductivity path between the coolant and the sensors.

Barrel cylinder	Radius (mm)	Full length (mm)	Module tilt angle (degrees)	Number of modules
3	284 (299)	1530 (1498)	11.00	384
4	355 (371)		11.00	480
5	427 (443)		11.25	576
6	498 (514)		11.25	672
Total				2112

Table 14. SCT barrel cylinder parameters and number of modules per cylinder. There are 12 modules per row. The quoted radii and length are those of the outer surface of the support cylinder. The average active sensor radii and overall length are shown in brackets. The tilt angle is with respect to the tangent to the support cylinder.

hybrid [77] with a carbon-fibre substrate bridges the sensors on each side. The two 770-strip (768 active) sensors on each side form a 128 mm long unit (126 mm active with a 2 mm dead space). High voltage is applied to the sensors via the conducting base-board.

2800 Precision alignment criteria were applied during assembly: the assembly tolerance as well as the achieved build accuracy are shown in Table 16. The important in-plane tolerance for positioning sensors within the back-to-back stereo pair was $< 8 \mu\text{m}$ and the achieved variance was $2 \mu\text{m}$. In the module plane, no additional distortions were measured after thermal cycling. Out-of-plane, the individual components and the assembly jiggling and gluing determine the module thickness and the intrinsic bow of the sensors determines the out-of-plane shape. A common distortion profile has
2805 been established for the sensors at the level of a few μm and a module thickness variation of $33 \mu\text{m}$ was maintained during fabrication. Following thermal cycling, the out-of plane distortions changed

Disk	1	2	3	4	5	6	7	8	9
$ z $ (mm)	853.8	934.0	1091.5	1299.9	1399.7	1771.4	2115.2	2505.0	2720.2
Outer	52								
Middle	40								None
Inner	None	40					None		

Table 15. The nominal z -position of the centre of each SCT end-cap disk and the number of modules on each disk (the total number of modules is 1976, summed over both end-caps). For geometrical acceptance reasons, disk 9 has only outer modules, while disks 1, 7 and 8 have no inner modules. The middle modules of disk 8 have only one sensor, again for geometrical acceptance reasons. For disks 1 to 8, the inner and outer module centres are displaced towards the interaction point by an average distance of 15.25 mm, while the middle modules are displaced away from the interaction point by 15.25 mm. For disk 9, the modules centres are displaced by 15.25 mm away from the interaction point. The modules are tiled in ϕ by ± 2.75 mm about these nominal centres.

by a few μm (RMS). When cooled from room to operating temperature, profile deviations did not exceed $20 \mu\text{m}$, even at the sensor corners not supported by the base-board.

Figure 35 shows the construction of an end-cap module [69]. There are three module types, as shown in Table 16. Each of the 1976 modules has two sets of sensors glued back-to-back around a central TPG spine with a relative rotation of ± 20 mrad to give the required space-point resolution in R - ϕ and R . The module thickness is defined by the individual components and variations are compensated by the glue thickness (nominally $90 \mu\text{m}$). The TPG spine conducts heat from the sensors to cooling and mounting points at the module ends and serves as the bias contact to the sensors. Glass fan-ins attach one end of the spine to a carbon base-plate with the polyimide flex-hybrid glued to it. The modules are arranged in tiled outer, middle and inner rings.

The precision alignment criteria applied to the end-cap modules were similar to those of barrel modules. The RMS spread of the module survey measurements after construction was $1.6 \mu\text{m}$ in the back-to-back position of the stereo pair, measured transverse to the strips, and $2.8 \mu\text{m}$ in the position of the mounting hole and slot measured transverse to the strips. In the module plane, no additional distortions were measured after thermal cycling. Out of the plane, the end-cap modules are less rigid, and are affected by variations of the spine thickness and bowing of the sensors. A common distortion profile has been established for the sensors at the level of a few μm and a module thickness variation of $15 \mu\text{m}$ was maintained during fabrication. Following thermal cycling, the out-of plane distortions changed by only a few μm (RMS). As for the pixels, the variations within the module are small compared to the subsequent module-placement accuracies.

The spatial resolution of individual SCT modules has been measured in a test beam, for both non-irradiated and fully irradiated modules [67]. At normal incidence, a combined spatial resolution of $\sim 16 \mu\text{m}$ is measured in R - ϕ ; the resolution is consistent with the binary readout of the two sensors with $80 \mu\text{m}$ strips, including a small fraction of multiple hits. The resolution is not significantly degraded after irradiation. The barrel staves are mounted at a tilt angle of $\sim 11^\circ$, defined by geometrical constraints and opposite in sign to that of the pixel barrel staves because of

Parameter	Description
Strips	2×768 active strips, ± 20 mrad stereo rotation
Nominal resolution	17 μm in-plane lateral ($R\text{-}\phi$) 580 μm in-plane longitudinal (z or R)
Module dimensions:	
-barrel	Active length 126.09 mm + 2.09 mm dead space between sensors
-outer end-cap	Active length 119.14 mm + 2.09 mm dead space, radius 438.77 – 560.00 mm
-middle end-cap	Active length 115.61 mm + 2.09 mm dead space, radius 337.60 – 455.30 mm
-short-middle end-cap	Active length 52.48 mm, radius 402.82 – 455.30 mm
-inner end-cap	Active length 59.1 mm, radius 275.00 – 334.10 mm
Specified build tolerance	Barrel back-to-back in plane: $< 8 \mu\text{m}$ (lateral) $< 20 \mu\text{m}$ (longitudinal) End-cap back-to-back in plane: $< 5 \mu\text{m}$ (lateral) $< 10 \mu\text{m}$ (longitudinal) Barrel out-of-plane (module thickness and sensor bowing): $< 70 \mu\text{m}$ Barrel envelope: $< 200 \mu\text{m}$ End-cap envelope: $< 115 \mu\text{m}$ Barrel module fixation points with respect to module centre: $< 40 \mu\text{m}$ End-cap module fixation points with respect to module centre: $< 20 \mu\text{m}$
Build accuracy of accepted modules (RMS)	Barrel back-to-back in plane: $\pm 2.1 \mu\text{m}$ (lateral) $\pm 2.7 \mu\text{m}$ (longitudinal) End-cap back-to-back in plane: $\pm 1.6 \mu\text{m}$ (lateral) $\pm 1.3 \mu\text{m}$ (longitudinal) Barrel module thickness: $\pm 33 \mu\text{m}$ Barrel out-of-plane (sensor bowing): $\pm 9 \mu\text{m}$ End-cap module thickness: $\pm 15 \mu\text{m}$ End-cap out-of-plane (sensor bowing): $\pm 20 \mu\text{m}$ Barrel module fixation points with respect to module centre $\pm 10 \mu\text{m}$ End-cap module fixation points with respect to module centre $\pm 6 \mu\text{m}$
Hybrid power consumption	5.5 – 7.5 W
Sensor power consumption	Up to 460 V bias, < 1 W at -7°C

Table 16. SCT barrel and end-cap module specifications and the RMS build accuracy for accepted modules. The barrel out-of-plane bowing specifications and the measured results are with respect to an average module shape. The quoted hybrid and sensor power consumption is more extensively described in Section 4.5.




Figure 35. The upper photograph shows the three SCT end-cap module types (outer, middle and inner from left to right). The lower schematic shows an exploded view of the different components for a middle module, including the high thermal conductivity spine, the polyimide hybrid and the ABCD readout ASIC's.

the different nature of the sensors used in each detector. The measured Lorentz angle, however, for a magnetic field of 2T, varies between 4.2° (before) and 2.7° (after) irradiation.

2835 The barrel and end-cap sensors are specified to operate at -7°C , with a maximum variation within and between modules of 5°C , to reduce the bulk leakage current after radiation damage. The hybrid power will be $5.5 - 7.5 \text{ W}$ per module, and the sensor load will reach $\sim 1 \text{ W}$ per module after ten years of operation. In addition, a convective load of $\sim 0.8 \text{ W}$ per module is expected with an additional convective load of $\sim 0.8 \text{ W}$ per module at the top of the barrel cylinders and outer
2840 disks. The heat is extracted by evaporating C_3F_8 at $\sim -25^\circ\text{C}$, circulating in cooling pipes attached to each module.

For the barrel, the sensor and hybrid heat leaves via the base-board and the hybrid substrate to the large beryllia facing on the base-board, which is interfaced to an aluminium block with a $\sim 100 \mu\text{m}$ layer of thermal grease and a copper-polyimide capacitive shunt shield. At full load
2845 for irradiated modules, the hybrid and sensor temperatures are expected to be approximately 14°C and 12°C above the cooling-pipe temperature, respectively. The block is itself soldered to a 3.6 mm

	$ z _{min}$ (mm)	$ z _{max}$ (mm)	R_{min} (mm)	R_{max} (mm)	Number of modules	Number of layers	Straws per module
Barrel (both sides)	0	780	554	1082	96	73	52544
Type-1 module (inner)	400	712.1	563	624	32	9	329
Type-1 module (outer)	7.5	712.1	625	694		10	
Type-2 module	7.5	712.1	697	860	32	24	520
Type-3 module	7.5	712.1	863	1066	32	30	793
End-cap (one side)	827	2744	615	1106	20	160	122880
Type-A wheels	848	1705	644	1004	12	8	6144
Type-B wheels	1740	2710	644	1004	8	8	6144

Table 17. Parameters of the TRT barrel modules and end-cap wheels (for one TRT end-cap). The quantities shown in bold are global parameters, including services and electronics. All other quantities are for individual modules and the active region. Type-1 barrel modules include two straw types, as described in the text.

diameter Cu/Ni cooling pipe. Each cooling loop serves 48 barrel modules.

For the end-cap, the sensor heat leaves via the spine, while the hybrid heat is transferred via the carbon-fibre hybrid substrate to a carbon-carbon cooling block, which is split to minimise heat transfer between the sensor and the hybrid. At full load, the ASIC and sensor temperatures are expected to be respectively $\sim 30^\circ\text{C}$ and $\sim 10\text{--}15^\circ\text{C}$ above the coolant temperature. A layer of thermal grease is applied between the modules and the cooling block. The blocks are soldered to a Cu/Ni cooling pipe that serves up to 33 modules.

All modules were tested electrically at room temperature and at the expected operating temperature of -7°C . Thermal cycling was performed on each module prior to completion of electrical testing. After module assembly and testing, 99.8% of the strips operate satisfactorily.

4.3.3 TRT modules

The TRT contains up to 73 layers of straws interleaved with fibres (barrel) and 160 straw planes interleaved with foils (end-cap), which provide transition radiation for electron identification. All charged tracks with $p_T > 0.5 \text{ GeV}$ and $|\eta| < 2.0$ will traverse at least 36 straws, except in the barrel-end-cap transition region ($0.8 < |\eta| < 1.0$), where this number decreases to a minimum of




Figure 36. Photograph of one quarter of the barrel TRT during integration. The shapes of one outer, one middle and one inner TRT module are highlighted. The barrel support structure space-frame can be seen with its triangular sub-structure.

22 crossed straws. Typically, seven to ten high-threshold hits from transition radiation are expected for electrons with energies above 2 GeV.

2875 The barrel TRT is divided into three
rings of 32 modules each, supported at each
end by a space frame, which is the main
component of the barrel support structure
(see Section 4.7). Each module consists
2880 of a carbon-fibre laminate shell and an in-
ternal array of straws embedded in a ma-
trix of 19 μm -diameter polypropylene fibres
serving as the transition radiation material.
The straws, described in Section 4.2, form
2885 a uniform axial array with a mean spacing
of ~ 7 mm. The module shells are non-
projective to reduce the dead region for high
 p_T tracks. The main barrel parameters are
shown in Table 17. Figure 36 shows a quad-
2890 rant of the TRT barrel during the integra-
tion of the modules at CERN. A ϕ -slice showing
one outer, one middle and one inner module
is highlighted.

The dimensional specifications are set
2895 by the intrinsic straw R - ϕ resolution of
130 μm , implying that each wire position
is constrained to within ± 50 μm . The mod-
ule shell, made of 400 μm thick carbon fibre
with high thermal conductivity and flat to
2900 within 250 μm , is measured to satisfy maxi-
mum distortions of < 40 μm under full load.

The module shells also serve as a gas
manifold for CO_2 which circulates outside
the straws to prevent high-voltage discharges
2905 and the accumulation of xenon due to possible gas leaks which would absorb the transition radiation
photons. The heat dissipated by the barrel straws is transferred to the module shell by conduction
through the CO_2 gas envelope. Each module shell is cooled by two cooling tubes located in the
acute corners. These tubes also serve as return pipes for the C_6F_{14} cooling circuits of the front-end
electronics.

2910 The module end with its components is shown in Fig. 38. The central element is the HV plate
which has stringent requirements on flatness and cleanliness to prevent discharges and on the straw
feed-through accuracy to ensure mechanical precision of the straw location. The HV plates were
individually surveyed after machining. The tension plate of Fig. 38 mounted on the HV plate is a
printed-circuit board holding the wire ends (and ensuring the wire tension) and providing electrical
2915 connections. It also closes the active gas volume and serves as a Faraday cage for the active module




Figure 37. Photograph of a four-plane TRT end-cap wheel during assembly. The inner and outer C-fibre rings can be seen, as well as the first layer of straws and the first stack of polypropylene radiator foils beneath it. Also visible are the plastic end-plugs which are used to position and fix the straws to the C-fibre outer ring. The high-voltage petals used to connect the straws to the high-voltage lines (see text) are laid back at this stage of the assembly and will only be folded vertically to push the petals into the straws at the next stage.




Figure 38. Detailed view of the end of the TRT barrel modules, showing the connection of the straw ends to high voltage through the high-voltage (HV) plate, of the wires to the front-end boards through the tension-plate and of the gas inlet to the individual straws through the active-gas manifold.




Figure 39. Schematic view of the inner and outer ends of the TRT end-cap wheels, showing the plastic end-plugs used to position and fix the straws in the inner and outer C-fibre rings, the crimping pins holding and positioning the wires, the inner and outer active-gas manifolds, and the flex-rigid printed-circuit board used to connect the straws to high voltage and the wires to the front-end electronics.

elements.

Mechanical and electrical tests on the modules and individual straws were made at successive assembly stages and following delivery of the modules to CERN. The module lengths were required to be in the range 1461.5–1462.9 mm and the twisting of individual modules was required to be < 1 mm. Stringent criteria were applied to the gas tightness, wire tension, straw straightness, and high voltage stability. The modules installed in the barrel had 98.5% operational channels.

The TRT end-caps each consist of two sets of independent wheels (see Fig. 37 and Table 17). The set closer to the interaction point contains 12 wheels, each with eight successive layers spaced 8 mm apart. The outer set of wheels contains eight wheels, also with eight straw layers but spaced 15 mm apart. Each layer contains 768 radially oriented straws of 37 cm length with uniform azimuthal spacing. The space between successive straw layers is filled with layers of 15 μm thick polypropylene radiator foils separated by a polypropylene net.

Each eight-plane wheel consists of two basic four-plane assembly units. To assemble a four-plane wheel, straws were inserted and glued into precisely drilled holes in grounded inner and outer carbon fibre rings. The rings and the straws constitute the main mechanical structure of the wheels. The successive straw layers, interleaved by the radiators, are rotated from one layer to the next by $3/8$ of the azimuthal straw spacing in a given layer. For high- p_T prompt tracks, this ensures optimal uniformity in the number of crossed straws, which varies radially from \sim six to \sim four straws across an eight-plane wheel.

Flex-rigid printed-circuit boards provide high-voltage and signal connections to the end-cap wheels through two separate layers (see Fig. 39). Each of the flexible layers has conducting paths on one side which connect to the rigid part of the board. To provide a reliable electrical connection, flexible "petals" in the high-voltage layer are forced into contact with the inner straw wall through the insertion of a plastic plug. A press-fit between similar but smaller petals in the signal layer and a metallic crimping pin positions and fixes the anode wires. There are 32 such boards per four-plane

wheel, each serving a ϕ -sector of 96 straws.

Each sector is further segmented into three groups of 32 readout channels and 12 high-voltage groups of eight straws sharing a common fuse and blocking capacitor. The carbon-fibre ring holding the straws and flex-rigid boards, together with a third carbon fibre ring and a simpler glass fibre board provide a rigid structure around the outer wheel perimeter that also serves as a gas manifold (see Fig. 39). The inner gas manifold is made from reinforced polyimide material.

The heat dissipated by the end-cap straws is evacuated through the CO₂ gas envelope which is forced to flow along the straws from the inner to the outer radius. Each group of wheels has its own CO₂ cooling circuit, passing the gas sequentially through all the wheels of the group. Heat exchangers cooled with C₆F₁₄ extract heat from the gas between adjacent wheels. The high flow rates required, 50 m³ per hour for type-A and 25 m³ per hour for type-B wheels, necessitates a closed-loop system capable of maintaining a small gas pressure between 0 and 5 mbar with a stability of ± 0.5 mbar inside the detector.

As for the barrel modules, all end-cap wheels passed quality control procedures during construction and after delivery to CERN, resulting in more than 99% of fully operational channels.

4.4 Readout of the inner detector

The readout architecture of the ID is optimised separately for each of the three sub-detectors, but is characterised by the following common elements:

- a) the reception of a 40.08 MHz clock signal synchronous with the LHC bunch-crossings used to time-stamp the signal generated in the low noise front-end electronics;
- b) signal generation in the front-end electronics and storage in binary or digital buffers for approximately $\sim 3.2 \mu\text{sec}$, compatible with the L1 trigger latency of $2.5 \mu\text{sec}$;
- c) following a L1 trigger, the subsequent transfer of the buffer content associated with the bunch-crossing or possibly several bunch-crossings to a readout driver (ROD) off the detector.

The readout of the pipelines is subject to the reception of a L1 trigger signal from the calorimeters or the muon detectors via the central trigger processor, as discussed in Section 8.2.3. The ID is not part of the L1 trigger. External power supplies provide regulated voltages to the front-end electronics and to the sensors (see Section 4.5).

4.4.1 Front-end electronics

4.4.1.1 Pixel front-end electronics Each front-end readout ASIC [64, 78, 79] of the pixel detector contains 2880 readout cells of $50 \times 400 \mu\text{m}^2$ size arranged in a 18×160 matrix. Sixteen ASIC's are bump-bonded to each sensor (module). The ASIC's are fabricated using commercial $0.25 \mu\text{m}$ CMOS technology. A high level of radiation tolerance is achieved by a combination of the process technology (thin gate oxide) and the use of special layout techniques (annular layout and guard rings for all NMOS transistors). The schematic and basic functionality of the pixel circuit is shown in Fig. 40.

Each readout cell contains an analogue block where the sensor charge signal is amplified and compared to a programmable discriminator threshold. The digital readout then transfers the hit




Figure 40. Layout and schematic description of the front-end readout ASIC for the pixel detector (see text).

2980 pixel address, a hit time stamp and a digitised amplitude (the time over threshold (ToT)) to buffers
 at the chip periphery. These hit buffers monitor each stored hit by inspecting the associated time
 stamp.

The charge sensitive amplifier uses a single-ended folded-cascode topology optimised for a
 nominal capacitive load of 400 fF and designed for the negative signal expected from DC-coupled
 2985 n^+ -on- n sensors. Attention has been paid to the preamplifier design because, following the irradiation
 expected at the LHC, the sensor leakage current (50 nA) is two orders of magnitude larger
 than the signal, that is itself reduced due to carrier trapping inside the silicon. The pre-amplifier has
 an approximate 5 fF DC feedback capacitance with a 15 ns risetime. The total analogue front-end
 (pre-amplifier, second stage amplifier and discriminator) has a bias current of only 24 μ A per pixel
 2990 for the default DAC settings. To ensure the separation of contiguous bunch crossings, a front-end
 time walk of <25 ns is required. To fulfill the requirements of sensor leakage current, a compensa-
 tion circuit is implemented that drains the leakage current and prevents any influence on the bias
 current of the fast feedback circuit used to discharge the feedback capacitor. Each pixel has several
 parameters that are tuned and stored in a 14-bit control register, for example the feedback trimming
 2995 and threshold levels. The digital circuitry in the readout cells generates the required hit information
 to measure the charge and associate the hit to the bunch crossing.

The readout is made using the column-based readout controller. The first task of the controller
 is the generation of the readout sequence to transfer the hit information from the pixel to the hit




Figure 41. Noise distribution for normal pixels of a non-irradiated module (upper left) and of a module irradiated with 24 GeV protons to a fluence F_{neq} of 10^{15}cm^{-2} (lower left), measured after retuning at an operating temperature of -4°C . The measured efficiency as a function of the incident beam particle arrival time is also shown for the non-irradiated (upper right) and irradiated modules (lower right). The arrows indicate the efficiency at the timing plateau.

3000 buffer at the end of the column. The second task is the digital processing of hit data in the front-end chip buffers. Hits are continuously compared to a delayed beam-crossing counter (see Fig. 40). When a match is found with the correct latency, the hit is flagged for readout if a L1 trigger signal is present, or deleted if there is no matching trigger. Flagged hits are transmitted to a serialiser and sent out of the chip. Hits older than the trigger latency are cleared from the front-end chip buffers.

3005 The module-control chip [64, 80] is a digital chip running with the same 40 MHz clock. It has three main system tasks: the loading of parameter and configuration data in the front-end chips and in the module-control chip itself, the distribution of timing signals such as bunch-crossing, L1 trigger and resets (TTC functions), and the front-end chip readout and event building. The design of the module-control chip reflects the required pixel performance during LHC operation: the association of signals to a bunch-crossing, the expected bandwidths at the highest luminosity, the maximum L1 trigger rate of 100 kHz (initially 75 kHz) and the number of front-end chips, which are controlled in a module. Because of the high radiation environment, particularly in layer-0 modules, special attention has been given to ensure a single-event upset (SEU) tolerant design.

3010 Extensive electrical tests have been made on ASIC chips, single front-end chip assemblies and full modules, before and after irradiation. Some production modules were irradiated to the end-of-life dose expected at LHC. The noise and hit efficiency are shown in Fig. 41 for a pixel




Figure 42. Schematic of the readout ASIC for the SCT detector, showing the successive signal processing steps.




Figure 43. The efficiency (circles) and noise occupancy (triangles) for SCT barrel modules measured in a test-beam before irradiation (left) and after exposure to a dose of $\sim 3 \times 10^{14}$ p/cm² in a 24 GeV proton test-beam (right). The nominal operating threshold is 1 fC. The dashed horizontal lines represent the nominal module performance specifications in terms of efficiency and noise. The vertical lines represent the range of thresholds over which these specifications are met after irradiation.

module before irradiation. Similar results are also shown for a module irradiated to a fluence F_{neq} of 10^{15}cm^{-2} (and to 500 kGy ionising dose) in a 24 GeV proton beam, at an operating temperature of -4°C and after threshold retuning. Both quantities are only slightly deteriorated by the irradiation and remain within the operating specifications. Furthermore, the difference in threshold dispersion after retuning is negligible.

4.4.1.2 SCT front-end electronics The readout hybrid of each SCT module (see Section 4.3) houses 12 identical 128-channel ASIC's [75] to read a total of 1536 sensor strips per module. The ASIC is fabricated in radiation tolerant bi-CMOS DMILL technology. The successive blocks of the ASIC are shown in the circuit schematic of Fig. 42. A pre-amplifier, shaper and tunable discriminator exists for each channel. A 132-length binary pipeline stores the hit information for each channel associated to the beam crossing for a period of $\sim 3.2 \mu\text{sec}$. Following a L1 trigger, the chip compresses the data pertinent to that beam crossing and serialises it for output. An 8-deep de-randomising buffer after the pipeline ensures that the dead-time is negligible for the expected data rates.

Two critical module performance specifications are the detection efficiency ($> 99\%$) and noise occupancy ($< 5 \times 10^{-4}$), for signals from the 12 cm long silicon strips with a capacitive load of $\sim 20 \text{pF}$. These have led to the choice of a front-end discriminator threshold of 1 fC. Extensive studies have been made using ASIC chips, single front-end chip assemblies and full modules, before and after irradiation. A sample of production modules was irradiated in a 24 GeV proton beam to a dose of $\sim 3 \times 10^{14}$ protons per cm^2 , equivalent to the damage expected for a fluence F_{neq} of $2 \times 10^{14} \text{cm}^{-2}$. The efficiency and noise occupancy measured in a test-beam are shown as a function of the discriminator threshold in Fig. 43. At the nominal operating threshold of 1 fC, the efficiency and noise-occupancy specifications are easily met before irradiation and almost met after irradiation.

The chips are daisy-chained so that all the data of one module are read out over two serial links. Several design features provide fault tolerance. For example, any faulty chip (except for the master on barrel modules) can be bypassed in the serial data path and if one link should fail, it is possible to send the data using the remaining active link. Likewise, if the primary clock and command lines to the ASIC's on a module fail, it is possible to instead use the clock and command signals from an adjacent module.

4.4.1.3 TRT front-end electronics The analogue signal processing and threshold discrimination to detect signals from both minimum-ionising particles and transition radiation, as well as the subsequent time digitisation and data-pipelining are implemented in two on-detector ASIC's. The signal is shown in Fig. 44 at each stage of the TRT signal readout chain which comprises:

- a) an eight-channel analogue ASIC [81], called the ASDBLR, fabricated in bi-CMOS radiation tolerant DMILL technology. It performs the amplification, shaping and baseline restoration. It includes two discriminators, one operating at low threshold (typically 250 eV) for minimum-ionising signal detection and one operating at high threshold (typically 6 keV) for transition radiation detection;
- b) a subsequent 16-channel ASIC fabricated in commercial radiation-tolerant $0.25 \mu\text{m}$ CMOS technology [82]. This ASIC performs the drift-time measurement ($\sim 3 \text{ns}$ binning). It in-




Figure 44. Schematic of the front-end readout of the TRT detector, showing the input signal shape and the signal shape after the amplification and shaping, the baseline restoration (BLR) and the dual-threshold discrimination which provides the ternary output corresponding to the low threshold set or both low and high thresholds set.

cludes a digital pipeline for holding the data during the L1 trigger latency, a derandomising buffer and a 40 Mbits/s serial interface. It also includes the necessary interface to the timing, trigger and control as well as DAC's to set the discriminator thresholds of the analogue ASIC and test-pulse circuitry for mimicking analogue inputs to the analogue ASIC.

3060

These ASIC's are housed on front-end boards attached to the detector. There are 12 different boards for the barrel and three different boards for the end-cap. The electronics are cooled by a liquid mono-phase fluorinert (C_6F_{14}) cooling system.

At the TRT operating low threshold used for tracking (equivalent to $\sim 15\%$ of the average signal expected from minimum-ionising particles), the mean straw noise occupancy is $\sim 2\%$, but a small fraction of 1% of channels have a noise occupancy exceeding 10%, which however remains small compared with the expected maximum straw occupancy of 40%. The full front-end electronics chain was exposed to a neutron dose of $\sim 4 \times 10^{14}/cm^2$ and to a γ -ray dose of 80 kGy. Changes of up to 25% were observed in the ASDBLR gain, but with no change in the effective thresholds and noise performance after a standard voltage compensation procedure.

3070

4.4.2 Data transmission and power-supply services and routing

The transmission of data from the ID modules to the off-detector electronics in the service cavern, as well as the digital transmission of the clock and control commands to the modules differs for the 3 sub-detectors. Figures 45, 46 and 47 summarise the layout and technology of the readout and control services for each of the pixel, SCT and TRT sub-detectors. The locations of key patch-panel connection boards are shown: PP0 close to the ends of the pixel detector, PP1 at the edges of the ID

3075




Figure 45. The routing of data links and power supply cables from each side of the pixel detector to respectively the off-detector electronics and power supplies in the service caverns, together with the number, type and utilisation of the cables and optical links.




Figure 46. The routing of data links and power supply cables from each side of the SCT to respectively the off-detector electronics and power supplies in the service caverns, together with the number, type and utilisation of the cables and optical links.

volume, PP2 in specifically designed parts of the muon spectrometer system and PP3 outside the ATLAS active detector volume. The numbers and lengths of lines for each module (pixel, SCT) or front-end board (TRT) are tabulated for both the barrel and end-caps. Similarly, the cables used for the readout electronic bias as well as the silicon sensor and TRT straw high-voltage lines are also shown.

3080




Figure 47. The routing of data links and power supply cables from each side of the TRT to respectively the off-detector electronics and power supplies in the service caverns, together with the number, type and utilisation of the cables and optical links.




Figure 48. Detailed schematic of the pixel optical link architecture. For each pixel module, one (layer-1, layer-2 and the disks) or two (layer-0) optical fibres transfer data to the ROD's, and one fibre transfers the control and clock signals to the module (see text). The SCT optical links have a similar design.

4.4.2.1 Pixel and SCT readout The digital transmission of the clock and control commands to the pixel and SCT module and data from the module to off-detector electronics in the service cavern is made via optical links [75, 83, 84, 85]. A detailed block diagram illustrating the pixel link structure is shown in Fig. 48.

As summarised in Table 18 each pixel module uses one or two up-link fibres according to the required bandwidth; the bit rate is either 40 or 80 Mbits/s per fibre link. One down-link fibre per module is used to transmit trigger, timing, clock signals and configuration data. Opto-boards on the detector side convert the electrical signals from the module-control chip to optical signals.

The down-link uses a bi-phase mark encoded format to transmit both the 40 MHz bunch-crossing

Sub-detector	Link	Speed (Mbits/s)	Format	Links per module	Number of links
Pixels	Data	40/80	NRZ	1 (layer-1, 2 and disks)	2812
		40/80	NRZ	2 (layer-0)	
	TTC	40	BPM	1	1744
SCT	Data	40	NRZ	2	8176
	TTC	40	BPM	1	4088

Table 18. Summary of the main characteristics of the optical links used for the pixel and SCT data readout and trigger, timing and control (TTC) signals (NRZ stands for Non-Return to Zero and BPM for Bi-Phase Mark).

clock and the data. The bi-phase mark light signal, detected by a PIN diode, is decoded by a digital opto-receiver integrated circuit (the DORIC chip). In the up-link the module-control chip output is converted to light by a driver chip coupled to a vertical-cavity surface-emitting laser diode (VCSEL). This latter chip and the corresponding SCT chips have four channels each. They have been made in the same 0.25 μm CMOS technology as the module-control chip and have been produced on the same silicon wafers. Opto-boards can service either six (half-staves or disk-sectors) or seven (half-staves) modules depending on which part of the detector they are connected to. The opto-boards have been placed on service quarter-panels near the internal surface of the pixel support tube (see Section 4.7), to reduce the radiation dose collected over ten years of LHC operation.

In the case of SCT modules, the service harness that provides power to 4–6 modules also includes two data fibres and one trigger/control fibre for each module, as well as VCSEL and pin-diode opto packages for electrical to optical conversion together with their control ASIC's. Each SCT barrel module is serviced by a polyimide-aluminium low-mass tape which provides redundancy and electrically isolates the modules. For the end-caps, polyimide-copper tapes are used for the control lines and copper-clad aluminium twisted pairs for the power lines, from the module to the edge of the disk. The module is then connected via polyimide-copper tape to the PP1 patch-panel outside the SCT and from there by conventional cables to the power supplies. The services include two data fibres and one trigger/control fibre for each module plus VCSEL and PIN diode opto packages for electrical to optical conversion together with their control ASIC's.

On the other side of the link, back-of-crate (BOC) cards interface the opto-signals with the electrical signals in the readout driver (ROD). These cards are the same for the pixel and SCT detectors and they contain two kinds of electro/optical converter plug-ins: a RX plug-in with an array of PIN diodes and a data receiver, and a TX plug-in with an array of VCSELs driven by a bi-phase mark encoding ASIC. The SCT uses a radiation-hard step index multi-mode fibre for the full length from the back-of-crate cards to the detector. The pixel detector uses a long length of radiation-tolerant graded-index fibre, spliced to 7 m lengths of the radiation-hard fibre inside the

detector.

3120 **4.4.2.2 TRT readout** As described in Section 4.3.3, the TRT readout is segmented in 32 ϕ -
sectors in order to simplify the data transmission for L2 triggers. The initial 40 Mbits/s LVDS data
readout uses small custom designed twisted pair lines from the ASIC to the boards at patch panel
PP2, located after the first muon chambers. The data are then serialised in a Gbit serialiser [86]
and an electrical-to-optical conversion is made. The complete readout of the TRT requires 768
3125 1.6 Gbits/s optical links. Timing and control signals are electrically transmitted to and from the
TRT-TTC module.

4.4.2.3 Pixel, SCT and TRT readout drivers As described in Section 8.3.5, all readout driver
modules (ROD's) receive serialised data from the detector upon reception of a L1 trigger, perform
3130 de-serialisation, error checking, data compression (in the case of the TRT), and local event-building
and data-monitoring tasks. The ROD's combine the data received into a single event packet asso-
ciated to the L1 trigger. The event packets are transferred in a standard data format over a high-
speed optical link, the ATLAS S-link, to the readout system, which is common to all ATLAS
sub-detectors.

3135 Each pixel and SCT ROD crate is a 9U VME crate with up to 16 ROD's per crate. Each ROD is
paired with a back-of-crate card, which is plugged into the crate back-plane and provides the optical
to electrical interfaces [84, 87, 88]. The back-of-crate card can accept both 40 Mbits/s optical links
from pixel and SCT modules and 80 Mbits/s from pixel modules. Each ROD card services up to
48 SCT modules (96 data links) or between 6 and 26 pixel modules, depending on their geometrical
3140 location. The lower number of pixel modules serviced per ROD is due to the larger quantity of data
received from pixel modules at the smallest radius, due to their higher occupancy. The number of
modules which can be processed by one ROD is limited by the output bandwidth of the S-link. The
pixel and SCT ROD's are identical units except for the data-treatment algorithms. The ROD crates
receive clock, trigger and fast commands from the ATLAS TTC system through a TTC interface
3145 module (TIM), located in each crate which transmits these signals across the crate, using a custom
back-plane, to the ROD's and back-of-crate cards. Clock and command signals are then transmitted
to the detector modules using a simpler protocol over the optical fibres mentioned above.

The TRT off-detector electronics also uses two custom-designed TRT-ROD and TRT-TTC
modules [89]. Each ROD receives serialised data from 1/16 of one barrel side or 1/32 of one end-
3150 cap side, using eight optical links operating at 1.6 Gbits/s. After the optical-to-electrical conver-
sion, the ROD's perform the standard tasks described above, including a data-compression scheme
which does not introduce any data losses. The TRT-TTC module provides an interface between
the standard ATLAS TTC system and the TRT front-end electronics. It also feeds the ROD with
all necessary L1 trigger information (event ID, bunch crossing ID, trigger type). Each TRT-TTC
3155 module interfaces to the front-end using 40 links and to two ROD's via a dedicated back-plane
using the VME 9U P3 connector.

4.5 Electronics and detector power supplies and services

Both the sensors and the front-end electronics of the inner detector require substantial power for

		Pixel	SCT	TRT
Bias or high-voltage supplies	Voltage (maximum)	700 V	500 V	2000 V
	Voltage (nominal)	150–600 V	150–350 V	1600 V
	Current (maximum)	4 mA	5 mA	3 mA
	Segmentation	One per module 1744 channels	One per module 4088 channels	One per ~ 200 straws 1984 channels
Front-end electronics low-voltage supplies	Voltages	1.7–2.1 V analogue 2.1–2.5 V digital 0.8, 2.5, 10 V opto-device	3.5 V analogue 4 V digital 5, 10 V opto-device	±3 V analogue 2.5 V digital
	Current	3.7 kA	6 kA	6.5 kA
	Segmentation of bulk supply	Configurable 6–7 modules	One per module 4088 channels	1/32 of end-cap side 1/32 of barrel side
	Segmentation of regulated supplies	One regulator per module	See above	One regulator per 1/32 of wheel or per barrel module
Power	Front-end electronics power	6 kW	22 kW	22 kW
	Cables plus regulators:			
	- in ID volume	~ 4 kW	~ 5 kW	~ 3 kW
	- outside ID volume	~ 14 kW	~ 9 kW	~ 20 kW
	- total cable loss	~ 18 kW (~ 6V drop)	~ 14 kW (1.5–4.6 V drop)	~ 23 kW (3.8 V drop)
	Total power (initial operation)	~ 24 kW	~ 36 kW	~ 44 kW

Table 19. Summary of the silicon sensor bias voltage, straw-tube high-voltage and front-end electronics requirements and granularity. The different front-end electronic supply voltages are described in the text. The power consumption expected during initial operation is also shown. Following irradiation, the power consumption increases significantly in the sensors, front-end electronics and cables (see text). Additional power requirements to ensure thermal isolation in the ID volume are described in the text.

initial operation and the power load will increase during the high-luminosity operation. Table 19, together with Figs. 45 to 47, summarise the requirements for each of the pixel, SCT and TRT sub-detectors, in terms of bias voltage for the silicon sensors or high voltage for the straws (HV) and of low voltage (LV) for the front-end electronics, as well as the segmentation and routing used for the power distribution. The power distribution lines are connected at PP0 and PP1 on the cryostat wall, at PP2, which is external to the inner-detector volume and located after the first muon chambers, and at PP3 on the outside of the whole detector.

The pixel power supply system has four main components: the LV and HV power supplies, the regulator stations, and the supply and control for the opto-links. Two commercial LV supplies provide the analogue and digital parts of the front-end readout electronics. To protect the front-end electronics against transients, remotely-programmable and radiation-tolerant regulator stations are

3170 installed at PP2. Separate HV supplies are able to power the sensors up to 700 V with a maximum
module current of 4 mA. The worst-case expected module operating condition is a 600 V depletion
voltage with a 2 mA leakage current. The LV and HV lines are connected to respectively the
low-voltage and high-voltage patch-panels which distribute the power and monitor the currents of
individual lines. The supply and control of the optical link is a complex link in itself, consisting of
3175 three voltage sources and a control signal. It delivers the adequate levels for the operation of the
on-detector part of the optical link.

The SCT [90] maintains independent electrical services to each of the 4088 modules. Each
module receives multi-voltage LV channels providing power and control signals to the readout
chips, the optical electronics, timing and control electronics, as well as HV to the silicon sensors.
3180 The SCT LV power module provides the analogue and digital voltages for the SCT front-end
ASIC's and optical-link components. Each LV power module also includes the hybrid temperature
readout as well as the digital control lines. The HV bias provides a stable and controlled voltage
of 0–500 V with a current limit of 5 mA. The LV and HV modules are based on DC-DC converters
with a single power line serving each module; twelve identical 4-channel LV and six identical
3185 8-channel HV boards are housed in each crate. Each board in the crate is connected through a
system-interlock card to the detector control system (see Section 4.8.4). The four crates of each
rack are powered by four commercial power units (including one spare) connected in parallel to
provide a 48V DC supply. As an example for the barrel SCT, the power distribution lines are in
four parts: low-mass polyimide tapes for the innermost region of the detector, which are connected
3190 to PPB1 on the cryostat wall, thin cables along the cryostat wall, which are spliced to medium-size
cables outside the cryostat, and, after PP3, thick cables going to the power-supply crates. Because
of the large voltage drop in the cables (up to 4.6 V), the supplied voltage exceeds the maximum
allowed voltage of the ASIC's. This could be problematic, for example in case of a drop in the
power consumption of the front-end system. To avoid such incidents, a voltage-limiter circuit has
3195 been introduced at PP3. To dampen common-mode noise, inductors on all power and low-current
control and monitoring lines are also housed at PP3.

The TRT front-end electronics require three low-voltage power supplies (+2.5 V digital and
 ± 3 V analogue). Commercial supplies deliver power to boards located at PP2. The PP2 boards
house radiation-tolerant voltage regulators delivering power to each front-end board. The size of
3200 the cables feeding the PP2 power boards and the front-end boards is a compromise between the
available space and the power dissipation in the cable trays allowed by the cooling system. The
TRT straw high voltage is nominally 1530 V. Sets of about 200 straws are powered by a single
commercial HV source able to deliver up to 3 mA at 2 kV. A total of 1984 HV channels are needed
for the whole detector. Standard multi-wire HV cables are used up to the level of PP2, then custom
3205 miniature HV cables are used to reach the detector.

The power dissipation during initial detector operation is also indicated in Table 19. The initial
power dissipation from the front-end electronics and cables is ~ 62 kW within the ID volume,
and ~ 104 kW in total. However, during high-luminosity operation of the TRT and after irradiation
of the pixel and SCT sensors, several kilowatts of power are dissipated by the sensors (this is
3210 sensitive to temperature for the silicon sensors). Together with additional losses in the cables and
increased front-end power consumption, the total power loss within the ID volume is expected to
exceed 85 kW after several years of operation.




Figure 49. Schematic top view of the inner-detector sub-systems with their vertical support points. All supports are on the plane of the detector axis and symmetric with respect to this axis.

Additional power supplies are required for the pad heaters located on the thermal enclosures surrounding the pixel and SCT detectors (see Section 4.7.2) and for the heaters which are an integral component of the evaporative cooling system for the pixel and SCT detectors (see Section 4.8.3). The thermal enclosures require a total power of almost 30 kW, including the heaters surrounding the beam-pipe (see Section 4.8.1). The heaters for the cooling system require a total power of 126 kW, which is provided by four racks situated in the service caverns. These power supplies are controlled by programmable logic controllers, which regulate the entire system and manage, in particular, transients caused by changes in the power load as the detector electronics are switched on or off.

4.6 Grounding and shielding of the inner detector

The three components of the inner detector, the pixels, the SCT and the TRT are designed to be electrically isolated both from earth and from each other, as is generally required for all ATLAS components (see Section 9.4.8). The detectors are therefore floating by design. This provides the possibility of a controlled ground connection at a location identified as IDGND between the tile calorimeter and the first layer of barrel muon chambers. A star connection of all the safety grounding cables for ID components, as well as common items (support structure, pipes for cooling and gas, heaters, and DCS sensors), ends in IDGND. The IDGND is connected to the ATLAS equipotential network by means of a dedicated 95 mm² conductor. To provide reasonable isolation of the sub-detectors and at the same time to complete the requirements of the safety grounding, the connection to IDGND from any sub-detector is made from its electromagnetic shield using a single-point connection. The power supplies are floating and are referenced in most cases on the detector modules (analogue and digital grounds) and at the PP0 and PP1 patch-panels (connection of the detector-module grounds to the electromagnetic shields).

Each pixel module is referenced at PP0, and all the PP0's are grounded with both the electromagnetic shield and the beam-pipe shield at PP1. The pixel detector itself is surrounded by a Faraday cage defined by an aluminium layer surrounding the beam-pipe, and an aluminium layer surrounding the outside of the pixel support tube. The PP1 end-plates close the Faraday cage at the two ends and serve as the common grounding point for the pixel detector. There is a single




Figure 50. The barrel support structure of the TRT, which serves as the support for the full barrel ID, shown during initial assembly and measurements. The SCT detector is supported on carbon-fibre rails inside the carbon-fibre inner cylinder. The two-piece outer carbon-fibre cylinder is not assembled yet at this stage of the integration. The space-frame geometry at each end was designed to support individual TRT barrel modules. A number of mirrors, visible as bright spots, are photo-grammetry targets used for survey measurements.

3240 grounding connection to IDGND on side-A of the detector. The SCT barrel and end-cap module
 assemblies are each enclosed by outer electromagnetic shield skins. The individual module supply
 tapes, outer supply cable shields and detector array shield skins are electrically bonded at PPB1
 (for the barrel) and PPF1 (for the end-caps). The common shield nodes of the barrel and of each
 3245 barrel modules and wheels are fully enclosed in electromagnetic shields which are connected to
 IDGND using individual grounding cables.

4.7 Structure and mechanical integration of the inner detector

The complexity of the individual ID sub-detectors, together with the integration and installation requirements, the different sub-detector operating temperatures and the tight stability requirements,

3250 have resulted in a complex support scheme illustrated in Fig. 49. Most of the 66 support points are in the axial horizontal plane [62, 91, 92].

The ID system is supported by rails fastened to the barrel cryostat inner wall. The ID sub-detectors rest on these cryostat rails and are basically supported at four points, implementing a quasi-kinematic support scheme, with simple support conditions in all directions except the vertical.

3255 The mechanical support for the barrel SCT system and for the barrel TRT modules is the barrel support structure, shown in Fig. 50. It is designed for high stiffness and stability, with $<10\ \mu\text{m}$ displacements under the expected temperature and humidity variations. It consists of two 21 mm thick carbon-fibre space frames, joined by inner and outer carbon-fibre cylinders. The SCT and TRT end-caps are each supported from a pair of girders, sliding on the cryostat rails. The separate 6.6 m long pixel and beam-pipe package includes the pixel support tube, which slides inside the SCT and is itself supported by the SCT barrel. Two end-plates provide the external supports for the beam-pipe and the pixel support tube.

3260 Prior to integration as part of the full ATLAS detector, the barrel and end-caps for each of the pixel, SCT and TRT sub-detectors were separately assembled and fully tested on the surface. This section describes the overall mechanical structure of each of the sub-detectors. The subsequent integration of the ID in ATLAS comprises four steps which are also described in this section: the integration of the barrel SCT and barrel TRT, the integration of the end-cap SCT and end-cap TRT (two end-caps), the integration of the barrel and end-cap pixel detectors with the beam-pipe, and finally the insertion of the pixel package.

4.7.1 Pixel structure and integration

The pixel detector and the pixel support tube (PST) within the ID are shown schematically in Fig. 28. The detector with its associated services and the vacuum inner detector beryllium beam-pipe (see Section 4.8.1) are precisely located inside the PST. The pixel services (cooling, power and monitoring) are routed to the ends of the PST.

3275 The active region of the pixel detector is shown in Fig. 51. The parameters of the pixel detector, with its three barrel layers and two end-caps, are listed in Table 12 of Section 4.3. The total active area of silicon is approximately $1.7\ \text{m}^2$, with 112 barrel staves and 48 end-cap sectors (eight sectors per disk).

3280 In the barrel region, the bi-staves are mounted in half-shells as illustrated in Fig. 52. Each half-shell is a thin carbon-fibre shell formed with facets to match the number of staves, with cut-outs to reduce the mass and with mounting rings that position the staves at five locations. The location of the staves in the half-shells was measured using a coordinate measuring machine to reference the position of modules to known features on the shells. The disks of the pixel end-cap detector are bolted with precision bushings to a carbon-composite support ring. The disks are then held with four mounts within a section of the octagonal support frame to form an end-cap. An end-cap during the final stages of assembly is shown in Fig. 53, after connection of the cooling circuits. Two fully-loaded half-shells form a barrel layer. The largest layer in size, layer 2, is shown after this step in Fig. 54. Each barrel layer is inserted into the supporting octagonal frame and connected to end-cone structures with fingers to mate precisely with mounting brackets on the barrel. Capillaries and outlet cooling-tube extensions are then added.

3290




Figure 51. A perspective cut-away view of the pixel detector. The view shows individual barrel and end-cap modules, supported with their associated services on staves and disks within an octagonal support frame.




Figure 52. A pixel barrel half-shell, with its cutouts, being loaded with barrel bi-staves and services.




Figure 53. A pixel end-cap at the last stage of assembly, after connection of its cooling circuits.

Following the integration of the beam-pipe, the barrel octagon and the two end-cap octagons, the barrel cooling pipes and cables are passed over the outside of the end-cap frame. All end-cap services are on the inside of the frame.

3295

The PST itself consists of three sections. Each section is a cylinder with external stiffening rings and precision rails. The barrel section is made from carbon-fibre composite. Each end-section is composed of carbon fibre and fibreglass composites. The rails are carbon-fibre composites and are accurately positioned within each cylinder. The cylinders are joined at each end of the barrel by bolted carbon-fibre flanges. Heater panels (copper-on-polyimide printed-circuit boards) are glued to the cylinders. These are activated if there is a failure in the dry environmental gas around the pixel detector.

3300

The barrel PST section is precisely located with respect to the barrel SCT structure. Mounting points on the barrel PST receive mounts on the pixel detector support frame described below and locate the pixel detector to approximately $100\ \mu\text{m}$ with respect to the barrel SCT. An octagonal carbon-composite frame supports and positions the barrel and the two end-caps. The two end-cap sections are joined with composite bolts to the barrel frame. Mounts which position the pixel detector within the PST are located on the end-cap end-plates.




Figure 54. Barrel pixel layer-2, loaded with bi-staves, viewed along the axis after the joining of the half-shells.

Each pixel module receives LV power directly from the LV regulators located at PP2 (see Section 4.5). There are in total 88 cooling circuits for the pixel detector. Certain monitoring functions require connections within the pixel detector volume itself. The electrical and cooling connections pass through the end regions of the PST. Since the pixel detector is inserted from one end, all services, including connectors, must fit within the 230 mm radius PST inner envelope. The electrical services and cooling pipes are contained within service quarter-panels, which deliver one-quarter of the required services at each end of the pixel detector. Power and monitoring wiring is routed by individual twisted pairs and soldered to printed circuit boards (PP0) near the active detector, which contain miniature connectors. The low-mass cables from each module are plugged into these connectors. At the other end (PP1) the wires are soldered to printed circuit boards to penetrate the sealing plate. On the outside of the PP1 region, other twisted pairs are soldered to the printed circuit boards and terminated in commercial connectors. Optical transceivers are also located in the PP0 region and convert electrical signals from the detector to light transmitted by fibres to connectors at PP1. Control signals from outside the detector are also transmitted by fibres to the transceivers for conversion to electrical signals (see Section 4.4).

Each cooling circuit includes a custom heat exchanger that consists of an inlet and an outlet tube glued together along the length of the service quarter-panel. These tubes penetrate the plate at PP1 and dry-gas integrity is maintained by a bellows seal that also allows for the 2–3 mm contraction of these aluminium pipes when the detector is operating. The capillaries and outlet extensions are attached at the PP0 end.

The service quarter-panels and the beam-pipe are supported by a composite beam-pipe support structure. The beam-pipe supports are adjustable from the ends of the PST to position the beam-pipe. The overall detector integration is illustrated in Fig. 55.

The pixel detector is sensitive to the high instantaneous rates that might occur during accidental beam losses [93]. For this reason, a set of small, fast and radiation-hard diamond detectors, called the ATLAS beam-conditions monitor (BCM) has been built and integrated into the pixel package to




Figure 55. The pixel detector during integration of the barrel, end-caps and their services: (a) the end-cap region; (b) the barrel detector region; (c) Patch Panel 1 (PP1) region; (d) Patch Panel 0 (PP0) region and (e) region of the optical transceivers on the service quarter-panels. See text for details.

monitor the beam conditions and to distinguish lost beam particles from proton-proton interactions.
 3345 The BCM, described in more detail in Section 3.4.1, is shown as installed near the beam-pipe
 in Fig. 24.

4.7.2 SCT structure and integration

As shown in Fig. 28, the SCT consists of 4 coaxial cylindrical layers in the barrel region and two
 end-caps each having 9 disk layers [61]. Tables 14 and 15 of Section 4.3 show the SCT detector
 3350 parameters in detail.

The low-mass barrel cylinders are designed to be extremely stable to both temperature and
 humidity variations, and to long-term creep [92]. They are made from three-layer (0° , $+60^\circ$, -60°)
 carbon-fibre skins of $\sim 200 \mu\text{m}$ total thickness over a carbon fibre/cyanate ester honeycomb core
 to form a 6 mm sandwich. The cylinder ends are closed with flanges, incorporating holes that
 3355 are machined to high precision. Pads for the precision mounting of module brackets are attached
 to the surface and both the surface and a precise mounting hole are machined to within $\pm 20 \mu\text{m}$
 accuracy. A similar precision is specified for the inserts of an alignment system mounted on each
 barrel, and for the machined holes on the end flanges. Because of poor adhesion for a few pads, all
 the pads were subsequently attached using small stainless steel screws and plastic bushes, with a
 3360 slight precision loss. The external radius of all cylinders has been maintained to within 1 mm.




Figure 56. The upper photograph shows a complete SCT barrel with all its modules mounted. A blown-up detail of some of these barrel SCT modules mounted on the support cylinder, together with the module services including the polyimide signal and power cables, and the cooling tubes is also shown (bottom left). A drawing of the mounting brackets, which are attached to the barrel SCT cylinders (in this case, the innermost barrel cylinder), and of the attachment of the module and cooling pipes to the bracket, is also shown (bottom right).

The barrel modules [68] are mounted in rows of 12, on individual carbon-fibre brackets, as depicted in Fig. 56. The module is rotated by ± 20 mrad, alternating from barrel to barrel, to align the strips of one side along the cylinder axis. The modules are attached to the support structure at three points, two on the beryllia facing (cooling side) and one on the far side [94]. When mounted on the barrel, the variance of the mounting precision in z as measured by the mounting robot is $60 \mu\text{m}$, but there has been no systematic survey of the variance in $R-\phi$. To avoid HV breakdown, a distance of more than 1 mm is maintained between the sensor edges and any ground potential. The centres of adjacent modules in each row are radially separated by 2.8 mm using the tiling arrangement.

The end-cap disks support end-cap modules, as illustrated in Fig. 57, with tight stability and accuracy requirements, together with their electrical, mechanical and alignment services [95]. The 8.7 mm thick disks consist of carbon-fibre face skins ($200 \mu\text{m}$ thick) with an aramid/phenolic honeycomb core. The orientation of the carbon-fibre skins and the choice of materials minimise the effects of thermal and humidity changes. Individual modules are attached to cooling blocks




Figure 57. End-cap SCT modules mounted on end-cap SCT disk with outer and inner modules (left) and middle modules (right).

3375 held by inserts glued to the disk. A large cooling block at the hybrid end (230 mm^2) defines the
 position of the module, while a slot at the end of the module defines the $\pm 20 \text{ mrad}$ rotation of
 the module when attached to the smaller block (78 mm^2), to within $\pm 1 \text{ mrad}$. The RMS spread
 of the surveyed module in-plane placement positions on each disk was $10 \text{ }\mu\text{m}$ and the placement
 3380 uncertainty of the disk within the support cylinder is estimated to be $50\text{--}100 \text{ }\mu\text{m}$ ($< 1 \text{ mm}$ in the
 z -direction).

The disks of an end-cap are supported at 12 points around their diameter by springs in the
 support cylinder. The springs are soft in the radial direction but otherwise stiff, allowing for the
 radial expansion of the disks and cylinder. The cylinder composition is similar to that of the disks.
 Services leave the disks through apertures in the cylinders and run along the cylinders before exiting
 3385 at the far end of the end-cap thermal enclosures. Each support cylinder is in turn supported by two
 flat composite panels (of similar construction to the cylinder). These panels rest on the same rails
 as those which support the TRT.

Prior to the mounting of SCT modules on the support structures, each barrel or disk was
 equipped with electrical services, optical services and cooling loops. In the barrel, each cooling
 3390 loop, with two inlets and one common exhaust, services four rows of 12 modules. The loops are
 mounted on the module mounting fixtures and connected to the modules using thermal grease. For
 the end-caps, each cooling loop is attached to the disk and traverses the module mounting blocks
 for all the modules within a disk quadrant.

Each service assembly and each cooling loop were tested and all modules of a given barrel
 3395 or disk were powered and read out to verify the full module functionality [96]. Less than 0.5%

of modules needed intervention. For a final test of completed barrels and disks, test systems were constructed to operate and read up to one million channels simultaneously [88, 97]. After assembly, individual barrels were transported to CERN for final integration.

3400 The integration of individual barrels and disks differ due to their different support structures. Both the barrel and end-cap are surrounded by low-mass outer and inner thermal enclosures (see Section 4.8.2). Their role is to prevent condensation during operation by maintaining a low-temperature and low-humidity N₂ environment, to prevent the out-flow of N₂ gas surrounding the SCT modules, which would affect the TRT gas-gain and performance, and to prevent the in-flow of the ID environmental gas (CO₂). They also provide a Faraday shield to protect the SCT from external electrical noise. The outer thermal enclosures and the end-surfaces are covered with resistive pad heaters to ensure thermal neutrality.

3410 To complete the SCT barrel from four individual barrels, the outer thermal enclosure was mounted into a support cradle and the four SCT barrels sequentially inserted and fastened together using eight radial interlinks at each end, having dowel pins matching the precisely machined holes in the barrel flanges. The deformation of the cylinders was measured to be below 90 μm (RMS). Within the measurement accuracy of 20 μm (lateral) and 40 μm (along the axis), the four barrel SCT cylinders are concentric with their axes aligned. All services were sealed into slots of the outer thermal enclosure feed-through on the outer circumference of the SCT barrel. Inlet cooling capillaries, outlet exhaust cooling pipes and ground-reference connections were added between individual barrels and thermal-enclosure bulkheads. Finally the outer and inner thermal enclosures were sealed and the air-tightness of the barrel enclosure was checked.

3420 For the end-cap, each assembled disk was inserted into the overall carbon-fibre support cylinder [95]. Polyimide power tapes, optical fibres and cooling pipes were connected at the outer disk circumference, and run along the outside of the cylinder to a patch-panel at the end. Each completed end-cap assembly was fully tested and characterised. After assembly, each of the end-cap cylinders was transported to CERN for final integration and testing.

3425 Finally, the SCT is equipped with a geodetic grid of 842 interferometers to monitor real-time deformations at the scale of a few μm. Distances between nodes attached to the structure are measured simultaneously using frequency-scanning interferometry [63] to a precision of < 1 μm. The three-dimensional grid node positions can be reconstructed to better than 5 μm in the critical direction [98], thus offering access to short time-scale, low spatial frequency detector deformations, which may otherwise only be weakly constrained with data.

4.7.3 TRT structure and integration

3430 As indicated in Section 4.3.3 and Fig. 28, the TRT occupies the outer radial regions of the inner detector. There are three module layers with axially oriented straws in the barrel region [99] and 20 wheels with radial straws in each of the end-cap regions [100]. The active regions of each detector are shown in Table 17 (see Section 4.3.3).

3435 Each of the 96 barrel modules is supported at each end by the barrel support structure, which also provides the required overall module stability. It is a 21 mm thick carbon-fibre disk, machined to a triangular strut array and attached to two thin inner and outer carbon-fibre cylinders (see Fig. 50 and Fig. 58).




Figure 58. End-view of the TRT barrel structure, during the final attachment of cooling and electrical services.




Figure 59. A completed TRT end-cap during the final service integration, showing (from the left) twelve type-A wheels and eight type-B wheels surrounded by their services and supporting rings.

As described in Section 4.3.3, each of the TRT end-caps consists of twelve type-A wheels with a straw-layer spacing of 8 mm and of eight type-B wheels with a spacing of 15 mm. The geometry of the lower-density type-B wheels maintains the required number of straws crossed by a particle from the interaction point as well as keeping the material of the active detector approximately constant as a function of η .

Each eight-plane end-cap wheel consists of two basic blocks of four-plane wheels mounted in an inner and outer electrically grounded carbon-fibre ring. The rings and the straws serve as the mechanical support structure of the wheels. Each eight-plane wheel is covered with a thin metal-clad polyimide membrane on each front side connected at the inner radius and providing a signal-return path with defined electronics ground. Openings at the outer radius of the electronic shield allow a path for the CO₂ cooling gas.

As shown in Fig. 59, the type-A and type-B wheels are assembled in two independent groups and the individual wheels are held together using tie-rods between two solid carbon-fibre membranes of 10 mm thickness. The end-membranes of a stack are supported on rails fixed to the ATLAS barrel cryostat. Each group is sealed at the inner radius by a glass-fibre cylinder covered with a thin copper-clad polyimide foil. This cylinder serves as both an electrical Faraday shield and a mechanical link keeping the end-membranes parallel.

The end-cap CO₂ gas circulating between and within sectors is cooled by heat exchangers positioned at the outer ends of the straws between each second eight-plane wheel of type-A and each eight-plane wheel of type-B before flowing to the next wheel of the end-cap. The leak-tightness of both end-caps for the CO₂ flow has been verified. To prevent pressure changes of the cooling gas which might cause mechanical stress and deformations (leading to discharges) of the straws, a set of passive safety valves has been installed on each group of wheels.

Supplies for high-voltage, active gas, electronics cooling and the routing of the detector signals to the back-end system are organised in sectors covering 1/32 of all the wheels of one end-cap.




Figure 60. Insertion of SCT barrel into the TRT barrel. The three module types of the TRT barrel are clearly identified. The SCT outer thermal enclosure is visible, together with the barrel services extending on support frames from each end.

Along a given sector, all pipes and cables are brought in cable trays from outside the ATLAS cryostat, to the PPF1 patch-panel and distributed from there to the individual wheels.

4.7.4 Integration and installation of the inner-detector components

3465 The completed SCT barrel and end-caps were finally inserted into the corresponding TRT sub-
 3470 detectors. The barrel procedure is described here: the end-cap integration followed a similar pro-
 cedure. The SCT was supported on a cantilever frame (foreground of Fig. 60). The completed
 TRT was transferred into the final support and lifting frame (background of Fig. 60). The TRT
 was guided on rails over the SCT. During the movement the mechanical alignment and electrical
 isolation of the sub-detectors were verified. After insertion, the SCT was positioned on rails inside
 the TRT with a precision of $\sim 250 \mu\text{m}$. The final survey of the SCT barrel with respect to the
 TRT barrel shows a displacement of their axes by -0.29 mm (horizontal) and 0.2 mm (vertical) on
 side A, and -0.55 mm (horizontal) and 0.45 mm (vertical) on side C. The precision of this survey
 was $\pm 0.15 \text{ mm}$. The barrel positions are in good agreement with the global alignment found in a
 later cosmic run.

The pixel barrel and end-cap integration into the pixel support tube has been described above. The following elements were finally installed on horizontal rails inside the inner warm vessel of the cryostat: the SCT/TRT barrel, the two SCT/TRT end-caps and finally the pixel package (see Fig. 49).

3480 The result of the mechanical survey of the components is shown in Table 20 with respect to the

Barrel	x (side A)	y (side A)	x (side C)	y (side C)
	(mm)	(mm)	(mm)	(mm)
Pixel	+0.3	-1.2	-0.4	-0.9
SCT	+0.53	-1.42	+0.3	-0.9
TRT	+0.2	-1.6	-0.2	-1.35
End-cap (side A)	x at z_{\max}	y at z_{\max}	x at z_{\min}	y at z_{\min}
	(mm)	(mm)	(mm)	(mm)
SCT	+0.34	-1.14	-0.21	-1.50
TRT	+0.84	-2.13	+0.27	-1.60
End-cap (side C)	x at z_{\max}	y at z_{\max}	x at z_{\min}	y at z_{\min}
	(mm)	(mm)	(mm)	(mm)
SCT	+0.28	-1.48	0.92	-2.04
TRT	-0.35	-0.97	-0.85	-1.00

Table 20. Mechanical position of the axes for the installed barrel and end-cap assemblies with respect to their nominal positions relative to the inner warm vessel. The positions do not include the expected rail sag of -0.35 mm at the interaction point and of ~ -0.2 mm at the position of the patch panel (PPF1). In addition, for reasons of space, each end-cap has been placed 5 mm further away from the interaction point than its nominal position.

inner warm vessel on which the rails are fixed. The z -positions of the pixel and barrel assemblies are well-centred. However, the z -positions of the end-caps are displaced away from the interaction point by 4.88 mm (side A) and 5.35 mm (side C) from their nominal positions. A more complete description of the survey strategy and measurements and of the positioning accuracies achieved for the major components of the ATLAS detector is given in Section 9.3.2.3.

4.8 Inner detector environment and cooling services

4.8.1 Beam-pipe interface and operational aspects

The vacuum inner detector (VI) beam-pipe has a total length of 7100 mm, an inner radius of 29 mm and a nominal outer radius of 34.3 mm. The interfaces between it and the other components of the beam-pipe in the ATLAS experiment are described in Section 9.8. To reduce the amount of material to an absolute minimum, the VI beam-pipe has been manufactured from beryllium with a thickness of 0.8 mm. Figure 61 shows the structure of the completed VI beam-pipe as it was inserted into the pixel detector. Several additional layers of material have been added around the beryllium wall to provide the possibility for bake-out in situ (see Section 9.8).

The VI beam-pipe is supported by the pixel detector (see Section 4.7.1) and adjusted after insertion to be straight and on the nominal beam axis to within 3 mm. Each end was surveyed,




Figure 61. Cross-section of the beryllium vacuum pipe with its various layers for bake-out in situ. Dimensions are in mm.

then positioned to this target, and with an optical device developed to straighten the pipe, the mid-supports, internal to the pixel package, were externally adjusted to straighten the beam-pipe between its two end-flanges. It was then surveyed again and the final values recorded show that the beam-pipe is within 1 mm of the centre of the pixel package.

Other than the envelope and mechanical interfaces, the beam-pipe has environmental interfaces to the pixel package. The beam-pipe has an electromagnetic interference shield, namely a 50 μm thick aluminium foil wrapped around its outer radius, which is also shown in Fig. 61. This foil is electrically terminated to the pixel-system ground, where the pixel gas envelope is sealed to the beam-pipe, at the flange of the PP1 bellows seal at each end of the pixel package. These bellows close both the pixel gas volume and the Faraday cage to the beam-pipe.

As shown in Fig. 61, the beam-pipe has integral bake-out heaters laminated to the outer skin of the beryllium tube inside ~ 4 mm of aerogel insulation. With the beam-pipe at bake-out temperatures of up to 230°C , a heat load of approximately 100 W/m is dissipated into the pixel volume from a surface with a temperature of up to 100°C . The beam-pipe will be baked out only when the pixel detector modules are not operating. During bake-out, it is expected that the complete pixel cooling system would be operational, in order to ensure that no portion of the detector becomes too hot. In addition, there is a dedicated interlock system which will monitor all of the layer-0 module temperatures, as well as cooling-system faults, and will switch off the beam-pipe bake-out heaters in case of any danger to the detector. The cooling of the vertexing layer remains essential to ensure that its modules remain cold, which will prevent damage, particularly after significant exposure to




Figure 62. Schematic breakdown of the environmental conditions inside the inner-detector volume: in yellow are shown the TRT volumes, in light blue the SCT volumes, and in green the pixel volume inside the pixel support tube. Also shown is the routing of the cold tubes bringing C_3F_8 coolant to the pixel and SCT volumes, as well as the points where CO_2 gas is flushed into the ID volume (see text).

irradiation.

4.8.2 Inner-detector environmental gas

The environmental conditions under which the ID sub-detectors operate are rather different both in terms of temperature and of environmental-gas requirements, as shown in Fig. 62.

The pixel and SCT silicon detectors operate at a low temperature of approximately $-7^\circ C$ to minimise the impact of radiation damage, while the TRT detector operates at room temperature. The cooling systems required to operate these detectors in a safe and stable way at LHC design luminosity are described in Section 4.8.3. This section describes the environmental gases required and the technical solution implemented to maintain a stable environment in the overall inner-detector volume.

The environmental gas for the silicon detectors is dry nitrogen at temperatures of between $-7^\circ C$ and $0^\circ C$ to avoid condensation on the cold surfaces. The nitrogen, chosen because it is inert, has a small flow and does not provide any cooling power. In contrast, the TRT detector requires an overall CO_2 envelope gas acting as a barrier between the active part of the detector and the environment, to prevent contamination from nitrogen and possibly other sources of pollution. As described in Sections 4.3.3 and 4.7.3, the TRT barrel modules are ventilated with a flow of $\sim 3 \text{ m}^3/\text{h}$ of CO_2 . The cooling of the end-cap straws requires on the other hand the operation of a dedicated cooling system providing CO_2 flows up to $50 \text{ m}^3/\text{h}$.




Figure 63. Details of the inner detector end-plate (see text).

3535 The different operating temperatures and gases, as well as the need for thermal and gas neutral-
 ity between the different sub-detectors, has resulted in the definition of the independent units shown
 in Fig. 62, with specific thermal enclosures and environmental gas flows wherever required. Details
 are given for each sub-system in Sections 4.7.1, 4.7.2 and 4.7.3. Because of space constraints, the
 thermal enclosures must achieve thermal neutrality with minimum material. They are covered by
 3540 pad heaters to maintain the outer temperature of the cold volumes above the dew point, and, over
 the area facing the TRT, to minimise heat losses and avoid cooling of the warm TRT structures.
 The humidity and temperature of the gases are constantly monitored as described in Section 4.8.4.

To avoid condensation, primarily around the cold pipes of the silicon-detector cooling system,
 and to minimise the amount of water vapour present in the overall ID volume, the latter is sealed at
 3545 each end by two end-plates and the gaps between the various ID sub-detectors are flushed with CO₂.

As shown in Fig. 63, the end-plates consist of a large segmented double-skin aluminium plate,
 extending up to a radius of 2.5 m, which is sealed around the beam-pipe and at the 32 service feed-
 throughs located at the outer radius of the cryostat flange. The ID services, exiting from the PPF1
 3550 region and from the pixels and the beam-pipe at smaller radius, are routed between the two skins
 and form the core of this sandwich-type structure.

The inner portion of these end-plates seals the ID up to the LAr cryostat bore diameter of 2.3 m;
 it consists of a structural cross which is embedded in the core and holds the support nose and
 the four beam-pipe supports. The nose in itself provides the fixation points for the pixel support

3555 tube as well as an independent sealed environment for the pixel patch-panels. Additional local supports have been added between the nose and the cryostat rails to provide the required stiffness of the assembly during pixel insertion. The plates on the back of this portion of the end-plates are equipped with cooling circuits to extract the heat load coming from the pixel services and to preserve the thermal neutrality of the ID environment. The whole surface of the front skin, which seals the cold pixel environment from the ambient air of the cavern, is covered by electrical heater pads to avoid possible condensation.

3560 The front skin of the outer portion of the end-plates, with a diameter between 2.3 m and 5 m, consists of the outer wall of the cryostat flange itself. In contrast, the outer skin is made by several panels sealed together for easier assembly. These panels are equipped with dedicated openings for local maintenance access to the evaporative heater and splice-box areas during short access periods to the detector.

The distribution of the CO₂ flushing points inside the inner-detector volume has been optimised to achieve a nominal dew point of -30°C in about three hours with a total CO₂ flushing rate of 5 m³/h. Safety valves are installed at the outer edges of the end-plate gas seal to prevent the differential pressure between the inside of the ID volume and the outside from exceeding ± 0.5 mbar.

3570 4.8.3 Inner-detector cooling

As described in Table 19 and Section 4.5, the cooling systems of the ID [101] must remove up to ~ 85 kW of heat at the LHC design luminosity from the ID volume. The pixel and SCT sensors operate at $\sim -7^{\circ}\text{C}$ while the TRT operates at room temperature. The temperature stability within the ID volume must be maintained within $\pm 2^{\circ}\text{C}$. In addition, the ID power cables must be cooled and maintained at approximately room temperature as they exit the active detector volume. The cooling systems must be robust and have a high reliability over the lifetime of the experiment.

3580 As explained in Section 4.3.3, the heat from the TRT straws is removed in the barrel by the same room-temperature C₆F₁₄ coolant as that used for the front-end electronics. For the end-cap wheels, where the geometry is quite different, the heat from the straws is removed by a forced flow of CO₂ gas. The heat exchangers between the TRT end-cap wheels and the front-end electronics boards are cooled by a standard mono-phase cooling system using a room-temperature C₆F₁₄ fluorinert coolant. An identical cooling system is also used to cool the pixel, SCT and TRT cables outside the ID volume (and also in part those inside the ID volume), as well as the electronic boards housed in the PP2 crates (see Section 4.4.2).

3585 The evaporative cooling system adopted for the pixel and SCT detectors operates in a way similar to those used in the refrigeration industry. Nevertheless, the safe operation of a large number of sensitive and expensive semiconductor detectors and front-end electronics in the difficult and high-radiation environment of the LHC imposes new and non-standard constraints on the cooling system [102]. The particular choice of coolant (C₃F₈ or octafluoropropane) [103] is motivated by the requirement that it be non-flammable, electrically non-conductive, chemically inert in the event of a leak, and very stable against the high radiation doses expected at the LHC. In addition, the absolute pressure of the two-phase fluid as it flows in the cooling structures on the detector is between 1.7 and 6 bar, which is reasonably well matched to the requirement of minimal material for the cooling structures.




Figure 64. Schematic of the evaporative cooling for the barrel SCT. The heaters are located just outside the ID volume and the pressure regulators are located in the cooling racks at the periphery of the ATLAS detector.

3595 While different in detail, the general features of the individual cooling circuits of the pixel, barrel and end-cap SCT are conceptually the same: as an example, Fig. 64 shows the layout of one barrel SCT circuit. Each of the 204 circuits serving the detector has a fixed flow. The fluid is delivered to the detector in liquid form at room temperature, at an absolute pressure between 11 and 14 bar, from four distribution racks situated on the detector platforms in the underground cavern. 3600 Each rack serves one quadrant of the inner detector. The delivery pressure is regulated at the distribution racks by pneumatically-controlled pressure regulators. The fluid travels approximately 30 m to the detector via small (4 mm inner diameter) copper tubes.

3605 Inside the detector volume, the liquid enters a small-diameter capillary and the pressure starts to fall. The fluid emerging from the capillary is in two-phase form and starts to boil, lowering the temperature of the thin-wall cooling structures connected to detector modules. The pressure of the coolant, which in turn determines the evaporation temperature, is set by a pneumatically-controlled back-pressure regulator in the exhaust line. The back-pressure regulator is also mounted on the distribution racks remote from the detector. The operating pressure on the cooling pipework is typically varied between 6.4 bar (absolute pressure for an evaporation temperature of 14°C) for warm commissioning operation to 1.7 bar (absolute pressure for an evaporation temperature of -25°C) for normal operation. 3610 The capillary size has been chosen to deliver a given mass flow which depends on the particular detector requirements. Typical mass flows for the different detectors vary between 2.9 g/s for pixel disks to 8.9 g/s for the barrel SCT.

3615 Under normal operation, the fluid in the pipes which emerge from the detectors into the environmentally-controlled inner-detector volume is cold, at a temperature of -25°C. This cold exhaust is used to pre-cool the incoming liquid to maximise the available cooling capacity on the detector structures. This is done in a counter-flow heat exchanger, which is mounted outside the detector structure but inside the cryostat, where the inlet and exhaust pipes are allowed to come

into direct contact. Before the residual two-phase fluid is allowed to exit the environmentally-
3620 controlled area via the large diameter (12–14 mm inner diameter) pipes, it must be heated above
the dew-point of the cavern to avoid condensation on the exhaust pipes as they return to the rack.
This is achieved by an in-line heater, where the fluid is boiled and super-heated. The heater element
is a hot wire which is in the flow of the C_3F_8 cooling fluid. The power-rating for the heaters varies
between 250 W for the pixel disks to 1 kW for the barrel SCT. The 110 V power to the heaters is
3625 pulsed (on–off) with a regulated feedback mechanism to maintain a constant 20°C temperature of
the exhaust gas. The temperature of the liquid is measured by one of a pair of redundant thermo-
couples mounted either in the liquid flow or directly on the exhaust pipe. A hardware protection
system (see Section 4.8.4) protects the heaters against the possibility of over-heating by switching
off the supplied power when the temperature measured on the body of the heater exceeds a pre-set
3630 value (typically 55°C).

4.8.4 Inner detector controls, safety and interlocks

The ID-specific detector control system (DCS) operates within the general context of the overall
ATLAS DCS, as described in Section 8.5. It controls the detector powering and monitors and
controls the environmental parameters, in particular the pixel and SCT cooling, and the TRT gas
3635 gain. The sub-system-specific environmental monitoring is similar for all three sub-systems and is
described briefly below.

The evaporative and mono-phase cooling systems are controlled and monitored by the ID DCS.
Critical temperature, pressure and flow parameters are monitored and controlled, including the
pressure regulators, the heater temperatures and the current. Additional sensors monitor and control
3640 the thermal enclosure heaters, the electronics rack temperatures and voltages, and the cable and
patch-panel temperatures.

Within the pixel-detector volume, nearly 400 thermistors are installed. They monitor the tem-
peratures of the cooling loops, the service panels and the gas volume. In addition, to protect against
moisture, several radiation-tolerant humidity sensors are mounted inside the pixel detector and read
3645 out via the pixel DCS. Finally, each detector module and other sensitive units such as the regulator
stations and opto-boards are equipped with thermistors, which are connected to the pixel hard-wired
interlock system. Thanks to the high granularity of this system, even small groups of channels can
be switched off in case of risky situations. The pixel-detector interlock includes a signal from the
beam-conditions monitor system (see Section 3.4.1).

3650 For the SCT [104], the environmental monitoring measures 735 temperature and humidity
sensors across the detector. Three types of temperatures are measured: sensors are located at the
exhaust of each of the 116 cooling loops; sensors are attached to the mechanical structure of the
detector to monitor possible deformations due to temperature changes; and sensors are located
next to the laser interferometric survey monitoring system to measure the gas temperature inside
3655 the detector volume. Radiation-hard humidity probes are installed in various locations and are
used to calculate the dew point and hence to avoid condensation on the modules. The monitored
cooling-pipe temperatures also trigger the interlock, protecting the SCT modules if the cooling
stops. It is implemented in hardware without the use of DCS micro-processors. The custom-built
hardware compares the temperature sensor values to a preset threshold and signals the appropriate




Figure 65. Photograph of the ID barrel set-up for cosmic-ray studies (left) and the highlighted configuration of module groups chosen for this test (right).

power supply cards to turn off in ~ 1 second if the threshold is reached.

More than 3000 temperature sensors are distributed in the TRT to monitor the temperature of the active detector, as well as that of the front-end electronics and of the cooling circuits. A hardware interlock system cuts the power delivered to the front-end electronics if NTC thermistors located on the front-end boards detect a temperature above a critical hard-wired threshold. In addition, approximately 200 parameters of the closed-loop active gas system are monitored. The gas gain is continuously measured by a set of reference straws located outside the ID volume, and an automated algorithm adjusts the high voltage on the detector to preserve the stability of the gas gain as the environmental parameters (temperature and/or pressure) vary.

4.9 Performance status of the integrated inner detector

4.9.1 Electrical performance of the integrated detector

The electrical performance of individual pixel, SCT and TRT modules was monitored throughout assembly and some test modules were irradiated to the total dose expected after ten years of LHC operation. Following delivery to CERN, the SCT barrels and end-caps were fully connected at the surface to the SCT DAQ. The equivalent noise charge, noise occupancy at the nominal 1 fC operating threshold and the number of defective channels were measured and compared with previous data. After the SCT barrels and end-caps were fully integrated into the TRT, connectivity checks were repeated.

Before installation in the ATLAS pit, two opposite sectors of the barrel SCT and barrel TRT were tested [97]. The connected sectors comprised 1/8 of the TRT and 468 modules of the SCT barrels as shown in Fig. 65. A partial test was also made of one end-cap (side C), in which one quadrant (247 modules) of the SCT and 1/16 of the TRT wheels were connected. The tests emulated

SCT	Total channels	Not bonded	Dead	Not reached	Partially bonded	Noisy	Other defects	Total defects
Barrel	3244032	803	3595	1664	400	70	152	0.206%
End-caps	3035136	811	6464	-	352	230	5	0.259%

Table 21. Numbers, types and percentages of defective channels in the SCT after integration with the TRT in the surface assembly building.

TRT		Total channels	Dead channels	Total defects
Barrel	Side A	52544	926	1.8%
	Side C	52544	1050	2.0%
End-cap	Side A	122880	2115	1.7%
	Side C	122880	1993	1.6%

Table 22. Numbers and percentages of defective channels in the TRT after installation in ATLAS.

Pixels	Defective channels (%)			
	Layer-0	Layer-1	Layer-2	Average
Barrel	0.07	0.40	0.29	0.28
End-cap side A	Disk 1	Disk 2	Disk 3	
	0.14	0.23	0.52	0.30
End-cap side C	Disk 1	Disk 2	Disk 3	
	0.12	2.19	0.31	0.87
Pixel average				0.33

Table 23. Percentages of defective pixel channels after the final integration of the detector in the surface assembly building.

the final pit configuration, in particular the service routing and detector grounding. One pixel end-cap was also operated under realistic conditions and cosmic-ray data were studied.

The goals of the combined tests for SCT and TRT included the commissioning of the DCS system and the operation of the cooling system, as well as the measurement of the noise performance of the combined detectors under a wide range of operating conditions, to ensure the absence of cross-talk, to study timing and synchronisation issues and to obtain experience in the combined operation as preparation for commissioning. Cosmic-ray tracks passing through the SCT/TRT barrels

3685




Figure 66. Distribution of average noise occupancy for all active module sides of the barrel and end-cap SCT (outer or middle end-cap modules on side C), as obtained at 1 fC threshold.




Figure 67. Distribution of the noise in electrons, normalised to a temperature of 0°C, for all active modules (or front-end chips) in the barrel and end-cap SCT (side C).




Figure 68. Average noise in electrons for each active module in the pixel end-cap cosmic ray test.




Figure 69. Pixel occupancy for active modules in one end-cap disk during the pixel end-cap cosmic ray test as described in the text. The value BCID = 5 corresponds to the peak of the cosmic-ray timing distribution and the value BCID = 4,6 corresponds to adjacent time bins. The occupancy for other BCID values represents a measurement of the random pixel noise occupancy.

and end-caps were also used to study their performance. Table 21 shows the rate of defective SCT channels, after full integration with the TRT at the surface. Table 22 shows the rate of defective TRT channels after installation in the experiment.

The fraction of defective channels in the barrel and end-cap pixel components after the final integration of the detector in the surface assembly building is 0.33%, as summarised in Table 23.

The noise levels of the SCT, TRT and pixel sub-detectors were measured in many different configurations from data collected using random triggers.

For the SCT detector, the noise performance and occupancy were compiled for the individual barrels and end-caps. The distribution of measured noise occupancies of all active module sides of the barrel and end-cap cosmic ray runs are shown in Fig. 66 and, at a threshold of 1 fC, has a mean of $< 5 \times 10^{-5}$. The timing of the L1 trigger was optimised with respect to the scintillator trigger

3700 by histogramming the hit coincidences between the two sides of each module as a function of the timing offset. The corresponding equivalent noise-charge values extracted from threshold scans are shown in Fig. 67 for all active barrel and end-cap (side C) modules. All numbers are within specifications.

3705 The SCT noise was measured in extreme conditions: varying the trigger rate from 5 Hz to 50 kHz, with and without TRT operation, with the thermal enclosure heater pads on, off and switching between the two states, with several grounding schemes, and also while the TRT was being read out. No increase of noise was observed in any of the tested configurations. Similar tests were also performed for the TRT barrel and the TRT end-caps: the noise was verified before and after the insertion of the SCT, in a configuration in which the TRT analogue ground was connected to the SCT power return, for different SCT thresholds, and also while the SCT was being readout. The straw noise occupancy was in all cases close to 2% at a threshold of 250 eV.




Figure 70. The mean and RMS (indicated by error bars) of all hybrid temperatures for a given module position (from the cooling inlet to the cooling outlet), averaged over all cooling loops in the combined SCT-TRT run. The exhaust cooling pipe temperature was $\sim 10^\circ\text{C}$.

3720 The pixel detector was not included as part of the combined SCT-TRT barrel cosmic test since it is installed as a complete unit into the ID. However, one pixel end-cap was oriented with the disks in the horizontal plan, such that a reasonable cosmic-ray rate of ~ 1 Hz could be obtained, using a simple scintillator trigger. A prototype of a fraction of the complete internal pixel services was used to bring power and cooling to the end-cap, and to provide optical readout. The external cables, power supplies and readout were close to those used in the final detector. The complete end-cap with its 144 modules was also operated and tested.

3730 The modules in the pixel end-cap were tuned for a threshold of $\sim 4000 e$. The average noise for each active module was $\sim 170 e$ and the distribution is shown in Fig. 68. The average noise seen in the cosmic ray test was about 10% lower than that measured during individual module testing under similar, but not identical, circumstances. The observed pixel occupancy per BCID (beam crossing ID) included fixed pattern noise from a small fraction of the channels. Most of these hot (noisy) pixels were previously identified during individual module tests using an ^{241}Am source. The pixel occupancy for the active modules in one of the end-cap disks is shown in Fig. 69 after masking of the hot pixels. About 5×10^{-5} of the active pixels were masked for the modules shown. BCID 5 corresponds to the peak of the cosmic ray timing distribution with small tails before (BCID 4) and after (BCID 6). The pixel occupancy per BCID for other BCIDs was found to be $10^{-9} - 10^{-10}$ and is indicative of the random noise occupancy for these operating conditions.

4.9.2 SCT and pixel cooling performance after integration

3740 A satisfactory cooling performance of the pixel and SCT detectors is crucial to their operation. The heat dissipated inside the ID detector volume is at least 40 kW and eventually 50 kW, with

a temperature stability requirement within the SCT and pixels of $\sim 2^\circ\text{C}$ (see Table 19). There is a real risk of thermal runaway once the sensor leakage currents become sizable. Data following the detector installation are not yet available. However, data collected following integration on the surface have been analysed.

The performance of the pixel end-cap sectors was measured during the cosmic ray tests. The evaporative cooling was operated with a base temperature (no power) of approximately -25°C . Temperatures on each module were monitored. The average increase in temperature on a module for nominal power (23W for the six modules on a sector) was $9\text{--}10^\circ\text{C}$ (operation at -16 to -15°C) for the eight sectors on a disk and projected to be $13\text{--}14^\circ\text{C}$ for end-of-lifetime power (about 38W). The barrel pixel system has not been measured in such a realistic way. However, extensive measurements have been made on single and bi-staves. The difference in temperature between the coolant (nominally -25°C) and a module on a bi-stave is expected to be $\sim 15^\circ\text{C}$. Due to some corrosion problems with the aluminium pipes, some staves in layer 2 had to be reworked and the temperature gradient may be as high as $\sim 23^\circ\text{C}$ in these cases.

The various phases of integration have allowed a preliminary study of the SCT cooling performance. For the combined barrel test, a cooling pipe exhaust temperature of approximately 10°C was used. The average temperature on the barrel module hybrids (read by a thermistor on each module side) was $\sim 10^\circ\text{C}$ above the coolant temperature (unpowered) and $\sim 17\text{--}19^\circ\text{C}$ (powered). When averaged over the operating cooling loops, the mean and spread of the hybrid temperature for each module along a barrel cooling loop is shown in Fig. 70. For the end-cap, this temperature difference was slightly higher, $\sim 15^\circ\text{C}$, with respect to the selected exhaust temperature of 1°C .

4.10 Material distribution of the inner detector

The performance requirements of the ATLAS inner detector are more stringent than any tracking detector built so far for operation at a hadron collider. The harsh environment and the pile-up from multiple interactions per bunch crossing make a high detector granularity mandatory, with electronics, readout services and cooling within a detector volume that must have good mechanical stability. The overall weight (~ 4.5 tonnes) and material budget of the ID (in terms of radiation length X_0 and interaction length λ) are therefore much larger than those of previous tracking detectors. The consequences of this are quite serious and are currently the focus of many studies (see Section 10.2):

- a) many electrons lose most of their energy through bremsstrahlung before reaching the electromagnetic calorimeter;
- b) approximately 40% of photons convert into an electron-positron pair before reaching the LAr cryostat and the electromagnetic calorimeter;
- c) even in the case of low-energy charged pions, a significant fraction will undergo an inelastic hadronic interaction inside the inner detector volume.

A detailed modelling of the ID material has been implemented in the simulation. Figure 71 shows a map of generated photon-conversion vertices in the ID volume, integrated over azimuth, for electrons with $p_T > 0.5$ GeV from photons originating at the primary vertex in minimum-bias




Figure 71. Mapping of photon conversions as a function of z and radius, integrated over ϕ , for the ID. The mapping has been made from 500,000 minimum bias events (~ 40 minutes of data-taking at 200 Hz), using $\sim 90,000$ conversion electrons of $p_T > 0.5$ GeV originating from photons from π^0/η decays.




Figure 72. Material distribution (X_0 , λ) at the exit of the ID envelope, including the services and thermal enclosures. The distribution is shown as a function of $|\eta|$ and averaged over ϕ . The breakdown indicates the contributions of external services and of individual sub-detectors, including services in their active volume.

events. The sample shown consists of 500,000 events; at a collection rate of 200 Hz, such a data sample can be accumulated in approximately one hour. It should be noted that many structural elements, for example the end-plates of both the barrel and end-cap, are azimuthally discrete and have been modelled faithfully in the simulation. Figures 72 and 73 show the integrated radiation length, X_0 , and interaction length, λ , traversed by a straight track as a function of $|\eta|$ at the exit of the ID envelope. The most striking feature is the onset of non-active service and structural material




Figure 73. Material distribution (X_0, λ) at the exit of the ID envelope, including the services and thermal enclosures. The distribution is shown as a function of $|\eta|$ and averaged over ϕ . The breakdown shows the contributions of different ID components, independent of the sub-detector.

	$ \eta = 0$		$ \eta = 1.8$	
	Radius (mm)	X_0	Radius (mm)	X_0
Exit beam-pipe	36	0.0045	36	0.014
Exit pixel layer-0	57	0.037	57	0.105
Exit pixel layer-2	172	0.108	172	0.442
Entry SCT	253	0.119	253	0.561
Entry TRT	552	0.205	621	0.907
Exit TRT	1081	0.469	907	1.126

Table 24. Integrated radiation length (X_0) from interaction point, estimated as a function of radius R for $|\eta| = 0$ and for $|\eta| = 1.8$, using the detailed description of the inner-detector material implemented in the simulation. The quoted radii are the maximum radii over which the integration is made. The data are averaged over ϕ .

at the interface of the barrel and end-cap regions. This includes cooling connections at the end of the SCT and TRT barrels, TRT electrical connections, and SCT and TRT barrel services extending radially to the cryostat, to the PPB1 patch-panel, and then along the cryostat wall. Another service contribution is from the pixel services at $|\eta| > 2.7$, which leave the detector along the beam-pipe; their extended range in $|\eta|$ can clearly be seen. A large fraction of the service and structural material is external to the active ID envelope, therefore deteriorating the calorimeter resolution but not the tracking performance. Table 24 lists the contribution to X_0 as a function of radius for different elements of the ID and for straight tracks at $|\eta| = 0$ and $|\eta| = 1.8$.

The material breakdown is particularly important at small radius. The pixel barrel radiation length for perpendicular incidence is approximately 10.7% for the three pixel layers. This can be broken down into: electronics+bump-bonds (1.4%), sensors (1.1%), hybrids (1%), local support structures with cooling (5.4%), cables (0.3%) and global supports (1.5%). The corresponding num-

Sub-detector	Measured weight (kg)	Estimated weight from component measurements (kg)		Weight in simulation (kg)
		Pre-integration	Post-integration	
SCT barrel	201±20	187±5	222±6	222
TRT barrel	707±20	677±3	703±3	700
SCT+TRT barrel	883±20	864±6	925±7	922
SCT end-cap A	207±10	177±5	225±10	225
SCT end-cap C	172±10	177±5	225±10	225
TRT end-cap A	1118±12	n/a	1129±10	1131
TRT end-cap C	1120±12	n/a	1129±10	1131
Pixel barrel	n/a	20.2	n/a	18.3
Pixel end-cap A	n/a	7	n/a	6.1
Pixel end-cap C	n/a	7	n/a	6.1
Pixel total	193.5±5	33	201	197

Table 25. The measured weight of individual ID sub-detectors (where available) together with bottom-up estimates of the ID weight, before and after detector integration/installation. Those items which were not measured are labelled as not available (n/a). The post-integration weight includes all services up to, but not including, the PP1 patch-panels. The services of the installed pixel layers have not been estimated independently for the barrel and end-caps. The measured weights should be compared to the estimated weights before integration except in the case of the TRT end-caps for which the weights were measured only after integration of all services. The quoted estimates from simulations should be compared to the post-integration weights.

ber for the SCT barrel layers is 11.8% when averaged over the active area. This amounts to 2.96%
3800 per layer, including 1.33% for modules, 1.15% for services and 0.48% for the support structure.
The module budget includes the sensors (0.61%), the base-boards and BeO facings (0.20%), and
the hybrids with their components and the carbon-carbon bridge (0.32%). The equivalent material
per disk is 3.75%: modules (1.81%), services (1.47%) and disk support structure (0.47%). Services
include the electrical services, cooling blocks and pipes. Off-barrel or off-disk services, especially
3805 at the interface between the barrels and end-caps and in the forward direction for pixel services,
and items such as the thermal enclosures and outer support cylinder of the end-cap SCT, are not
included.

An attempt has been made to check the assumed material distribution in the ID by making
a comparison with the measured weight of different ID sub-detector units, and with a bottom-up
3810 inventory of the ID components. Table 25 shows the status of that comparison. There is good
agreement between the measured and estimated weights of the component sub-detectors and their
services. There remains a discrepancy between the measured weights of the two SCT end-caps.

5. Calorimetry

An overview of the ATLAS calorimetry system [105, 106] is given in Section 1.4. The overall system is depicted in Fig. 3, its general performance goals are listed in Table 1, and its main parameters are given in Table 3.




Figure 74. Cumulative amounts of material, in units of radiation length X_0 and as a function of $|\eta|$, in front of and in the electromagnetic calorimeters. The top left-hand plot shows separately the total amount of material in front of the presampler layer and in front of the accordion itself over the full η -coverage. The top right-hand plot shows the details of the crack region between the barrel and end-cap cryostats, both in terms of material in front of the active layers (including the crack scintillator) and of the total thickness of the active calorimeter. The two bottom figures show, in contrast, separately for the barrel (left) and end-cap (right), the thicknesses of each accordion layer as well as the amount of material in front of the accordion.

The ATLAS calorimeters consist of a number of sampling detectors with full ϕ -symmetry and coverage around the beam axis. The calorimeters closest to the beam-line are housed in three cryostats, one barrel and two end-caps. The barrel cryostat contains the electromagnetic barrel calorimeter, whereas the two end-cap cryostats each contain an electromagnetic end-cap calorimeter (EMEC), a hadronic end-cap calorimeter (HEC), located behind the EMEC, and a forward calorimeter (FCal) to cover the region closest to the beam. All these calorimeters use liquid argon as the active detector medium; liquid argon has been chosen for its intrinsic linear behaviour, its stability of response over time and its intrinsic radiation-hardness.

The precision electromagnetic calorimeters are lead-liquid argon detectors with accordion-shape absorbers and electrodes. This geometry allows the calorimeters to have several active layers

in depth, three in the precision-measurement region ($0 < |\eta| < 2.5$) and two in the higher- η region ($2.5 < |\eta| < 3.2$) and in the overlap region between the barrel and the EMEC. In the precision-measurement region, an accurate position measurement is obtained by finely segmenting the first compartment in η . The η -direction of photons is determined by the position of the photon cluster in the first and the second compartments. The calorimeter system also has electromagnetic coverage at higher η ($3.1 < |\eta| < 4.9$) provided by the FCal. Furthermore in the region ($0 < |\eta| < 1.8$) the electromagnetic calorimeters are complemented by presamplers, an instrumented argon layer, which provides a measurement of the energy lost in front of the electromagnetic calorimeters.

For the outer hadronic calorimeter, the sampling medium consists of scintillator tiles and the absorber medium is steel. The tile calorimeter is composed of three parts, one central barrel and two extended barrels. The choice of this technology provides maximum radial depth for the least cost for ATLAS. The tile calorimeter covers the range $0 < |\eta| < 1.7$ (central barrel and extended barrels). The hadronic calorimetry is extended to larger pseudorapidities by the HEC, a copper/liquid-argon detector, and the FCal, a copper-tungsten/liquid-argon detector. The hadronic calorimetry thus reaches one of its main design goals, namely coverage over $|\eta| < 4.9$.

The numbers of radiation and interaction lengths in front of and in the electromagnetic and hadronic calorimeters are shown in Figs. 74 and 75.

Section 5.1 describes the LAr cryostats and feed-throughs. Sections 5.2 and 5.3 are devoted to the description of the electromagnetic and hadronic calorimetry, respectively. The instrumentation in the gaps between the cryostats is described in Section 5.4. The front-end readout electronics, back-end electronics and services are described in Section 5.5. Finally, test-beam measurements obtained with production modules of the different calorimeters are shown in Section 5.6.

5.1 Cryostats and associated feed-throughs

5.1.1 Cryostat description

As mentioned above, the liquid-argon calorimeters are located in three different cryostats. Figure 3 shows the general features and location of the barrel and end-cap cryostats in the overall calorimeter system. Each cryostat is composed of two concentric aluminium vessels, an inner cold vessel and an outer warm vessel. Each vessel forms a cylindrical torus centred on the beam axis and the space in between the vessels is under vacuum. As described in Section 2.1.1, the central solenoid is housed in the insulating vacuum of the barrel cryostat and supported by the inner




Figure 75. Cumulative amount of material, in units of interaction length, as a function of $|\eta|$, in front of the electromagnetic calorimeters, in the electromagnetic calorimeters themselves, in each hadronic compartment, and the total amount at the end of the active calorimetry. Also shown for completeness is the total amount of material in front of the first active layer of the muon spectrometer (up to $|\eta| < 3.0$).




Figure 76. Cut-away view of an end-cap cryostat showing the positions of the three end-cap calorimeters. The outer radius of the cylindrical cryostat vessel is 2.25 m and the length of the cryostat is 3.17 m.

cylinder of the warm barrel vessel. The two end-cap cryostats each house one electromagnetic end-cap calorimeter (EMEC), one hadronic end-cap calorimeter (HEC) and one forward calorimeter (FCal). Figure 76 shows a more detailed cut-away view of one end-cap cryostat with the three different calorimeter systems it contains.

The different vessels are closed at both ends by bulkheads. All bulkheads are circular-shaped plates with passage holes. The barrel passage holes allow the insertion of the inner detector and of its associated beam-pipe into the inner warm vessel bore. The end-cap passage holes allow the passage of the next section of the beam-pipe and the insertion of its associated pump station (see Section 9.8). The barrel bulkheads and the inner cold vessel of the barrel cryostat have tapered walls to minimise the material and dead space between the interaction region and the electromagnetic calorimeters. To further minimise the amount of material, the cold and warm front end-cap bulkheads are flat, and insulating spacers allow the warm bulkhead to sit on the cold one.

The cold vessels rest on four feet made of a glass-fibre composite material, which provide electrical and thermal insulation to the warm vessels. The barrel cryostat load is transferred directly to the ATLAS main rail system through the tile-calorimeter support saddles. For the support on the side of the end-cap cryostat pointing towards the interaction point, this is not possible and the

load is transferred through the inner radius of the extended calorimeter on a vertically adjustable
3885 support.

5.1.2 Signal feed-throughs

The signal feed-throughs bring all the signal, monitoring, calibration and spare lines through the
insulating vacuum from the liquid-argon cold volume to the front-end crates located at room tem-
perature around and on the outside of the barrel and end-cap cryostats. A total of 64 feed-throughs
3890 serve the 122,800 lines of the barrel electromagnetic calorimeter, whereas a total of 2×25 feed-
throughs serve the $2 \times 48,000$ lines of the two end-caps. In each end-cap, 20 feed-throughs are
used by the EMEC, four by the HEC and one by the FCal (the EMEC uses also part of the four
HEC feed-throughs).

A feed-through [107] consists primarily of a warm flange and a cold flange, with a flexible
3895 bellows welded between them. The volume between the two flanges is under vacuum. Each flange
houses four gold-plated pin carriers, providing a total of 1920 signal connections per feed-through.
Both pin carriers are interconnected with flat polyimide strip-line flexible cables. To minimise the
lengths of the connection lines, the signal feed-throughs are distributed radially around each end of
the barrel cryostat and at the back of each of the end-cap cryostats.

5.1.3 High-voltage feed-throughs

Two high-voltage feed-through ports [108] are mounted on each cryostat. They are located close to
the highest point at each end of the barrel cryostat and at the back of each end-cap cryostat. A port
consists of a single warm bulkhead mounted on the warm vessel, from which the wire bundle leads
to the calorimeter via a tube connected to the cold vessel by stainless-steel bellows. The liquid
3905 argon rises up to the bellows and is followed by a 70 cm high column of argon gas at 1.25 bar.
The high-voltage wire consists of a solid 0.41 mm diameter constantan (Cu/Ni alloy) conductor
with a 0.30 mm thick radiation-hard plastic (polyether-ether-ketone) cladding. Each port holds
approximately 840 HV lines. These lines serve the standard high-voltage lines and the spares may
be used to supply single electrodes or split sectors which do not hold the nominal high voltage
3910 (see Section 5.2.5).

5.2 Electromagnetic calorimetry

5.2.1 Accordion geometry

An accordion geometry has been chosen for the absorbers and the electrodes of the barrel and end-
cap electromagnetic calorimeters (see Fig. 77). Such a geometry provides naturally a full coverage
3915 in ϕ without any cracks, and a fast extraction of the signal at the rear or at the front of the electrodes.
In the barrel, the accordion waves are axial and run in ϕ , and the folding angles of the waves vary
with radius to keep the liquid-argon gap constant (see Figs. 78 and 79). In the end-caps, the waves
are parallel to the radial direction and run axially. Since the liquid-argon gap increases with radius
in the end-caps, the wave amplitude and the folding angle of the absorbers and electrodes vary
3920 with radius (see Fig. 80). All these features of the accordion geometry lead to a very uniform
performance in terms of linearity and resolution as a function of ϕ .




Figure 77. Layout of the signal layer for the four different types of electrodes before folding. The two top electrodes are for the barrel and the two bottom electrodes are for the end-cap inner (left) and outer (right) wheels. Dimensions are in millimetres. The drawings are all at the same scale. The two or three different layers in depth are clearly visible. The first layer is finely segmented along η , as for example in the barrel where there are eight strips in front of a middle cell. One can note however the coarser granularity of the first layer in the edge zones of the barrel (top right) and end-caps, as explicitly described in Table 3. The second layer collects the largest fraction of the energy of the electromagnetic shower, and the third layer collects only the tail of the electromagnetic shower and is therefore less segmented in η .

The absorbers are made of lead plates, to which two stainless-steel sheets (0.2 mm thick) are glued using a resin-impregnated glass-fibre fabric to provide mechanical strength. The absorbers in the barrel have a thickness of 1.53 mm for $|\eta| < 0.8$ and of 1.13 mm for $|\eta| > 0.8$. The change in lead thickness at $|\eta| = 0.8$ limits the decrease of the sampling fraction as $|\eta|$ increases. In the end-cap calorimeters, the plates have a thickness of 1.7 mm for $|\eta| < 2.5$ and of 2.2 mm for $|\eta| > 2.5$.

The readout electrodes [109] consist of three conductive copper layers separated by insulating polyimide sheets. The two outer layers are at the high-voltage potential and the inner one is used for reading out the signal via capacitive coupling. The segmentation of the calorimeter in η and in depth is obtained by etched patterns on the different layers, as shown in Fig. 77. The ϕ -segmentation is obtained by ganging together the appropriate number of electrodes (see Section 5.2.6). Each barrel gap between two absorbers is equipped with two electrodes, one type for $|\eta| < 0.8$ and another for $|\eta| > 0.8$. Similarly, each end-cap gap between two absorbers is equipped with one type of electrode for $|\eta| < 2.5$ and with another for $|\eta| > 2.5$.

5.2.2 Barrel geometry

The barrel electromagnetic calorimeter [110] is made of two half-barrels, centred around the z -




Figure 78. Sketch of a barrel module where the different layers are clearly visible with the ganging of electrodes in ϕ . The granularity in η and ϕ of the cells of each of the three layers and of the trigger towers is also shown.

axis. One half-barrel covers the region with $z > 0$ ($0 < \eta < 1.475$) and the other one the region with $z < 0$ ($-1.475 < \eta < 0$). The length of each half-barrel is 3.2 m, their inner and outer diameters are 2.8 m and 4 m respectively, and each half-barrel weighs 57 tonnes. As mentioned above, the barrel calorimeter is complemented with a liquid-argon presampler detector, placed in front of its inner surface, over the full η -range.

A half-barrel is made of 1024 accordion-shaped absorbers, interleaved with readout electrodes. The electrodes are positioned in the middle of the gap by honeycomb spacers. The size of the drift gap on each side of the electrode is 2.1 mm, which corresponds to a total drift time of about 450 ns for an operating voltage of 2000 V. Once assembled, a half-barrel presents no discontinuity along the azimuthal angle ϕ ; however, for ease of construction, each half-barrel has been divided into 16 modules, each covering a $\Delta\phi = 22.5^\circ$. The total thickness of a module is at least 22 radiation lengths (X_0), increasing from 22 X_0 to 30 X_0 between $|\eta| = 0$ and $|\eta| = 0.8$ and from 24 X_0 to 33 X_0

between $|\eta| = 0.8$ and $|\eta| = 1.3$.

3950 At the inner and outer edges, each absorber is encased in the groove of a precision-machined glass-fibre composite bar. The purpose of these bars is to accurately position each absorber with respect to its neighbours and also to provide space for the connectors of the electrodes. The stacking of these bars defines the cylindrical geometry of the half-barrel.

3955 Seven stainless-steel outer rings support a half-barrel and provide it with the required rigidity. Each ring is made of 16 ring-pieces corresponding to the 16 modules. All ring-pieces are identical with an I-beam cross-section except for the two ring-pieces at the level of the cryostat rails. Similarly, eight composite inner rings define the inner geometry of a half-barrel. Each inner ring is also made of 16 identical ring-pieces. The absorber bars are screwed into these ring-pieces.

3960 A module, as depicted in Figs. 78 and 79, has three layers or compartments in depth (front, middle and back as viewed from the interaction point). The front compartment is read out at the low-radius side of the electrode, whereas the middle and back compartments are read out at the high-radius side of the electrode. The readout granularity of the different compartments is shown in Table 3. In total, there are 3424 readout cells per module, including the presampler cells. The amount of dead material in front of the presampler and between the presampler and the first calorimeter compartment as well as the thickness of each calorimeter compartment are shown in Fig. 74 in units of X_0 .

3970 The presampler [111] is a separate thin liquid-argon layer (11 mm in depth), which provides shower sampling in front of the active electromagnetic calorimeter and inside the barrel cryostat. This presampler layer is made of 64 identical azimuthal sectors (32 per half-barrel). Each sector is 3.1 m long and 0.28 m wide, thus covering the half-barrel length and providing a coverage in $\Delta\eta \times \Delta\phi$ of 1.52×0.2 . It is composed of eight modules of different size, with a length increasing with $|\eta|$ to obtain a constant η -granularity of $\Delta\eta = 0.2$ for each module, except for the module at the end of the barrel, for which the η -coverage is reduced to 0.12.

3975 The presampler modules are made of interleaved cathode and anode electrodes glued between glass-fibre composite plates. The electrode spacing varies slightly, from 1.9 to 2.0 mm, with the presampler module type. The cathodes are double-sided printed-circuit boards while the anodes have three conductive layers separated by glass-fibre composite layers. The required segmentation, $\Delta\eta \sim 0.025$ and $\Delta\phi = 0.1$, for each module is obtained by ganging the appropriate number of anodes in the η direction and by subdividing (i.e. etching) each anode into two halves in the ϕ -direction. A high voltage potential of +2 kV is applied to the outer layers of the anodes and the signal is read out through capacitive coupling to the central layer at ground potential.

5.2.3 End-cap geometry

3985 The EMEC calorimeters [112] consist of two wheels, one on each side of the electromagnetic barrel. Each wheel is 63 cm thick and weighs 27 tonnes, with external and internal radii at ambient temperature of 2098 mm and 330 mm, respectively. It covers the region $1.375 < |\eta| < 3.2$. In the transition region between the barrel and the end-cap calorimeters, the material in front of the calorimeter amounts to several X_0 , as shown in Fig. 74. In order to improve the energy measurement in this region, a liquid-argon presampler is implemented in front of the end-cap calorimeter, covering the range $1.5 < |\eta| < 1.8$.



Figure 79. Photograph of a partly stacked barrel electromagnetic LAr module. A total of six out of seven outer support rings into which the absorbers can be seen. The backbone behind the outer support rings and the assembly bench below the stacked modules are also visible.



Figure 80. Photograph showing a side view of an electromagnetic end-cap LAr module (the beam axis is vertical). The first accordion absorber of each wheel is clearly visible, as well as the summing boards, the mother-boards and the cables.

3990 Each end-cap calorimeter consists itself of two co-axial wheels. The boundary between the
 inner and the outer wheel, which is 3 mm wide and located at $|\eta| = 2.5$, is mostly filled with
 low-density material. This boundary is approximately projective and matches the acceptance of
 the inner detector (see Table 2). Each end-cap wheel is further divided into eight wedge-shaped
 3995 modules without introducing any discontinuity along the azimuthal angle owing to the accordion
 geometry. A view of a module is shown in Fig. 80. Each end-cap contains 768 absorbers interleaved
 with readout electrodes in the outer wheel and 256 absorbers in the inner wheel. The electrodes
 are positioned in the middle of the gaps by honeycomb spacers. The total active thickness of an
 end-cap calorimeter is greater than $24 X_0$ except for $|\eta| < 1.475$. The thickness increases from 24
 to $38 X_0$ as $|\eta|$ increases from 1.475 to 2.5 (outer wheel) and from 26 to $36 X_0$ as $|\eta|$ increases
 4000 from 2.5 to 3.2 (inner wheel). The calorimeter support frame is composed of six support rings:
 three on the front (at three different radii) and three on the back. The total weight of the detector is
 supported mainly by the rings at the largest radius.

In the outer wheel, signals from the different pads are read out from both sides of the electrode,
 as in the case of the barrel electromagnetic calorimeter. In the inner wheel, because of the higher
 4005 radiation levels, the signals are all read out from the back side.

As for the barrel electromagnetic calorimeter, the precision region in the end-cap electromag-
 netic calorimeters ($1.5 < |\eta| < 2.5$) is divided in depth into three longitudinal compartments.
 The front compartment, about $4.4 X_0$ thick, is segmented with strips along the η direction. The
 transverse size of the projective cell in the middle compartment is the same as defined in the bar-
 4010 rel electromagnetic calorimeter, $\Delta\eta \times \Delta\phi = 0.025 \times 0.025$. The back compartment has a twice
 coarser granularity in η . The outermost region $|\eta| < 1.5$ of the outer wheel and the inner wheel
 ($2.5 < |\eta| < 3.2$) are segmented in only two longitudinal compartments and have a coarser trans-
 verse granularity. Table 3 summarises the longitudinal and transverse readout granularities of the
 electromagnetic end-cap calorimeter as a function of $|\eta|$. The η -granularity in the front compart-
 4015 ment varies with η in order to keep the copper strip width larger than a few mm (see Table 3). The

ϕ -granularity is obtained by ganging the signals from adjacent electrodes (see Section 5.2.6). Each module contains 3984 readout channels, including the 96 channels in the presampler.

Each end-cap presampler consists of 32 identical azimuthal sectors or modules. These are placed in a 5 mm deep cavity in the back of the cryostat cold wall. The granularity of the presampler is $\Delta\eta \times \Delta\phi = 0.025 \times 0.1$. One end-cap presampler module consists of two, 2 mm thick, active liquid argon layers, formed by three electrodes parallel to the front face of the EMEC calorimeter. The electrodes, made from double-sided printed-circuit boards, are separated by honeycomb spacers and glued together at the inner and outer radius with 2 mm thick bars. A negative high voltage is applied to the external electrodes and the signals are read out from the central electrode which is segmented into pads. The same signal, calibration and high-voltage cables as for the end-cap calorimeter are used. Two high-voltage cables feed separately the left and right side of a presampler module.

5.2.4 Shape and placement of the electromagnetic calorimeters

After insertion into the cryostat, the circularity of the barrel electromagnetic calorimeter was measured. The deformation due to its weight can reach up to 3 mm at the top and bottom of each half-barrel and is in fair agreement with the corresponding finite-element calculations [110]. In the cold, the deformation is reduced due to the Archimedes thrust caused by the liquid argon which has a relative density of 1.4. The centres of the fitted inner and outer circles coincide to better than 0.5 mm. Finally, the sagging of the absorbers was measured as a function of the azimuthal angle and a maximum value of 2.5 mm near the horizontal plane was found. The same set of measurements has been performed for both end-caps, once they were in vertical position. The measured deformation was found to be up to 2 mm at the top. The sagging of the absorber plates and the small deformation of the barrel electromagnetic calorimeter will induce some distortions in the position measurements but their impact on the energy resolution is expected to be very small.

Due to difficulties in the positioning of the half-barrels inside the cryostat, both half-barrels when cold are 4 mm too low and laterally displaced by 2 mm with respect to the cryostat axis [110]. Part of the vertical misalignment has been compensated by positioning the cryostat approximately 2 mm higher than originally planned inside the experiment (see Section 9.3.2.3). To make room for the inner-detector services, the two end-cap cryostats have been positioned 40 mm away from their nominal position in the z -direction. The projectivity of the geometry of these calorimeters has therefore been slightly degraded.

5.2.5 High-voltage distribution

To provide a measure of redundancy, the two high-voltage sides of each electrode are fed independently. The nominal high-voltage for the barrel electromagnetic calorimeter is 2000 V (see Table 26). If one side of an electrode is not powered, only half of the signal will be collected. The variation of the signal [113] with high voltage is mainly due to the variation of the drift velocity with high voltage and is shown in Fig. 81. The variation of the signal amplitude with high voltage is moderate: for example, 77% of the signal is collected when decreasing the high voltage to half of its nominal value. The high-voltage granularity of the barrel electromagnetic calorimeter, including the presampler, is $\Delta\phi \times \Delta\eta = 0.2 \times 0.2$, meaning that 32 electrode sectors are powered simultaneously.




Figure 81. Measured electromagnetic cluster energy as a function of the applied high voltage. The results are shown for a barrel module (left), for 245 GeV electrons (open circles), 100 GeV electrons (open diamonds) and for the 100 GeV results at the nominal voltage of 2 kV scaled to the corresponding result at 245 GeV (stars). The results obtained with an end-cap module (right) are shown for 193 GeV electrons. The curves correspond to fits with a functional form $E_{\text{tot}} = a \times V^b$.

In contrast to the barrel electromagnetic calorimeter, the drift gap on each side of the electrodes is not constant for the EMEC, but is a function of radius and varies from 2.8 mm to 0.9 mm in the outer wheel and from 3.1 mm to 1.8 mm in the inner wheel. To obtain a uniform η -independent detector response, one would therefore need a high voltage continuously varying with η , as illustrated in Fig. 82. In practice, this ideal behaviour is approximated with a variation in steps, also illustrated in Fig. 82. Table 26 provides explicitly these nominal high-voltage values as a function of the η -range, thereby defining seven (two) high-voltage sectors for the outer (inner) wheel. The high-voltage granularity is thus 0.2 in ϕ as in the barrel and approximately 0.2 in η , as indicated in Table 26.

Different solutions have been implemented in case of high-voltage problems. One electrode sector may be powered individually, if the problem was identified at room temperature. On the other hand, if the problem has been identified only in the cryostat when testing individual modules at liquid-argon temperature, one high-voltage sector may be divided into two halves, for example $\Delta\phi \times \Delta\eta = 0.1 \times 0.2$ in the barrel, and the two halves are powered separately. The effect of not supplying high voltage at all to one side of the electrode was studied in the test-beam to verify the expected factor of two loss in signal. A slight deviation from this expectation was observed for the measured signal amplitude at the transition region with a good




Figure 82. High-voltage distribution as a function of $|\eta|$ for the EMEC. A uniform calorimeter response requires a high voltage which varies continuously as a function of $|\eta|$, as shown by the open circles. This has been approximated by a set of discrete values shown as full triangles.

Component	η -range and voltage (V)						
Electromagnetic barrel	0 – 1.475						
Voltage	+2000						
Electromagnetic end-cap (EMEC)	1.375 – 1.5	1.5 – 1.6	1.6 – 1.8	1.8 – 2.0	2.0 – 2.1	2.1 – 2.3	2.3 – 2.5
Voltage	+2500	+2300	+2100	+1700	+1500	+1250	+1000
Electromagnetic end-cap (EMEC)	2.5 – 2.8				2.8 – 3.2		
Voltage	+2300				+1800		
Barrel presampler	0 – 1.52						
Voltage	+2000						
End-cap presampler	1.5 – 1.8						
Voltage	–2000						
Hadronic end-cap (HEC)	1.5 – 3.2						
Voltage	+1800						
Forward calorimeter	3.1 – 4.9						
FCal1 voltage	+250						
FCal2 voltage	+375						
FCal3 voltage	+500						

Table 26. Nominal high-voltage settings for the various liquid-argon calorimeter components and as a function of $|\eta|$ in the case of the EMEC outer and inner wheels.

sector and this was understood as a geometrical effect, which is easy to correct for since it is well reproduced by simulation. An alternative for sector sides with a permanent short is to supply them with high-voltage modules permitting large DC current draws of up to 3 mA, which is in general sufficient to operate the faulty sector at 1000 V or above.

5.2.6 Electronic boards and cables inside the cryostats

The first layer of each calorimeter is read out from the front (inner radius for the barrel electromagnetic calorimeter and small- $|z|$ face for the EMEC), whereas the second and third layers are read out from the back (outer radius for the barrel electromagnetic calorimeter and large- $|z|$ face for the EMEC). The readout is accomplished by several interconnected boards. These include summing boards which generate analogue sums from adjacent calorimeter gaps and mother-boards for analogue sum readout and for distribution of the calibration pulses.

The summing boards are connected to the electrode connectors grouping the signals in ϕ . For the barrel electromagnetic calorimeter and the outer EMEC wheel, each summing board services respectively 16 and 12 electrodes (corresponding to $\Delta\phi = 0.1$). For the first layer, all electrodes are summed into one readout cell, while for the second and third layers they are summed in four adjacent readout cells ($\Delta\phi = 0.025$). For the inner EMEC wheel, each summing board services eight electrodes ($\Delta\phi = 0.2$) and for each of the two layers, four adjacent electrodes are summed to form readout cells ($\Delta\phi = 0.1$).

A mother-board provides readout for a $\Delta\eta \times \Delta\phi$ region of size 0.2×0.2 . The receptacles at the output of the summing boards and the input of the mother-boards are connected via pins mounted on two thin printed-circuit boards. Each channel is connected twice to provide maximum robust-

Layer	Number of channels	Maximum rate of faults accepted	Maximum number of faults accepted	Number of faulty channels	
				at room temperature	in liquid argon
Electromagnetic barrel calorimeter					
Presampler	7808	0.05%	4	0	0
Layer 1	57216	2 (per module)	64	9	11
Layer 2	28672	0.05%	14	2	5
Layer 3	13824	0.05%	6	0	4
Barrel end	2048	0.05%	1	3	3
Calibration	8192	0.05%	4	1	1
Electromagnetic end-cap calorimeter					
Presampler	1536	0.05%	1	0	1
Layer 1	28544	2 (per module)	32	3	4
Layer 2	23424	0.05%	12	2	7
Layer 3	10240	0.05%	5	0	2
Calibration	5952	0.05%	3	2	2

Table 27. Number of signal channels in each electromagnetic calorimeter layer and total number of calibration channels, pre-defined acceptable maximum rate and maximum number of faults, and numbers of faulty channels found after insertion at room temperature and later after cooling down and filling with liquid argon. Faulty channels correspond to either missing signal channels or signal channels with a bad signal shape [110, 112].

ness and redundancy. For the second and third calorimeter layers, every other printed-circuit board
4105 also houses a low-capacitance transient voltage suppressor [110], referenced to ground, with a turn-
on at 6 V to protect against accidental discharges, which might damage the calibration resistors
mounted on the mother-board. The mother-boards route the outputs to the readout cables through
a connector designed to minimise cross-talk between readout channels. The mother-boards also in-
clude surface-mounted 0.1% tolerance tantalum-nitride calibration resistors (non-hygroscopic with
4110 a temperature coefficient of 70 ± 5 ppm/ $^{\circ}$ C) to distribute calibration pulses to all readout channels.
The values of the calibration resistors in various parts of the detector are chosen to match the ex-
pected currents from particle showers [110, 112]. The calibration pulse is injected at the analogue
input to the mother-board, which is the point closest to the origin of the analogue signals from the
electrodes. Each presampler module is also equipped with a mother-board which collects signals
4115 from the readout cells and distributes the calibration pulses through a set of precision calibration

resistors as mentioned above.

5.2.7 Quality-assurance tests

At each stage of the construction of the modules and assembly of the calorimeters, a number of quality-assurance tests were carried out. The most important tests were those performed after
4120 insertion of the calorimeters inside their cryostats and after the first cool-down and liquid-argon filling. All the signal and calibrations channels have been tested. Table 27 shows that the number of missing or bad channels is below the specified acceptable number of faults.

A high-voltage test at the nominal settings was also performed at this stage. After this first cold test, the calorimeters were heated back to room temperature before transportation and installation
4125 in the ATLAS pit. High-voltage tests have also been performed on all calorimeters in their final position in ATLAS. In the case of the barrel presampler and of the EMEC, most of the shorts can be eliminated by applying a high-voltage discharge [112, 114]. As a result, approximately 3% of the barrel electromagnetic channels and approximately 1% of the barrel and end-cap presampler and EMEC channels remain with shorts, and therefore will be powered at a reduced voltage, as
4130 explained in Section 5.2.5.

5.3 Hadronic calorimeters

This section describes the ATLAS hadronic calorimeters: the tile calorimeter, the liquid-argon hadronic end-cap calorimeter (HEC) and the liquid-argon forward calorimeter (FCal).

5.3.1 Tile calorimeter

5.3.1.1 Overview The tile calorimeter [106] is a sampling calorimeter using steel as the absorber and scintillator as the active medium. It is located in the region, $|\eta| < 1.7$, behind the liquid argon electromagnetic calorimeter and is subdivided into a central barrel, 5.8 m in length, and two extended barrels, 2.6 m in length and each having an inner radius of 2.28 m and an outer radius of 4.25 m, as shown in Fig. 3. The radial depth of the tile calorimeter is approximately 7.4λ
4135 (interaction lengths). Each barrel consists of 64 modules or wedges of size $\Delta\phi \sim 0.1$, made of steel plates and scintillating tiles [115].

The assembled module forms an almost-periodic steel-scintillator structure with a ratio by volume of approximately 4.7:1. The geometry is sketched in Fig.83. The orientation of the scintillator tiles radially and normal to the beam line, in combination with wavelength-shifting fibre readout
4145 on the tile edges, allows for almost seamless azimuthal calorimeter coverage. The grouping of the readout fibres into the readout photomultiplier tubes (PMT's) provides an approximately projective geometry in pseudorapidity (see Table 3). The gap region between the barrel and the extended barrel is instrumented with special modules, made of steel-scintillator sandwiches with the same sampling fraction as the rest of the tile calorimeter and with thin scintillator counters in the sectors
4150 where the available space in the gaps is even more limited. These devices allow to partially recover the energy lost in the crack regions of the detector and are described in more detail in Section 5.4.

The electronics and readout of the tile calorimeter are highly integrated with the mechanical structure. The photomultiplier tubes and all the front-end electronics are mounted in 1.4 m long aluminium units, called drawers, which are inserted inside the support girder at the rear of each

4155 module (see Section 5.5.1). The front-end electronics also provide analogue sums of subsets of
 the channels, forming trigger towers, for the L1 trigger (see Section 5.5). The low-voltage power
 supplies which power the readout are mounted in an external steel box, which has the cross-section
 of the support girder and which also contains the external connections for power and other services
 for the electronics (see Section 5.5.3.1). Finally, the calorimeter is equipped with three calibration
 4160 systems: charge injection, laser and a ^{137}Cs radioactive source. These systems test the optical
 and digitised signals at various stages and are used to set the PMT gains to a uniformity of $\pm 3\%$
 (see Section 5.5.2) .

5.3.1.2 Mechanical structure

The mechanical structure of the tile calorimeter is designed as a self-supporting, segmented
 4165 structure comprising 64 modules, each subtending 5.625 degrees in azimuth, for each
 of the three sections of the calorimeter [115]. The module sub-assembly is shown in Fig. 84.
 4170 Each module contains a precision-machined strong-back steel girder, the edges of which
 are used to establish a module-to-module gap of 1.5 mm at the inner radius. To
 maximise the use of radial space, the girder provides both the volume in which the tile
 calorimeter readout electronics are contained and the flux return for the solenoid field. The
 4175 readout fibres, suitably bundled, penetrate the edges of the girders through machined
 holes, into which plastic rings have been precisely mounted. These rings are matched
 to the position of photomultipliers. The fundamental element of the absorber structure
 consists of a 5 mm thick master plate, onto which 4 mm thick spacer plates are glued
 4180 in a staggered fashion to form the pockets in which the scintillator tiles are located [116].
 The master plate was fabricated by high-precision die stamping to obtain the dimensional tolerances
 required to meet the specification for the module-to-module gap [117]. At the module edges,
 the spacer plates are aligned into recessed slots, in which the readout fibres run. Holes in the
 4185 master and spacer plates allow the insertion of stainless-steel tubes for the radioactive source
 calibration system.

Each module is constructed by gluing the structures described above into sub-modules on a
 custom stacking fixture. These are then bolted onto the girder to form modules, with care being
 taken to ensure that the azimuthal alignment meets the specifications. The calorimeter is assembled
 4190 by mounting and bolting modules to each other in sequence. Shims are inserted at the inner
 and outer radius load-bearing surfaces to control the overall geometry and yield a nominal module-




Figure 83. Schematic showing how the mechanical assembly and the optical readout of the tile calorimeter are integrated together. The various components of the optical readout, namely the tiles, the fibres and the photomultipliers, are shown.




Figure 84. Azimuthal view of the tile-calorimeter module-to-module interface showing the bearing locations at the inner and outer radii, and the azimuthal gap with a nominal width at the inner radius of 1.5 mm, in which the readout fibres are routed to the photomultipliers inside the girder.




Figure 85. Radial view, looking inwards towards the interaction point, showing the fibre routing in the barrel modules. The fibre shown outside the channel accepts light for the tile onto which it is pressed. Further along the channel, this fibre is routed through a slot in the channel, into the interior region, while one of the other fibres is routed outside to view the next series of scintillator tiles in depth.

to-module azimuthal gap of 1.5 mm and a radial envelope which is generally within 5 mm of the nominal one [115, 118].

5.3.1.3 Optical components Eleven sizes of scintillating tiles (one for each depth in radius) of 3 mm thickness and with radial lengths ranging from 97 mm to 187 mm and azimuthal lengths ranging from 200 mm to 400 mm form the active medium of the tile calorimeter. Ionising particles crossing the tiles induce the production of ultraviolet scintillation light in the base material (polystyrene) and this light is subsequently converted to visible light by wavelength-shifting fluors (the polystyrene is doped with 1.5% PTP as the primary fluor and with 0.044% POPOP as the secondary fluor). Over 460,000 scintillating tiles were produced for the tile calorimeter by injection molding of individual tiles: this eliminated the need for machining to form the trapezoidal shapes and drilling to cut the holes through which the calibration tubes must pass. The tolerance for all dimensions was held to ± 0.10 mm. Approximately 5% of the tile production was tested with a ^{90}Sr radioactive source and the results were used to characterise the light output of each small group of approximately twenty tiles in terms of maximum intensity and attenuation length [119]. Two sources of raw polystyrene were used for tile fabrication; during assembly, the groups of tiles were sorted so that tiles with similar response were inserted in contiguous areas of the detector.

Irradiation tests of tile/fibre assemblies indicated that in the first longitudinal sampling, for an integrated dose corresponding to ten years of operation at the LHC design luminosity, a light loss of less than 10% is expected. Smaller losses will occur in the other samplings, where the radiation dose is less.

Prior to insertion into the calorimeter, the tiles are inserted into a plastic sleeve, which both




Figure 86. Segmentation in depth and η of the tile-calorimeter modules in the central (left) and extended (right) barrels. The bottom of the picture corresponds to the inner radius of the tile calorimeter. The tile calorimeter is symmetric about the interaction point at the origin.

protects the tile and improves the scintillation light yield due to its high reflectivity of 95%. A mask pattern is printed on the sleeve to improve the optical uniformity. The resulting non-uniformity over the surface of a tile is generally below 5% for the sum of signals on both sides of the tile.

Wavelength-shifting fibres placed in contact with the tile edges collect the scintillation light produced in the scintillators and convert it to a longer wavelength. Each fibre collects light from tiles located at one or two radial depths in the calorimeter and transmits it to the PMT's located inside the girder. The fibres used have a diameter of 1 mm, are equipped with a double cladding and are characterised by an emission peak at 476 nm with a decay time of ~ 6 ns. As part of the quality control, light output and attenuation were measured for sample fibres from each of the 65 pre-forms used for fibre production [120]. The fibres had a spread of $\sim 4\%$ in attenuation length and of $\sim 3\%$ in relative light output. To improve the light output, the fibres are aluminised at the end opposite to the PMT. The aluminium mirrors were deposited using magnetron sputtering on bundles of 1261 fibres. The reflectivity was required to be 75% and the spread of its variation is $\sim 7\%$, based on measurements from a sample of fibres in each bundle [119]. The light output spread, at a fixed distance from the aluminised end, relative to a reference fibre, was found to be $\sim 3\%$ for each fibre length [121].

To facilitate assembly, a novel approach was developed to manage the 540,000 fibres required to read out the scintillator tiles and form calorimeter cells. The fibres (four for the central barrel and three for the extended barrel calorimeters) are inserted into a plastic channel much like a straw (see Figs. 83 and 85), which illustrate the scheme whereby depth segmentation is achieved in the readout. The channels are opaque to avoid light cross talk, and painted with a high-reflectivity ink on the side facing the tiles to maximise the light collection. Slots in the channels allow fibres to pass from a location in which they face a tile to a location in which they do not, in order to decouple the fibres at different depths in the calorimeter. The fibres were inserted in these channels using a robot [122] and are glued to the channel at a spot near the aluminised end of the fibre.

5.3.1.4 Instrumentation with scintillator and fibres The light produced in the scintillating material is collected at the edges of each tile using two wavelength-shifting fibres as illustrated



Figure 87. Glued fibre bundle in girder insertion tube (left) and fibre routing (right) for tile-calorimeter module.

4245 in Fig. 83. The fibres are grouped together and coupled to the PMT's which are housed at the outer edge of each module. The fibre grouping is used to define a three-dimensional cell structure in such a way as to form three radial sampling depths, approximately $1.5, 4.1$ and 1.8λ thick at $\eta = 0$. These cells have dimensions $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ in the first two layers and 0.2×0.1 in the last layer. The depth and η -segmentation of the barrel and extended barrel modules are shown
 4250 in Fig. 86. The fibres coupled to each edge of the scintillating tiles are read out by two different PMT's to provide redundancy and sufficient information to partially equalise signals produced by particles entering the scintillating tiles at different impact positions.

Module instrumentation [115] comprised many steps, which ranged from cleaning the slots in the absorber structure to a set of final
 4255 module quality-control measurements based on the individual tile fibre response. Following insertion of tiles and channel/fibre assemblies into a module, the fibres are grouped together to form cells and are glued as a bundle into the fibre-insertion tube, as shown in Fig. 87. These tubes are then fixed into the girder plastic rings mentioned above, to obtain a precise match to the position of the photo-
 4260 multipliers. The tubes and fibres are then cut and polished inside the girder to give the optical interface to the PMT. This interface requires that these fibres be physically present at the time of module instrumentation. However, the gap and crack scintillators described
 4265 in Section 5.4 are mounted only following calorimeter assembly in the cavern. An optical connector is used, therefore, to couple the light from their readout fibres to the already glued and polished optical fibres which penetrate the girder.




Figure 88. Average cell response uniformity, measured using the cesium calibration system for the barrel and the two extended barrel tile calorimeters.




Figure 89. Schematic view of a HEC module, with a cut-away showing the readout structure and the active-pad electronics.




Figure 90. Schematic of the arrangement of the HEC readout structure in the 8.5 mm inter-plate gap. All dimensions are in mm.

Quality-control checks have been made at several moments during the instrumentation process: during fibre bundling and routing, during fibre gluing, cutting and polishing, during tile-fibre optical coupling when the tile was excited by either a blue LED or a ^{137}Cs γ -source. Tile-fibre pairs with a response below 75% of the average response of the tile row for the cell under consideration were repaired in most cases (typically by re-insertion of the plastic channel to improve tile-fibre optical coupling or replacement of the fibre). The average cell response uniformity for production modules, measured using the cesium source, is shown in Fig. 88 and meets the design specification of 10%.

5.3.2 Hadronic end-cap calorimeters (HEC)

5.3.2.1 Overview Figure 91 shows a schematic view of the design of a HEC module. The hadronic end-cap calorimeter (HEC) [123, 105] is a copper/liquid-argon sampling calorimeter with a flat-plate design, which covers the range $1.5 < |\eta| < 3.2$. The HEC shares each of the two liquid-argon end-cap cryostats with the electromagnetic end-cap (EMEC) and forward (FCal) calorimeters (see Fig. 76). The HEC consists of two wheels in each end-cap cryostat: a front wheel (HEC1) and a rear wheel (HEC2), each wheel containing two longitudinal sections. The wheels are cylindrical with an outer radius of 2030 mm. Each of the four HEC wheels is constructed of 32 identical wedge-shaped modules, as illustrated in Fig. 89. A stainless-steel connecting-bar system at the outer wheel perimeter guarantees the mechanical integrity of the wheel structure. At the inner radius, small copper connecting bars link the plates of each set of neighbouring modules. Two sliding

HEC module and section	Total number of channels	Number of faulty channels (ambient)	Number of faulty channels (in liquid)	Percentage (%)
HEC1 (first section)	1536	0	0	0
HEC1 (second section)	1472	0	3	0.20
HEC2 (first section)	1344	0	0	0
HEC2 (second section)	1280	1	2	0.16
Calibration	1024	1	1	0.1

Table 28. Table showing the number of faulty signal channels in each segment and the number of faulty calibration lines for both hadronic end-caps after insertion at room temperature and afterwards in liquid argon.

rails support the wheels inside the cryostat and are an integral part of the mechanical structure. The final vertical deformation of the wheel structure has been measured for the four wheels to represent a sag of 0.3 mm on average. The wheels remain perpendicular to their axis within ± 1.0 mm.

The modules of the front wheels are made of 24 copper plates, each 25 mm thick, plus a 12.5 mm thick front plate. In the rear wheels, the sampling fraction is coarser with modules made of 16 copper plates, each 50 mm thick, plus a 25 mm thick front plate. The gaps in between the plates all have a thickness of 8.5 mm. The resulting sampling fractions for HEC1 and for HEC2 are 4.4% and 2.2% respectively. The wheels have an inner radius of 372 mm for the first nine plates of HEC1 and of 475 mm for the remaining plates of HEC1 and for all 17 plates of HEC2. The structural strength of the modules is provided by seven stainless-steel tie-rods with 12 mm (16 mm) diameter for the front (rear) modules. Annular high-precision spacers on the tie-rods maintain the 8.5 mm gaps between the copper plates.

Three electrodes divide the 8.5 mm gaps into four separate LAr drift zones of 1.8 mm width each. Figure 90 shows schematically the arrangement of the readout structure. The space between the electrodes is maintained using a honeycomb sheet which is also held in place by the seven tie-rods. Each drift zone is individually supplied with high voltage (see Table 26). The middle electrode carries a pad structure covered by a high-resistivity layer, serving as the readout electrode and defining the lateral segmentation of the calorimeter. The two other electrodes carry surfaces of high resistivity to which high voltage is applied. These high-voltage planes form an electrostatic transformer (EST). Electrically, this structure corresponds to a double gap of 3.6 mm each. The EST structure has the advantage of a lower high-voltage for operation, of a double high-voltage safety and thus of reduced problems with space-charge effects due to ion density at large values of $|\eta|$. For the nominal high voltage of 1800 V, the typical drift time for electrons in the drift zone is 430 ns.

The readout cells are defined by pads etched on the central foil in each gap. The arrangement of these pads provides a semi-pointing geometry (see Fig. 91). The size of the readout cells is $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ in the region $|\eta| < 2.5$ and 0.2×0.2 for larger values of η .

The signals from the pads of this electrode structure are amplified and summed employing the concept of active pads [124]: the signals from two consecutive pads are fed into a separate




Figure 91. Schematic $R - \phi$ (left) and $R - z$ (right) views of the hadronic end-cap calorimeter. The semi-pointing layout of the readout electrodes is indicated by the dashed lines. Dimensions are in mm.

preamplifier (based on GaAs electronics). The use of GaAs preamplifiers in the cold provides the optimum signal-to-noise ratio for the HEC. An important aspect of the HEC is its ability to detect muons and to measure any radiative energy loss. The density of the electronics on the HEC wheels with their rather modest number of read-out channels (5632 in total) and the power consumption of the GaAs integrated circuits (15 mW for one preamplifier channel and approximately 200 mW for the entire chip) are sufficiently low that the heating effect of the electronics on the liquid argon does not produce bubbling. The outputs of the preamplifiers are summed on the same GaAs integrated circuit to produce one signal from each cell. The signal sent to the feed-through for each cell is thus comprised of the amplified and summed signals of the eight or sixteen pads with the same η and ϕ within a readout section.

5.3.2.2 Wheel assembly and installation Fig. 92 shows a HEC wheel fully assembled on its assembly table. The geometrical precision of the wheel is given by 32 datum pins on the assembly table. During the wheel assembly, each module had to pass a series of quality-assurance tests: high-voltage reliability, capacitance control, electronic cabling and signal reconstruction verification using the calibration procedures in warm and in cold. These tests were repeated after the wheel




Figure 92. Photograph of a fully assembled HEC wheel on the assembly table. The active-pad electronic boards located on the outer circumference are clearly visible.

assembly, after the wheel rotation, after the wheel insertion, after the full cabling of the HEC1 and HEC2 wheels inside the cryostat and finally through the feed-throughs of the cryostat.

4340 After closing the end-cap cryostat, each end-cap has been cooled down, filled with LAr and the final cold tests prior to the movement to the ATLAS cavern have been performed. For the HEC, these tests included quality-assurance measurements and cabling checks of all signal and calibration lines, full calibration and delay scans of all signal channels, full pulse-shape analysis of all signal channels, detailed noise measurements of all channels including coherent noise and a long-term high-voltage test of all high-voltage lines. Typically, three signal channels per end-cap were
 4345 found to be not fully operational, corresponding to 0.1 % of all channels, as shown in Table 28. The fraction of high-voltage lines with shorts was $\sim 4\%$ in liquid argon, but, most importantly, with four individual high-voltage lines per individual HEC gap, all HEC regions are operational (see also Table 26 which lists the nominal high-voltage settings for the HEC).

5.3.3 Forward calorimeters

4350 **5.3.3.1 Overview** The forward calorimeters (FCal) are located in the same cryostats as the end-cap calorimeters and provide coverage over $3.1 < |\eta| < 4.9$. The close vicinity and coupling between these systems result in a quite hermetic design, which minimises energy losses in cracks




Figure 93. Schematic diagram showing the three FCal modules located in the end-cap cryostat. The material in front of the FCal and the shielding plug behind it are also shown. The black regions are structural parts of the cryostat. The diagram has a larger vertical scale for clarity.




Figure 94. Electrode structure of FCal1 with the matrix of copper plates and the copper tubes and rods with the LAr gap for the electrodes. The Molière radius, R_M , is represented by the solid disk.

between the calorimeter systems and also limits the backgrounds which reach the muon system. As the FCal modules are located at high η , at a distance of approximately 4.7 m from the interaction point, they are exposed to high particle fluxes. This has resulted in a design with very small liquid-argon gaps, which have been obtained by using an electrode structure of small-diameter rods, centred in tubes which are oriented parallel to the beam direction. The liquid-argon gaps are smaller than the usual 2 mm gap of the electromagnetic barrel calorimeter to avoid ion build-up problems and to provide at the same time the highest possible density. These smaller gaps also lead to a faster signal with roughly the same instantaneous current but smaller integrated current. In the electromagnetic compartment (FCal1), the triangular current pulse at the electrode has a full drift time of 60 ns. For FCal2 and FCal3, the full drift time scales with the gap size since the field in the gaps is similar for all three modules.

5.3.3.2 Module description Each FCal is split into three 45 cm deep modules: one electromagnetic module (FCal1) and two hadronic modules (FCal2 and FCal3), as illustrated in Fig. 93. To optimise the resolution and the heat removal, copper was chosen as the absorber for FCal1, while mainly tungsten was used in FCal2 and FCal3, to provide containment and minimise the lateral spread of hadronic showers. A shielding plug made of a copper alloy has been mounted behind FCal3 to reduce backgrounds in the end-cap muon system. A full description of the FCal detector can be found in [125].

The basic parameters of the FCal modules are listed in Table 29. The FCal1 compartment is made of copper plates stacked one behind the other. The plates have 12,260 holes drilled in them through which the electrode structures are inserted. An electrode consists of a co-axial copper rod and copper tube separated by a precision, radiation-hard plastic fibre wound around the rod. The arrangement of electrodes and the effective Molière radius for the modules can be seen in Fig. 94. Mechanical integrity is achieved by a set of four tie-rods which are bolted through the structure. The electrode tubes are swaged at the signal end to provide a good electrical contact.




Figure 95. View of the FCal hadronic module absorber matrix, including a set of tungsten rods and copper tubes surrounded by 1 cm long tungsten slugs.




Figure 96. Schematic of the FCal1 module cabling from the electrodes to the cryogenic feed-through. The other modules differ only by the number of rods grouped together on the interconnect board (six for FCal2 and nine for FCal3).



Figure 97. Assembly of FCal modules: from left to right, the three modules plus the copper alloy plug can be seen on the support mandrel with most of the cabling in place.



Figure 98. Completed FCal assembly with its bulkhead and cone attached, just before insertion into the end-cap cryostat.

The hadronic modules FCal2 and FCal3 are optimised for a high absorption length. This is achieved by maximising the amount of tungsten in the modules. These modules consist of two copper end-plates, each 2.35 cm thick, which are spanned by electrode structures, similar to the ones used in FCal1, except for the use of tungsten rods instead of copper rods. Swaging of the copper tubes to the end-plates is used to provide rigidity for the overall structure and good electrical contact. The space between the end-plates and the tubes is filled with small tungsten slugs, as shown in Fig. 95. The inner and outer radii of the absorber structure formed by the rods, tubes and slugs are enclosed in copper shells.

Signals are read out from the side of FCal1 nearer to the interaction point and from the sides of FCal2 and FCal3 further from the interaction point. This arrangement keeps the cables and

	FCal1	FCal2	FCal3
Function	Electromagnetic	Hadronic	Hadronic
Mass of module (kg)	2119	3826	3695
Main absorber material	Copper	Tungsten	Tungsten
LAr gap width (mm)	0.269	0.376	0.508
Radiation length (X_0)	27.6	91.3	89.2
Absorption length (λ)	2.66	3.68	3.60
Number of electrodes	12260	10200	8224
Number of readout channels	1008	500	254

Table 29. Table of parameters for the three FCal modules

Module	Channels		
	Total number of channels	Number of dead channels	Number of faulty channels
FCal1 (side A)	1008	0	26
FCal2 (side A)	500	1	4
FCal3 (side A)	254	1	59
FCal1 (side C)	1008	5	51
FCal2 (side C)	500	0	8
FCal3 (side C)	254	1	4

Table 30. Summary of faults found in the FCal modules after installation in ATLAS. Shown for each module are the total number of channels, the number of dead channels and the number of faulty channels which will require special treatment in the reconstruction of the signal (see text). Faulty channels correspond to either missing signal channels or signal channels with a bad signal shape [125].

connectors away from the region of maximum radiation damage which is near the back of FCal1. Readout electrodes are hard-wired together with small interconnect boards on the faces of the modules in groups of four, six and nine for FCal1, FCal2 and FCal3 respectively. The signals are then routed using miniature polyimide co-axial cables along the periphery of the FCal modules to summing boards which are mounted on the back of the HEC calorimeter. The summing boards are equipped with transmission-line transformers which sum four inputs. High voltage (see Table 26) is also distributed on the summing boards via a set of current-limiting resistors, as shown in Fig. 96 for the specific case of FCal1. The signal summings at the inner and outer radii of the modules are

in general different due to geometric constraints and higher counting rates at the inner radius [125].

5.3.3.3 Integration into the cryostat There are two sets of FCal modules, one on each side of the interaction point. Each set of modules, along with a shielding plug, are mounted inside a cylindrical tube, known as the FCal support tube, which is shown as the black horizontal line between the FCal and the HEC in Fig. 93. This support tube, along with an upstream conical nose and a downstream flange are the structural components at the inner radius of the end-cap cryostat. To load the modules into the support tube, a mandrel was used as a temporary support for the modules. This stage of the assembly is illustrated in Fig. 97. This arrangement allowed the readout cables to be dressed in the cable trays along the sides of the modules. The support tube was then slid over the modules, the module weights were transferred to the support tube and the cone and flange were attached to the support tube, as shown in Fig. 98. The complete unit was then installed into the end-cap cryostat. At this point, the cables were connected to the FCal summing boards.

5.3.3.4 Placement of FCal There is only a small clearance between the FCal modules and the FCal support tube, such that the $x - y$ position of each module is fixed relative to the tube. The central axis of the detectors is about 1 mm below the central axis of the support tube. The distance between modules is constrained by bumpers attached to the end-plates of the modules. In this manner, the modules can move in the z direction within the tubes, but will be within approximately 1 mm of their design location. The survey of the FCal is then a matter of a survey of the FCal support tube. This was carried out after the installation of the FCal into the cryostat. The central axis of the support tube is within about 2 mm of the central axis of the cryostat outer warm vessel and has been surveyed to a precision of about 0.3 mm.

5.3.3.5 Tests during installation/integration A series of quality-control checks were carried out on the modules. These included high-voltage testing, capacitance measurements and reflection testing. The results of the capacitance measurements provided a measure of the uniformity of the electrodes. High-voltage checks were carried out at each stage of assembly and integration of the detectors both at warm and at cold. It was found that the number of shorts did vary during this process. The reflection tests were used to look for other detector faults, such as disconnected channels and broken ground braids.

A summary of the faults observed in the FCal modules after the cryostats were installed in the ATLAS cavern and filled with liquid argon can be found in Table 30. The right-hand column shows the number of readout channels per module which have a fault leading to a distorted signal, a reduction in signal amplitude or to both of these effects. It is expected that the output from these channels can be adjusted in the signal reconstruction to partially recover the true amplitude, including when the fault arises from high-voltage problems. The 1762 channels in each FCal are served by 112 high-voltage lines. Two of these were lost during the FCal integration into the end-cap cryostat, one for FCal1 on side C and the other for FCal3 on side A. This accounts for the large number of faulty channels shown for these two modules in Table 30. There are sixteen high-voltage shorts on the FCal on side A and nine on side C. These account for roughly half of the remaining faults listed in Table 30. The other half are dominated by damaged blocking capacitors (see Fig. 96). A more detailed description of the faults observed in the modules can be found in Ref. [125].




Figure 99. Schematic of the transition region between the barrel and end-cap cryostats, where additional scintillator elements are installed to provide corrections for energy lost in inactive material (not shown), such as the liquid-argon cryostats and the inner-detector services. The plug tile calorimeter is fully integrated into the extended barrel tile calorimeter. The gap and cryostat scintillators are read out together with the other tile-calorimeter channels (see text).

5.4 Instrumentation in gap between cryostats

The transition region or gap between the barrel and end-cap calorimeters is filled with cables and services for the inner detector as well as power supplies and services for the barrel liquid-argon calorimeter. One of the standard gap regions is shown in Figs. 86 and 99. At the outer radius of the detector, a reduced section of a standard tile-calorimeter sub-module, the plug, provides additional coverage in this region and significantly reduces the neutron flux from the inner-detector volume into the muon system.

In the region of the tile-calorimeter module where part of the calorimeter structure has been removed, two scintillators (E1 and E2 shown in Fig. 86) with independent readout (the gap scintillators) are attached to the module surface. The gap scintillators cover the region $1.0 < |\eta| < 1.2$ and provide signals which can be used to correct for energy losses in the inactive material in the gap.

Two scintillators (E3 and E4 shown in Fig. 86), the cryostat scintillators, penetrate the region between the outer walls of the barrel and end-cap cryostats. These are physically attached to the tile calorimeter and read out through its electronics. The cryostat scintillators cover the region $1.2 < |\eta| < 1.6$ and provide a signal which can be used to correct for energy losses in the outer wall of the barrel cryostat and in the inner-detector services. Both gap and crack scintillators have full coverage in ϕ , with a segmentation of $\Delta\phi = 0.1$.

For initial data-taking at low luminosities, two sets of sixteen scintillator counters installed on the front face of the end-cap calorimeter cryostats will be used to trigger on minimum-bias events. Each set of counters is segmented in eight units in ϕ and two units in η . They are located




Figure 100. Expected electronic noise in individual cells of the various sampling layers of the calorimeters as a function of $|\eta|$. Note that the presampler noise is corrected for by the appropriate sampling fractions as discussed in Section 5.5.2.1.

at $|z| = 3560$ mm, the innermost set covers radii between 153 mm and 426 mm, corresponding to the region $2.82 < |\eta| < 3.84$ and the outermost set covers radii between 426 mm and 890 mm, corresponding to the region $2.09 < |\eta| < 2.82$. The readout for these counters will be provided by that of sixteen of the cryostat scintillators covering the region $1.2 < |\eta| < 1.4$. Eight of these cryostat scintillators had to be removed on each side for the routing of the signals from the minimum-bias scintillators. They will be reinserted for operations at full luminosity.

There are several locations around the circumference of the calorimeter, in which the standard gap instrumentation is not possible, for example the region opposite the support feet for the barrel cryostat [115]. In these locations, the plug calorimeter is reduced further in its length. In the location of the exit of the barrel cryogenic line, no plug calorimeter can be accommodated and it is replaced entirely by scintillators. Finally, the inner support of the end-cap cryostat sits on the tile calorimeter at its inner radius. In these regions, the tile-calorimeter modules themselves are cut to provide the support surface on which the cryostat support jacks sit. The gap scintillators are also missing in this region.

5.5 Calorimeter readout electronics, calibration and services

5.5.1 Readout electronics

5.5.1.1 Overview This section describes the on-detector (front-end) and off-detector (back-end) electronics of the ATLAS calorimetry. The major tasks required of the readout electronics are to provide the first-level (L1) trigger system with the energy deposited in trigger towers of size $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ and to measure, for L1-triggered beam crossings, the energy deposit in each calorimeter cell to the specifications set by the performance requirements. The dynamic range to be covered ranges from a maximum of 3 TeV down to a lower limit of ~ 10 MeV set by thermal noise in the calorimeters. The thermal noise from the front-end electronics varies widely across the various calorimeter systems, as illustrated in Fig. 100. Coherent noise over the many cells used to measure electron/photon energies in the electromagnetic calorimeters and jet energies and E_T^{miss} in the overall calorimetry should be kept below 5% of the incoherent noise. The readout should proceed without any appreciable dead-time up to a trigger rate of 75 kHz.

As already described for the overall experiment in Section 3.1, all electronics components situated in the cavern have to be radiation-tolerant to specifications including large safety factors. However, the necessity for a large dynamic range and for very low coherent noise favours a readout architecture with analogue processing very close to the detectors. These quite stringent requirements have led to solutions where the analogue signals are amplified, processed and digitised by radiation-tolerant front-end electronics mounted very near the detectors (just above the signal feed-throughs for the LAr calorimeters and in drawers incorporated into the girders for the tile calorimeters) and followed by subsequent digital processing of the signals in back-end systems which are located off-detector. For example, the overall readout architecture of the LAr calorimeters is sketched out in Fig. 101. This overall electronics readout system [126, 127] is used for all the LAr calorimeters and, as shown in Fig. 101, provides also the necessary input for the L1 calorimeter trigger (see Section 8.2).

The radiation levels in the vicinity of the LAr cryostats and tile drawers (fluences of 10^{12} neutrons per cm^2 per year and ionising doses of 20 Gy per year) have required the development and use of a large number of specifically designed radiation-hard ASIC's for many key components of the front-end electronics of the calorimeters. The back-end systems, located in the main services cavern (USA15), where there are no radiation issues, use commercial components. The back-end electronics systems for all LAr calorimeters and for the tile calorimeter have been designed with the same architecture, although some important differences exist in the functional implementation.

As discussed in Section 9.4.8, an important effort has been invested in implementing the grounding and electromagnetic compatibility rules for ATLAS during the design and installation phases of the calorimeters and of their electronics. This has been particularly demanding for the LAr cryostats, since all services, such as cryogenic lines, cooling pipes, etc., and all mechanical connections to the surrounding detector systems had to be installed in a way which preserved the electrical isolation of the cryostats. The most difficult challenge has been related to the large amount of pipes which run along the cryostat walls and enter the inner bore of the barrel cryostat to service the inner detector. The cryostat surfaces are wrapped in insulating foils or covered by thin glass-fibre sheets to prevent any accidental electrical contact. Concerning the calorimeter electrical connections, the DC isolation of the power to the front-end boards is taken care of by the




Figure 101. Block diagram of the LAr readout electronics. The cold electrical circuit is depicted at the bottom, followed above by the on-detector front-end electronics crate and at the top (left) by a schematic view of the readout crate with its ROD boards and TTC modules. Also indicated at the middle and top (right) are the LAr front-end tower builder electronics and the interfaces to the L1 trigger system with its central trigger processor (CTP).

DC-DC converters of the low-voltage power supplies (see Section 5.5.3.1). The shields of all cables connected to the cryostats are interrupted at one end. For example, the shield of the L1 cables (the only copper links in the calorimeter readout) is disconnected at the entrance of the receiver boards. The cryostats themselves are connected to the safety ground [128] by long cables with a large copper cross-section to a single point close to the L1 receiver racks. This choice minimises the potential difference between the analogue sending and receiving ends of the trigger signals. The tile calorimeter itself is structurally supported on rails which are grounded to the cavern safety network. Additionally, dedicated grounding cable connections for each of the three tile calorimeter barrels are made to the cavern network [129].

5.5.1.2 LAr calorimeter front-end electronics The sensitive analogue electronics [126] are located on the detector. Inside the cryostat, the calorimeter electrodes are grouped to form readout cells and small coaxial cables bring the cell signals to the cold-to-warm feed-throughs (see Section 5.1), except for the HEC (see Section 5.3.2). Custom front-end crates are mounted near the feed-throughs and contain several types of boards:

1. The calibration board [130] injects precisely known current pulses through high-precision resistors to simulate as accurately as feasible energy deposits in the calorimeters. These resistors are mounted on the detector in the cold for the barrel and end-cap electromagnetic calorimeters and for the HEC, and mounted on the base-plane of the front-end crates for the FCal (see Section 5.5.2).
2. Front-end boards [131] (FEB) receive the raw calorimeter signals and, as described in more detail below, perform the analogue processing, digitisation and transmission off-detector of the calorimeter signals.
3. Tower-builder boards perform the final level of analogue summation for the barrel and end-cap electromagnetic calorimeters to form trigger-tower signals and transmit the analogue signals to the L1 trigger processor, where they are digitised and processed. The corresponding boards which transmit the HEC and FCal trigger towers are called tower-driver boards, since no summation is performed in this case.
4. Front-end crate-controller boards receive and distribute the 40 MHz clock, the L1 trigger accept signal, as well as other fast synchronous signals and information to configure and control the front-end electronics boards (SPAC protocol).
5. Additional boards transmit the information from the mechanical sensors monitoring possible stresses of the detector, from the temperature gauges situated in the liquid argon, and from the purity monitors.

A total of 58 front-end crates are needed to read out the complete LAr calorimeter system. The configuration of such a crate is typically the following (e.g. for the barrel electromagnetic calorimeter): 28 FEB's and two each of the calibration, trigger-tower builder and front-end crate controller boards described above.




Figure 102. Block diagram of the FEB architecture, depicting the dataflow for four of the 128 channels (see text for details).

The analogue processing and digitisation of the calorimeter signals is performed by the FEB's. Each FEB processes up to 128 calorimeter channels, and a total of 1524 FEB's are used to read out the 182,468 channels of LAr calorimetry. A block diagram indicating the main features of the FEB architecture is shown in Fig. 102. The raw signals from the calorimeter are mapped onto the FEB inputs as they emerge from the cryostat feed-throughs. Four-channel preamplifier hybrids amplify these raw signals (in the specific case of the HEC, preamplifiers mounted on the detector and operating in the cold provide some amplification of the signals before they reach the FEB's).

The preamplifier outputs are AC-coupled into a four-channel shaper ASIC, which splits each signal into three overlapping, linear gain scales in the ratio 1/10/100, and applies a bipolar $CR - (RC)^2$ analogue filter to each scale to optimise the signal-to-noise ratio. The single differentiation step serves

to remove the long tail from the detector response, while the two integrations limit the bandwidth in order to reduce the noise. The time constant is chosen to be 15 ns to minimise the overall noise level, as a compromise between the two main contributions, namely electronic noise and pile-up noise, which respectively decrease and increase with the shaping time. This value is optimised for a luminosity of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. The triangular input current pulse from the detector and the shaped output pulse from the FEB are depicted for the case of a barrel electromagnetic cell in Figure 103.

The shaped signals are then sampled at the LHC bunch-crossing frequency of 40 MHz by four-channel switched-capacitor array (SCA) analogue pipeline ASIC's. The SCA stores the analogue signals during the L1 trigger latency in pipelines of 144 cells, and also serves as a multiplexer and de-randomising buffer in front of the ADC for triggered events.

For events accepted by the L1 trigger, typically five samples per channel and only one of the three gain scales are read out from the SCA. A gain-selector chip (GSEL) is used to select the optimal readout gain individually for each calorimeter channel and separately for each L1 trigger. Two




Figure 103. Amplitude versus time for triangular pulse of the current in a LAr barrel electromagnetic cell and of the FEB output signal after bi-polar shaping. Also indicated are the sampling points every 25 ns.




Figure 104. Block diagram of tile-calorimeter readout electronics.

commercial dual op-amp chips couple the SCA outputs to commercial 12-bit ADC's used for the digitisation of the analogue signals. The signals are formatted by the GSEL, multiplexed (SMUX), serialised at 1.6 Gbits/s (GLINK), and then transmitted via a single optical transmitter (OTx) based on vertical-cavity surface-emitting laser diodes (VCSEL).

Each shaper ASIC contains an analogue summing circuit, which sums its four channels as the first step in producing sums for the L1 trigger system. Channels can be individually enabled or disabled during FEB configuration in order to mask noisy channels. The four-channel sums are routed on the FEB to plug-in boards, which contain the next stage of the trigger analogue summing tree and which drive their output sums through the crate back-plane to the tower-builder (or tower-driver) boards, where the sums are completed before transmission off-detector to the L1 trigger system.

5.5.1.3 Tile calorimeter front-end electronics The front-end electronics of the tile calorimeter are housed inside drawers located within the steel girders which constitute the external support frame of the tile-calorimeter modules (see Section 5.3.1.1). A block diagram of the tile-calorimeter front-end electronics and readout components inside the drawer is shown in Fig. 104.

A key element in the readout is the photomultiplier (PMT) block [132]. It is a mechanical structure comprising a steel cylinder and mu-metal shield for magnetic shielding, which contains a light mixer, a photomultiplier tube, a voltage divider and the so-called 3-in-1 card. The light

mixer is an optical plastic insert which mixes the light from the readout fibres to ensure uniform illumination of the photo-cathode. The PMT's with their compact 8-dynode structure are used to measure the scintillation light [133]. All PMT's were burned in and tested for linearity, stability, 4605 dark current and operating voltage for a nominal gain of 10^5 [115]. The average operating voltage for nominal gain is 680 V. The assembled PMT blocks are inserted inside precision slots inside the aluminium structure of the drawers, which ensure accurate placement of the light mixer relative to the fibre bundle for each readout cell.

The high voltage for each PMT is regulated with an accuracy of better than 0.25 V by micro- 4610 controllers inside the drawer. A single micro-controller services up to 48 PMT channels and has a memory into which the nominal and set values of the voltage for each PMT are loaded. Each high-voltage distributor system drives seven temperature probes located at key locations inside each drawer.

Most analogue functions of the front-end electronics are contained on the 3-in-1 card [134]. 4615 This board provides unipolar signal shaping to yield a pulsewidth of 50 ns and two linear outputs with a relative gain of 64 to achieve the overall 16-bit dynamic range required. The full scale on the low-gain channel corresponds to 800 pC (the PMT gain is adjusted to provide approximately 1 pC/GeV). Charge-averaging amplifiers on the 3-in-1 card are used for cell calibration and monitoring and a charge-injection system is incorporated to calibrate the readout over its full 4620 dynamic range. The 3-in-1 card also produces the analogue signals for the L1 trigger.

The digitiser system [135] samples the incoming data from the 3-in-1 cards every 25 ns using two dedicated 10 bit ADCs (one for each of the low-gain and high-gain channels) and stores the information in a fixed-length pipeline, while waiting for a L1 trigger accept (L1A). Each digitiser services up to six PMT channels. Each triggered event is recorded over an extended time frame, 4625 with a programmable length of up to 16 samples. In normal data-taking mode, following a L1A signal, seven samples are kept: one close to the peak, four before the peak (with the two first samples providing a measurement of the signal pedestal) and two after the peak.

There is one interface board [136] per pair of drawers. It receives and distributes the timing, trigger and control (TTC) signals, collects and formats data from the digitiser cards, and transmits 4630 the digitised data via an optical link. The interface board also performs cyclic redundancy checks on the input and output datastreams.

A system of four linked mother-boards form the base of the readout system. They distribute power to the 3-in-1 cards, the ADC integrator, the mezzanine card and the adders. Each mother- 4635 board also contains circuitry which generates a precise reference voltage used by the 3-in-1 cards for charge injection and integrator calibration. The mezzanine card is mounted on the first mother-board in the drawer. It decodes control commands sent via the TTC system and sends them on differential serial lines to the 3-in-1 cards. These commands control the charge injection, integrator gain setting and calibration, trigger-output gating and integrator bus access for each 3-in-1 card. The mezzanine card can also receive commands via the ADC integrator card, which is mounted on 4640 the first mother-board. It sends and receives messages via a dedicated CANbus daisy chain. The micro-processor in the ADC integrator card queues commands to the mezzanine card. It sets the integrator gain on a 3-in-1 card, multiplexes the 3-in-1 card onto the integrator readout bus and reads out the charge induced by the cesium source, by minimum-bias events, or by the on-board calibration DAC. The ADC integrator can also read back 3-in-1 card status and control information.

4645 Adder boards [137] are distributed along the drawer mounted on the mother-boards. Each
 adder board receives the analogue trigger outputs from up to six 3-in-1 cards. The inputs are
 connected to form trigger towers. The adders perform an analogue sum of the input signals and
 send two output signals via long cables to the L1 trigger system. One of these signals (tower
 signal) comprises the sum of all samples, the second one (muon output) contains only the last
 4650 tile-calorimeter sampling. In the case of the gap and crack scintillators (see Section 5.4), the
 adder board provides amplification to match the small signals from these scintillators to the signal
 amplitudes required by the L1 trigger boards. It also provides the sum of the signal in all four of
 the gap and crack scintillators. The muon output from these adder boards (see Fig. 104) provides
 the signal from the scintillator covering the region $1.2 < |\eta| < 1.4$. As mentioned in Section 5.4,
 4655 sixteen of these channels will be used for the trigger from the minimum-bias trigger scintillators
 during initial data-taking.

5.5.1.4 Back-end electronics

Overview of architecture The back-end system [127, 138] for all calorimeters is housed
 in the main services cavern (USA15), located 70 m away from the detector, and contains three
 4660 different sub-systems: the readout driver (ROD) system, which constitutes the core of the back-
 end, the trigger, timing and control (TTC) system, and the L1 trigger receiver. The back-end system
 reads the data from the front-end electronics, and digitally processes the data before transmitting it
 to the data acquisition system at a L1 trigger rate of up to 75 kHz for the physics data-taking mode
 and at a 10 kHz rate for the calibration mode. The back-end system also receives the TTC signals,
 4665 distributes the timing clock and trigger to the front-end electronics and the ROD's, and configures
 and controls the front-end crates. All of these tasks must be performed fast enough to cope with
 the above trigger rates.

The ROD's are responsible for receiving, digitally processing and formatting the data from
 the front-end electronics and for performing various data-integrity checks and higher-level moni-
 4670 toring tasks. The ROD module comprises a ROD mother-board on which are plugged up to four
 processing-unit mezzanine cards. Each processing unit houses two commercial Digital Signal Pro-
 cessor (DSP) chips operating at 720 MHz. Each DSP is preceded by an input buffer used to perform
 data-integrity checks as well as to reformat the input data to suit the DSP architecture. Following
 each DSP is an output buffer used to buffer processed events before transfer to the data acquisition
 4675 system. For each trigger, data from the TTC stream and from the front-end electronics are pushed
 into the ROD modules, where they are checked, processed, formatted and sent on a readout link to
 the ATLAS data acquisition system.

Optimal filtering The most important task of the DSP is to apply an optimal-filtering method
 to the signals, including the relevant calibration constants, in order to provide the deposited en-
 ergy (in GeV) and the signal timing (in ns) for each cell [139]. This optimal-filtering method is at
 the core of the back-end processing role and has been implemented in all the calorimeter ROD's.
 It calculates the energy E , the timing τ and a simplified χ^2 for the shaped and digitised signal from
 each cell, by means of a weighted sum over the relevant sample amplitudes, s_i , expressed as:

$$E = \sum_{i=1}^n a_i(s_i - ped) \quad E\tau = \sum_{i=1}^n b_i(s_i - ped) \quad \chi^2 = \sum_{i=1}^n (s_i - ped - E(g_i - \tau g'_i))^2, \quad (5.1)$$

where n is the number of samples and ped is the pedestal value of the corresponding readout channel. The optimal filtering weights, a_i and b_i , are evaluated while minimising the dispersion in E and $E\tau$ arising from electronics and pile-up noise, taking into account the noise auto-correlation matrix and the pulse shape at the input of the digitiser.

Trigger, timing and control Four VME crates are used to implement the timing trigger and control (TTC) system, where up to two partitions are implemented in a single crate. Each TTC partition contains a controller and a local trigger processor (see Section 8.2.3 and Fig. 171).

The trigger latency, which is the delay between the bunch-crossing time and the time when the L1 accept signal (L1A) arrives to the front-end electronics, has been minimised to a value below $2.5 \mu\text{s}$. The TTC rack location has been chosen to minimise the length of the TTC fibres to the front-end crates and the associated contribution to the trigger latency. In addition, the programmable delay lines of the calibration boards are preset to reproduce the timing of signals generated by particles originating from the interaction point.

Receiver system The receiver system is part of the trigger-sum chain and interfaces the tower-builder and tower-driver boards to the calorimeter L1 processor. One important function of the trigger-sum chain is to convert the signal from energy to transverse energy. The final gain adjustment is left to the receiver. Each receiver crate contains 16 receiver modules, two monitoring modules and one controller module. The LAr system consists of six receiver crates: two for the barrel electromagnetic, two for the EMEC's and two for the HEC and FCal. Two additional crates are for the tile calorimeters.

LAr back-end electronics As described above, the optimal-filtering method calculates the energy for each cell while minimising the noise and pile-up contributions. For cells with an energy above a certain threshold (approximately 5% to 10% of all cells), the precise timing of the signal as well as a χ^2 -like quality factor, allowing to flag cells with a large pile-up contribution, are determined and transmitted as well. For an even lower fraction of cells, the individual samples are also transmitted.

The pulse after the bipolar shaping is sampled every 25 ns (see Fig. 103). For normal data-taking, five samples are typically used, whereas for calibration and commissioning, up to 32 samples can be taken. The energy E , time τ and simplified χ^2 of the signal pulse are calculated as in Eq. 5.1. In the weighted sum, g_i and g'_i are respectively the normalised ionisation pulse shape and its




Figure 105. Plot of the electronic noise in the electromagnetic barrel as a function of the number of readout samples. The circles correspond to a middle-layer cell and the triangles to a strip-layer cell. For the case of five readout samples, as planned for normal data-taking, the electronic noise is reduced by a factor of ~ 1.7 compared to the case of only one readout sample.

4720 derivative with time. They are estimated in different ways for the various calorimeters, using a number of methods: the pulse-shapes measured in test-beam, the measured calibration pulse, and a detailed modeling of the front-end electronics and readout chain [127]. These methods must take into account the different signal injection points and the differences between the shapes of the ionisation and calibration signals.

4725 As an illustration, Fig. 105 shows the variation of the noise with the number of samples, as obtained in situ for the strip and middle compartments of the barrel electromagnetic calorimeter. For the chosen number of five samples used in the back-end for normal LHC operation, the electronic noise is reduced by a factor ~ 1.7 with respect to a readout which would use only one sample.

4730 To increase modularity and allow for concurrent running of the various parts, the LAr calorimetry is split into six partitions: one for each half-barrel, one for each EMEC, and one for HEC and FCal on a given side. Each partition is associated to a computer, which controls and monitors the system, and to a TTC sub-system. A typical barrel partition (see [127] for the configuration of the other partitions) is composed of four ROD crates and the associated front-end electronics. Each ROD crate contains 14 ROD boards and drives and controls four front-end crates.

4735 **Tile calorimeter back-end electronics** In addition to the ROD system described above, modified slightly to have two rather than four processing units per ROD board, the tile-calorimeter back-end electronics system also includes an optical multiplexer board [140]. This board has the responsibility of checking data coming from the front-end electronics. There are two fibres per front-end drawer carrying the same data and the board selects the error-free one in real time by checking the integrity of the data received using a cyclic redundancy check. It can also act as data injector for the ROD system.

4740 The ROD system of the tile calorimeter is physically divided into four independent partitions, two for the readout of the barrel calorimeter and one for each extended barrel calorimeter. Each partition is composed of a 6U VME TTC crate and a 9U VME ROD crate (containing eight ROD processing-unit boards), which read out up to 64 front-end drawers. The ROD processing unit performs various data-integrity checks and reconstructs the amplitudes of the digitised signals using the optimal-filtering method described above. Several algorithms based on optimal filtering are being developed to deal with detector and pile-up noise. The optimal-filtering method is also used to reconstruct the signal in charge-injection calibration runs.

5.5.2 Calorimeter calibration systems

4750 The LAr and tile calorimeters have very precise charge-injection systems and calibration boards to calibrate the response of the front-end electronics boards. The tile calorimeter is equipped with two additional calibration systems to monitor the behaviour with time of the photomultipliers and of the optical components.

4755 **5.5.2.1 Calibration of LAr calorimeters** The aim of the electronic calibration is to determine the conversion factor, expressed as the ratio between the measured signal in ADC counts to the signal current in μA , for each individual channel and to monitor this factor with time.

The use of fast shaping for the calorimeter readout requires that the distribution of the calibration signals be done via injection resistors directly placed at the input of the detector cell and




Figure 106. Schematic diagram of the calibration system for the LAr electromagnetic calorimeters. The components in the left part of the diagram are located on the calibration board itself at ambient temperature in the front-end crates, whereas the right part of the diagram depicts the distribution of the calibration signal into the calorimeter cells. This is achieved through the precision resistors (R_{inj}), which are located on printed-circuit boards in the liquid argon. The impedance of the cable between the calibration board and the printed-circuit boards in the liquid argon is denoted Z_c .

not at the input of the preamplifier located outside the cryostat. This is shown in Fig. 106 for the barrel and end-cap electromagnetic calorimeters for which the greatest precision is required. In the case of the HEC, the calibration pulse is injected in front of the cold preamplifiers. In the case of the FCal, however, the calibration pulses are applied at the base-plane of the front-end crates, where the signals are split into two. One of the split pulses goes directly into the front-end boards and the other goes through the cold electronics chain, reflects off the electrodes and is observed as a delayed pulse. The direct pulses are used to calibrate the electronics in the front-end crates and the delayed pulses are used to examine problems related to the FCal and its cold electronics.

The conversion of the signal in μA to a signal in GeV has been obtained for all LAr calorimeters from test-beam measurements. Figure 107 shows these conversions factors for the electromagnetic calorimeters. This factor only changes once in the barrel when the absorber thickness changes at $|\eta| = 0.8$, whereas it changes almost continuously in the end-cap calorimeter because of the geometry and of the varying high-voltage. For the presampler cell energy scale, an effective

sampling fraction has been applied with a value of 1/20 for the barrel and of 1/60 for the EMEC. These values correspond approximately to the weights which have to be applied to the presampler signal to correct for the energy loss of electrons before reaching the calorimeter. The conversion factors in Fig. 107 agree with predictions based on Geant 4 simulations and knowledge of the LAr ionisation energy and drift velocity [141].

In the HEC1 wheel, this conversion factor is $0.306 \mu\text{A}/\text{GeV}$ and it is half that value for HEC2 with its twice thicker absorber plates. The corresponding conversion factors for the FCal are 1.49, 0.69 and $0.64 \mu\text{A}/\text{GeV}$, respectively for FCal1, FCal2 and FCal3.

The calibration boards [130] are built to deliver a uniform, stable, linear and precise voltage-driven signal with an amplitude and shape as close as possible to the calorimeter ionisation current signal over a 16-bit dynamic range (with currents up to 100 mA over 50Ω). The triangular ionisation pulse is in practice approximated by an exponential shape. The chosen distribution scheme, typically one calibration channel every four signal channels, also allows the measurement of cross-talk [141]. The largest cross-talk is between neighbouring strips in the first compartment of the electromagnetic calorimeter. It amounts to 5 – 8% and is corrected for in the

strip energy measurement. The calibration boards are equipped with programmable delays which are used to compensate the timing differences due to the cable lengths. Apart from pedestal and gain measurements, a measurement of the pulse shape is performed by precisely delaying the trigger, using the TTCrx chip located on the calibration board. A total of 132 calibration boards have been produced. They fulfill the required performance of an integral linearity better than 0.1%, a uniformity better than 0.2%, and a stability as a function of time better than 0.1%.

In the computation of the calibration constants for the electromagnetic calorimeter, however, a small difference in the response between physics and calibration pulses for the same initial current should also be taken into account. This arises from the different injection points of the current (see Fig. 106) and from the difference between the exponential current as a function of time for the calibration and the triangular pulse shape of the physics pulse (see Fig. 103). This difference depends on the electrical properties of the readout cells and on pseudorapidity and never exceeds 4%. The procedure to compute this correction has been validated during the study of the response uniformity with test-beam data of production modules. For this reason, several calibration methods have been used to determine the optimal-filtering coefficients (see Section 5.5.1.4).

5.5.2.2 Calibration of tile calorimeter Each tile-calorimeter cell can be divided into three sections for calibration purposes; the optical part consisting of scintillator and fibres, the photomultipliers, and the front-end electronics which shape and digitise the light signals. A calibration and monitoring system is used to certify each of these parts independently [142].




Figure 107. Initial ionisation current per deposited energy from electromagnetic showers in the barrel and end-cap electromagnetic calorimeters. These values have been derived from electron test-beam data.

The charge-injection system is designed to calibrate the front-end shaping and digitising circuits to an accuracy of 1%. It uses a precision DAC on each of the 3-in-1 cards, in conjunction with solid-state switches and precision capacitors. Dedicated charge-injection runs are taken scanning the full range of both gains. Fits to the measured pulses convert ADC counts to pC and yield a typical non-linearity of about two ADC counts over the full dynamic range, with a stability of better than 1% over many months of monitoring.

The laser system is used to calibrate and monitor the photomultiplier response with a target accuracy of better than $\pm 1\%$. This system sees both the front-end electronics and the photomultipliers. A Nd:YVO4 laser, synchronised to the TTC clock generates pulses of ~ 10 ns length at a wavelength of 532 nm. The laser light is split in the main services cavern (USA15) and sent via clear plastic fibres to each calorimeter finger where it is again split to each photocathode. The pulse-to-pulse variation is monitored at the laser by several pin-diodes and, by means of a filter wheel, the laser light intensity can span a dynamic range of 10^4 . The global photomultiplier non-linearity is found to be $< 0.5\%$ above 80 pC and up to where saturation sets in. In addition, the system can be used to set the global calorimeter timing and to investigate saturation-recovery techniques.

To calibrate and monitor the tile-calorimeter scintillator and optical system, a ^{137}Cs γ -source of ~ 10 mCi strength is moved hydraulically inside the calorimeter body. This system sees the full chain from the optical part to the front-end electronics and to the photomultiplier. It enables the setting of the high voltage of each channel and inter-calibration of all the 10000 tile channels. These calibrations are performed during dedicated runs, with the source traversing each of the 463,000 tiles in the detector [143]. The photomultiplier current is integrated by the 3-in-1 cards as the source passes within the corresponding cell. From the digitised current, the response of each scintillator tile can be extracted. The accuracy of a single tile response is better than $\pm 2\%$, and the mean response accuracy for a calorimeter cell is $\pm 0.3\%$. Based on monitoring a single module over a period of four months, the long-term stability of the system was determined to be 1%. A procedure to set the gains of each photomultiplier based on the integrated cell-source current allows intercalibration of cells and modules to better than $\pm 3\%$ [142]. The absolute energy calibration (pC/GeV) is set with electrons, as described in Section 5.6.4.1.

5.5.3 Calorimeter power supplies and services

This section describes briefly the LAr and tile calorimeter power supplies, cooling systems and controls.

5.5.3.1 Low-voltage power supplies The low-voltage power supplies for the LAr and tile calorimeters are located in the main cavern, as close as feasible to the on-detector electronics boards. To this end, they had to be custom built and validated to operate in a high radiation environment, in a significant residual magnetic field of up to 100 mT, and in a very limited volume.

Each front-end crate of the LAr calorimeter is powered by a low-voltage power supply located close to it in the gap between the cryostats. This power supply delivers with multiple, redundant DC-DC converters, the seven DC voltages necessary to the front-end crate. There are a total of 58 low-voltage power supplies, 2×16 for the barrel electromagnetic calorimeter and 2×13 for the end-cap calorimeters. Each low-voltage power supply is powered by a 280 V DC power supply located in the underground counting room (USA15). The low-voltage power supplies and associ-

4855 ated 280 V power supplies are controlled and monitored using a custom interface board, ELMB,
described in Section 8.5.

Low-voltage power to the tile-calorimeter drawer electronics is supplied by a custom power
supply located in a mechanical extension on the ends of the module. These extensions are located
on either side of the central barrel and on the outer surface of each extended barrel at the support
4860 girder. Each water-cooled power supply consists of eight isolated switching DC-to-DC convertors
(bricks) which are powered by a customised external 200 V DC bulk power supply. The bricks
themselves are controlled and monitored at two levels externally through CANbus using a custom
interface board based on the ELMB described in Section 8.5 and using current loops. An independ-
ent external reset from USA15 is available in the case of a failure in the voltage control network
4865 (for example due to a single event upset). A second auxiliary board, also based on the ELMB and
located in USA15, provides overall control.

5.5.3.2 High-voltage power supplies High voltage is delivered to the LAr and tile calorimeters
through commercial or custom-built external power supplies located in the underground counting
room.

4870 Several different versions of the same industrial power supply are used to power the different
LAr detectors, according to the voltage and current needed. The maximum voltage and current
of each version are given in Table 31. Furthermore, the barrel electromagnetic, EMEC or HEC
channels showing a problem, such as a resistive short, are powered by high-current (maximum
of 3 mA) supplies. A power supply module contains two boards with 8 (FCal) or 16 (all detectors
4875 but FCal) channels each. All high-voltage channel parameters can be controlled remotely and in
particular the voltage itself, the ramp speed and the trip levels for the current drawn. In total, there
are 157 modules housed in 20 crates containing eight modules each for the whole LAr system.

High voltage is delivered to the tile-calorimeter drawers from external bulk power supplies.
The supply can deliver two voltages to the input of the drawer (-830 V and -950 V) to match the
4880 PMT operating requirements and provides a maximum current of 20 mA. The voltage is distributed
inside the drawer and regulated to provide the desired voltage for each PMT. A single bulk supply
provides high voltage to 16 drawers.

5.5.3.3 Cryogenics and cooling The cryogenics services for the LAr calorimeters are described
in some detail in Section 9.4.5 and shown in Figure 180. The cooling systems for the LAr and
4885 tile calorimeter electronics are both based on a leak-less system for the on-detector implementa-
tion (see Section 9.4.3 for an overview of the off-detector systems). The cooling system supplies
demineralised water at 18°C and works with sub-atmospheric pressure to prevent the cooling liquid
from leaking out of the system should a break or holes occur in a cooling pipe.

The total consumption of the LAr front-end readout electronics exceeds 200 kW (each FEB
4890 channel has a consumption of ~ 0.7 W). Therefore, each electronic board in a front-end crate
is conductively cooled using two aluminium plates placed on each side of the board. These alu-
minium plates are implemented as part of the leak-less cooling system. The temperature of a FEB
is maintained constant at $\pm 1^\circ\text{C}$. Under these conditions, no temperature dependence is observed on
the pedestals, the gains and the delays of the readout electronics. The low-voltage power supplies
4895 are also operated using the same cooling system.

Detector	Nominal detector voltage (V)	Maximum module voltage (V)	Maximum module current (μA)
Electromagnetic barrel	2000	2500	75
Barrel presampler	2000	2500	75
Electromagnetic end-cap (EMEC)	1000 to 1700	2000	200
	1800 to 2500	2500	200
End-cap presampler	-2000	-2500	75
Hadronic end-cap (HEC)	1800	2500	75
Forward calorimeter (FCal)	250, 375, 500	600	6000
Purity monitors	± 2500	± 2500	75

Table 31. Nominal voltages, maximum voltage and maximum current of the different LAr calorimeter high-voltage power supplies.

A heat dissipation of 300 W is expected from a pair of tile-calorimeter drawers and a leak-less cooling system [144] provides cooling inside the drawers. The performance requirements were determined with a prototype system [145] which was used during the calibration of production modules in a test-beam. By varying the input temperature and flow-rate it was possible to change the temperature inside the drawer while measuring the response of the calorimeter to a high-energy pion beam. The fractional energy variation as a function of temperature changes of the PMT block was found to be $0.2\%/^{\circ}\text{C}$ and the variation in response of the PMT gain itself as a function of temperature changes of the cooling water is $0.15\%/^{\circ}\text{C}$. To ensure a PMT gain stability of $\pm 0.5\%$, the temperature variations of the PMT block must be smaller than $\pm 2.5^{\circ}\text{C}$. Long-term tests carried out during calibration of production modules showed that the system could maintain the temperature in a drawer to within $\pm 0.3^{\circ}\text{C}$.

5.5.3.4 Detector control systems Within the context of the overall ATLAS detector safety (DSS) and control (DCS) systems described in Section 8.5, this Section describes the main operating parameters monitored by the specific DCS implementations of the different calorimeter systems.

LAr controls The LAr DCS controls and monitors two important operational parameters: the liquid-argon temperature and purity.

- **Temperature measurements**

The measured sensitivity [141] of the LAr signal to temperature is -2% per degree. Numerous calibrated ($\Delta T = 10$ mK) temperature probes (PT100 platinum resistors) have therefore been installed in each calorimeter to measure precisely the liquid-argon temperature. A total

of 192 temperature probes are glued on the absorbers of the barrel electromagnetic calorimeter at two radii (inner and outer) and regularly distributed in $z \times \phi$ (6×32). Similarly, a total of 32 (respectively 18) temperature probes are located at the front and back sides of the external (respectively internal) EMEC wheels, close to their outer radius. These probes are also uniformly distributed in ϕ .

A total of 192 HEC temperature probes fixed on copper blocks at three radii (inner, middle and outer) are located near the HEC-wheel rear surface and uniformly distributed in ϕ . Finally, the three FCal modules are equipped with a total of 14 temperature probes, four at the outer radius and ten at the inner radius, where heating from beam interactions and heat transfer through the cryostat wall is highest. These sets of temperature probes are completed by other sets, used to control and monitor the temperature during the cooling or heating of the cryostats.

- **Purity monitors**

To control the argon purity, ten purity monitors have been installed on the outer radius of the barrel calorimeter between the support rings. They are located in the median azimuthal plane and at the top and the bottom of the half-barrels. Ten purity monitors are also installed in each end-cap cryostat, two in the EMEC region, six in the HEC region and two in the FCal region. The impurity measurement is based on the energy deposition by radioactive sources in the liquid argon. The ionisation charge is collected by an electric field and measured by a cold pre-amplifier. Each device consists of two radioactive sources: an ^{241}Am source, which emits 5.5 MeV α -particles, and a ^{207}Bi source, which emits 1 MeV conversion electrons. The ratio of the measured charges between the Bi and Am sources is used to extract the absolute oxygen content in the liquid argon [110].

The other quantities controlled and monitored by the LAr DCS are mainly concerned with the readout electronics. On the detectors, DCS controls also the position monitors, the voltage, the current and the temperature of the HEC front-end power supplies. On the high-voltage system, DCS controls the voltage, the intensity and the pre-set value for switching off each high-voltage channel. The temperature of each high-voltage module is also monitored. On the front-end electronics, DCS controls the power supplies and monitors the voltage, the current and the status of the 280 V power supplies, which powers the front-end crates, and of the low-voltage power supplies in the front-end crates (the temperature of each FEB is monitored by the DAQ system). Finally, on the back-end electronics, DCS monitors the voltage, the current and the temperature of each crate.

Tile controls The tile-calorimeter DCS controls and monitors mainly the on-detector low-voltage and high-voltage power supplies and electronics cooling, as well as the off-detector high-voltage and bulk 200 V DC power supplies, auxiliary boards and electronics racks. The communication with the high-voltage bulk power-supply crates is done using RS422. CANbus is used for the communication with the micro-controller cards located on the high-voltage side of the drawers. A client-server system has been developed and commissioned for the control and monitoring of the photomultiplier high voltages. Control and monitoring of the on-detector low-voltage power supplies is implemented via a custom auxiliary board, which multiplexes currents and temperatures

from sensors inside the drawers, and communicates via CANbus. The 200 V DC power supplies are controlled and monitored using Modbus.

A total of more than 30,000 parameters are monitored in this system. Configuration data such as system structure (lists and hierarchies of devices), device properties (configuration of archiving, smoothing, etc.) and settings (output values, alert limits) are stored in the configuration database. The data produced by PVSS is stored in the DCS ORACLE-based archive. Due to the large amount of data monitored by DCS, smoothing is applied to reduce the amount of data stored. A subset of these data (for example photomultiplier high-voltage settings) is also available to offline reconstruction.

5.6 Test-beam measurements and results

5.6.1 Electromagnetic module performance

Four (of 32) barrel and three (of 16) end-cap production modules have been exposed to test-beams with a set-up which includes a cryostat on a moveable table, enabling detailed position scans. These tests have been carried out in the CERN SPS H8 and H6 beam lines, using electron and positron beams of energies between 1 and 250 GeV. A comprehensive analysis of module performance has been reported [146, 147, 141, 148].

The electron energy reconstruction procedure is similar to that described in Section 10.4: for example, for the barrel electromagnetic calorimeter, an electron cluster is first constructed using the measurements in the middle accordion compartment, where cells within a square window of 3×3 cells around the cell with the highest energy are clustered together. For the other accordion compartments, all cells intersecting the geometrical projection of this square window are included. The energy in each compartment, E_0 for the presampler and E_1, E_2, E_3 for the strip, middle and back compartments, is the sum of the energies of the selected cells in the cluster. The energy of the electron is then given by the algorithm:

$$E = \text{offset} + w_0 \times E_0 + w_{01} \times \sqrt{E_0 E_1} + \lambda (E_1 + E_2 + E_3) + w_3 \times E_3,$$

where the offset and w_0 correct for the energy lost upstream of and in the presampler, w_{01} corrects for the energy deposited in the dead material between the presampler and the strip compartment, λ corrects for the energy dependence of the sampling fraction and for lateral leakage outside the electron cluster, and w_3 corrects for the energy deposited downstream of the calorimeter. These coefficients depend somewhat on the electron energy and the η value. An iterative procedure is therefore needed to reconstruct the electron energy. This procedure has been shown to converge after two iterations. The correction coefficients have been determined by simulating the test-beam set-up using GEANT4. In the case of the EMEC, where the presampler is present only in the region $1.5 < |\eta| < 1.8$, where the material in front of the calorimeter was independent of η , and where the beam energy did not exceed 180 GeV, a similar but simplified scheme is used [141].

For the barrel electromagnetic calorimeter, the linearity of the response and the energy resolution have been studied as a function of energy in the range from 10 to 245 GeV at $\eta = 0.687$. At 245 GeV, the beam energy was not known well enough and the linearity of response could not be assessed accurately. The experimental measurements, after noise subtraction, have been fitted




Figure 108. Linearity of response as a function of the electron beam energy, E_{beam} , for a barrel LAr electromagnetic module at $|\eta| = 0.687$. All points are normalised to the value measured at $E_{beam} = 100$ GeV. The band represents the total uncertainty on the beam energy measurement.




Figure 109. Fractional energy resolution as a function of the electron beam energy, E_{beam} , for a barrel LAr electromagnetic module at $|\eta| = 0.687$. Electronic noise was subtracted from the data before plotting the results. The curve represents the results of a fit to the data using Eq. 5.2.

with the expression [147]:

$$\frac{\sigma(E)}{E} = \frac{a}{\sqrt{E(\text{GeV})}} \oplus b, \quad (5.2)$$

where a is the stochastic term and b the constant term reflecting local non-uniformities in the response of the calorimeter. In the energy range 15 – 180 GeV, the reconstructed energy response is linear within $\pm 0.1\%$, as shown in Fig. 108. For the lowest-energy point at 10 GeV, a non-linearity of 0.8% is measured. A stochastic term of $10\% \cdot \sqrt{\text{GeV}}$ and a constant term of 0.17% have been obtained from the fit to the fractional energy resolution, as illustrated in Fig. 109. These results are in agreement with dedicated Monte-Carlo simulations of the test-beam set-up. Similar results for the energy resolution have been obtained for the EMEC [113].

At low energy, the calorimeter response to minimum-ionising particles has provided a detailed exploration of the structure of the active material of the detector. In particular, the muon energy deposition is much more localised than the electron energy depositions. Typically, muons only cross one middle cell in η and one or two in ϕ . They can therefore be used to study the fine structure of the calorimeter, without having to deconvolute the effects of the showering process. The signal-to-noise ratio, evaluated as the ratio of the most probable energy deposit divided by the RMS spread of the pedestal, varies between seven and twelve for middle barrel cells and between six and seven for middle end-cap cells [148]. These results have shown that cosmic muons can be used to commission the favourably oriented part of the electromagnetic calorimeters in terms of timing and response uniformity, and may serve to verify the integrity of the detector and of the whole readout chain before the start-up of the LHC [149].

The response uniformity at high energy as a function of η has been measured using an electron beam of 245 GeV for the barrel and of 119 GeV for the end-cap [141]. The goal for ATLAS is to achieve a constant term of 0.7% or smaller over the full calorimeter acceptance. Non-uniformities of the response on the tested modules do not exceed 0.7% and do not exceed even 0.5% in the case




Figure 110. Distribution of the average energies measured in all cells of all tested modules as a function of the cell η , normalised to the mean energy measured in the modules. In the barrel, this mean energy was ~ 245 GeV, while it was ~ 120 GeV in the EMEC. For each bin in η , the distribution is normalised to the number of middle cells in that bin (design value). This normalisation is only used to define the colour of each bin in the plot.




Figure 111. Linearity of response as a function of the beam energy, E_{beam} , at $|\eta| = 0.687$, for a barrel LAr electromagnetic module in the combined test-beam set-up exposed to electron beams with different amounts of material placed upstream of the active calorimeter.




Figure 112. Fractional energy resolution as a function of the electron beam energy, E_{beam} , for a barrel LAr electromagnetic module in the combined test-beam. Electronic noise has been subtracted from the data. The results are shown for an amount of upstream material of $2.4 X_0$, which is that expected in ATLAS at $\eta = 0.4$. The curves represent the results of fits to the data and the simulation using Eq. 5.2.

of the barrel modules, as shown in Fig. 110. The overall constant term in the energy resolution, using the above formula, ranges between 0.5% and 0.7% and therefore meets well the calorimeter design performance goals.

The performance of the barrel electromagnetic calorimeter in terms of its finely segmented

first sampling has been studied by using electron, photon and pion beams [146]. The position resolution along η was measured to be about 1.5×10^{-4} and 3.3×10^{-4} (in units of pseudorapidity) at 245 GeV for the front and middle compartments, allowing to achieve a polar angle resolution in the range $50\text{--}60(\text{mrad})/\sqrt{E}$ (GeV) over the whole coverage (barrel and end-caps). The π^0 rejection was measured from real data, using a photon beam and mixing together different events to mimic photon pairs from π^0 decays, and found to be 3.54 ± 0.12 (statistical) for π^0 with $p_T = 50$ GeV and for a single photon efficiency of 90%.

A spare electromagnetic barrel module, identical to the series modules, was built for the combined test-beam period described in Section 10.1.2 and was exposed to electron, photon, pion and muon beams with energies between 1 and 350 GeV. The amount of material in front of this module was very close to the material expected in ATLAS and therefore great importance was given to the task of verifying that the linearity, energy resolution and uniformity [150] are well understood in terms of the detector description and the simulations using GEANT 4. The performance of the electromagnetic calibration scheme, using longitudinal weights similar to the ones described above, was tested by adding in a controlled way different amounts of material just in front of the electromagnetic calorimeter. This extra material corresponded to $2.4\text{--}3.3 X_0$ and emulated in this way the material in front of the LAr barrel calorimeter in ATLAS in the region $|\eta| < 1.2$. A linearity of 0.5%, independent of the upstream material, is obtained over the energy range $20\text{--}250$ GeV, as shown in Fig. 111. Figure 112 shows the noise-subtracted ($\sigma_{\text{noise}} = 170$ MeV) resolution at $\eta = 0.4$ for an amount of upstream material of $2.4 X_0$. The data points have been corrected for the beam momentum spread. A stochastic term of $(10.1 \pm 0.4)\% \times \sqrt{\text{GeV}}$ and a constant term of $(0.2 \pm 0.1)\%$ are obtained and the data are well described by the simulation, which incorporates the most detailed and up-to-date description of the detector and GEANT 4 (version 4.8.2) model of electromagnetic showers. More detailed results from this last test-beam period for the barrel LAr electromagnetic calorimetry can be found in [151].

5.6.2 Hadronic end-cap performance

About 25% of the series production modules were exposed to beams of muons, electrons and pions with energies up to 200 GeV at the CERN SPS [152]. Two partial HEC wheels, consisting of three HEC1 and three HEC2 modules were used. The goal was not only to prove the uniformity of the production modules as defined by the hardware tolerances, but also to determine their performance and calibration, as obtained from muons, electrons and pions. The data taken with electrons yield a fractional energy resolution after noise subtraction and a fit to the functional form in Eq. 5.2 with a stochastic term, $a = (21.4 \pm 0.1)\% \cdot \sqrt{\text{GeV}}$ and a constant term, b , compatible with zero, in good agreement with predictions from simulation [152]. Horizontal and vertical scans with beams across the surface of the calorimeter showed a homogeneity of the electron signal of $\pm 1\%$ without corrections.

The pion measurements are of particular importance for the prediction of the final performance of the calorimeter for jets. A stochastic term of $(70.6 \pm 1.5)\% \cdot \sqrt{\text{GeV}}$ and a constant term of $(5.8 \pm 0.2)\%$ were obtained from a fit to the functional form in Eq. 5.2. The data have been compared in detail to simulations (GFLUKA, GCALOR and GHEISHA) with satisfactory results [153].

From 2002 to 2004, a new set of combined test-beam measurements was launched. The main purposes were to define calibration procedures and constants for initial operation in ATLAS and




Figure 113. Fractional energy resolution as a function of reconstructed energy for π^- and π^+ data taken during the 2002 EMEC/HEC combined test-beam period compared to different predictions from simulation using GEANT 4. The analysis employs the signal-weighting technique described in the text. The data are plotted after noise subtraction and the curve represents as an example a fit to the π^+ data using Eq. 5.2.




Figure 114. Energy response on the electromagnetic scale for 200 GeV pions when performing a vertical scan across the transition region between the EMEC/HEC and FCal calorimeters. Shown is the total energy response together with the individual responses in the different electromagnetic and hadronic components of the calorimetry. The data (full symbols) are compared to GEANT 4 predictions (open symbols).

5050 to operate the EMEC, the HEC and the FCal together in conditions as close as feasible to those expected in ATLAS. The first of these combined test-beam periods took place in 2002 and was devoted to the region $1.6 < |\eta| < 1.8$ [154, 155]. The second period took place in 2004 and was dedicated to a scan of the transition region around $|\eta| = 3.2$.

5055 A three-dimensional clustering algorithm and a signal-weighting approach (see Section 10.5.2), used already in previous experiments, have been tested and the first results yield good energy resolution for pions. The signal-weighting technique exploits the fact that local energy depositions of high density are mainly due to electromagnetic interactions, whereas for hadronic interactions, the corresponding density is substantially lower. Thus, for a segmented calorimeter, the energy deposited in individual readout cells can, on a statistical basis, be identified to be of electromagnetic
5060 or hadronic origin. For ATLAS, these weights are derived from simulations of single particles and jets. In test-beam, the volume of the related clusters in EMEC and HEC has been used to obtain the cluster energy density. The weighting function has been derived from the data directly, but after correcting for leakage (for details see [154]).

5065 Figure 113 shows the energy dependence of the fractional energy resolution separately for the π^- and π^+ data. For energies up to 80 GeV, a differential Cerenkov counter has been used to separate π^+ and protons. The proton contamination in the beam increases with energy and its contribution to the π^+ data in Fig. 113 is the dominant source of systematics when comparing these data to the π^- data and to simulation. Fits to the data using Eq. 5.2 yield stochastic terms of $(84.6 \pm 0.3)\% \sqrt{\text{GeV}}$ and of $(81.7 \pm 0.4)\% \sqrt{\text{GeV}}$ for π^- and π^+ respectively, and constant
5070 terms of zero within errors. The energy range available is not extensive enough to avoid any correlation between the stochastic and constant terms. Nevertheless, the results give some indication of the effectiveness of the signal-weighting technique in achieving a good level of compensation.

The GEANT 4 simulations (version 5.0) with two different hadronic physics lists (LHEP 3.3 and QGSP 2.3) shown in Fig. 113 are in reasonable agreement with the data, but neither model yields an optimal description of the data [153].

The scan of the transition region around $|\eta| = 3.2$ was meant to assess the performance in the complex region of overlap of the three end-cap calorimeters. The EMEC, HEC and FCal modules were positioned as in ATLAS, including mock-ups of the details of cryostat walls and supports (dead material). One quarter of the full HEC1 and HEC2 wheels were assembled with dedicated small modules having reduced η -coverage and encompassing only the forward region. Similarly, one EMEC inner-wheel module (one-eighth of the full EMEC wheel) and one quarter of the full FCal1 and FCal2 detectors were assembled. The results of this scan for the performance of the EMEC, HEC and FCal on the electromagnetic scale are described in [156]. These results demonstrate that the response and resolution of electrons and pions across most of the EMEC/HEC and FCal acceptance, including the crack region around $\eta = 3.2$, show reasonable agreement with expectations from detailed simulation.

As an example, Fig. 114 shows the response on the electromagnetic scale for 200 GeV pions, when moving from the FCal region (left) across the crack to the EMEC/HEC region (right). Shown is the total energy response together with the individual responses in the different electromagnetic (FCal1 and EMEC) and hadronic (FCal2 and HEC) components of the calorimeter. For the energy reconstruction, a cone size of $\Delta R \times \Delta\phi = 0.30$ has been used. The data are compared with simulation (GEANT 4, version 7.1 with physics list QGSP-GN 2.6). The prediction for the total electromagnetic energy measured in the FCal region shows some deviations from the data. As already known from previous studies with the QGSP physics list, the simulation tends to predict too compact hadronic showers in comparison to the data. But the general loss of signal (on the electromagnetic scale), when traversing the crack region is not more than $\sim 20\%$. The details of the complex material structure in the crack region are rather well described in the simulation.

5.6.3 FCal performance

There have been two FCal stand-alone test-beam periods at the CERN SPS, one in 1998 and one in 2003. The 1998 period used engineering modules which were full-depth segments of FCal1 and FCal2. The electron response of FCal1 and an estimate of the electromagnetic scale for FCal2 were extracted from these measurements [157]. The 2003 test-beam was carried out with the final production modules and used electronics which were very similar to the final electronics implemented for ATLAS. Data were taken with electron and pion beams with momenta from 10 to 200 GeV. The data were analysed using the standard LAr technique of optimal filtering (see Section 5.5.1.4). Electronic noise was calculated from the data and was subtracted in quadrature before energy resolution plots, such as those shown in Figs. 115 and 116, were produced.

Figures 115 and 116 illustrate the results obtained after subtraction of the electronic noise from the data. The points have been fitted with the function in Eq. 5.2. The fit to the electron data in Fig. 115 yielded stochastic and constant terms of $(28.5 \pm 1.0)\% \sqrt{\text{GeV}}$ and $(3.5 \pm 0.1)\%$, respectively. For pions, the energy resolution depends on the technique used to correct the reconstructed energies from the electromagnetic scale to the hadronic scale. Two techniques were employed: the first one uses a single weight per module (flat weights) and yields stochastic and constant terms of $(94.2 \pm 1.6)\% \sqrt{\text{GeV}}$ and $(7.5 \pm 0.4)\%$, respectively. A more sophisticated technique, using ra-




Figure 115. Fractional energy resolution obtained for electrons measured in the first module of the forward calorimeter as a function of the beam energy, E_{beam} . The curve corresponds to the result of a fit to the data points using Eq. 5.2.




Figure 116. Fractional energy resolution obtained for pions, measured in all three modules of the forward calorimeter, as a function of the beam energy, E_{beam} . The data are shown for two cell-weighting schemes and the curves correspond to the result of a fit to the data points using Eq. 5.2.

5115 dial weights, which exploits the fine transverse segmentation of the FCal, improves the sampling
term from 94% to 70% and yields a constant term of 3.0%. The measured energy resolution for
pions with an energy of 200 GeV is improved from 10.1% to 5.8% as shown in Fig. 116.

The resolutions obtained in these FCal beam tests meet well the design specifications, which
were expressed as stochastic and constant terms of 100% and 10% respectively for hadrons. Further
5120 details of the results of the 2003 test-beam measurements can be found in [158].

5.6.4 Tile-calorimeter performance

5.6.4.1 Stand-alone performance Approximately 12% of all production modules of the tile
calorimeter have been measured extensively in dedicated test-beam periods at the CERN SPS.

Muons with an energy of 180 GeV and incident at 90° to the module symmetry plane were
5125 used to study the photo-electron yield for all eleven tile sizes in the detector (monitoring the photo-
electron yield checks for any deterioration in time of the optical response of the calorimeter). The
photo-electron yield was calibrated using two independent techniques. The first one uses the spread
in the difference of signals from the two photomultiplier tubes for muons passing through the centre
of the cell. The second one uses the spread in the laser-generated light to calibrate the photomul-
5130 tiplier response in photo-electrons per pC and the measured signal for muons and electrons. The
results of the two methods are in excellent agreement: the light yield is independent of the size of
the tiles within 1–2% uncertainties. These results also reproduced the batch-to-batch light output
variation seen during tile production, together with the systematic difference associated with the
two sources of raw polystyrene (the photo-electron yield is about 80 per GeV for one source and
5135 about 100 per GeV for the other source). In the extended-barrel modules, the yield is 10% to 20%
lower, reflecting the choice which was made to use batches with higher light yields for the barrel
modules. The muon data at 90° incidence also provided the means to study the local and global
features of the calorimeter response. The spread of the mean signals from the eleven tile sizes

is 2.0% (respectively 2.8%) for the barrel (respectively extended-barrel) modules and the spread within each tile size is about 3%. The average response of a module varies between modules with a spread of 1.2%, while the variation within a given module is typically 4–5%.

The muon signal-to-noise ratios are very large, typically 44 when summing over a tower (three radial compartments and over $\Delta\eta \times \Delta\phi \approx 0.1 \times 0.1$) and 18 in the last radial compartment. This feature is expected to provide reliable identification of isolated muons down to energies of ~ 2 GeV even in the presence of pile-up noise, which is negligible in the last compartment. The response to 180 GeV muons, averaging over angles in the range $-1.49 < |\eta| < 1.35$ has a spread of 1.9% (rms), excluding angles for which the coverage is not complete.

The modules were exposed to electron beams with energies between 10 and 180 GeV, to set the energy conversion scale (i.e. the photomultiplier high voltage to be supplied, as discussed in Section 5.5.1.3) for a significant fraction of the cells of the entire module. The cesium system will be used to transfer this basic calibration to the majority of modules which were not calibrated using high-energy beams. Electron showers are mostly contained in the first radial compartment, therefore precise response values can only be extracted from this first layer of cells. The spread of these responses over the entire exposed module sample is up to 3%. The electron response is linear with deviations of about 1%. The electron response is observed to vary with the angle of incidence, as expected because of the variation in the effective calorimeter granularity with angle. Between extremes, corresponding to angles to the tile/steel plate planes from 0° to 90° , the response increases by approximately 8%.

Systematic measurements were performed using hadron beams with energies in the 20–180 GeV range, with a few additional runs at 350 GeV. The set-up consisted of a vertical stack of three modules, in which the production module under test occupied the central position. For hadrons incident on this module, transverse leakage of hadron showers is approximately 1%, while longitudinal leakage significantly affects the resolution particularly at higher energies and lower impact angles. The beams typically contained a mixture of pions, kaons and protons. The beam-line included a Cerenkov counter which was used to separate pions and protons in the case of positive beam energies (between 50 GeV and 180 GeV).

The fractional energy resolution, σ_E/E , for isolated pions was studied as a function of beam energy and impact angle. The results are summed at the electromagnetic scale over all cells and the resulting energy resolutions for $\eta = 0.35$ are shown in Fig. 117. The parametrisation of Eq. 5.2 was used to fit the results and the best fit yields $a = (56.4 \pm 0.4)\%$ and $b = (5.5 \pm 0.1)\%$. The statistical errors quoted here do not display the correlation between the two terms. These results are in good agreement with earlier stand-alone measurements made with prototype modules, when accounting for the fact that the latter were radially longer by 1.5 interaction lengths in order to represent the




Figure 117. Fractional energy resolution obtained for pions as a function of the inverse square root of the beam energy at an angle of incidence equivalent to $|\eta| = 0.35$.




Figure 118. Linearity of response as a function of the pion beam energy, E_{beam} , for combined LAr and tile calorimetry at $|\eta| = 0.25$.




Figure 119. Fractional energy resolution obtained for pions as a function of the inverse square root of the beam energy, E_{beam} , for combined LAr and tile calorimetry at $|\eta| = 0.25$. The curve corresponds to the result of a fit to the data points with the functional form as shown.

total depth of electromagnetic and hadronic calorimetry.

The fractional energy resolution for production modules displays a significant dependence on η , mostly as a result of the increase in effective depth and decrease of longitudinal leakage as η increases. As representative examples, σ_E/E at $\eta = 0.25$ is $(14.2 \pm 0.1)\%$ at 20 GeV and $(6.6 \pm 0.1)\%$ at 350 GeV, whereas, at $\eta = 0.55$, σ_E/E is $(13.0 \pm 0.1)\%$ at 20 GeV and $(5.9 \pm 0.1)\%$ at 350 GeV. The module-to-module uniformity has been studied with 180 GeV pions entering the calorimeter under various impact points and incidence angles. The uniformity in the mean response was found to be independent of η over all modules measured in the test-beam and shows an average spread of 1.4% [159].

5.6.4.2 Combined LAr and tile calorimeter test-beam measurements The combined performance of the barrel LAr electromagnetic and tile calorimeters was measured in 1996 in the H8 beam at the CERN SPS. The set-up used prototype modules of the two calorimeters. The LAr stack consisted of two modules, each spanning nine degrees in the azimuthal direction. The modules were longitudinally segmented into three compartments, of $9 X_0$, $9 X_0$ and $7 X_0$ each at $\eta = 0$, for a total of 25 radiation lengths (1.22 interaction lengths). The segmentation was 0.018×0.02 in $\Delta\eta \times \Delta\phi$ for the first two longitudinal compartments and 0.036×0.02 for the third compartment.

Five prototype modules of the tile calorimeter, each covering $\Delta\phi = 0.1$, were stacked vertically downstream of the LAr cryostat and as close as possible to it; the distance between the active parts of the two detectors was nevertheless about 55 cm, roughly twice as much as in ATLAS. The tile calorimeter modules had the same steel and scintillator plate geometry as the production modules, but, in the longitudinal direction, the active portion of the calorimeter measured 180 cm, rather than 152 cm as in the production modules. These modules were segmented into four longitudinal compartments, of about 1.5, 2.0, 2.5 and 3 interaction lengths; in the η -direction, each module was segmented into five equal cells of size $\Delta\eta = 0.2$.

The hadron beam consisted of pions with an energy-dependent proton component and impinged on the combined calorimeter assembly at an angle of 12° . Data were taken in the energy

range between 10 and 300 GeV. The pion energy was reconstructed by a cell-weighting technique, following which the response of cells with relatively small signals is corrected upwards, to equalise it to that of cells with large, typically electromagnetic, signals. This method had been successfully tested in a combined LAr and tile calorimeter test-beam run, as described in [160]. The total energy is expressed as:

$E = \sum_{\text{em.cells}} W_{\text{em}}(E_{\text{cell}}, E_{\text{beam}}) \cdot E_{\text{cell}} + \sum_{\text{had.cells}} W_{\text{had}}(E_{\text{cell}}, E_{\text{beam}}) \cdot E_{\text{cell}} + E_{\text{cryo}}$, where the last term accounts for the energy lost in the dead region between the electromagnetic and the hadronic calorimeters and is taken to be proportional to the geometric mean of the energies deposited in the last electromagnetic compartment and the first hadronic compartment: $E_{\text{cryo}} = W_{\text{cryo}} \cdot \sqrt{E_{\text{em}_3} \cdot E_{\text{had}_1}}$

As a first step, the optimal weights for each readout cell and each beam energy are determined by minimising the energy resolution in a dozen cell energy bins, under the constraint that the mean reconstructed energy reproduces the nominal beam energy. A substantial reduction in the number of energy reconstruction parameters is then obtained by expressing the weights for any bin as simple functions of the cell energies: $w_{\text{em}} = a_E + b_E/E_{\text{cell}}$, $w_{\text{had}} = a_H + b_H/E_{\text{cell}}$, where the a and b coefficients are obtained specifically for each beam energy. As a second step, the dependence of the a and b coefficients on the beam energy is parametrised. At this point, all weights are expressed as a function of only eight parameters which include the weight corresponding to the cryostat term. A further step implements an energy reconstruction method which does not rely on a priori knowledge of the incident particle energy: a first-pass value of the total energy is calculated and used as the nominal value in minimising the resolution for the first-pass determination of the weights, according to the functional forms just described. The procedure is iterated until the reconstructed energy is stable.

In the procedure described here, the only information gained using a priori knowledge of the beam energy is the functional form of the weights, but their values are boot-strapped from the raw data. Using this algorithm, the linearity and resolution results in Figs. 118 and 119 are obtained [161]. The hadron response deviation from linearity is within $\pm 2\%$ over the entire beam-energy range. This is a consequence of the constraint applied in the first pass of the weight determination. The fractional energy resolution is well fitted by the sum in quadrature of a stochastic term of $(52.0 \pm 1.0)\%$, a constant term of $(3.0 \pm 0.1)\%$ and an electronic-noise term of $1.6 \text{ GeV} \pm 0.1\%$.

6. Muon spectrometer

6.1 Overview

The muon spectrometer forms the outer part of the ATLAS detector and is designed to detect charged particles exiting the barrel and end-cap calorimeters and to measure their momentum in the pseudorapidity range $|\eta| < 2.7$. It is also designed to trigger on these particles in the region $|\eta| < 2.4$. The driving performance goal is a stand-alone transverse momentum resolution of approximately 10% for 1 TeV tracks, which translates into a sagitta along the z (beam) axis of about 500 μm , to be measured with a resolution of $\leq 50 \mu\text{m}$. Muon momenta down to a few GeV (~ 3 GeV, due to energy loss in the calorimeters) may be measured by the spectrometer alone. Even at the high end of the accessible range (~ 3 TeV), the stand-alone measurements still provide adequate momentum resolution and excellent charge identification.

Precision-tracking chambers in the barrel region are located between and on the eight coils of the superconducting barrel toroid magnet, while the end-cap chambers are in front and behind the two end-cap toroid magnets. The ϕ symmetry of the toroids is reflected in the symmetric structure of the muon chamber system, consisting of eight octants. Each octant is subdivided in two sectors with slightly different lateral extensions, leading to a small region of overlap in ϕ . This overlap of the chamber boundaries minimises gaps in detector coverage and also allows for the relative alignment of adjacent sectors using tracks recorded by both a large and a small chamber (see Section 10.3.2).

The chambers in the barrel are arranged in three concentric cylindrical shells around the beam axis at radii of approximately 5 m, 7.5 m, and 10 m. In the two end-cap regions, muon chambers form large wheels, perpendicular to the z -axis and located at distances of $|z| \approx 7.4$ m, 10.8 m, 14 m, and 21.5 m from the interaction point. Figure 120 [162] shows the overall layout of the muon system while Figs. 121 and 122 give cross-sections in the planes transverse to, and containing, the beam axis. In the centre of the detector ($|\eta| \approx 0$), a gap in chamber coverage has been left open to allow for services to the solenoid magnet, the calorimeters and the inner detector. The size of the gap varies from sector to sector depending on the service necessities, the biggest gaps of 1-2 m being located in the large sectors. The angular range, seen from the interaction point, where a high momentum (straight) track is not recorded in all three muon layers due to the gaps is about $\pm 4.8^\circ$ ($|\eta| \leq 0.08$) in the large and $\pm 2.3^\circ$ ($|\eta| \leq 0.04$) in the small sectors. Additional gaps in the acceptance occur in sectors 12 and 14 due to the detector support structure (feet). The consequences of the acceptance gaps on tracking efficiency and momentum resolution are shown in Figs. 227 and 224, respectively. A detailed discussion is given in Section 10.3.4.

Table 32 gives an overview of the main parameters of the four sub-systems of the muon detector. An overview on the performance of the muon system is given in [163]. The precision momentum measurement is performed by the Monitored Drift Tube chambers (MDT's), which combine high measurement accuracy, predictability of mechanical deformations and simplicity of construction (see Section 6.3). They cover the pseudorapidity range $|\eta| < 2.7$ (except in the innermost end-cap layer where their coverage is limited to $|\eta| < 2.0$). These chambers consist of three to eight layers of drift tubes, operated at an absolute pressure of 3 bar, which achieve an average resolution of 80 μm per tube, or about 35 μm per chamber. An illustration of a 4 GeV and a




Figure 120. Initial configuration of the muon spectrometer with its four chamber sub-systems: the precision-measurement tracking chambers (MDT's and CSC's) and the trigger chambers (RPC's and TGC's). In the end-cap, the first TGC layer (I) is located in front of the innermost tracking layer; the next three layers stand in front (M1) and behind (M2 and M3) the second MDT wheel. The first letter (B and E) of the MDT naming scheme refers to barrel and end-cap chambers, respectively. The second and third letters refer to layer (inner, middle, and outer) and sector (large and small) types, respectively (see Section 6.3.1).

20 GeV muon track traversing the barrel region of the muon spectrometer is shown in Fig. 123.

In the forward region ($2 < |\eta| < 2.7$), Cathode-Strip Chambers (CSC) are used in the innermost tracking layer due to their higher rate capability and time resolution (see Section 6.4). The CSC's are multiwire proportional chambers with cathode planes segmented into strips in orthogonal directions. This allows both coordinates to be measured from the induced-charge distribution. The resolution of a chamber is $40 \mu\text{m}$ in the bending plane and about 5 mm in the transverse plane. The difference in resolution between the bending and non-bending planes is due to the different readout pitch, and to the fact that the azimuthal readout runs parallel to the anode wires. An illustration of a track passing through the forward region with $|\eta| > 2$ is shown in Fig. 124.

To achieve the sagitta resolution quoted above, the locations of MDT wires and CSC strips along a muon trajectory must be known to better than $30 \mu\text{m}$. To this effect, a high-precision optical alignment system, described in Section 6.3.4, monitors the positions and internal deformations of the MDT chambers; it is complemented by track-based alignment algorithms briefly discussed in Section 10.3.2.

An essential design criterion of the muon system was the capability to trigger on muon tracks. The precision-tracking chambers have therefore been complemented by a system of fast trigger chambers capable of delivering track information within a few tens of nanoseconds after the passage of the particle. In the barrel region ($|\eta| < 1.05$), Resistive Plate Chambers (RPC) were selected




Figure 121. Cross-section of the barrel muon system perpendicular to the beam axis (non-bending plane), showing three concentric cylindrical layers of eight large and eight small chambers. The outer diameter is about 20 m.




Figure 122. Cross-section of the muon system in a plane containing the beam axis (bending plane). Infinite-momentum muons would propagate along straight trajectories which are illustrated by the dashed lines and typically traverse three muon stations.




Figure 123. Trajectories of muons with momenta of 4 GeV and 20 GeV in the bending plane of the barrel muon spectrometer. In general, the tracks cross 2×4 inner, 2×3 middle, and 2×3 outer layers of MDT tubes. The cyan and dark blue areas in each MDT layer illustrate the granularity of the mezzanine cards (Section 6.3.3).

for this purpose (Section 6.6), while in the end-cap ($1.05 < |\eta| < 2.4$) Thin Gap Chambers (TGC) were chosen (Section 6.7). Table 32 gives the intrinsic time resolution of the detectors, to which contributions from signal propagation and electronics have to be added. The design goal was to




Figure 124. Structure of the barrel and end-cap regions with a track at large η , passing through a CSC in the inner wheel and through MDT's in the middle and outer wheels. For $|\eta| > 2.0$, the 2×4 hits in the inner MDT, as explained in Fig. 123, are replaced by four CSC hits.

Type	Function	Chamber resolution (RMS) in			Measurements/track		Number of	
		z/R	ϕ	time	barrel	end-cap	chambers	channels
MDT	tracking	$35 \mu\text{m} (z)$	–	–	20	20	1088 (1150)	339k (354k)
CSC	tracking	$40 \mu\text{m} (R)$	5 mm	7 ns	–	4	32	30.7k
RPC	trigger	10 mm (z)	10 mm	1.5 ns	6	–	544 (606)	359k (373k)
TGC	trigger	2–6 mm (R)	3–7 mm	4 ns	–	9	3588	318k

Table 32. Parameters of the four sub-systems of the muon detector. The quoted spatial resolution (columns 3, 4) does not include chamber-alignment uncertainties. Column 5 lists the intrinsic time resolution of each chamber type, to which contributions from signal-propagation and electronics contributions need to be added. Numbers in brackets refer to the complete detector configuration as planned for 2009.

5300 keep these contributions low enough for reliable beam-crossing identification with $\geq 99\%$ prob-
 ability. Both chamber types deliver signals with a spread of 15–25 ns, thus providing the ability
 to tag the beam-crossing. The trigger chambers measure both coordinates of the track, one in the
 bending (η) plane and one in the non-bending (ϕ) plane.

5305 The purpose of the precision-tracking chambers is to determine the coordinate of the track in
 the bending plane. After matching of the MDT and trigger chamber hits in the bending plane, the
 trigger chamber's coordinate in the non-bending plane is adopted as the second coordinate of the
 MDT measurement. This method assumes that in any MDT/trigger chamber pair a maximum of
 one track per event be present, since with two or more tracks the η and ϕ hits cannot be combined in
 an unambiguous way. Simulations have shown that the probability of a track in the muon spectrom-




Figure 125. Overall three-dimensional view of the passive material in the muon system, which consists of such items as the barrel and end-cap toroid coils and vacuum vessels, as well as the support structures of the calorimeters, muon chambers, and toroid magnets.

eter with $p_T > 6$ GeV is about 6×10^{-3} per beam-crossing, corresponding to about 1.5×10^{-5} per
 5310 chamber. Assuming uncorrelated tracks, this leads to a negligible probability to find more than one
 track in any MDT/trigger chamber pair. When *correlated* close-by muon tracks do occur, caused
 for example by two-body-decays of low-mass particles, the ambiguity in η and ϕ -assignment will
 be resolved by matching the muon track candidates with tracks from the inner detector.

The muon system was designed to tolerate the radiation levels in the experimental hall, which
 5315 may induce ageing effects in the detectors and damage the electronics. All components were there-
 fore tested to withstand at least five times the radiation levels predicted by the simulation studies
 (see Section 3.3.2).

Due to budget constraints, parts of the muon spectrometer have been staged. This includes the
 readout electronics of certain RPC trigger chambers in the feet region, which will only be installed
 5320 in 2009. It also affects the MDT chambers in the transition region between barrel and end-cap (EE
 chambers), the construction of which has been completed and their installation expected in 2009.
 In Table 32, the numbers in parentheses correspond to the completed detector.

6.2 Precision-tracking chambers

The overall layout of the MDT's is projective: the layer dimensions and the chamber sizes increase
 5325 in proportion of their distance from the interaction point. The MDT chambers cover the region
 $|\eta| < 2.7$, except for the innermost layer where they are replaced by CSC's for $2 < |\eta| < 2.7$.
 In the initial configuration of the spectrometer at the start-up of the LHC, there are 1088 MDT

chambers with 18 main types in the detector for a total area of 5500 m². Besides the main type of chambers, a considerable number of chambers with special shapes have been built to minimise acceptance losses in the regions around the magnet coils and support structures, as illustrated by the many different MDT chamber names in Fig. 120.

The particle fluxes and muon-track density are highest in the forward direction ($2 < |\eta| < 2.7$). These issues, which directly impact the pattern recognition and the muon-reconstruction efficiency as well as the radiation ageing of the detectors, are particularly severe in the first forward muon-chamber layer, just downstream of the end-cap calorimeter. This layer is therefore equipped with CSC's, which measure both track coordinates simultaneously; multitrack ambiguities are resolved by correlating the magnitude of the charge collected on the two planes of orthogonal strips. Each track is measured in four consecutive CSC planes. The number of CSC measurements along a track is less than in the MDT's; however, the faster time response helps restore the tracking efficiency. The CSC's cover an area of about 65 m² in total.

The air-core magnet concept for the muon spectrometer minimises the amount of material traversed by the muons after exiting the calorimeters. However, the muons also encounter the muon chambers themselves and their supports, as well as other passive materials such as the toroid coils, vacuum vessels and magnet support structures (Fig. 125). In the barrel, chambers are mounted on aluminium structures supported by the coils of the toroid magnet. In the end-cap, special wheel-like aluminium support structures have been built to carry the MDT's, TGC's and their respective services. Fig. 126 shows the material to be transversed by a muon in one octant in units of radiation lengths (X_0). A barrel-toroid coil is visible in this figure at $\phi = 22.5^\circ$ and two half-end-cap coils at $\phi = 0^\circ$ and 45° . The increase in radiative thickness above $\eta \sim 1.4$ is caused primarily by the heavy mechanical structure of the end-cap toroid. Over a large fraction of the acceptance in the barrel toroid, the cumulative amount of material traversed by a muon is about $1.3 X_0$. As a consequence, multiple scattering effects dominate the momentum resolution between 30 GeV and 200 GeV.




Figure 126. Amount of material in units of radiation lengths (X_0) traversed by muons after exiting the calorimeters, as a function of η and ϕ .

6.3 Monitored drift tube chambers

6.3.1 Structure and function of the drift tube

The basic element of the monitored drift tube chambers is a pressurised drift tube with a diameter of 29.970 mm, operating with Ar/CO₂ gas (93/7) at 3 bar (see Fig. 127). The electrons resulting from ionisation are collected at the central tungsten-rhenium wire with a diameter of 50 μ m, at a potential




Figure 127. Cross-section of a MDT tube.




Figure 128. Longitudinal cut through a MDT tube.

of 3080 V. The wire is held in position at the tube ends by a cylindrical end-plug (Fig. 128) which
 5370 guarantees the concentricity of the wire with respect to the tube with an accuracy of $\sigma < 10 \mu\text{m}$.
 The central conductor holding the wire also serves for the gas transfer in and out of the tube. Signal
 transmission to the electronics and connection to the HV supply system are at opposite ends. The
 choice of the operating parameters of the drift tubes is discussed in [164, 165, 166].

Building the precision-tracking chambers out of individual tubes offers several advantages.
 5375 The stiffness of the tube assembly allows to combine high mechanical precision with robustness
 of the chambers. A high level of operational reliability can be expected because the failure of
 a single tube does not affect the operation of most of the others. The tube concept lends itself
 to high-pressure operation, thereby reducing the deleterious effect of diffusion on the single-wire
 resolution. One more advantage is that the cylindrical geometry results in a radial electric field: the
 5380 measurement accuracy, therefore, depends only weakly on the angle of incidence of the track onto
 the chamber plane, as the coordinate of the track is determined by the radius of the circle around the
 wire to which the track is tangential (see Fig. 127). This is important as the angle of incidence of
 infinite momentum tracks onto the chamber plane extends up to 45° . In a rectangular drift geometry,
 the field lines are such that the path length travelled by a electron resulting from ionisation would
 5385 vary significantly depending on where along the muon track that electron is produced, thereby
 degrading the drift-time resolution.

A disadvantage of the radial drift geometry is the long pulse train caused by the track segments
 far from the tangential point, which may produce several threshold crossings (hits) per track. The
 maximum drift time from the wall to the wire is about 700 ns. A track passing close to the wire
 5390 thus generates a pulse train with a duration of this order, while only the arrival time of the central
 track part, closest to the wire, is relevant for the track coordinate (r_{min} in Fig. 127). To prevent an
 inflation of the data volume by multiple track hits, an adjustable dead-time has been implemented
 in the front-end of the readout chain, see Section 6.3.3.

The operating gas was selected because of the good ageing properties. Deposits on the wires
 5395 have never been observed in clean samples of this gas mixture, the formation of polymers not being
 possible in the absence of hydrogen. A high degree of cleanliness is however required in all parts
 of the gas system as traces of hydrocarbons or silicon oil may lead to deposits on the wires and
 a reduction of the signal pulse height [167, 168, 169, 170]. A disadvantage of this gas mixture

is the non-linear space-drift time relation and the drift time of about 700 ns, which is about 50% longer than is typical for linear gases such as Ar/CH₄. The non-linearity of the Ar/CO₂ gas leads to a reduction of spatial resolution at high counting rates due to the distortion of the electric field created by the positive ions. At full LHC luminosity, counting rates of up to 30 kHz per tube will be expected due to the conversion of background photons and neutrons [35, 37, 171]. The corresponding degradation of the average resolution has been determined in tests at high gamma backgrounds and is expected to be 60-80 μm per tube at the expected background levels [172, 173, 174, 168]. Detailed results are given in Section 6.3.5. An additional complication for tracking comes from the fact that the detailed shape of the space drift-time relation in ArCO₂ depends on environmental parameters like temperature and pressure as well as on the local magnetic field due to the Lorenz force. In order to maintain the high spatial resolution under varying environmental conditions, an online calibration system based on measured tracks is foreseen [175, 176].

A small water admixture to the gas of about 300 ppm is foreseen to improve HV stability. The effect of this admixture on the drift behaviour is expected to be negligible [177].

6.3.2 Mechanical structure

The main parameters of the MDT chambers are listed in Table 33. The chambers are rectangular in the barrel and trapezoidal in the end-cap. Their shapes and dimensions were chosen to optimise solid angle coverage, while respecting the envelopes of the magnet coils, support structures and access ducts. A detailed presentation of MDT chamber construction and installation is given in [178, 179, 180].

The naming of chambers is based on their location in the barrel or end-cap (B,E), their assignment to inner, middle, or outer chamber layer (I, M, O) and their belonging to a large or a small sector (L,S). The sector number (1–16) and the sequence number of the chamber in a row of chambers in a sector are added to completely specify a MDT chamber. A BOS chamber, for example, is located in a small sector of the barrel, outer layer, while an EML lies in the large sector of the middle layer of the end-cap. In sectors 12 and 14 (see Fig. 121) of the barrel, special chambers were designed to keep the acceptance losses due to the ATLAS support structure (“feet”) to a minimum. Modified BOS chambers are therefore called BOF or BOG, and the modified BMS chambers BMF. The BEE (Barrel End-cap Extra) are special chambers located in the castellations of the end-cap toroid cryostats. Although labelled barrel chambers, the BEE are used in the measurement of tracks passing from the barrel to the end-cap. The BEE chambers consist of a single multi-layer of four tube layers.

As can be seen in Fig. 122, the outer part (in the radial direction) of the EI wheel does not project into the EO wheel. In order to allow for momentum measurement in this region, an intermediate ring of chambers, the EES and EEL (the “E” is derived from extra) chambers, have been introduced with an offset of about 3-3.6 m with respect to the corresponding EI wheel chambers.

All regular MDT chambers consist of two groups of tube layers, called multi-layers, separated by a mechanical spacer. In the innermost layer of the muon detector, each multi-layer consists of four tube layers to enhance the pattern-recognition performance; in the middle and outer layer of the muon detector, each multi-layer consists of three tube layers only. Fig. 129 shows the structure of a barrel chamber with 2 × 3 tube layers. The height of the support beam between the multi-

Parameter	Design value
Tube material	Al
Outer tube diameter	29.970 mm
Tube wall thickness	0.4 mm
Wire material	gold-plated W/Re (97/3)
Wire diameter	50 μm
Gas mixture	Ar/CO ₂ /H ₂ O (93/7/ \leq 1000 ppm)
Gas pressure	3 bar (absolute)
Gas gain	2×10^4
Wire potential	3080 V
Maximum drift time	\sim 700 ns
Average resolution per tube	\sim 80 μm

Table 33. Main MDT chamber parameters.

5440 layers (spacer) depends on the chamber type, varying from 6.5 mm to 317 mm (Tables 34 and 35). Detailed information about chamber dimensions and other parameters is available in [181].

During chamber construction, a high precision of tube placement and a high level of mechanical strength had to be achieved in order to maintain the inherent resolution of the drift tubes. The two multi-layers of a MDT chamber are mounted on a support frame of solid aluminium beams, providing mechanical rigidity to the structure. The lateral support beams (cross-plates), designated by RO for readout and HV for the high voltage supply side, are interconnected by two longitudinal beams. Three kinematic mounting points, not shown in Fig. 129, are attached to the frame for installation onto the rail system of ATLAS.

5450 The frames also carry most of the interfaces to gas supplies, and to the electrical, monitoring and alignment services. The 3 (4) tube layers of a multi-layer are joined together with epoxy glue, layer by layer, after having been carefully aligned on a granite table. The reference surface for the tubes is the precisely-milled end-plug, which also serves as reference for wire positioning. This method ensures a high precision of relative wire positioning at construction time.

5455 The straightness of the tubes is required to be better than 100 μm . The relative positioning of wires reached during production, has been verified to be better than 20 μm . The gap between adjacent tubes filled by glue is 60 μm . A detailed account of MDT chamber construction and quality assurance is given in [182, 183, 184, 185, 186, 187].

5460 In spite of the solid construction of the MDT chambers, deformations are expected to occur in the various mounting positions in ATLAS and may change in time when thermal gradients are present. Therefore, an internal chamber alignment system was implemented, which continuously monitors potential deformations of the frame. The alignment system consists of a set of four




Figure 129. Mechanical structure of a MDT chamber. Three spacer bars connected by longitudinal beams form an aluminium space frame, carrying two multi-layers of three or four drift tube layers. Four optical alignment rays, two parallel and two diagonal, allow for monitoring of the internal geometry of the chamber. RO and HV designate the location of the readout electronics and high voltage supplies, respectively.

optical alignment rays, two running parallel to the tube direction and two in the diagonal direction as shown in Fig. 129. The lenses for the light rays are housed in the middle, while LED's and CCD sensors are located in the outer spacers. This system can record deformations of a few μm and is designed to operate during production, installation, and operation of ATLAS. Details of the in-plane alignment system of the MDT chambers are given in Section 6.3.4.

Due to gravitational forces, chambers are not perfectly straight but suffer a certain elastic deformation. The BOS chambers for example, with a tube length of 3.77 m, have a gravitational sag of about $800 \mu\text{m}$ when supported at the two ends in the horizontal position. The wires in the tubes have only $200 \mu\text{m}$ sag at their nominal tension of 350 g. In order to re-establish the centricity of the wires, the sag of the multi-layers can be corrected by the sag-adjustment system, which applies an adjustable force to the central cross-plate. Using the in-plane alignment system as reference, deformations can be corrected with a precision of about $10 \mu\text{m}$. Thus, for each angle of installation in the ATLAS detector, the sag of drift tubes and wires can be matched, leading to wire centricity and circular symmetry of the drift field. This technique, however, is only applicable for the barrel chambers, while for the outermost end-cap chambers (EOL, EOS, EEL, and EES) they have been corrected for by construction.

The precision achieved in construction has been verified in cosmic ray tests and in the X-ray tomography facility operated at CERN. The result was an average deviation of the wire positions of $\leq 20 \mu\text{m}$ from the nominal over the full area of the chamber [188, 189].

Due to the tight construction tolerances and to the continued monitoring of global chamber deformations, the relative positions of MDT wires are sufficiently well known for the accuracy of a track segment in the 6 (8) tube layers to be limited only by the single-tube resolution (about $80 \mu\text{m}$). Therefore, the resolution on the central point of a track segment in a 3 (4)-tube multi-layer is 50

Name	Layer	Standard chambers	Special chambers	Tube layers	Location in R	Tubes/layer	Tubes total	Length along z	Width along ϕ	Spacer height
BIS	inner	96		2×4	4550	30	240	916	1820	6.5
BIS.7	inner		16	2×4	4550	30	240	916	1820	6.5
BIS.8	inner		16	1×3	4620	16	48	496	1000	–
BIL	inner	72		2×4	4949	36	288	1096	2820	170
BIM	inner		20	2×4	5373	36	288	1096	1685	170
BIR	inner		24	2×4	6056	36	288	916	1685	170
BMS	middle	72		2×3	8095	48	288	1497	3220	170
BMF	middle		12	2×3	8095	64	384	1937	3220	170
BML	middle	94		2×3	7139	56	336	1697	3700	317
BEE	middle		32	1×4	4415	48	192	1457	1060	–
BOS	outer	72		2×3	10569	72	432	2177	3920	317
BOF	outer		16	2×3	10675	64	384	2177	3920	317
BOG	outer		18	2×3	10675	40	240	1216	3920	317
BOL	outer	96		2×3	9500	72	432	2177	5110	317
Total		502	154							

Table 34. Parameters of MDT barrel chambers. Numbers and dimensions (mm) of the tubes refer to the most common values. The special chambers have non-standard shapes to fit between magnet coils, support structures etc. The BIS.8 and BEE chambers have only one multi-layer. The BEE are located at the outer circumference of the end-cap cryostat and record tracks passing from the barrel to the end-cap.

5485 (40) μm ; combining the two multi-layers into a chamber yields an accuracy of 35 (30) μm . The relative positions of the chambers in radially-consecutive stations are monitored by the projective alignment system with a design accuracy of about 30 μm (Section 6.3.4). For a track crossing three MDT chambers, a sagitta resolution of $\Delta S = 45 \mu\text{m}$ is thus expected, corresponding to a momentum resolution $\delta p/p = \Delta S \times p/500 \mu\text{m}$, where p is given in units of TeV.

5490 To monitor thermal deformation, each chamber carries from 3 to 28 temperature sensors, depending on chamber type, while the local magnetic field is monitored by two to four B-field sensors per chamber. Both environmental parameters are read out via the Detector Control System (DCS) (see Section 8.5), in which each chamber represents a node.

5495 An important design criterion in chamber construction is electrical integrity: low-impedance connections have to be provided between all metallic components. Because the bonds between most chamber components are epoxy glue, the gas bars, support structures and Faraday cages are connected with solid wire straps. The end-plugs of the tubes are screwed to a continuous metal plane (ground plane) which in turn is connected to the support structure. The Faraday cage which surrounds the electronics is made of thin aluminium sheets which have undergone a chemical treatment on the surface (chromatisation) to assure a conductive, unalterable surface.

5500 To avoid ground loops among chambers, the kinematic mounts which slide on the rails have

Name	Layer	Chambers	Tube layers	Location in $ z $	Radial chambers	Tubes/layer	Tubes total	Radial length	Width along ϕ	Spacer height
EIS	inner	32	2 × 4	7261	2	36	288	1096	1745	121
EIL	inner	48	2 × 4	7675	4	54	432	1637	3295	121
EES	extra	32	2 × 3	10276	2	40	240	1216	2951	121
EEL	extra	30	2 × 3	11322	2	40	240	1216	4703	121
EMS	middle	80	2 × 3	13878	5	64	384	1937	3860	170
EML	middle	80	2 × 3	14294	5	64	384	1937	6008	170
EOS	outer	96	2 × 3	21424	6	48	288	1457	4202	170
EOL	outer	96	2 × 3	21840	6	48	288	1457	6503	170
Total		494								

Table 35. MDT chambers in the end-cap. In the region $1.3 < |\eta| < 2.4$, the EI, EM and EO chambers provide the three points for momentum measurement while in the transition region $1.0 < |\eta| < 1.3$ not covered by the EO wheel, the second and third points are provided by the EE and EM chambers (see Fig. 122). The EE chambers will only be installed in 2009. Dimensions are in mm. Values in columns 9–10 apply to the outermost trapezoidal chambers in each radial row.

a ceramic inner surface, while non-conductive transition pieces are used to connect the gas bars to the external supplies. Care is also taken to isolate the readout of the in-plane alignment (RAS-NIK [190]) and the DCS from the chamber ground. After installation, the resistance between individual chambers and the ATLAS structure is typically $> 1 \text{ G}\Omega$. To assure a low-impedance connection to the safety ground, rows of chambers are connected to *one* common ground line, which, in turn, is connected to the ATLAS structure at a single point. A detailed account of final quality controls and chamber installation in the barrel is given in [191].

6.3.3 Signal path and readout electronics

The architecture of the MDT readout electronics chain is shown in Fig. 130. In the first stage, the raw tube signals are amplified, shaped and discriminated, eight tubes being served by one custom-designed monolithic ASD (Amplifier/Shaper/Discriminator) chip. The binary differential signals output by the ASD are routed to the Time-to-Digital Converter (TDC), where the arrival times of leading and trailing edges are stored in a large buffer memory together with an identifier of the corresponding tube and a flag indicating whether the edge is leading or trailing. The time is measured in units of the Timing, Trigger and Control (TTC) clock of 40.08 MHz, 12 bits being used for the beam-crossing identification and 5 bits for the subdivision of the 25 ns-period (fine time). The fine time period of 0.78 ns therefore leads to a RMS timing error of 0.23 ns corresponding to an average position error of about $5 \mu\text{m}$. A detailed description of the MDT readout and of its components is given in [192, 193, 194].

An additional feature of the ASD is the measurement of the pulse height of the signal, which allows the monitoring of the gas gain as well as pulse height dependent corrections (slewing corrections) [195]. The pulse height is measured by an Analogue-to-Digital Converter (ADC) and encoded as the delay between leading and trailing edge, a delay of 150 ns corresponding to the maximum pulse height to be recorded by the ADC. Thus the *trailing* edge of the ASD output is




Figure 130. Schematic diagram of the MDT readout electronics. See text for details.

not determined by the discriminator status and does not contain any timing information. An internal dead-time disables the discriminator until after the pulse height encoding has safely elapsed. Beyond this minimum value, the dead-time is programmable and can be increased up to 540 ns in order to mask multiple hits from the same track. After the end of the dead-time, the first low-to-high transition of the discriminator will create the next leading edge of the ASD output. The standard dead-time setting will be 540 ns.

Programming of the ASD and TDC is done following the JTAG protocol via a shift register, allowing for a string of 54 and 196 bits for the ASD and the TDC, respectively. This way, many parameters and functions can be controlled, like the setting of discriminator threshold and dead-time as well as triggering of test pulses for calibration or deactivation of noisy channels. The TDC with three ASD's, serving 24 tubes, is implemented on a $9 \times 11 \text{ cm}^2$ printed circuit-board (mezzanine card). The mezzanine card is the basic readout element of the MDT's and is connected to the tubes via a signal distribution card (hedgehog board). There are about 15000 mezzanine cards to serve the 354000 MDT tubes in the system.

MDT chambers contain up to 18 mezzanine cards, which are controlled by a local processor, the Chamber Service Module (CSM), see Fig 130. The CSM broadcasts the TTC signals to the TDC's and collects the hits they sent back for each L1 trigger (see Section 8.2). Subsequently, the data are formatted, stored in a large derandomising buffer, and sent via optical link to the MDT Readout Driver (ROD) in the USA15 service cavern. The ROD is a VME module, serving up to six CSM's. Its main task is to format the front-end data for rapid transfer to the Readout Buffer (ROB), where data are stored until the event has been either validated or discarded by the L2 trigger logic. Table 36 summarises the modularity of the readout system.

An important feature of the ROD is large processing power, which allows it to monitor the incoming data. The data stream received from the CSM's can be sampled for tube and chamber occupancies and their deviations from nominal values, possibly pointing to a malfunction. As the

Region	Tubes	ASD's	Mezzanine cards	CSM's	Chambers	ROD's
Barrel	191904	24162	8054	624	656	110
Two end-caps	162240	20280	6760	494	494	94
Total	354144	44442	14814	1118	1150	204

Table 36. Modularity of the MDT readout electronics. The difference between the number of chambers and CSM's in the barrel is due to each pair of BEE and BIS.8 chambers being served by only one CSM. The allocation of ROD's to the barrel or end-cap is approximate, some ROD's serving CSM's from both regions. See the text for definition of the acronyms.

ROD's see the full L1 event rate, they can accumulate significant statistics in a short time, allowing for an early recognition of errors. Depending on average event size and L1 rate, a variable fraction of the events will be monitored in order to not slow down data transfer to the ROB's, which has priority. A detailed description of the MDT readout is given in [196] and [197].

5555 6.3.4 Alignment system

In the ATLAS spectrometer, the MDT chambers and their auxiliary alignment objects (alignment bars and reference system plates, see below) are installed with a precision of about 5 mm and 2 mrad with respect to their nominal positions. In contrast, the actual chamber locations must be known to a precision of $\leq 30 \mu\text{m}$ in order to achieve the required momentum resolution: this is more than two orders of magnitude tighter than the positioning accuracy. To reach this precision goal, an alignment system was built which relates the position of each chamber to that of its neighbours, both within an MDT layer and along $R - z$ trajectories within MDT towers. With the internal accuracy of the MDT's guaranteed by construction, and their planarity monitored by optical in-plane alignment sensors (see Section 6.3.2), this alignment system forms a dense, stable grid for monitoring (in the barrel) or reconstructing (in the end-cap) the position of each MDT wire in the muon system. A schematic representation of the alignment system in a large sector is shown in Fig. 131. Due to different geometrical constraints, the alignment strategies are somewhat different in the barrel and end-cap regions (Section 6.3.4.1) and several specialised optical alignment sensors have been developed (Section 6.3.4.2). The overall system layout is detailed in Section 6.3.4.3.

5570 The optical alignment techniques described below are insufficient to reconstruct, on their own, the absolute positions of the MDT barrel chambers: only variations in relative position can be determined with the required precision. Track-based alignment algorithms must therefore be used in combination with the optical system to achieve the desired sagitta accuracy, and also to determine the global positions of the barrel and end-cap muon-chamber systems with respect to each other and to the inner detector, in the overall ATLAS reference frame (Section 10.3.2). A more extensive presentation of the alignment strategy and technical implementation can be found in [198, 199].

6.3.4.1 Alignment strategies In the barrel, the chambers inside a MDT row are referenced to each other by chamber-to-chamber alignment sensors (praxial and axial systems, see Fig. 131),




Figure 131. Principle of the alignment of the ATLAS muon spectrometer. See text for details.

while the projective system connects the inner, middle, and outer MDT layers, simulating the trajectories of an infinite-momentum track originating from the interaction point. This optical network is able to reliably detect relative changes in chamber position at the $20 \mu\text{m}$ level. However, because of its essentially projective configuration, it remains insensitive to some specific geometrical distortions, and therefore unable to provide absolute chamber positions on its own. Additional optical lines (Table 37) link MDT chambers to the barrel toroid cryostat, or relate adjoining chambers in large and small sectors. Once fully functional, the optical alignment system will provide a slightly overconstrained network of optical sensors which can determine the absolute chamber positions in the barrel region to an accuracy of a few hundred microns, and yield a systematic uncertainty of $100 - 200 \mu\text{m}$ on the residual sagitta. Because of mechanical positioning uncertainties on some optical sensors, the ultimate $30 \mu\text{m}$ precision on the sagitta (but not necessarily on the absolute chamber positions), and in particular the exact reference values with respect to which relative changes are monitored, can only be reached, in the barrel region, by combining optical and track-based alignment algorithms.

In the end-caps, the light path between the inner and the middle layer is obstructed by the cryostat vessels of the toroid magnets. The polar alignment therefore proceeds in two steps using high-precision rulers (alignment bars), as an intermediate reference [200]. The polar (quasi-projective) optical lines connect the alignment bars installed in each EI, EM and EO layer; the chambers, in turn, are optically connected to their neighbouring bars by proximity sensors. Additional sensor families, listed in Table 37, determine the remaining degrees of freedom. Simulations have demonstrated that this network is sufficiently precise and overconstrained to determine absolute chamber positions within each end-cap region to a $300 \mu\text{m}$ accuracy, and relative chamber positions along muon trajectories to a $40 \mu\text{m}$ accuracy on the sagitta. Test-beam results and preliminary tests after installation of the EM layers in the ATLAS cavern have already demonstrated that an alignment accuracy at about this level is reached.

The alignment bars in the EI wheel are also used for the alignment of the CSC's. The positions




Figure 132. The RASNIK alignment system. The image sensor (RasCam) is a CMOS sensor in the barrel and a CCD sensor in the end-cap. An infrared filter is placed in front of the sensors to avoid stray light. A RasMux multiplexer is installed on each chamber, servicing up to eight RasCam sensors. A MasterMux can multiplex up to 16 RasMux's, sending the data to the USA15 hall for processing.

5605 of the large CSC's relative to these bars and the position of the small CSC's relative to the large ones are measured with optical sensors, in a similar way as is done for the MDT's. The alignment precision of the CSC's is thus expected to be at the same level as for the MDT's.

The relative positions of the barrel MDT system and of the forward end-cap chamber systems must also be known to a few tens of microns to ensure accurate momentum measurement for muons
 5610 crossing from one region to the other. Track-based alignment algorithms are being developed to provide the required precision.

6.3.4.2 Optical alignment sensors Three different optical systems are used in the ATLAS alignment: RASNIK's, BCAM's and SaCam's. All three systems are based on the principle of a three point straightness monitor: an optoelectronic image sensor (CCD or CMOS) monitors the position
 5615 of an illuminated target through a lens. The sensor image is analysed online and converted into four parameters characterising the deviation from the nominal geometry: two translations perpendicular to the optical axis, the rotation around and the longitudinal position along the axis, the latter being derived from the optical magnification.

Fig. 132 shows a schematic of the RASNIK system, the most frequently used element of the muon alignment system. An image sensor monitors through a lens an illuminated target carrying a chequer-board pattern imprinted on a mask. For a symmetric RASNIK with the lens halfway between image sensor and mask, a transverse resolution of $1 \mu\text{m}$ and a magnification resolution of 2×10^{-5} have been obtained. A description of the RASNIK alignment system can be found in [190] and [201].
 5620

The BCAM is a camera consisting of a CCD image sensor which monitors the position of a laser diode through a lens. The BCAM housing also contains a set of two or four laser diodes which can be used as targets by another BCAM. BCAM's are mostly used in pairs, facing each other. The image sensor is placed close to the focal plane of the lens (76 mm), while the target with the laser diodes is placed at a distance of 0.5 m to 16 m. The image of a laser diode on the CCD appears as a circular spot.
 5625
 5630

BCAM's achieve a resolution of $5 \mu\text{rad}$ for a target distance of 16 m. If the target consists of two or more laser diodes with a known separation, relative angles can be extracted, which yield




Figure 133. Layout of the optical-alignment lines (red) for three adjacent barrel sectors. See text and Table 37 for details. The Chamber-to-Chamber Connector sensors (CCC) connect chambers in a small sector to those in an adjacent large sector.

magnification and rotation around the optical axis. If BCAM's are used in pairs, the absolute angular position of the partner can be determined with $50 \mu\text{rad}$ accuracy.

5635 Like the BCAM, the SaCam consists of a camera with an image sensor and a lens at fixed distance. The target consists of four back-illuminated holes, covered by a light diffuser. The lens is mounted at distances ranging from 25 mm to 80 mm, depending on the distance between image sensor and target. A relative resolution of $5 \mu\text{rad}$ has been achieved with the SaCam, a result similar to that of the BCAM.

5640 Another specialised sensor type is the 3D sensor, a pin-hole camera similar to the BCAM, which monitors the position of four laser diodes placed at two different distances along the optical axis. This sensor type is used for the alignment of the CSC's.

6.3.4.3 Layout of the alignment system The alignment rays in three adjacent barrel sectors are shown in Fig. 133. The praxial systems (two short, crossed RASNIK's connecting adjacent chamber corners) and the axial systems (interleaved long RASNIK systems along the chamber) monitor the geometry of each row of six chambers. The relative position of the chamber layers perpendicular to the polar direction is monitored by the projective rays. The projective system works with RASNIK's, where BIL chambers hold the mask, BML's the lenses and BOL's the image sensor. Chambers in the small sectors (BIS, BMS and BOS) are not equipped with projective lines and are referenced to the adjacent large chambers by Chamber-Chamber Connection sensor (CCC).
5645
5650

The layout of the alignment rays in the end-cap is shown in Fig. 134. The presence of the cryostat of the end-cap toroid magnet allows only for three polar lines per small sector, one of them passing through a horizontal stay tube connecting the two outer flanges of the cryostat. To cope




Figure 134. Layout of the alignment of two MDT's and CSC's in the end-cap. Only alignment sensors belonging to these sectors are shown, and the EE chambers and bars have been omitted. See text and Table 37 for details.

Barrel alignment				End-cap alignment			
Name	Type	Number	Function	Name	Type	Number	Function
In-plane	R	2110	MDT deformation	In-plane	R	1984	MDT deformation
Praxial	R	2006	MDT-MDT distance	In-bar	R	352	Bar deformation
Axial	R	1036	MDT-MDT planarity	Radial	B	96	Bar deformation
Projective	R	117	Tower alignment	Polar	B	208	Bar-bar link
Reference	S	256	Link to toroid	Azimuthal	B	736	Bar-bar link
CCC	S	260	Small-large link	Proximity	R	2384	MDT-bar and MDT-MDT link
BIR-BIM	R	32	BIR/BIM link	Saloon door	B	584	MDT-bar link
				3D sensor	B	192	CSC-bar and CSC-CSC link
Total:		5817		Total:		6536	

Table 37. The alignment components in barrel and end-cap. Three different types of imaging systems are used, adapted to the special application: RASNIK (R), BCAM (B) and SaCam (S). The BIR and BIM chambers are shown in Fig. 121. See text for details.

5655 with this limitation, a grid of alignment bars using a quasi-projective alignment system was built. The chamber positions are referenced to this grid of aligned precision rulers. Each MDT wheel carries eight bars, with lengths of up to 9.6 m, placed at the border between large and small sectors. The bar straightness and elongation are monitored by a set of internal, interleaved RASNIK's and by a number of temperature sensors. Prior to installation, the geometry of each bar is carefully measured using a large coordinate measuring machine.

5660 The bar positions in the grid are determined by a network of BCAM sensors measuring the bearing angle of light sources located on adjacent bars. Radial BCAM's monitor the bar deformation, azimuthal BCAM's connect adjacent bars within a wheel, and polar BCAM's connect bars of different wheels. MDT chambers are referenced to the alignment grid using RASNIK proximity sensors, where image sensor and lens are placed on the chamber, while the mask is on the neighbouring bar. In order to monitor chamber rotations around the bar axis, azimuthal BCAM's also look at targets (the saloon-door sources) on the opposite edge of the adjacent MDT chamber. The positions of CSC chambers relative to the bars and to each other are measured by 3D sensors.

6.3.5 Performance of the MDT chambers

MDT chambers have been tested under various conditions, trying to match the environment of the final experiment as closely as possible.

5675 A set of 12 MDT chambers were tested in a high-energy muon beam in 2003 and 2004 in the H8 area at CERN. Six barrel MDT's (two of each type of BIL, BML, BMS chambers) were arranged as a barrel projective tower part of a large sector. Six end-cap MDT's were arranged to reproduce a projective region formed by one large and one small adjacent chambers in each of the EI, EM and EO stations. The main purpose was to verify the functionality and performance of the readout electronics and the projective alignment system. The resolution dependence on thresholds and gas gain was tested as well as efficiency and noise rates. The alignment system was used to monitor the chamber positions relative to each other recording, for example, day-night displacements of the chambers due to temperature. The optical alignment and the track alignment with the MDT's demonstrated the expected correlation with an accuracy $20 \mu\text{m}$. These test-beam data were also used to develop the calibration algorithms measuring the space-time relation and to correct for the time offsets of the individual tubes. The calibration procedure is described in [176].

5685 In 2004, tests were performed with the muon trigger chambers and with other sub-systems like the pixel detector, SCT, TRT, and the calorimeters. The main result was a confirmation of the performance with respect to the resolution parameters, readout functionality and stability of the electronics as well as the correct functioning of the alignment system. The results of the tests in 2003 and 2004 are presented in [185, 202, 203].

5690 Additional tests were done in 2003/2004 at the Gamma Irradiation Facility (GIF) at CERN to test the behaviour of the chambers in a high background environment as is expected in the ATLAS hall. In the GIF test area, a high-energy muon beam was available as well as an adjustable, intense γ -source. The muon tracks were defined by a silicon telescope and the resolution of the MDT tubes could be studied accurately as a function of the γ -intensity and of the track's distance from the central wire. The expectation was that electrons from Compton scattering in the tube walls would lead to high occupancies, while positive ions, slowly drifting towards the wall, would locally distort the drift field, leading to a radius-dependent degradation of the resolution. Tests done in 5695 1999 [172] on a few tubes, and in 2003 and 2004 on full chambers using final electronics [173, 174], demonstrated the feasibility of using Ar/CO₂ as the drift gas. Figure 135 shows the most recent results. The degradation of resolution with gamma intensity is clearly visible and is highest for the longest drift path close to the wall ($R = 15 \text{ mm}$). The highest irradiation level expected for MDT chambers is $\leq 150 \text{ hits}/(\text{cm}^2\text{s})$. In the centre of the tube ($R = 7.5 \text{ mm}$) this would degrade the




Figure 135. Resolution as a function of the impact parameter of the track with respect to the tube wire at various levels of γ -irradiation with a threshold of 16 photoelectrons (p.e.). The maximum rate expected for the MDT's is ≤ 150 hits/cm²s.

resolution from 60 μm to about 80 μm , which is acceptable.

6.4 Cathode-strip chambers

6.4.1 Layout of the CSC system

The limit for safe operation of the MDT's is at counting rates of about 150 Hz/cm², which will be exceeded in the region $|\eta| > 2$ in the first layer of the end-cap. In this η region of the first layer, the MDT's are replaced by cathode-strip chambers, which combines high spatial, time and double track resolution with high-rate capability and low neutron sensitivity. Operation is considered safe up to counting rates of about 1000 Hz/cm², which is sufficient up to the forward boundary of the muon system at $|\eta| = 2.7$.

As in the case of the MDT's, the CSC's are segmented into large and small chambers in ϕ . The whole CSC system consists of two disks with eight chambers each (eight small and eight large) as shown in Fig. 136. Each chamber contains four CSC planes resulting in four independent measurements in η and ϕ along each track.

The CSC's are multiwire proportional chambers with both cathodes segmented, one with the strips perpendicular to the wires (providing the precision coordinate) and the other parallel to the wires providing the transverse coordinate. The position of the track is obtained by interpolation between the charges induced on neighbouring cathode strips. The CSC wire signals are not read out.

The resolution achieved with this procedure depends on the signal-to-noise ratio and the read-out pitch, the latter being the main cost-driving factor for the readout electronics. With a readout pitch of 5.08 mm in the bending direction, the CSC reaches a resolution of 60 μm per CSC plane,

Parameter	Value
Operating voltage	1900 V
Anode wire diameter	30 μm
Gas gain	6×10^4
Gas mixture	Ar/CO ₂ (80/20)
Total ionisation (normal track)	90 ion pairs

Table 38. Operating parameters of the CSC's.

to be compared with the 80 μm resolution of a MDT tube layer. In the non-bending direction the cathode segmentation is coarser leading to a resolution of 5 mm.

Apart from the precision and relative simplicity of the coordinate determination, there are a number of other characteristics which make the CSC's suitable for regions of high particle densities:

- a) Good two-track resolution.
- b) Pairing of the measurements in the two coordinates via the pulse height to resolve the ambiguities if more than one track is present.
- c) Electron drift times of less than 40 ns resulting in a timing resolution of about 7 ns per plane.
- d) Low neutron sensitivity because of the small gas volume and the absence of hydrogen in the chamber gas (Ar/CO₂).

A detailed description of the CSC system is given in [204]. The operating parameters of the CSC are shown in Table 38.

6.4.2 Spatial and time resolution

Fig. 137 shows the symmetric cell structure of the CSC, the anode-cathode spacing being equal to the anode-wire pitch. The measured charge distribution as induced onto the cathode strips by the avalanche on the wire is given in Fig. 138.

The segmentation of the cathode aims to sample the induced charge distribution as precisely as possible while limiting the number of electronic readout channels. The following scheme has been chosen: between two neighbouring readout strips (connected to amplifiers) there are two intermediate (floating) strips capacitatively coupling the induced signal to the readout strips. The width of readout and intermediate strips was selected to minimise the differential non-linearity of the position measurement to $\leq 1\%$. Fig. 139 shows the segmentation of the CSC cathode where (for the




Figure 136. Layout of a CSC end-cap with eight small and eight large chambers.




Figure 137. Left: structure of the CSC cells looking down the wires. The wire pitch s is equal to the anode-cathode spacing $d = 2.54$ mm. Right: view in the perpendicular direction (bending plane), down the readout and intermediate strips. The spacing of the readout strips w is 5.08 mm. The induction of the avalanche is spread out over 3–5 readout strips (see Figs. 138 and 139).




Figure 138. Charge distribution on the CSC cathode induced by the avalanche on the wire.

large chambers) readout and intermediate strips are 1.13 mm and 1.72 mm wide, respectively, with a gap between strips of 0.25 mm. The interstrip capacitance is about 10 times the strip capacitance to ground. The intermediate strips are connected to ground via a high-resistance path to define the DC potential.

5755

As the precision coordinate in a CSC is obtained by the *relative* measurement of charges induced on adjacent strips, the performance is immune to the variation of conditions encountered by the whole chamber, like gas gain, temperature variations or pressure. The primary limiting factor for the spatial resolution of the CSC's is electronic noise of the pre-amplifiers, and therefore only a small number of strips around the centre is used in the clustering algorithm. In this geometry, the best results are obtained with 3 to 5 strips around the peak of the distribution.

5760

The second most significant contribution to the spatial resolution of the CSC is the effect of inclined tracks and the Lorentz angle. The charge interpolation is optimal when the avalanche is formed on a single point along the wire, while a significant extent of the anode charge results in a degradation. The spreading of the charge deposition can be caused by delta electrons, inclined tracks and a Lorentz force along the wire, as the magnetic field is not collinear with the electric field around the wire. While the effect of the Lorentz force cannot be corrected in a CSC, it does not lead to a systematic shift of the measured coordinate, but only to a slight reduction in resolution. The effect of inclined tracks is minimised by tilting the chamber towards the interaction point so that infinite momentum tracks are normal at the centre of the chamber (Fig. 140).

5765

5770

The strips in the second cathode plane are parallel to the wires and measure the position of the avalanche in the non-bending ϕ -plane. The spacing of the readout strips is larger than the one in





Figure 139. The spacing x of the readout strips is 5.08 mm. The intermediate strips lead to a gaussian-like distribution of the signal by capacitive coupling to the adjacent readout strips. The intermediate strips are not connected to readout electronics.




Figure 140. The inclination of the CSC's towards the interaction point. Dimensions are given in mm.

the bending η -plane (15.7 mm) leading to a resolution of about 5 mm.

If more than one track is present in a given event, the pulse height is used to match the measurements in the two projections. Because of the broad pulse height distribution in a gas chamber and the high resolution of the ADC's, both tracks will have, in most of the cases, a different pulse height value, allowing for an unambiguous η/ϕ assignment of the tracks from the separate measurements in η and ϕ . In the rare cases of near equality, the information from the other three chamber layers will be used. In using the information of all four chamber layers, the η/ϕ assignment for two or more tracks seems to be sufficiently robust.

The time resolution of the CSC is an important parameter for the correct tagging of the beam-crossing. The maximum drift path of a primary electron to the closest wire being only 1.25 mm, one would expect a maximum drift time of about 20 ns, given the average drift velocity of 60 $\mu\text{m}/\text{ns}$. However, the vanishing drift field at the boundary between two cells creates regions of very low drift velocity leading to long tails in the distribution of the arrival times, which prevents reliable tagging of the beam-crossing. By OR-ing the signals of corresponding cathode strips in the four chamber planes, the tails disappear, as virtually no track can pass close to a cell boundary in all four layers. The time of arrival distribution thus obtained is symmetric with an RMS value of 3.6 ns, resulting in a reliable tagging of the beam-crossing.

6.4.3 Mechanical design

The CSC's are located at a distance of about 7 m from the interaction point. They are mounted together with the MDT's and TGC's on the shielding disk (see Section 3.2) forming the so-called




Figure 141. Shapes of large and small CSC's. Dimensions are given in mm.

small wheel and occupy the radial space between 881 mm and 2081 mm, corresponding to the pseudorapidity range $|\eta| = 2 - 2.7$.

5795 The CSC's come in two types of chambers as shown in Fig. 141 and are arranged in two disks of eight chambers each, as shown in Fig. 136. Each chamber consists of four wire planes leading to a similar configuration like in the multi-layer of the MDT system, but with much finer granularity. The modularity of the CSC system is summarised in Table 39.

5800 The CSC design utilises low-mass materials to minimise multiple scattering and detector weight. A four layer chamber is formed by five flat rigid panels, each made of an 18.75 mm thick sheet of polyurethane foam and two 0.82 mm thick copper-clad laminates, where the 17 μm thick copper cladding forms the cathodes. Precision strips glued on the panels provide the 2.54 mm step for the anode wire plane. The anode wires have a diameter of 30 μm and are made of gold-plated tungsten with 3% rhenium.

5805 There are 250 and 402 wires per chamber plane along the bending direction in the small and large chambers, respectively (Table 39). The readout pitch in the bending plane is 5.56 mm and 5.31 mm for the small and large sectors respectively with a total of 192 channels per plane for both. The odd strip width is the result of the different radial chamber lengths and the desire to fix the number of channels in multiples of 192. The transverse coordinate strips have widths of 12.52 mm and 21.00 mm for the small and large chambers respectively with a total of 48 strips per plane. 5810 The cathode strips are lithographically etched for highest precision. The five panels are precisely positioned with respect to each other with the aid of locating pins. The outer copper-clad laminates of each chamber form an electromagnetic shield for the detector. A cutout view of one gap formed by two panels is shown in Fig. 142. Signals from the precision cathode strips are directly connected 5815 to the front-end (ASM-I) boards located on the chamber edges. The transverse coordinate signals




Figure 142. Structure of the CSC.




Figure 143. Model of a CSC chamber with four planes showing the location of the readout electronics.

are transferred via short ribbon cables to similar front-end boards mounted on the chamber frame (Fig. 143).

The 16 chambers of an end-cap are mounted onto precisely manufactured stainless steel frames which, in turn, are mounted on adjustable brackets of the small wheel structure.

5820 **6.4.4 Readout electronics**

The CSC's have to operate at the innermost part of the small wheel, a region which is charac-




Figure 144. Schematics of the CSC front-end electronics.

		Wires	Strips	Readout	Strips	Readout	Total	Front-end	Front-end
			channels	channels	across wires	channels			
		along wires		across wires		channels			
Plane	small	250	48	48	574	192	240	10	
	large	402	48	48	574	192	240	10	
Chamber	small	1000	192	192	2296	768	960	40	10
	large	1608	192	192	2296	768	960	40	10
One side	16 chambers	20864	3072	3072	36736	12288	15360	640	160
System	32 chambers	41728	6144	6144	73472	24576	30720	1280	320

Table 39. Modularity of the CSC system and readout.

terised by a high density of detector components, difficult accessibility and high radiation levels. As standard components like CMOS-processors, FPGA's or commercial ADC's are not expected to operate reliably in this environment, the readout electronics had to consist of custom-designed building blocks, fabricated in radiation hard technologies, combining performance with simplicity and reliability. All other commercial components were radiation certified according to ATLAS requirements. To limit the number and complexity of custom components, the front-end functionality had to be reduced to a minimum. The design therefore aimed at shifting the data as early as possible, i.e. after the L1 trigger latency to the ROD's in the shielded area in USA15.

The readout architecture of the CSC's is shown in Fig. 144. At the first stage, the chamber signals are amplified and shaped (ASM-I). At the second stage, the pulse train is sampled at a rate of 40.08 MHz and is stored in a Switched Capacitor Array (SCA) while awaiting the L1 trigger decision. Upon arrival of a L1 trigger accept signal, the stored data are digitised, multiplexed and transferred to the ROD.

A particularly demanding aspect of the centroid-finding system is to achieve an intercalibration of $< 0.5\%$ between adjacent channels. This requires high stability of the front-end amplifiers and small leakage in the analogue storage system, as well as an appropriate calibration system.

Signal processing in the front-end chip, crucial for the system performance, is as follows. Subsequent to the charge-sensing pre-amplifier, the pulse-shaping stage consists of a seventh order filter providing a bipolar pulse with a 70 ns peaking time, to reject low-frequency noise and to

assure a prompt return of the signal to the baseline. The choice of these parameters was based on an assumed maximum rate of about 600 kHz per strip at full LHC luminosity, including the usual safety factor of 5. Pre-amplifier and shaper are realised as monolithics with 24 channels per chip. The 0.5 μm CMOS technology was used because of known reliability and high radiation tolerance.

5845 Each CSC chamber has 768 readout channels in the precision (bending) and 192 in the transverse direction, so 32 and 8 chips (24 channels) are needed for the readout of both directions, respectively. A group of four chips is mounted on a printed circuit-board (ASM-I), such that the 40 chips of a chamber are served by 10 ASM-I boards, which are mounted on the lateral frame of the CSC, see Fig. 143. Two ASM-I boards in turn are served by one ASM-II board which sends the
5850 signals to the off-detector electronics via optical fibres. A breakdown of components used in the CSC readout is given in Table 39.

The signals from the ASM-II boards are routed to the CSC ROD's in the USA15 control room. The ROD's are VME modules which perform zero suppression, discard signals with incorrect timing, apply calibration constants and build event fragments to be sent to the TDAQ system. They
5855 also provide control signals for the on-chamber electronics. ROD's are interfaced to the front-end and to the TDAQ via transition modules which handle all data traffic via optical fibres.

Each ROD services two CSC's and consists of 13 DSP-based generic processing units, which communicate with the ASM-II's (10), build event fragments (2) and handle the operation of the ROD and the interface to the TDAQ (1). More details on the CSC readout are given in [205, 206].

5860 An additional constraint for the construction of the readout system came from the fact that natural air convection is practically absent around the CSC's, and so liquid cooling had to be provided to evacuate the dissipated heat.

6.4.5 Performance of the CSC

The CSC's have been tested under conditions as close as possible to those in the final experiment.
5865 For this, the Gamma Irradiation Facility at CERN has been used. In any tested chamber of four CSC planes, two layers were used as reference for the track location while the other two were considered the devices under investigation. The intensity of the gamma background could be adjusted up to 2000 Hz/cm². The results are shown in Fig. 145. Compared to zero background, the resolution degrades from $\sim 45 \mu\text{m}$ to $65 \mu\text{m}$ at 1 kHz/cm², while the inefficiency of a single CSC layer
5870 increases from $\sim 4\%$ to 10%. The 10% inefficiency at 1 kHz/cm² is acceptable, as the probability for two or more layers to be inefficient in the same event is very low.

6.5 Trigger chambers

The trigger chambers of the muon system provide fast information on muon tracks traversing the detector, allowing the L1 trigger logic to recognise their multiplicity and approximate energy range.
5875 The main requirements for the trigger system are:

- discrimination on muon transverse momentum;
- bunch-crossing identification;
- fast and coarse tracking information to be used in the high-level trigger stages;




Figure 145. CSC resolution and efficiency in a high rate test.

- second coordinate measurement in the non-bending ϕ -projection to complement the MDT measurement, see Section 6.2;
- robustness towards random hits due to n/γ -background in the experimental hall.

The trigger detectors must provide acceptance in the range $|\eta| \leq 2.4$ and over the full ϕ -range. This poses a considerable challenge to the design of the trigger system as resolution requirements in barrel and end-cap are quite different, an obvious reason being that muon momenta, corresponding to a given p_T , are strongly increasing with η . At $|\eta| = 2.4$, for example, p is about 5.8 times larger than p_T , while the integrated bending power is only about twice the value as at $\eta = 0$. This leads to the necessity of an increased and η -dependent granularity in the end-cap trigger system, if the p_T -resolution is to match the one in the barrel. The fact that the three trigger layers in the end-cap are outside the magnetic field, seeing no curvature, and that their respective distances are smaller than the ones in the barrel (Fig. 146) also calls for a finer granularity of the end-cap trigger readout. Furthermore, radiation levels in the end-cap region reach a factor of 10 higher than in the barrel. Another difficulty for end-cap triggering comes from the strong inhomogeneities of the magnetic field in the region $1.3 \leq |\eta| \leq 1.65$ as can be seen in Fig. 14. In this transition region, the superposition of the fields of barrel and end-cap toroids leads to a complex field geometry with large field components in ϕ (the non-bending plane) and strong inhomogeneities of the integrated bending power, which in two locations in the η and ϕ plane is close to zero. In this angular region, all tracks are nearly straight, similar to tracks with very high momentum. In order to avoid high fake trigger rates, this region can be excluded from the trigger by a masking algorithm, which again calls for a fine readout granularity to keep the resulting trigger losses to a minimum.

Taking these constraints into account, two different technologies have been selected for barrel ($|\eta| \leq 1.05$) and end-cap ($1.05 \leq |\eta| \leq 2.4$) regions. In the barrel, Resistive Plate Chambers (RPC's) are used due to good spatial and time resolution as well as adequate rate capability. A RPC has no wires, which simplifies its construction and makes chambers less sensitive to small deviations from planarity if appropriate spacers are used to keep the gap width constant. Being located in the comparatively homogeneous field of the barrel toroid and having sufficient spacing

between the three trigger layers (see Table 40), RPC's give sufficient trigger selectivity even with moderate channel count, i.e. spatial resolution.

In the end-cap region, Thin Gap Chambers (TGC's) have been selected: they operate on the same principle as multi-wire proportional chambers, and they provide good time resolution and high rate capability. Their spatial resolution is mainly determined by the readout channel granularity, which can be adjusted to the needs by wire ganging. TGC's have demonstrated a high level of reliability and robustness in previous experiments.

To reduce the probability of accidental triggers caused by random combinations of converted γ 's, the coincidence condition in both types of trigger chambers is established separately in the η and ϕ -projection, a valid trigger requiring a coincidence of both. This also suppresses fake triggers from curling tracks, i.e. multi-MeV electrons from γ -conversions, spiralling in the magnetic field,

potentially creating correlated hits in the trigger chambers. In the barrel, three layers of trigger chambers are implemented while in the end-cap a fourth layer is added to increase the trigger robustness in case of higher backgrounds. The schematic layout of the trigger system is shown in Fig. 146. In the barrel, two layers (RPC1 and RPC2) sandwich the MDT's of the middle layer, while the third one (RPC3) is located close to the outer MDT layer. In the end-cap, the three layers are in front (TGC1) and behind (TGC2 and TGC3) the second MDT wheel, while the fourth layer is located in front of the innermost tracking layer. The trigger information is generated by a system of fast coincidences between the three last layers along the trajectory of the muon particle. Each coincidence pattern corresponds to a certain deviation from straightness, i.e. curvature of the track, which is used as a criterion for the track to have passed a predefined momentum threshold. The deviation from straightness is the deviation of the slope of the track segment between two trigger chambers from the slope of a straight line between the interaction point and the hit in a reference layer called the pivot plane, which is the second layer in the barrel (RPC2) and the last layer in the end-cap (TGC3), as illustrated in Fig. 146. For the low (high)- p_T trigger in the barrel, for example, the slope between RPC2 and RPC1 (RPC3) is compared to the slope between the interaction point and RPC2.

To assure full acceptance down to the low-momentum limit, the trigger chambers have regions of overlap with adjacent chambers and between the barrel and end-cap regions. As this may cause double counting of tracks leading to fake two-muon triggers, algorithms are in place to treat these




Figure 146. Schematics of the muon trigger system. RPC2 and TGC3 are the reference (pivot) planes for barrel and end-cap, respectively.




Figure 147. Cross-section through the upper part of the barrel with the RPC's marked in colour. In the middle chamber layer, RPC1 and RPC2 are below and above their respective MDT partner. In the outer layer, the RPC3 is above the MDT in the large and below the MDT in the small sectors. All dimensions are in mm.

5950 overlap regions, either within the barrel or end-cap trigger logic, or as a part of the muon interface to the central trigger processor, see Section 8.2.2.3. An overview of the muon trigger chambers is given in [207, 208] and details about the coincidence system and readout logic are given in [209] and [210].

6.6 Resistive plate chambers

5955 The trigger system in the barrel consists of three concentric cylindrical layers around the beam axis, referred to as the three trigger stations. Fig. 147 shows a standard barrel sector and the location of the RPC's (shaded) relative to the MDT's. The large lever arm between inner and outer RPC's permits the trigger to select high momentum tracks in the range 9–35 GeV (high- p_T trigger), while the two inner chambers provide the low- p_T trigger in the range 6–9 GeV. The chamber count and radial position of the RPC's are given in Table 40. Each station consists of two independent
 5960 detector layers, each measuring η and ϕ . A track going through all three stations thus delivers six measurements in η and ϕ . This redundancy in the track measurement allows the use of a 3-out-of-4 coincidence in both projections for the low- p_T trigger (RPC1 and RPC2 stations) and a 1-out-of-2 OR for the high- p_T trigger (RPC3 station). This coincidence scheme rejects fake tracks from noise hits and greatly improves the trigger efficiency in the presence of small chamber inefficiencies.

5965 The naming scheme of the RPC's is identical to the one in the MDT's, a RPC in a small sector

Name	Small sector					Large sector				
	Units	Chambers	Location in R	To pivot	Maximum z	Units	Chambers	Location in R	To pivot	Maximum z
RPC1	148	84	7820	545	9362	149	94	6800	678	9147
RPC2	148	84	8365		9362	149	94	7478		9660
RPC3	176	92	10229	1864	12847	192	96	9832	2354	12267
Total	472	260				490	284			

Table 40. Segmentation of the RPC system in 544 chambers and 962 units (126 chambers are made of only one unit). Columns 4 and 9 give the radial positions of the RPC stations, while columns 5 and 10 give the distance to the pivot station (RPC2), relevant for the p_T resolution. All dimensions are in mm.

of the middle layer thus being called a BMS. To denote a RPC/MDT pair in the outer layer the term station is used, while for the RPC/MDT/RPC packages in the middle layer the term superstations is used.

6.6.1 Principle of operation

5970 The RPC is a gaseous parallel electrode-plate (i.e. no wire) detector. Two resistive plates, made
of phenolic-melaminic plastic laminate, are kept parallel to each other at a distance of 2 mm by
insulating spacers. The electric field between the plates of about 4.9 kV/mm allows avalanches to
form along the ionising tracks towards the anode. The signal is read out via capacitive coupling
to metallic strips, which are mounted on the outer faces of the resistive plates. The gas used is a
5975 mixture of $C_2H_2F_4$ /Iso- C_4H_{10} / SF_6 (94.7/5/0.3) which combines relatively low operating voltage
(due to the low SF_6 concentration), non-flammability and low cost, while providing a comfortable
plateau for safe avalanche operation. The main parameters and performance of the RPC's are given
in Table 41. A detailed presentation of the RPC system is given in [211].

5980 RPC's can be operated both in avalanche and streamer mode. In the high background envi-
ronment encountered at the LHC, the avalanche mode offers the benefit of higher rate capability
and rate-independent time resolution and has therefore been selected as the operation mode. At the
nominal operating voltage of 9.8 kV, a signal with a width of about 5 ns is generated by the track
with a streamer probability of less than 1%.

5985 The small jitter of the RPC signal in avalanche mode with respect to the passage of the particle
is due to the primary electrons not having to drift to a region of amplification, as is the case in
all types of wire chambers. In the strong and uniform electric field inside a RPC cell, all primary
electron clusters form avalanches simultaneously, producing *one* single signal instantaneously after
the passage of the particle. The charge multiplication in each avalanche continues until its arrival
at the anode plane and, therefore, the gas gain of each avalanche depends on the distance of the
5990 primary cluster from the anode. The total signal charge is thus dominated by the few clusters
produced at the largest distances from the anode. The RPC operating parameters are discussed in
[212, 213, 214].

6.6.2 Mechanical structure

A RPC trigger chamber is made of two rectangular detectors, contiguous to each other, called units.

Parameter	Design value
E-field in gap	4.9 kV/mm
Gas gap	2 mm
Gas mixture	C ₂ H ₂ F ₄ /Iso-C ₄ H ₁₀ /SF ₆ (94.7/5/0.3)
Readout pitch of η and ϕ -strips	23–35 mm
Detection efficiency per layer	$\geq 98.5\%$
Efficiency including spacers and frames	$\geq 97\%$
Intrinsic time jitter	≤ 1.5 ns
Jitter including strip propagation time	≤ 10 ns
Local rate capability	~ 1 kHz/cm ²
Streamer probability	$\leq 1\%$

Table 41. RPC parameters and performance.

5995 Each unit has two independent resistive plate structures, called gas volumes, which are each read out by two orthogonal sets of pick-up strips.

The structure of the gas volumes (i.e. 1/2 unit) is identical for all RPC's (Fig. 148): two resistive plates (plastic laminate, 2 mm) with a volume resistivity of 10^{10} Ωcm , enclose a gas gap of 2 mm, the correct distance being assured by a series of insulating spacers. The outside surface of the resistive plates is coated with a thin layer of graphite paint (100 $\text{k}\Omega/\square$) to assure the HV and ground connection of the resistive electrodes. The graphite electrodes are insulated from the pick-up strips by means of PET films (190 μm), glued to the graphite surfaces. The pick-up strips outside the PET layers are bonded on polystyrene plates (3 mm) and connected to the front-end electronics. The outside surface of the polystyrene plates carries a copper sheet for grounding. A readout signal is induced on the strips by the drift motion of the avalanche electrons. The graphite electrode interposed between the gas gap and the strips does not shield the induction in a significant way due to the graphite electrode's high resistivity and the fast rise-time of the signal.

Each RPC unit is thus made of two detector layers (i.e. gas volumes) and four readout strip panels. The detector layers are interleaved with three support panels made of light-weight paper honeycomb (40 kg/m^3) and are held in position by a solid frame of aluminium profiles. The two external support panels interconnected by the aluminium profiles give the required stiffness to the chamber. The BOL chambers being the largest size ones have a reinforced structure using aluminium plates (2 mm) and aluminium honeycomb. The total thickness of a RPC unit with two gas volumes, support panels and aluminium covers is 96 mm (106 mm for the BOL) and increases to 112 mm (122 mm for the BOL) if the lateral profiles are included. The two units forming a chamber have an overlap region of 65 mm to avoid dead areas for curved tracks. The BMS gas volumes are made of a single piece covering the chamber in full length. All other standard chambers, whose size exceeds the maximum length (3200 mm) of the available plastic laminates




Figure 148. Cross-section through a RPC, where two units are joined to form a chamber. Each unit has two gas volumes supported by spacers (the distance between successive spacers is 100 mm), four resistive electrodes and four readout planes, reading the transverse and longitudinal direction. The sandwich structure (hashed) is made of paper honeycomb. The ϕ -strips are in the plane of the figure and the η -strips are perpendicular to it. Dimensions are given in mm.

6020 have gas volumes divided in two segments with a 9 + 9 mm inefficient region in between due to the edge frames. The readout-strip panels are also segmented in the longitudinal direction, including the case of the BMS, in order to get an homogeneous trigger scheme for all chamber types. This reduces the η -strips time jitter by a factor of two. Detailed information on the assembly procedure is available in Ref. [211]. RPC's have been certified for high rate operation and ageing at the GIF facility at CERN [215, 216].

6025 All standard RPC's, as listed in Table 40, are assembled together with a MDT of equal dimensions in a common mechanical support structure. In order to avoid constraining forces amongst the two chamber types, light-weight kinematical supports were used, holding the chambers with minimum clearance as required by the confined radial space.

6030 A number of small RPC chambers (special RPC's) not paired with MDT's are used around the magnet ribs and in the feet region, where MDT's cannot be installed because of lack of space. RPC's, requiring less space, are used in these regions to keep the trigger acceptance loss to a minimum.

6.6.3 Signal path and readout electronics

6035 A RPC operating in the avalanche mode produces signals of 5 ns full width at half-maximum with a time jitter of 1.5 ns. To preserve this high inherent precision, the pick-up strips must be high-quality transmission lines with low attenuation, terminated at both ends with the characteristic impedance.

The layout of a readout strip plane is shown in Fig. 149. The strips with a width of 25–35 mm consist of 17 μm copper on a 190 μm PET foil glued on a 3 mm plate of rigid polystyrene, which

is covered, on the outside, by 50 μm PET and 17 μm copper as ground reference. The strips are separated by a 2 mm gap with a 0.3 mm ground strip at the centre for improved decoupling. This sandwich structure creates an impedance of about 25 Ω for the strips, slightly depending on the width.

The front-end circuit is a three-stage shaping amplifier followed by a comparator. It is implemented in a eight-channel GaAs die (in order to avoid the encapsulation cost) which is bonded with 20 μm -diameter gold wire on a printed board serving eight strips. The boards are mounted along the edges of the readout panels as shown in Fig 149 and soldered directly to the strips thus avoiding the introduction any further insensitive area in addition to those already contained in the gas volume. While this prevents exchange of faulty amplifiers in the field, the known reliability of GaAs-ASIC's minimises the risk of losing readout channels during the experiment.




Figure 149. Layout of a RPC readout strip plane.

The frequency response of the GaAs-amplifiers has a maximum at 100 MHz and a 60 MHz bandwidth, well adapted to the rise time of the chamber signals. The front-end electronics input is coupled to the strips with a transformer integrated in the printed circuit matched to the signal polarities of η and ϕ -strips, which are opposite to each other. The voltages defining the thresholds of the comparators are supplied by external digital-to-analogue converter units, located in racks in the experimental hall controlled via the DCS. Details of the RPC front-end electronics are given in [217, 218].

The front-end output signals of the ϕ -strips are wired OR-ed such that two physical strips form a logical strip of twice the length but with the same time jitter as a single strip. This is done to avoid unnecessary granularity and to adapt the RPC readout segmentation to the trigger sector segmentation as discussed in Section 8.2.2. The ϕ -strip length and the relative time jitter is on average about half that of an η -strip, giving the ϕ -strips a more precise timing.

In order to be used for the trigger, a signal from a RPC has to be compared with those in the two other RPC's along the path of the particle, i.e. in the same sector and tower, a task which is accomplished by a system of fast coincidence units close to the chambers. Thus, coincidences between strips in RPC1 and RPC2 are used to create the low- p_T trigger. A high- p_T trigger requires hits in all three trigger stations: RPC1 and RPC2 must fulfil the low- p_T condition *and* a confirming hit must be found in RPC3. The inclusion of the low- p_T trigger condition in the high- p_T trigger definition leads to more robustness against fake triggers compared to a simple coincidence between RPC3 and only one inner RPC station. A detailed discussion of the RPC trigger system is given in Section 8.2.2.1.

	EM big wheel			I layer	
	M1 triplet	M2 doublet	M3 doublet	FI (small wheel) doublet	EI doublet
Overall centre in z	13440	14733	15153	6983	7370
Inner radial extent	1900	2526	2610	2170	4700
Outer radial extent	10682	11917	11917	4434	6190

Table 42. Location and radial extension of the TGC wheels. All dimensions are in mm.

6.7 Thin gap chambers

6.7.1 Introduction

Thin Gap Chambers (TGC's) provide two functions in the end-cap muon spectrometer: the muon trigger capability and the determination of the second, azimuthal coordinate to complement the measurement of the MDT's in the bending direction. The middle layer of the MDT's in the end-cap (EM-wheel) is complemented by seven layers of TGC's, while the inner (I) layer is complemented by only two layers. The inner layer is segmented radially into two non-overlapping regions: end-cap (EI) and forward (FI, also known as the small wheel). EI TGC's are mounted on support structures of the barrel toroid coils.

The azimuthal coordinate in the outer MDT wheel (EO) is obtained by the extrapolation of the track from the middle layer, which can be done accurately due to the lack of magnetic field between EM and EO.

The radial, bending coordinate is measured by the TGC wire groups, the azimuthal coordinate by the radial strips. The TGC's need good time resolution to tag the beam-crossing with high efficiency ($\geq 99\%$) and fine granularity to provide a sufficiently sharp cut-off in the momentum of the triggering muon. To match the granularity to the required momentum resolution, the size of the wire groups varies from 6 to 31 as a function of η , corresponding to a variation in width from 10.8 mm to 55.8 mm. The alignment of wire groups in consecutive layers is staggered to optimise the position resolution for a given number of electronics channels. The radial strips are staggered in a similar way to achieve an azimuthal granularity of 2–3 mrad, as seen from the interaction point.




Figure 150. TGC structure showing anode wires, graphite cathodes, G-10 layers and a pick-up strip, orthogonal to the wires.

Figure 168 shows a longitudinal cut through the end-cap. TGC's are located in the innermost layer (marked I) and in the middle layers (EM-wheels marked M1–M3, corresponding to TGC1–3). The location of the MDT in a small (S) and large sector (L) are shown for reference. The location along z and the radial extension of the TGC wheels are given in Table 42.

6115 A detailed, up-to-date description of the TGC chamber system is given in [219] and [210]. A detailed listing of all relevant construction parameters is given in the TGC parameter book [220].

6.7.2 Principle of operation

The main operational parameters of the TGC's are summarised in Table 43.

6120 TGC's are multi-wire proportional chambers with the characteristic that the wire-to-cathode distance of 1.4 mm is smaller than the wire-to-wire distance of 1.8 mm, as shown in Fig. 150. With a highly quenching gas mixture of CO_2 and $n\text{-C}_5\text{H}_{12}$ (n -pentane), this cell geometry allows for operation in a quasi-saturated mode, i.e. with a gas gain of $\sim 3 \times 10^5$. This relatively low gas gain, compared to previous uses does not allow to make full use of its independence from the primary ionisation. Some of its characteristics are still kept, even at such a low gas gain. In particular:

- 6125 • The highly quenching gas prevents the occurrence of streamers in all operating conditions.
- The pulse height observed in the interaction of low energy neutrons (1–10 MeV) is only a factor 30 larger than for a minimum ionising particle.

6130 The high electric field around the TGC wires and the small wire-to-wire distance lead to very good time resolution for the large majority of the tracks. Only tracks at normal incidence passing midway between two wires have much longer drift times due to the vanishing drift field in this region. This effect was already discussed in the context of the CSC's which have a similar cell geometry, see Section 6.4.2. In the TGC wheels, however, the angle of incidence for tracks emerging from the interaction point will always be greater than 10° , thus a part of the track will be outside of the low field region. Including the variation of the propagation time on wires and strips, signals arrive with
6135 99% probability inside a time window of 25 ns.

6.7.3 Mechanical structure

The seven detector layers in the middle layers (big wheels, see Section 9.5) are arranged in one triplet and two doublets. The triplet is to cope with false coincidences from background hits, which are more likely in the end-cap region than in the barrel.

6140 The trigger detectors, forming circular disks, are mounted in two concentric rings, an outer or end-cap one covering the rapidity range $1.05 \leq |\eta| \leq 1.92$ and an inner or forward one covering the rapidity range $1.92 \leq |\eta| \leq 2.4$. Figure 151 shows the cross-section of a TGC triplet and doublet. The structure consists of wire planes (anode), cathode planes, strip planes, shields and honeycomb support structures. The cathode planes consist of 1.6 mm thick FR4 (Flame Resistant 4) plates, graphite coated on the inside, i.e. facing the wires, and with copper cladding on the other side.
6145 The resistivity of the graphite is of $1 \text{ M}\Omega/\square$ in the end-cap and of $0.5 \text{ M}\Omega/\square$ in the inner/forward region. Two of the copper layers in the triplet and doublets are segmented into readout strips to read the azimuthal coordinate (marked "Cu strips" in Fig. 151). The solid copper layers serve as DC grounds, while the segmented ones, being connected to amplifiers, are virtual grounds for

Parameter	Design value
Gas gap	2.8 ± 0.10 mm
Wire pitch	1.8 ± 0.05 mm
Wire diameter	50 μ m
Wire potential	2900 ± 100 V
Operating plateau	200 V
Gas mixture	CO ₂ /n-pentane (55/45)
Gas amplification	3×10^5

Table 43. Main TGC operational parameters.

Modularity	EM big wheel						I layer		Total
	M1 triplet		M2 doublet		M3 doublet		I doublet		
	Inner	Outer	Inner	Outer	Inner	Outer	Inner	Outer	
Modules/sector	2	4	2	4	2	4			
Units/module	1	4	1	5	1	5			
Chambers/unit	3	3	2	2	2	2			
Units/sector	2	16	2	20	2	20			
Units/side	24	192	24	240	24	240	24	21	789
Units/system	48	384	48	480	48	480	48	42	1578
Chambers/sector	6	48	4	40	4	40			
Chambers/side	72	576	48	480	48	480	48	42	1794
Chambers/system	144	1152	96	960	96	960	96	84	3588

Table 44. TGC modularity. Each wheel consists of 12 sectors, each sector containing an inner (forward) and an outer (end-cap) part, having a different azimuthal segmentation. A module covers 15° in azimuth in the inner and 7.5° in the outer part.

6150 the wire signals. On the outside of the triplet and doublet chambers the honeycomb stiffeners are covered by 0.5 mm thick FR4 plates for rigidity and mechanical protection. The adjacent chambers in a module (three in a triplet, two in a doublet) are separated by a 20 mm thick honeycomb paper stiffener.

6155 A gas volume containing a wire plane and two cathodes is called a chamber, while the entirety of three or two chambers in a triplet or doublet arrangement is called a unit. In the outer ring four or five chambers in triplet and doublets, respectively, are mounted in the way of a ladder forming modules. A module covers 7.5° in azimuth in the outer and 15° in the inner ring. Altogether there are 744 units in the big wheels, corresponding to 1704 chambers. The TGC wheel in the innermost layer has a slightly different geometrical structure, containing 45 units and 90 chambers on each




Figure 151. Cross-section of a TGC triplet and doublet module. The triplet has three wire layers but only two strip layers. The dimensions of the gas gaps are enlarged with respect to the other elements.

6160 side. All TGC units are enclosed on their periphery by a gas-tight envelope which is continuously
 flushed by CO₂. This is done to keep a dry atmosphere in the region where the HV elements are
 located as well as to dilute any potential leak of the operating gas (n-pentane). If traces of this
 flammable gas are detected in the CO₂ stream at the output of the chambers, HV and LV as well
 as gas supplies are automatically switched off, and an alarm is activated. A summary of the TGC
 6165 segmentation is given in Table 44.

6.7.4 Signal path, readout, and detector controls

The data flow starting with the primary wire and strip signals is as follows. After amplification in
 the front-end amplifiers, signals are time-aligned and synchronised to the beam-crossing frequency.
 The subsequent signal processing makes use of the redundancy of the track measurement in the
 6170 TGC layers to mitigate the effects of chamber inefficiencies and improve rejection of fake triggers
 caused by background or noise hits. Thus the three signals of the triplet are combined to form a 2-
 out-of-3 coincidence, while the four signals of the two EM-doublets form a 3-out-of-4 coincidence.
 Track segments which have passed this early filter are forwarded to the coincidence matrix, which
 performs a fast search for signal coincidences corresponding to muon tracks in a certain momentum
 6175 range. A detailed description of the algorithm used and of the signal processing in the higher levels
 of the trigger system is given in Chapter 8. The TGC readout is discussed in more detail in [210].

The detector control system (DCS) is fully integrated into the trigger electronics. For the TGC,
 its role is not only of monitoring operating parameters like temperatures, voltages and currents as
 well as the relative alignment of the detectors (down to 0.2 mm), but it can also initialise its various
 6180 elements. Furthermore, by combination with a special trigger signal of the patch-panel board, it
 can monitor the charge left by an ionising particle in every chamber, and therefore act as a gas
 monitor for every single detector.

6.8 Commonalities in the muon system

6.8.1 The gas supplies in the muon system

6185 While gas mixtures and operating conditions in the sub-systems of the muon spectrometer are quite different, being adapted to the requirements of the four chamber types, there are also many commonalities amongst them from an operational point of view.

- 6190 • The gases must be circulated in a closed loop system (mainly for cost reasons), except for the CSC's. This requires a complex recuperation and re-mixing system, using specialised purifiers to clean the gas before recirculation into the supply system. For safety and space reasons, all recuperation systems are housed on the surface, requiring tight pressure and flow control for the transfer to the underground experimental area.
- 6195 • Distribution of the gas in the experimental cavern needs a high degree of modularity to be immuned to leaks in any given supply line. Because of the considerable hydrostatic pressure variations along the height of the hall (25 m), pressure regulation has to be segmented accordingly. Therefore, gas distributors (racks) are located at three different levels in the galleries and wheels of the detector each one supplying chambers in a ± 3 m range of relative elevation.
- 6200 • The flow rate in the gas distribution branches must be supervised by a correspondingly large number of flow and pressure meters for control of gas circulation (purity), operating pressure and leak detection. Precise pressure control is most important for CSC, RPC and TGC chambers, where operation is defined at ~ 2 mb above the outside pressure at the corresponding height and where a significant overpressure might lead to the destruction of the chamber. While this is not the case for the MDT's, they also need precise pressure control for reasons
6205 of drift time and gain stability.

Besides these commonalities there are many differences due to the particularities of the gases used in each system. The high flammability of the TGC gas mixture, for example, requires stringent safety precautions, the outside of the TGC's being constantly flushed with CO₂ in a separate, closed gas circuit. The return gas is monitored for traces on n-pentane, which could point to leaks in the
6210 primary distribution system.

The main characteristics of the four gas supply systems are summarised in Table 45 and in more detail in Table 56.

6.8.2 Electronics services and power consumption

The segmentation of the electrical services of the muon system follows the structure of the detector
6215 which is characterised by a high level of modularity, consisting of electrically independent chambers. The service architecture aims to match this modularity in order to limit data loss in case a given supply channel should fail. Power supply channels are kept floating i.e. the ground outputs of the power supplies are not connected amongst each other nor to the power supply housing or to any local safety ground. This way, the return current of a given chamber has only one path to return to
6220 its power supply. Ground loops (multiple ground returns) are therefore not possible. Ground loops, if present, might create noise coupling between near-by chambers and would expose the readout to

Characteristic	MDT	CSC	RPC	TGC
Gas mixture	Ar/CO ₂ /H ₂ O	Ar/CO ₂	C ₂ H ₂ F ₄ /Iso-C ₄ H ₁₀ /SF ₆	CO ₂ /n-pentane
Composition	93/7/(≤ 1000 ppm)	80/20	94.7/5/0.3	55/45
Gas gain	2×10^4	6×10^4	10^7	3×10^5
Flammability	–	–	low	high
Operating pressure (bar)	3	1	1	1
Volume at operating pressure (m ³)	710	0.5	14	16
Gas exchanges/day	1	10	20	4
Flow at operating pressure (m ³ /h)	30	0.2	12	2.7

Table 45. Main characteristics of the muon gas supply systems.

noise pick-up from sources in the ATLAS electrical environment. To further prevent pick-up, all low-voltage (LV) supply cables are shielded by a copper mesh enclosing the cable wires. The cable shield is connected only at one end (chamber side) to the return wire.

6225 For the definition of the DC level and for safety reasons all chambers are connected from one point of the chamber to the ATLAS structure ground. This star grounding does not allow currents to circulate, i.e. does not create ground loops.

6230 For cost reasons, the maximum supply modularity could not be maintained. Thus, every two MDT chambers are supplied from one LV supply channel. The segmentation of the MDT high-voltage (HV) supplies is finer, with each multi-layer of three (four) tube layers being supplied by a separate channel. Close to the chamber, a passive HV-split is installed (splitter box) to supply each tube layer with an independent cable. In case of shorted tubes (broken wire), the corresponding tube layer can be disconnected, while the other layers would continue to work. Under normal LHC conditions and at an operating voltage of 3080 V, typical chamber currents are around 100 μ A, 6235 yielding a negligible heat dissipation of about 0.3 W/chamber. The total LV power consumption of the MDT system is about 38.5 kW. Only one 4.5 V supply line is needed for each MDT chamber with an average consumption of about 25 W.

6240 The CSC LV system supplies one voltage of 9 V. The operating voltages of 3.3 V and 5 V are derived with radiation tolerant regulators close to the chamber. Each of the 16 chambers of an end-cap consumes about 40 A, leading to a total consumption of about 11.5 kW for both end-caps. The power dissipated in the readout electronics, is evacuated by a water cooling system. The RPC's have separate LV supplies for the analogue front-end and the digital (trigger) electronics, operating at -5 V and 3.3 V, respectively. A third voltage of -2 V is needed to supply the wired-ORs (pull down voltage), combining two aligned physical wires into a "logical" wire. A number of reference 6245 voltages are supplied to each chamber to define the discriminator thresholds for the ASD boards. The total power consumption of the RPC system is about 42 kW. The HV for the RPC's (9.8 kV) requires one channel for each gas gap. To limit the required number of primary supply channels, splitter boxes are used, which are able to safely handle voltages up to 15 kV.

6250 All power supply channels for the MDT barrel, the CSC's, RPC's and TGC's in the small (EI) wheel are housed in racks in the UX15 experimental hall and have been certified for tolerance to

	Voltage at power supply	Voltage on circuit	Supply channels	Average current per channel	Average power per channel	Total power (both sides)	
	V	V	Number	A	W	consumed	supplied
						kW	kW
MDT							
- mezzanine cards & CSM	4.5	3.3	1140	7.5	25	28.2	38.5
CSC (ASM-I & ASM-II)							
- analogue front-end	9	3.3	32	5	17	0.5	1.4
- digital read-out	9	5	32	35	175	5.6	10.1
RPC							
- analogue front-end	-6.7	-5	528	4	20	10.6	14.2
- wired-OR (pull-down)	-2.3	-2	408	0.2	0.4	0.2	0.2
- splitters & PAD boards	4.3	3.3	400	16	53	21.1	27.5
TGC EM wheel							
- analogue front-end	3.3	3	24	217	651	15.6	17.2
- analogue front-end	-3.3	-3	24	56	168	4.0	4.4
- readout (PS-boards)	3.4	3.3	24	55	182	4.4	4.5
- readout (HS-crates)	3.3	3.3	24	34	112	2.7	2.7
TGC EI wheel							
- analogue front-end	3.5	3	2	80	240	0.5	0.6
- analogue front-end	-3.5	-3	2	18	54	0.1	0.1
- readout (PS-boards)	3.3	3.3	2	25	83	0.2	0.2
Total power muon spectrometer						94	122

Table 46. Power consumption of the four sub-systems of the muon spectrometer. Voltages delivered by the power supplies (column 2) are approximate, depending on cable length and conductor cross-section. The differences between power *supplied* and *consumed* are due to cable losses as well as losses in voltage regulators, used to stabilise voltages on the board. The aim of the cabling layout was to limit cable losses to about 15–20% of the power used by the circuits. This could not be achieved for the CSC’s due to limitations on cross-sections for supply cables. The PS-boards and HS-crates are components of the TGC digital readout (see Section 8.2.2).

radiation and magnetic field. The TGC chambers and the MDT in the big wheel (EM wheel) are supplied from racks directly mounted on the wheel structure (non-standard “mini-racks”, due to space limitations at the circumference of the EM wheel) in order to limit cable lengths and power loss on the LV cables.

6255 The TGC’s, like the RPC’s, have separate LV supplies for analogue and digital readout electronics. The segmentation of the supplies in the EM wheel follows that of the TGC’s, with 12 sectors per side. For the small wheel there is only one supply channel per side for each voltage. The HV modularity provides one supply channel per chamber with the exception of the 144 inner chambers in the M1 triplet which have two supply channels per chamber. A total of 3732 HV channels
6260 are therefore used to serve the 3588 chambers of the TGC system.

Services for the small wheel, containing MDT’s, CSC’s and TGC’s, are routed through a flex-

ible support structure (cable chains). When the small wheel is moved along the z -direction to give access to the calorimeters and inner detector, most service channels will not require interruption, reducing access time and minimising the risk of damage to the small wheel chamber operation.

6265 An overview of the segmentation and power consumption of the four muon sub-systems is given in Table 46.

7. Forward detectors

In addition to the main ATLAS detector systems described in the previous chapters, three smaller sets of detectors are being built to provide good coverage in the very forward region [221].

6270 Ordered according to their distance from the interaction point as shown in Fig. 152, the first system is a Cerenkov detector called LUCID (LUminosity measurement using Cerenkov Integrating Detector, see Section 7.1). LUCID is the main relative luminosity monitor in ATLAS and is located at a distance of ± 17 m from the interaction point, near the TAS (Target Absorber Secondaries) collimator. The second system is the Zero-Degree Calorimeter (ZDC), located at a distance of ± 140 m from the interaction point (see Section 7.3). This corresponds to the location where the LHC beam-pipe is divided into two separate pipes. The ZDC is embedded in the TAN (Target Absorber Neutral), located between the beam-pipes just after the split. The ZDC's primary purpose is to detect forward neutrons in heavy-ion collisions. The most remote detector is the absolute luminosity detector ALFA (Absolute Luminosity For ATLAS, see Section 7.2). ALFA consists of scintillating-fibre trackers located inside Roman pots at a distance of approximately ± 240 m from the interaction point. Additional proton-tagging detectors are being considered at an even greater distance of about ± 420 m [222] from the interaction point. These detectors, together with new radiation hard-detectors at ± 220 m [223] are part of a possible upgrade programme.




Figure 152. Placement of the forward detectors along the beam-line around the ATLAS interaction point (IP). See text for details.

7.1 The LUCID detector

6285 Luminosity can be measured by several detectors in ATLAS, but LUCID (LUminosity measurement using Cerenkov Integrating Detector) [221] will be the only detector which is primarily dedicated to online luminosity monitoring. LUCID is a relative luminosity detector. Its main purpose

is to detect inelastic p - p scattering in the forward direction, in order to both measure the integrated luminosity and to provide online monitoring of the instantaneous luminosity and beam conditions. Potentially, LUCID could also be used for diffractive studies, for example as a rapidity-gap veto or as a tag for a diffractive signal [224].

During the initial period of LHC operation, the absolute calibration procedure for LUCID will use calculations based on LHC machine parameters, initially determining luminosity to a precision of $\sim 20 - 30\%$. Physics processes can also be used, such as W and Z counting [225] and exclusive muon pair production via two-photon interactions [226]. The most accurate calibration will be obtained with the ALFA detectors which measure elastic Coulomb scattering. The goal is thus to measure the luminosity with an uncertainty of better than 5%.

At the LHC design luminosity, most bunch-crossings will result in multiple p - p interactions. The number of interactions per beam-crossing must be known in order to determine luminosity. LUCID is based on the principle that the number of interactions in a bunch-crossing is proportional to the number of particles detected in this detector. This holds true even when most of the detected particles originate from secondary interactions.

The main requirements for the LUCID detector are:

- good acceptance for minimum-bias events;
- sufficient time resolution to measure individual bunch-crossings which occur every 25 ns;
- resistance to very high radiation levels;
- the ability to count individual charged particles;
- pointing capability in order to suppress signals from tracks which do not point to the interaction point.

The concept of a luminosity detector consisting of an array of Cerenkov tubes (cones) was developed by the CDF collaboration [227]. The main challenges in using this type of detector at the LHC are the much higher radiation and background levels. ATLAS has approved the construction of a detector with a reduced number of readout channels for the initial low-luminosity phase of the LHC operation (up to few $10^{33} \text{ cm}^{-2} \text{ s}^{-1}$), with a view to then possibly upgrade the detector to a larger coverage if it can be demonstrated through further studies that the detector would work well at the LHC design luminosity.

Since there is no Landau tail in the measurement of Cerenkov light, it is in principle possible to determine the number of particles which have entered a Cerenkov tube by using pulse-height measurements (particle-counting mode). The probability that several particles enter the same tube is quite low for the LUCID design, typically a few percent at a luminosity of $\sim 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$. Simulations show that a simpler measurement consisting in counting only the number of tubes with a signal above a preset threshold (hit-counting mode) provides a measurement for online monitoring with an acceptable systematic error. These measurements will be done with different scalers for each bunch in the LHC, so that the luminosity of the individual bunches can be monitored online. A pre-scaled LUCID trigger will make it possible to record a sub-set of the LUCID events for offline analysis which may be done in particle-counting mode and thereby obtaining a more precise luminosity determination using charge and time-to-digital converters.

7.1.1 Detector design

The LUCID detector is depicted in Fig. 153. It consists of twenty aluminium tubes which surround the beam-pipe and point toward the interaction point. The 1.5 m long mechanically polished tubes with a diameter of 15 mm are placed in a light-weight aluminium gas vessel which ensures that the tubes are filled with C₄F₁₀ at a constant pressure of 1.2–1.4 bar, providing a Cerenkov threshold of 2.8 GeV for pions and 10 MeV for electrons. There are two detectors installed, one in each end-cap region of ATLAS, at a distance of approximately ± 17 m from the interaction point, and placed at a radial distance of approximately 10 cm from the beam-line ($|\eta| \approx 5.8$).

The Cerenkov light emitted by a particle traversing the tube has a half-angle of 3° and is reflected on average three times before the light is measured by photomultiplier tubes (PMT's) which match the size of the Cerenkov tubes. The signal amplitude from these PMT's can be used to distinguish the number of particles per tube, and the fast timing response provides unambiguous measurements of individual bunch-crossings. LUCID is situated in a high radiation area and, although the PMT's have quartz windows and are radiation hard, their performance is expected to degrade after several years of operation at nominal LHC luminosity. It is therefore foreseen to upgrade the detector after some years of operation. The Cerenkov light from the upgraded detector could be transmitted through radiation-hard quartz fibres to a region outside the forward shielding (see Fig. 19) with a low level of background radiation, in which the fibres would be read out with multi-anode PMT's. In this readout scheme, a Winston cone at the end of each Cerenkov tube focuses the light onto a bundle of 37 fused silica fibres with 0.8 mm diameter. In order to gain experience with this type of readout, four of the twenty Cerenkov tubes are read out by fibres in this manner. The initial LUCID detector will not have a re-circulating gas system as it is not expected to be needed for the low-luminosity phase of the LHC. However, such a system could be envisaged for high-luminosity operations, should it prove to be necessary.

The 16 electrical signals from the PMT readout are fed into a custom-designed front-end card containing a fast amplifier and a differential line-driver. These analogue signals pass through a constant-fraction discriminator, which registers a hit each time the PMT pulse height is above a preset threshold. A charge-to-digital converter allows offline analysis of the signal amplitudes in a sub-sample of the triggered events. One of the digital outputs of the front-end card is fed into a multi-hit time-to-digital converter for offline timing analysis, with a second output going to a custom-designed 9U VME readout card. The readout card houses FPGA's used to calculate the luminosity for each bunch-crossing using different algorithms and stores one scaler value per bunch-crossing per luminosity algorithm. It is thus at the core of the LUCID online processing: after receiving a L1 trigger accept signal, it sends the pattern of hits to the overall readout system. It also provides an independent LUCID trigger. The LUCID readout buffer is sent to the trigger and data acquisition system via the S-link interface for inclusion in the overall event record.

7.1.2 Experimental results from test-beam and radiation tests

A six-tube prototype detector, dedicated purely to test-beam measurements, has been validated at DESY. The prototype was used to study the light-collection efficiency of both readout schemes, either with PMT's directly connected to the Cerenkov tubes or with bundles of optical fibres coupled to Winston cones attached to the Cerenkov tubes. The light output was studied as a function of




Figure 153. a) Picture of the two LUCID vessels fully assembled and ready to be installed in ATLAS. b) Sketch of LUCID integrated in the cone supporting the beam-pipe. c) Design of the gas vessel. d) Expanded view of the readout area showing the coupling between the 15 mm diameter Cerenkov tubes and the photomultiplier tubes. Also shown is the coupling between the Cerenkov tubes and the readout fibre bundle through a Winston cone (see text).

angle, track position in the tube, pressure, Cerenkov-tube polish, PMT type, fibre type, number of
6370 fibres, and of the Winston cone design parameters. Fig. 154 (left) shows an example of the number
of photoelectrons collected in the PMT, when the Cerenkov tube was placed in a coaxial electron
beam with an energy of 6 GeV. This test was performed using a Cerenkov tube with a larger diam-
eter than the PMT window resulting in two peaks depending on whether the beam particles missed
the PMT window or went through it. On average, Cerenkov light produced in the gas yielded about
6375 70 photoelectrons and light produced in the PMT window about 50 photoelectrons. These values
are being determined from test-beams for each tube and will be used to set their thresholds. One
of the main objectives of the test-beam measurements was to validate the GEANT4-based Monte
Carlo [228] simulations of the detector. Fig. 154 (right) shows a comparison of a measurement
6380 of the average number of photoelectrons as a function of radiator gas pressure with the simulated
detector response. The agreement observed shows that the light collection is correctly simulated.

The radiation hardness of the detector and its readout is an important issue, although for the
initial operation of the LHC, the total dose per year seen by LUCID will be less than one tenth of
the dose at design luminosity (7 Mrad per year). The body of the detector is made of aluminium




Figure 154. Left: distribution of the number of photoelectrons for an aluminised mylar Cerenkov tube as obtained at the electron test facility at DESY. The peak coming from the Cerenkov light in the radiator gas alone (left peak) and the one produced by the radiator gas and the quartz window of the readout PMT (right peak) are clearly visible. The curves superimposed on the experimental data represent the results of the fit with two gaussian distributions for the signal plus a linearly decreasing shape for the background. Right: comparison of the photoelectron yield from the Cerenkov light in the radiator gas as a function of the radiator gas pressure for the test-beam data (red circles) and Monte Carlo data (blue squares).

and is extremely radiation hard. However damage to other components may be a concern (for instance, the PMT's, the optical fibres, the signal cables, the power cables and certain seals and bonds in the detector which were made using epoxy). Radiation testing of these components has been performed using both a strong gamma source and a reactor. So far, no serious problems have been observed in any of the materials tested.

7.1.3 LUCID installation in ATLAS

The LUCID detector is fully assembled and mechanically tested for gas leaks and light tightness. The front-end and readout scheme has been validated in the beam test at DESY. The calibration using LED light injected into each Cerenkov tube has been successfully tested in the fully assembled detectors. The detector is ready to be installed together with the beam-pipe in 2008.

7.2 The ALFA detector

Traditionally, the absolute luminosity at hadron colliders has been determined via elastic scattering at small angles. This is also the approach pursued by ATLAS with the ALFA (Absolute Luminosity For ATLAS) [221] detector. The optical theorem connects the elastic-scattering amplitude in the forward direction to the total cross-section and can be used to extract luminosity. In addition, other interesting physics topics may be explored with such scattering data [229].

The extremely small scattering angles ($3 \mu\text{rad}$) needed to make these measurements are smaller than the nominal beam divergence. These measurements can therefore only be performed with specially prepared beam conditions, which are also used for calibration of the LUCID detector. High-beta (β^*) optics in combination with reduced beam emittance are required. Moreover the detectors




Figure 155. Schematic view of the support mechanics for one of the ALFA detectors and of its location at a distance of ± 240 m from the ATLAS interaction point (left). One of the as-built structures, which will house the scintillating-fibre trackers (right).

have to be placed far away from the interaction point and as close as possible to the beam. The
 6405 Roman-pot technique has been successfully used in the past [230] for measurements very close to
 circulating beams and has therefore been adopted for ATLAS. The ATLAS Roman pots have been
 designed to move the detectors as close as 1 mm to the beam, but only from above and below, due
 to the mechanical constraints imposed by the two horizontal beam-pipes of the LHC, as shown in
 Fig. 155. The Roman pots will be located ± 240 m away from the interaction point, and on each
 6410 side there will be two Roman-pot stations separated by four metres.

The main requirements on the detectors which will be housed in the Roman pots are: a spatial
 resolution of about $30 \mu\text{m}$, no significant inactive edge region, minimal sensitivity to the radio-
 frequency noise from the LHC beams and compatibility with the vacuum in the Roman pots.

Given these requirements, ATLAS has chosen a scintillating-fibre tracker. The high- β^* op-
 6415 eration of the LHC corresponds to a very low instantaneous luminosity, and for this reason no
 radiation-hard technology is required for this specific set of detectors. Prototypes of the ALFA
 detectors have been validated in beam tests at DESY [231] and CERN [232]. In order to fulfil the
 spatial resolution requirements, the detector is built of ten double-sided modules, each with 64 fi-
 bres arranged in stereo $u - v$ geometry on both sides. The square fibres have a width of 0.5 mm,
 6420 resulting in an effective detection area of $32 \times 32 \text{ mm}^2$. The modules are staggered in depth by
 multiples of a tenth of the effective fibre pitch, i.e. $70 \mu\text{m}$. The fibres are aligned and glued on a
 precisely machined support structure made of titanium. Before assembly, the fibres are aluminised




Figure 156. Schematic layout of the ALFA detector in the Roman pot, showing the scintillating-fibre stack, the fibre connectors, the multi-anode photomultipliers, and the front-end boards. The Roman pots (labelled as upper and lower) approach the beam-line from above and below (left). Details of the scintillating-fibre stack (right).

to reduce light losses and optical cross-talk. The fibre positions are measured by means of optical metrology at various stages of the manufacturing, resulting in a set of equations describing the location of each fibre in a given detector.

6425

The overall layout with the front-end electronics is shown in Fig. 156. The module support structure is connected to a base-plate, which is mounted on the Roman-pot flange. The base-plates also house optical feed-through connectors coupling the fibres to the pixels of multi-anode phototubes. The multi-anode phototubes with 64 channels are aligned to the fibres and shielded against residual stray magnetic fields. Anode signals from these phototubes are fed into the front-end boards and further processed by the front-end multi-anode readout chips.

6430

To achieve an accurate alignment of the fibre tracker with respect to the centre of the beam, the Roman pots are equipped with special interleaved extrusions, yielding a vertical overlap area when in beam position. These extrusions are instrumented with three layers of vertically staggered overlap detectors consisting of 30 horizontally oriented fibres in each layer. The overlap detectors use the same type of fibres and are read out in the same way as the other detectors.

6435

The central crossing area is further instrumented with two optically separated trigger scintillator tiles providing a fast L1 trigger signal. Single trigger tiles are used in each of the overlap




Figure 157. The spatial resolution as function of the beam energy compared to GEANT4 simulations with and without the contributions from multiple scattering.




Figure 158. Comparison between the edges of ALFA and a reference trigger counter of the y coordinates measured with the silicon detector. Dotted lines represent edge fits smeared by detector resolution for 6 GeV electrons.

extrusions. The trigger tiles are coupled to clear light-guide fibres routed to standard PMT's with
 6440 8 mm diameter. The PMT signals are amplified and processed before being sent to the central
 trigger processor (see Section 8.2.3).

7.2.1 Experimental results from test-beam

Two test-beam campaigns were carried out for the ALFA detectors; the first one at DESY [231]
 to validate the detector design, and the second one at CERN [232] to test the overlap, trigger,
 6445 and electronics concepts. Prototypes with ten modules but only a fraction of nominal fibres per
 module were built and exposed to the beam. The test-beam facilities were partly equipped with
 high-precision silicon-strip detectors for precise resolution studies. The photoelectric yield of the
 scintillating fibres was extracted from charge-amplitude measurements of the multi-anode photo-
 tubes. The charge response of the multi-anode phototube to single photons was calibrated using a
 6450 pulsed blue LED. The average number of photoelectrons was 4.2 ± 0.5 , with a relative cross-talk
 contribution from adjacent fibres of about 3%.

The spatial detector resolution, obtained after deconvolution of the intrinsic contribution of the
 silicon reference detector of $30 \mu\text{m}$, was determined to be $36 \mu\text{m}$ in both the x and y directions.
 For the low-energy electron beam used at DESY, a sizable resolution degradation is expected from
 6455 multiple scattering, as shown in an energy scan in Fig. 157. Using a high-energy hadron beam at
 CERN, where multiple scattering is irrelevant, a spatial resolution for ALFA of $25 \pm 3 \mu\text{m}$ was
 determined, in agreement with GEANT4 [228] Monte-Carlo expectations.

The sensitivity of ALFA at the edge closest to the beam, which is essential to cover the re-
 quired phase-space for the measurements, was studied with an additional scintillator touching the
 6460 fibres and defining the physical edge. Coordinates of tracks passing either through ALFA or the
 scintillator were recorded with the silicon detector and fit by edge functions, smeared for spatial

resolution effects. As shown in Fig. 158, the edge functions cross very close to the physical ALFA edge and no insensitive region was found. The fitted smearing is consistent with the spatial resolution of the silicon detector. All measured performance parameters were thus found to fulfil the
6465 ATLAS requirements [221].

7.2.2 ALFA installation in ATLAS

The Roman pots (without detectors) with their precision support mechanics will be produced at the beginning of 2008 and the installation will take place when the LHC schedule will allow it. The ALFA detector itself is still at the prototype stage. A prototype, representing a complete
6470 detector for one Roman pot (1/8 of the full system) is under construction. After validation, the series production for the entire system will begin in 2008, aiming for final installation and first measurements in 2009.

7.3 The zero-degree calorimeters

The primary purpose of the ATLAS Zero-Degree Calorimeters (ZDC) [221] is to detect forward
6475 neutrons with $|\eta| > 8.3$ in heavy-ion collisions. The ZDC's play a key role in determining the centrality of such collisions, which is strongly correlated to the number of very forward (spectator) neutrons. During the start-up phase of the LHC (p - p collisions with luminosities well below $10^{33} \text{ cm}^{-2} \text{ s}^{-1}$), the ZDC's will enhance the acceptance of ATLAS central and forward detectors for diffractive processes and provide an additional minimum-bias trigger for ATLAS.

6480 Significant backgrounds in hadron-collider experiments are created by beam-gas and beam-halo effects. These can be greatly reduced by requiring a tight coincidence from the two arms of the ZDC's, located symmetrically about the interaction point. For p - p scattering, this coincidence, as determined from PYTHIA [233] simulations, is found to be 9% of the total inelastic rate and will therefore be important for initial beam tuning at the beginning of LHC operation. The time
6485 resolution of the ZDC's is roughly 100 ps, which is sufficient to locate the interaction point to about 3 cm in z (along the beam axis). This would therefore provide some knowledge of the vertex location without using the inner detector information.

The ZDC's reside in a slot in the TAN (Target Absorber Neutral) absorber, which would otherwise contain inert copper bars as shielding. The TAN is located at ± 140 m from the interaction
6490 point, at the place where the straight-section of the beam-pipe is divided back into two independent beam-pipes.

7.3.1 ZDC module description

There will be four ZDC modules installed per arm: one electromagnetic (EM) module (about 29 radiation lengths, X_0 , deep) and 3 hadronic modules (each about 1.14 interaction lengths, λ , deep).

6495 Table 47 summarises the mechanical properties of the ZDC system. The EM module of one arm is shown in Fig.159a). It consists of 11 tungsten plates, with their faces perpendicular to the beam direction. The height of these plates is extended in the vertical direction with 290 mm long steel plates. A total of 96 quartz rods of 1.0 mm diameter penetrate the tungsten plates parallel to the beam in a 8×12 matrix, which is transverse to the beam. Along the face of the module the rods are bent upwards at an angle of 90° and are viewed by multi-anode phototubes capturing the Cerenkov
6500

light from shower products of incident particles. The position of the rods containing a signal thus corresponds to the position of the incident particle transverse to the beam, and the intensity of the light reflects the energy of the particle. Only one of the arms, at -140 m, is equipped with the position-sensing capability for its EM module.

6505 Between the tungsten and steel plates are placed 1.5 mm quartz strips which are routed vertically and are viewed from above by photomultiplier tubes via air light-guides. These strips are actually rows of quartz rods called strips, in order to distinguish them from the position-measuring rods. These strips are depicted in Fig.159b). The purpose of the strips is to obtain a second, improved measurement of the incident particle energy. In the position-sensing EM module there are
6510 four such light-guides, side by side, since more than one incident particle per event is expected.

The hadronic modules are similar, but while the EM module maps each of the 96 position-measuring rods onto one pixel of the multi-anode phototube, the hadronic modules map clusters of four rods into individual phototubes and have only one light-guide. Not all hadronic modules have position-sensing rods (there is only one position-sensing module per arm).

6515 **7.3.2 Calibration and monitoring**

7.3.2.1 Flashers Measurements over long periods of time and limited access to the tunnel require that the PMT gains be remotely monitored to an accuracy of about 1%. This can be achieved with a system based on blue LED's. A controllable generator located in the underground side-cavern USA15 will excite the LED in the tunnel. Each PMT will view a fibre connected to the LED, and
6520 one fibre will illuminate each large single channel PMT.

The LED is mounted opposite a pin-diode which tracks its stability. Around the pin-diode are eight primary optical fibres. Light in these fibres will travel 5 m to the modules. The light from three of these primary fibres is split further into eight secondary fibres. Light from each of these secondary fibres impinges on each of the 24 phototubes in the hadronic position-sensing module.
6525 Four more of the eight primary fibres will calibrate the phototubes in the hadronic position-sensing and energy-measuring modules. The eighth primary fibre is a spare.

The LED monitoring system is not subjected to the large doses of radiation seen in the ZDC's and so deterioration of the system due to radiation damage is not expected.

7.3.2.2 Particles Full calibration of the ZDC detector begins with the determination of the gains
6530 in all rod/phototube chains, and continues with the measurement of shower-shape distribution functions, such as the dependence of single-rod amplitudes on the distance to the shower centre. These parameters will be permanently monitored during operation. This preliminary calibration is followed by one using π^0 decays in situ in p - p collisions. During heavy-ion operations, peripheral ion collisions will yield samples of mono-energetic neutrons with the beam energy of 2.75 TeV,
6535 providing a useful redundancy check of the calibration constants.

7.3.3 ZDC installation in ATLAS

The ZDC is being installed in two stages because it shares the TAN location with the LHCf detector [234]. The LHCf experiment will run during early 43-bunch operation of the LHC and will then be removed. Since the LHCf detector occupies the space of the first ZDC module in each
6540 arm, only the three downstream modules in each arm will be installed during this phase. The full




Figure 159. a) The electromagnetic ZDC module. The incident particles impinge on tungsten plates at the bottom of the module and produces showers of particles. The quartz rods pick up the Cerenkov light generated by the shower and transmit it to multi-anode phototubes at the top of the module. The phototubes measure light from the quartz strips through four air light-guides. b) Details showing the placement of the quartz strips.

ZDC will be in place for the luminosity ramp-up of the machine beyond 2008. As soon as the LHC reaches luminosities above $10^{33} \text{ cm}^{-2} \text{ s}^{-1}$, the ZDC modules will be removed for p - p operation and reinstalled only for the heavy-ion periods. This will be necessary in order to minimise radiation damage to the quartz radiator in the ZDC.

	Electromagnetic (at -140 m)	Electromagnetic (at +140 m)	Hadronic (per arm)
Number of modules	1	1	3
Number of position-sensing modules	1	none	1
Number of energy-measuring light-guides	4	1	3
Total depth	29 X ₀	29 X ₀	4.6 λ (includes EM module)
Module size (mm):			
width	93.8	93.8	93.8
height	738	738	738
depth	150	150	150
Number of plates/module	11	11	11
Tungsten plate size (mm):			
width	91.4	91.4	91.4
height	180	180	180
depth	10	10	10
Module weight (kg)	80	80	80
Number of 1.5 mm quartz strips	648	648	648
Number of 1 mm pixel rods	96	none	96
Number of pixel readout channels	96	none	24

Table 47. Summary of ZDC mechanical properties.

6545 8. Trigger, data acquisition, and controls

8.1 Introduction to event selection and data acquisition

As described in Chapter 1, the trigger consists of three levels of event selection: Level-1 (L1), Level-2 (L2), and event filter. The L2 and event filter together form the High-Level Trigger (HLT). The L1 trigger is implemented using custom-made electronics, while the HLT is almost entirely
6550 based on commercially available computers and networking hardware. A block diagram of the trigger and data acquisition systems is shown in Fig. 160.

The L1 trigger searches for signatures from high- p_T muons, electrons/photons, jets, and τ -leptons decaying into hadrons. It also selects events with large missing transverse energy (E_T^{miss}) and large total transverse energy. The L1 trigger uses reduced-granularity information from a
6555 subset of detectors: the Resistive Plate Chambers (RPC) and Thin-Gap Chambers (TGC) for high- p_T muons, and all the calorimeter sub-systems for electromagnetic clusters, jets, τ -leptons, E_T^{miss} , and large total transverse energy. The maximum L1 accept rate which the detector readout systems can handle is 75 kHz (upgradeable to 100 kHz), and the L1 decision must reach the front-end electronics within $2.5 \mu\text{s}$ after the bunch-crossing with which it is associated.

The L2 trigger is seeded by Regions-of-Interest (RoI's). These are regions of the detector
6560 where the L1 trigger has identified possible trigger objects within the event. The L2 trigger uses RoI information on coordinates, energy, and type of signatures to limit the amount of data which must be transferred from the detector readout. The L2 trigger reduces the event rate to below 3.5 kHz, with an average event processing time of approximately 40 ms.

The event filter uses offline analysis procedures on fully-built events to further select events
6565 down to a rate which can be recorded for subsequent offline analysis. It reduces the event rate to approximately 200 Hz, with an average event processing time of order four seconds.

The HLT algorithms use the full granularity and precision of calorimeter and muon chamber data, as well as the data from the inner detector, to refine the trigger selections. Better information
6570 on energy deposition improves the threshold cuts, while track reconstruction in the inner detector significantly enhances the particle identification (for example distinguishing between electrons and photons). The event selection at both L1 and L2 primarily uses inclusive criteria, for example high- E_T objects above defined thresholds. One exception is the L2 selection of events containing the decay of a B -hadron, which requires the reconstruction of exclusive decays into particles with low
6575 momentum.

The data acquisition system (DAQ) receives and buffers the event data from the detector-specific readout electronics at the L1 trigger rate. The data transmission is performed over point-to-point Readout Links (ROL's). It transmits to the L2 trigger any data requested by the trigger
(typically the data corresponding to RoI's) and, for those events fulfilling the L2 selection criteria,
6580 event-building is performed. The assembled events are then moved by the data acquisition system to the event filter, and the events selected there are moved to permanent event storage.

In addition to controlling movement of data down the trigger selection chain, the data acquisition system also provides for the configuration, control and monitoring of the ATLAS detector during data-taking. Supervision of the detector hardware (gas systems, power-supply voltages, etc.)
6585 is provided by the Detector Control System (DCS).




Figure 161. Block diagram of the L1 trigger. The overall L1 accept decision is made by the central trigger processor, taking input from calorimeter and muon trigger results. The paths to the detector front-ends, L2 trigger, and data acquisition system are shown from left to right in red, blue and black, respectively.

rel and TGC's in the end-caps. The trigger searches for patterns of hits consistent with high- p_T muons originating from the interaction region. The logic provides six independently-programmable p_T thresholds. The information for each bunch-crossing used in the L1 trigger decision is the multiplicity of muons for each of the p_T thresholds. Muons are not double-counted across the different thresholds.

The overall L1 accept decision is made by the Central Trigger Processor (CTP), which combines the information for different object types. Trigger menus can be programmed with up to 256 distinct items, each item being a combination of requirements on the input data. The trigger decision, together with the 40.08 MHz clock and other signals, is distributed to the detector front-end and readout systems via the Timing, Trigger and Control (TTC) system, using an optical-broadcast network.

While the L1 trigger decision is based only on the multiplicity of trigger objects (or flags indicating which thresholds were passed, for global quantities), information about the geometric location of trigger objects is retained in the muon and calorimeter trigger processors. Upon the event being accepted by the L1 trigger, this information is sent as RoI's to the L2 trigger (see Section 8.3.6), where it is used to seed the selection performed by the HLT.

An essential function of the L1 trigger is unambiguous identification of the bunch-crossing of interest. The very short (25 ns) bunch-crossing interval makes this a challenging task. In the case

of the muon trigger, the physical size of the muon spectrometer implies times-of-flight exceeding the bunch-crossing interval. For the calorimeter trigger, a serious complication is that the width of the calorimeter signals extends over many (typically four) bunch-crossings.

6625 While the trigger decision is being formed, the information for all detector channels has to be retained in pipeline memories. These memories are contained in custom electronics placed on or near the detector, where often radiation levels are high and access is difficult. In the interest of cost and reliability, it is desirable to keep the pipeline length as short as possible. The L1 latency, which is the time from the proton-proton collision until the L1 trigger decision, must therefore be kept as short as possible. The design of the trigger and front-end systems requires the L1 latency
6630 to be less than $2.5 \mu\text{s}$, with a target latency of $2.0 \mu\text{s}$, leaving $0.5 \mu\text{s}$ contingency. About $1 \mu\text{s}$ of this time is accounted for by cable-propagation delays alone. To achieve this aim, the L1 trigger is implemented as a system of purpose-built hardware processors, which are described in more detail below.

8.2.1 Calorimeter trigger

6635 L1Calo [235] is a pipelined digital system designed to work with about 7000 analogue trigger towers of reduced granularity (0.1×0.1 in $\Delta\eta \times \Delta\phi$ in most parts, but larger at higher $|\eta|$) from the electromagnetic and hadronic calorimeters. It sends the results for each LHC bunch-crossing to the CTP approximately $1.5 \mu\text{s}$ after the event occurs, resulting in a total latency for the L1Calo chain of about $2.1 \mu\text{s}$, well within the allowed envelope.

6640 The L1Calo system is located off-detector in the service cavern USA15. Its architecture, shown in Fig. 162, consists of three main sub-systems. The pre-processor digitises the analogue input signals, then uses a digital filter to associate them with specific bunch-crossings. It uses a look-up table to produce the transverse-energy values used for the trigger algorithms. The data are then transmitted to both the Cluster Processor (CP) and Jet/Energy-sum Processor (JEP) sub-
6645 systems in parallel. The CP sub-system identifies electron/photon and τ -lepton candidates with E_T above the corresponding programmable threshold and satisfying, if required, certain isolation criteria. The JEP receives jet trigger elements, which are 0.2×0.2 sums in $\Delta\eta \times \Delta\phi$, and uses these to identify jets and to produce global sums of scalar and missing transverse energy. Both processors count the multiplicities of the different types of trigger objects. The CP and JEP send these
6650 feature multiplicities, as well as transverse-energy threshold information, to the CTP for every bunch-crossing.

When there is a L1 Accept (L1A) decision from the CTP, the stored data from the L1Calo sub-systems are read out to the data acquisition system: this includes input data, intermediate calculations and trigger results in order to allow full monitoring and verification of the L1 trigger func-
6655 tionality. These data can also provide useful diagnostics for the LHC machine (see Section 9.10) and the ATLAS sub-detectors. The types and positions of jet, τ -lepton and electromagnetic cluster candidates are also collected and sent to the RoI builder (see Section 8.3.6) for use by the L2 trigger.

6660 The L1Calo architecture is relatively compact, with a minimal number of crates and cable links. This helps in reducing the latency. Some of the hardware modules were designed to fulfil several different roles in the system, in order to reduce hardware costs and design efforts, as well




Figure 162. Architecture of the L1 calorimeter trigger. Analogue data from the calorimeters are digitised and associated with the correct bunch-crossing in the pre-processor and then sent to two algorithmic processors, the jet/energy-sum processor and the cluster processor. The resulting hit counts and energy sums are sent to the central trigger processor.

as to reduce the number of spares required.

8.2.1.1 The analogue front-end Analogue signals from trigger towers in all the calorimeters are sent to the USA15 cavern on 16-way twisted-pair cables. These cables are specially routed to minimise their length, and hence the trigger latency; they range from about 30 m to 70 m in length. Liquid-argon electromagnetic calorimeter signals are converted from energy to transverse energy in the tower builder boards located on the detector, but all hadronic calorimeter signals are transmitted proportional to energy. All the trigger-tower signals arrive at 64-channel receiver modules. The main function of the receiver modules is to adjust the gains, in order to provide transverse energy rather than energy for hadronic calorimeter signals, and to compensate for differences in energy calibration and signal attenuation in the long cables. The receivers reshape the signals, and

contain linear, variable-gain amplifiers controlled by DAC's. Receiver outputs are sent as differential signals on short twisted-pair cables to the pre-processor. A further function of the receivers is to monitor a small, programmable selection of analogue input signals.

6675 **8.2.1.2 The pre-processor** The pre-processor consists of eight 9U VMEbus crates. Four crates
process electromagnetic trigger towers and four process hadronic towers. Each crate contains 16
Pre-Processor Modules (PPM's) which each receive four analogue cables on the front panel and
process 64 analogue input signals. The granularity of these signals is reduced compared to the
full calorimeter data. This is done by analogue summing at the detector of variable numbers of
6680 calorimeter cells, ranging from a few up to 60. The main signal processing is performed by 16
multi-chip modules, each of which processes four trigger towers. Ten-bit Flash ADC's (FADC's)
digitise the signals with a sampling frequency of 40.08 MHz. Fine adjustment of the timing of each
digitisation strobe is performed by a four-channel ASIC, which provides programmable delays in
steps of 1 ns across the 25 ns LHC clock period. The digitised values are then sent to a custom
6685 pre-processor ASIC.

The pre-processor ASIC synchronises the timing of the four inputs, to compensate for different
times-of-flight and signal path-lengths. It then assigns signals to the correct bunch-crossing, as
detailed below. A look-up table is used to carry out pedestal subtraction, apply a noise threshold,
and do a final transverse-energy calibration, resulting in 8-bit trigger-tower energies. Finally, it
6690 performs bunch-crossing multiplexing (see below) for the CP and it sums the four values into
 0.2×0.2 jet elements (2×2 sum) for the JEP. Two 10-bit low-voltage differential signalling (LVDS)
serialisers operating at 400 Mbit/s transmit the processed trigger-tower data to the CP, while a third
serialiser sends the summed 9-bit jet elements to the JEP.

The pre-processor ASIC reads out data to the data acquisition system upon receiving a L1A
6695 signal. The readout data are taken from pipeline memories at two stages: the raw digitised values
from the FADC's, and the 8-bit processed trigger-tower data from the look-up tables. Data from
the bunch-crossing of interest, as well as a programmable number of bunch-crossings around it
(typically up to five in all), allow monitoring of pulse shapes at the FADC's, and of the bunch-
crossing identification and energy calibration at the look-up table outputs. These readout data are
6700 serialised and sent to the data acquisition readout over an optical fibre. In addition, various rates
based on the input signals are monitored and histogrammed automatically in the ASIC and are read
out by VMEbus.

Bunch-crossing identification The analogue pulses from the electromagnetic and hadronic
calorimeters have widths of several bunch-crossings, so it is essential that trigger-tower signals are
6705 associated with the correct bunch-crossing. This is a crucial requirement not only for normal-size
pulses, but also for saturated pulses (above about 250 GeV) and for pulses as small as possible
(down to 2–3 GeV, just above the noise level). The pre-processor ASIC is capable of identifying
a signal's bunch-crossing using three different methods, which provides ample redundancy for
consistency checks during commissioning.

6710 For normal, unsaturated signals, a digital pipelined finite-impulse-response filter processes five
consecutive FADC samples. A subsequent peak-finder attributes the maximum value of this sum

to the corresponding bunch-crossing. The working range of the method spans from small signals (energy depositions of a few GeV) up to the near-saturation level of around 250 GeV.

For saturated signals, two consecutive samples are compared to a low and a high threshold, making use of the finite peaking-time (approximately 50 ns) of an analogue input signal. Thus, 6715 detection of a leading edge allows attribution of the virtual peak to a specific bunch-crossing. This method is valid from around 200 GeV up to the maximum energy range of the calorimeters.

A third method uses comparators with programmable thresholds on the analogue input signals to present a rising-edge signal to the multi-chip modules. Given the known peaking time, 6720 bunch-crossing identification can be performed using an appropriate programmed delay in the pre-processor ASIC. The validity of this method begins well above the comparator threshold and extends up to the full energy range. There is thus a large overlap with the two previous methods, allowing consistency checks between the methods to be performed.

The finite-impulse-response filter output is presented to the look-up table to extract a calibrated 6725 E_T value for the trigger tower. If the bunch-crossing identification criteria are met, this value is sent to the CP outputs. In the case of saturation, the tower is assigned the maximum 8-bit value of 255 GeV. For the JEP outputs, any 0.2×0.2 -sum jet element which contains a saturated trigger tower, or which has a 9-bit sum in overflow, is assigned the maximum 9-bit value of 512 GeV. A tower or jet element with a maximum value is understood to be saturated by the CP and/or JEP 6730 sub-systems. The trigger menu will be set up so that any event where a saturation condition occurs will produce a L1A signal, and the RoI's sent to the L2 trigger will be flagged by saturation bits.

Bunch-crossing multiplexing The data from the pre-processor modules consist of four 8-bit trigger towers per multi-chip module to the CP, and one 9-bit 0.2×0.2 -sum jet element to the JEP. To economise on the number of links needed, it was noted that the bunch-crossing identification 6735 algorithm is essentially a peak-finding scheme. This means that an occupied bunch-crossing will always be followed by an empty (zero) one. This allows two trigger towers being sent to the CP to share a single serial link. Trigger towers are paired at the pre-processor ASIC output stage, and the scheme is called Bunch-Crossing Multiplexing (BC-mux). By using it, data transmission to the CP sub-system is achieved with only two links per multi-chip module instead of four. However, for the 6740 JEP, where a sum of four towers is transmitted, this cannot be done.

Output signal fan-out and pre-compensation The high-speed serial outputs to the CP and JEP are fanned out in order to provide the trigger algorithms with overlapping data between detector quadrants in azimuth. The data pass through the back-plane to 11 m long shielded parallel-pair cables. RC pre-compensation is done to improve signal-driving capabilities, since observed signal 6745 attenuation and distortion from the cables may compromise data integrity. Bit-error rates of less than 10^{-14} have been achieved.

8.2.1.3 The cluster and jet/energy-sum processors The CP and JEP sub-systems share many architectural features and some common hardware. The jet algorithm in the JEP and the electron/photon and τ cluster algorithms in the CP both perform feature searches in overlapping, sliding 6750 windows. Therefore, a large amount of data duplication between processor modules is required, and this is done as follows. Both sub-systems divide the calorimeters into four azimuthal quadrants, with each processor module within a quadrant covering a slice in pseudorapidity and 90° in

azimuth. Overlapping data from neighbouring azimuthal quadrants are provided by duplicated serial links from the pre-processor. Within each quadrant, modules only need to share input data with their nearest neighbours, over short (roughly 2 cm) point-to-point back-plane links. This architecture minimises the number of cable links from the pre-processor, and the back-plane fan-out is simplified.

The CP is a four-crate system, with 14 Cluster Processor Modules (CPM's) in each crate covering one calorimeter quadrant. The JEP is contained in two crates, each containing eight Jet/Energy Modules (JEM's) from two opposing quadrants in azimuth (16 JEM's total). Results from the processor modules are brought to two Common Merger Modules (CMM's) in each crate: these sum the data to produce crate-level results. The CMM's also perform the system-level summation of data from the different crates, and transmit the final results to the CTP.

The electron/photon and τ triggers extend out to $|\eta| = 2.5$, which is the fiducial limit for precision measurements with the inner detector and electromagnetic calorimetry. The jet trigger extends out to $|\eta| = 3.2$. The E_T^{miss} and total transverse-energy triggers include the forward calorimetry, in particular to provide adequate E_T^{miss} performance, which means that they extend to $|\eta| = 4.9$. This also allows the FCal to be used for forward-jet triggers.

8.2.1.4 The cluster processor module The electron/photon trigger algorithm [236], shown in Fig. 163, identifies 2×2 clusters of trigger towers in which at least one of the four possible two-tower sums (1×2 or 2×1) of nearest-neighbour electromagnetic towers exceeds a pre-defined threshold. Isolation-veto thresholds are set for the 12-tower surrounding ring in the electromagnetic calorimeter, as well as for the 2×2 hadronic-tower core sum behind the cluster and the 12-tower hadronic ring around it. All these thresholds are programmable.

The τ algorithm uses the same basic elements to select narrow hadronic jets. Each of the four possible two-tower sums of nearest-neighbour electromagnetic towers is added to the 2×2 hadronic-tower core sum directly behind, and the result is compared to a pre-defined threshold. Isolation veto thresholds are set separately for each of the surrounding 12-tower rings in both the electromagnetic and hadronic calorimeters.

The isolation thresholds for both algorithms are absolute values, rather than ratios of isolation energy to cluster energy. This simpler approach was chosen on the basis of studies, which showed that the expected isolation sums are relatively insensitive to shower energies. In practice, high-energy clusters will generally have looser isolation criteria to maximise the efficiency for possible low-rate exotic signal processes, while lower-energy clusters will have stricter isolation criteria in order to minimise the rates at the expense of a limited loss of signal.

These algorithms are run over all possible 4×4 windows, which means that the windows overlap and slide by steps of 0.1 in both η and ϕ . This implies that an electron/photon or τ cluster can satisfy the algorithm in two or more neighbouring windows. Multiple-counting of clusters is avoided by requiring the sum of the four central electromagnetic plus the sum of the four central hadronic towers to be a local maximum with respect to its eight nearest overlapping neighbours. In order to avoid problems in comparing digital sums with identical values, four of the eight comparisons are 'greater than' while the other four are 'greater than or equal to', as shown in Fig. 164. The location of this 2×2 local maximum also defines the coordinates of the electron/photon or τ RoI.




Figure 163. Electron/photon and τ trigger algorithms, as described in the text.




Figure 164. E_T local-maximum test for a cluster/ROI candidate. The η -axis runs from left to right, and the ϕ -axis from bottom to top. The symbol R refers to the candidate 2×2 region being tested.

6795 The CPM identifies and counts clusters satisfying sets of threshold and isolation criteria. Eight threshold sets are reserved for electron/photon triggers, while eight further threshold sets can each be used for either electron/photon or τ triggers.

Each CPM receives and deserialises input data on 80 LVDS cables from the pre-processor modules, brought in to the rear of the module through back-plane connectors. The data are then shared between neighbouring modules via the back-plane, and finally fanned out to eight CP FPGA's, which perform the clustering algorithms. The serialiser FPGA's also store the input data in pipelines for eventual readout to the data acquisition system upon reception of a L1A signal.

6805 The eight CP FPGA's each service eight overlapping 4×4 windows. Pipelines implemented in each one of them save output data for readout to the data acquisition system, and also save cluster types and coordinates for readout as ROI's to the L2 trigger. Two hit-multiplicity FPGA's collect and then sum the 3-bit cluster multiplicities from the CP FPGA's, for reporting to the crate-level merging of CP results. These multiplicities are transmitted via the back-plane. If more than seven instances of a cluster type are identified (a very rare occurrence, given that the mean occupancy is less than one), the multiplicity is reported as seven. Two additional FPGA's collect input data from the serialiser FPGA's, ROI data from the CP FPGA's, and output data from the hit-multiplicity FPGA's upon reception of a L1A signal, and transmit them to readout driver modules serving the data acquisition system and the L2 trigger on two optical fibres from the front panel of the module.

6815 **8.2.1.5 The jet/energy module** The Jet/Energy Module (JEM) works with jet elements which are the sums of 2×2 trigger towers in the electromagnetic calorimeters added to 2×2 trigger towers in the hadronic calorimeters. The jet algorithm identifies E_T sums within overlapping windows




Figure 165. Jet trigger algorithms, based on 0.2×0.2 jet elements and showing RoI's (shaded). In the 0.6×0.6 case there are four possible windows containing a given RoI. In the 0.8×0.8 case the RoI is required to be in the centre position, in order to avoid the possibility of two jets per window.

consisting of 2×2 , 3×3 , or 4×4 jet elements, corresponding to window sizes of 0.4, 0.6, or 0.8 in η and ϕ , as shown in Fig. 165. These sums are then compared to pre-defined jet energy thresholds. Multiple-counting of jet candidates is avoided by requiring the window to surround 2×2 jet elements whose sum is a local maximum, with the same definition as for electron/photon and τ clusters. The location of this 2×2 local maximum also defines the coordinates of the jet RoI. Eight independent combinations of jet E_T threshold and window size are available for trigger menus.

The energy-summation algorithm produces sums of E_T , E_x and E_y , and uses the system-level sums of these to report on four total- E_T and eight E_T^{miss} thresholds to the CTP.

Each JEM receives and deserialises data from 88 LVDS links, corresponding to 44 jet elements for both the electromagnetic and hadronic calorimeters. Four input FPGA's receive the data, sum the electromagnetic and hadronic parts of each jet element to 10-bit values, and send these sums to the main processor FPGA's on the same and neighbouring modules. Pipelines in each input FPGA save input data for readout to the data acquisition system upon reception of a L1A signal.

The jet and energy-summation algorithms are implemented in two large main-processor FPGA's per JEM. The main processors are also responsible for reporting results to the crate-level merging, as well as pipelining of data acquisition and RoI information for readout. The jet output of each JEM is a data stream consisting of eight 3-bit jet multiplicities. The energy output is also a data stream containing the values of E_x and E_y , each compressed from 12 bits to an 8-bit quad-linear scale (6-bit mantissa plus two multiplier bits).

A single readout-controller FPGA collects input data from the input FPGA's, and output and RoI data from the main processor FPGA's, for readout to readout driver modules serving data acquisition and the L2 trigger on two optical fibres from the front panel of the module.

8.2.1.6 The common merger module Two modules in each CP and JEP crate carry out crate-level merging of results received from the crate's processor modules. In the CP crates, each merger module is responsible for calculating 3-bit cluster multiplicities for eight of the 16 electron/photon and τ cluster definitions. In the JEP crates, one merger module produces 3-bit multiplicities for the eight jet definitions, while the other produces sums of E_T , E_x and E_y . Each Common Merger Module (CMM) receives data from the crate's 14 CPM's or 16 JEM's, over point-to-point links on the crate back-plane.

The CMM carries out all of these merging functions by using different firmware versions. Each CMM receives up to 400 bits of data per bunch-crossing from the crate's CPM's or JEM's. A large FPGA performs crate-level merging. Parallel LVDS cable links between the sub-system crates bring all crate-level results to one CMM of each type, which is designated as the system-merger CMM. A second FPGA on the CMM carries out the system-level merging.

At the system level, the CMM carries out the logic to provide global trigger results. Three-bit overall multiplicities for each of the electron/photon, τ , and jet thresholds are formed and sent to the CTP. The overall sums of E_x and E_y are applied together as the address to a look-up table. In one operation this works out whether the resulting vector sum, i.e. missing E_T , is above or below eight programmable missing- E_T thresholds and codes the result in an eight-bit word. For total scalar E_T , the global sum is compared to four threshold values. Finally, a rough approximation of the total E_T in jets, based on the numbers of jets passing each of the eight jet thresholds, is compared to four threshold values.

All of these calorimeter trigger results are passed to the CTP by cable. As with other processor modules, input and output data on each CMM are stored in FIFO's and read out to the data acquisition system over an optical fibre. RoI data on the missing and total E_T values are sent to L2.

8.2.1.7 The processor back-plane The CP and JEP use a common, custom processor back-plane. It has 16 CPM/JEM positions flanked by two CMM positions. At the left it has a slot for a commercial VMEbus CPU. At the right is a slot for a timing control module, which interfaces to the TTC (e.g. to distribute clock signals) and to the detector control system, which uses CANbus to monitor voltages and temperatures on all trigger modules.

The processor back-plane is a monolithic printed-circuit board of 9U height. It is populated almost entirely with 2 mm hard-metric connectors, with 1148 signal and ground pins in each JEM/CPM and CMM position. There are point-to-point links between neighbouring processor modules for input data fan-in/fan-out. Connections are provided from each CPM or JEM to the two CMM's at the right and left of the processor modules. To conserve pins, a non-standard VMEbus with the minimum possible number of pins (43 signals with 16 data bits and 24 address bits) is used.

The LVDS serial-input and merger-interconnect cables are connected to the rear of the processor back-plane and passed through it to the modules in front. This results in a system with fewer cables on the front panels of the modules and, as a consequence, hopefully fewer recabling errors and less cable damage over the lifetime of the experiment.

8.2.1.8 The readout driver The trigger system has two separate readout systems. Input, output, and some intermediate data from each module are read out to the data acquisition system and at the same time the CP and JEP sub-systems report feature types and coordinates as RoI data to the L2 trigger.

The readout system has been designed to handle one bunch-crossing of RoI data and up to five bunch-crossings of data acquisition data per event at a L1A rate of up to 75 kHz. A common approach has been adopted in all L1Calo sub-systems for data acquisition and RoI readout.

On each module to be read out, readout FIFO's on each processor FPGA or ASIC are read out as serial streams to a readout controller FPGA for timing alignment. This passes the serial streams




Figure 166. Schema (left) and segmentation (right) of the L1 muon barrel trigger. Left: The RPC's are arranged in three stations: RPC1, RPC2, and RPC3. Also shown are the low- p_T and high- p_T roads. See text for details. Right: areas covered by η and ϕ coincidence-matrix (CM) boards, by an RoI, by a Pad logic board, and by sector logic boards.

in parallel to the inputs of a G-link transmitter, which transmits them serially at 800 Mbit/s over optical fibres to a Readout Driver (ROD).

A common ROD module is used by both the data acquisition and RoI readout sub-systems to gather and report data from the pre-processor modules, CPM's, JEM's, and CMM's, using different firmware configurations for different readout tasks and modules. The ROD is a 9U-module residing in a standard VME64x crate. It has 18 G-link receivers, which pass their parallel outputs to the FPGA's for data compression, zero suppression, and some data monitoring. The ROD also contains four S-link transmitters on a rear-transition module for passing compressed event data to the data acquisition and RoI readout buffers. Routing of data to the different outputs is carried out by a switch controller FPGA, whose settings depend on the type and source of data being read out. In addition, a further large FPGA provides monitoring capability on a sample of readout data.

8.2.2 Muon trigger

The L1 muon trigger is based on dedicated finely segmented detectors (the RPC's in the barrel and the TGC's in the end-caps, as described in detail in Section 6.5) with a sufficient timing accuracy to provide unambiguous identification of the bunch-crossing containing the muon candidate.

The trigger in both the barrel and the end-cap regions is based on three trigger stations each. The basic principle of the algorithm is to require a coincidence of hits in the different trigger stations within a road, which tracks the path of a muon from the interaction point through the detector. The width of the road is related to the p_T threshold to be applied. A system of programmable coincidence logic allows concurrent operation with a total of six thresholds, three associated with the low- p_T trigger (threshold range approximately 6–9 GeV) and three associated with the high- p_T trigger (threshold range approximately 9–35 GeV). The trigger signals from the barrel and the muon end-cap trigger are combined into one set of six threshold multiplicities for each bunch-crossing in the muon to CTP interface, before being passed on to the CTP itself.

8.2.2.1 Muon barrel trigger

Trigger signals The muon trigger for the barrel regions [237] ($|\eta| < 1.05$) makes use of dedicated RPC detectors. The RPC is a gaseous detector providing a typical space-time resolution of $1 \text{ cm} \times 1 \text{ ns}$ and a rate capability of about 1 kHz/cm^2 . As shown on the left side of Fig. 166, the RPC's are arranged in three stations. The two Barrel Middle (BM) stations, RPC1 and RPC2, are arranged on either side of the Monitored Drift Tube (MDT) BM stations at approximately 7.5 m radial distance from the interaction point (see Chapter 6). The RPC3 Barrel Outer (BO) station, mounted on the inside (large sectors) or outside (small sectors) of the MDT BO stations, is located at a radial distance of about 10 m. Each station is made of one RPC doublet, i.e. two independent detector layers, each measuring η and ϕ . Both planes are used in the trigger. The η -strips are parallel to the MDT wires and provide the bending view of the trigger detector. The ϕ -strips are orthogonal to the MDT wires and provide the second coordinate measurement. These strips are also needed for the pattern recognition. The RPC's are organised in several modules, and their dimensions have been chosen to match those of the corresponding MDT chambers. In most stations the RPC's are composed of two units along the beam direction. To avoid dead areas between adjacent units, the active zones of neighbouring RPC's are partially overlapped in η .

Trigger algorithm The trigger algorithm operates in the following way: if a track hit is generated in the second RPC doublet (the pivot plane), a search for a corresponding hit is made in the first RPC doublet, within a road whose centre is defined by the line of conjunction of the hit in the pivot plane with the interaction point. The width of the road is a function of the desired cut on p_T : the smaller the road, the higher the cut on p_T . The system is designed so that three such low- p_T thresholds in each projection can be applied simultaneously. The algorithm is performed in both the η and the ϕ projections to reduce accidental triggers from low-energy particles in the cavern. A 3-out-of-4 coincidence of the four layers of the two doublets is required, which ensures excellent rejection of fake tracks from noise hits and greatly improves the stability of the trigger performance in the presence of small chamber inefficiencies.

The high- p_T algorithm makes use of the low- p_T trigger built from hits in RPC1 and RPC2, and of the information generated in the RPC3 station. The algorithm operates in a similar way to the low- p_T one. The centre of the road is determined in the same way as for the low- p_T trigger, and in addition to the low- p_T trigger pattern result, 1-out-of-2 possible hits of the RPC3 doublet is required. As with the low- p_T trigger, three p_T thresholds operate simultaneously, resulting in a total of six thresholds reported to the central trigger logic for each event. For both low and high- p_T triggers, trigger information in η and ϕ is combined to form RoI's to be sent to the L2 trigger.

System implementation The trigger scheme for the barrel muon trigger is implemented in custom-built electronics, mounted either directly on the RPC detectors or located outside the main experimental cavern. A schema of the trigger signal and readout chain is shown in Fig. 167. Signals from the RPC detectors are processed in Amplifier-Shaper-Discriminator (ASD) boards (shown as triangles in Fig. 167) attached to the chambers at the end of the RPC strips. In the low- p_T trigger, for each of the η and the ϕ projections the RPC signals of the RPC1 and RPC2 doublets are sent to a coincidence matrix board containing a coincidence matrix chip. This chip performs most of the functions of the trigger algorithm and of the readout. At this stage the signals are aligned, the




Figure 167. Schema of the trigger signal and readout chain of the L1 barrel muon trigger.

coincidence operations are performed, and the three p_T thresholds are applied. The coincidence matrix board produces an output pattern containing the low- p_T trigger results for each pair of RPC doublets in the η and ϕ projections. The information of the two adjacent coincidence matrix boards in the η projection, and similarly in the ϕ projection, are combined in the low- p_T Pad Logic board (low- p_T Pad in Fig. 167) board. The four low- p_T coincidence matrix boards and the corresponding Pad board are mounted on top of the RPC2 detector stations. The low- p_T Pad board generates the low- p_T trigger result and the associated RoI information. This information is transmitted to the corresponding high- p_T Pad board, which collects the overall results for low- p_T and high- p_T .

In the high- p_T trigger, for each of the η and ϕ projections the signals from the RPC3 doublet, and the corresponding pattern result of the low- p_T trigger, are sent, via dedicated LVDS links, to a coincidence matrix board very similar to the one used in the low- p_T trigger. This board contains the same coincidence matrix chip as the low- p_T board, programmed for the high- p_T algorithm. The high- p_T board produces an output pattern containing the high- p_T trigger results for a given RPC doublet in the η and ϕ projection. The information of two adjacent coincidence matrix boards in the η projection, and similarly in the ϕ projection, are combined in the high- p_T Pad logic board (high- p_T Pad in Fig. 167). The four high- p_T coincidence matrix boards and the corresponding Pad board are mounted on top of the RPC3 detector.

The high- p_T Pad board combines the low- p_T and high- p_T trigger results. The combined information for each bunch-crossing is sent via optical links to sector logic boards located in the USA15 counting room. Each sector logic board receives inputs from seven (six) low- p_T (high- p_T) Pad boards, combining and encoding the trigger results of one trigger sector. The sector logic board sends the trigger data for each bunch-crossing to the Muon to Central Trigger Processor Interface (MUCTPI, see Section 8.2.2.3), located in the USA15 counting room.

6975 For events which are selected by the L1 trigger, data are read out from both the low- p_T and the
high- p_T Pad boards. These data include the RPC strip pattern and some additional information used
in the L2 trigger. The readout data for events accepted by the L1 trigger are sent asynchronously to
ROD's located in the USA15 counting room and from there to Readout Buffers (ROB's). The data
links for the readout data are independent of the ones used to transfer partial trigger results to the
6980 sector logic boards.

System segmentation and latency From the trigger point of view the barrel is divided into
two halves, $\eta < 0$ and $\eta > 0$, and within each half-barrel 32 logically identical sectors are defined.
The correspondence between these logical sectors and physical chambers is indicated in the dia-
gram on the right of Fig. 166. The barrel large chambers and the barrel small chambers of both
6985 middle and outer RPC stations are logically divided in two in azimuth to produce two large sectors
and two small sectors per half-barrel octant. Inside a sector, the trigger is segmented in Pads and
RoI's.

A large sector contains seven Pad regions, while a small sector contains six Pad regions. The
region covered by a Pad is about 0.2×0.2 in $\Delta\eta \times \Delta\phi$. Inside the Pad the trigger is segmented into
6990 RoI's. A RoI is a region given by the overlap of an η coincidence-matrix and a ϕ coincidence-
matrix. The dimensions of the RoI's are about 0.1×0.1 in $\Delta\eta \times \Delta\phi$. The total number of Pads is
 $7 \times 2 \times 32$ for the small sectors and $6 \times 2 \times 32$ for the large ones, giving 832 Pads altogether. Since
one Pad covers four RoI's, the total number of RoI's is 3328.

To avoid losing efficiency due to uncovered regions in the trigger system, different parts of
6995 the system overlap. However, this overlap can cause double-counting of muon candidates. In
the barrel trigger system, overlap is treated and solved at three different levels. Within a Pad
region the Pad logic removes double-counting of tracks between the four RoI's of the region. In
addition, if it is found that a trigger was generated in a zone of overlap with another Pad region,
this trigger is flagged as 'border' trigger and any overlap will be solved later on. The sector logic
7000 then prevents double-counting of triggers within a sector. Triggers generated in zones of overlap
between different sectors are flagged by the sector logic and sent to the MUCTPI, which prevents
double-counting between sectors.

The latency of the muon barrel trigger is about 2.1 μ s, well within the allowed envelope.

8.2.2.2 Muon end-cap trigger

7005 **Trigger signals** The muon trigger for the end-cap regions [238] is based on signals provided
by TGC detectors. The time resolution is not as good as for RPC's, but good enough to provide an
efficiency greater than 99% for bunch-crossing identification for the 25 ns gate of ATLAS. Crucial
for the end-cap region of ATLAS is their larger rate capability of more than 20 kHz/cm². The
TGC's are arranged in nine layers of gas volumes grouped into four planes in z (Fig 168 (left) and
7010 see also Chapter 6). The TGC inner station (I) at $|z| \sim 7$ m consists of one plane of doublet units. At
 $|z| \sim 14$ m seven layers are arranged in one plane of triplet chambers (M1, closest to the interaction
point) and two planes of doublet chambers (M2, M3). The doublet forming the plane farthest from
the interaction point in each end-cap (M3) is referred to as the pivot plane, and its chamber layout
and electronics are arranged such that, to a good approximation, there are no overlaps or holes
7015 in this plane. For triggering, the TGC's cover a pseudorapidity range $1.05 < |\eta| < 2.4$, except




Figure 168. Schema (left) and segmentation (right) of the L1 muon end-cap trigger. See text for details.

for the innermost plane which covers a range $1.05 < |\eta| < 1.92$. Each trigger plane consists of a wheel of eight octants of chambers symmetric in ϕ . Each octant is divided radially into the ‘forward region’ and the ‘end-cap region’. Anode wires of TGC’s are arranged in the azimuthal direction and provide signals for R information, while readout strips orthogonal to these wires provide signals for ϕ information. Both wire and strip signals are used for the muon trigger. Signals from two wire-planes and two strip-planes are read out from the doublet chambers, and signals of three wire-planes but only two strip-planes are read out from the triplet chambers. Anode wires are grouped and fed to a common readout channel for input to the trigger electronics, resulting in wire-group widths in the range between 10.8 mm and 55.8 mm. Wire groups are staggered by half a wire group between the two planes of a doublet station, and by one third of a wire group between each of the planes of a triplet station. Each chamber has 32 radial strips, and thus the width of a strip is 4 mrad (8 mrad for the forward region). Strips are also staggered by half a strip-width between the two strip-planes in a triplet or a doublet chamber.

Trigger algorithm The scheme of the L1 muon end-cap trigger is shown on the left hand side of Fig. 168. The trigger algorithm extrapolates pivot-plane hits to the interaction point, to construct roads following the apparent infinite-momentum path of the track. Deviations from this path of hits in the trigger planes closer to the interaction point are related to the momentum of the track. Coincidence signals are generated independently for R and ϕ . A 3-out-of-4 coincidence is required for the doublet pair planes of M2 and M3, for both wires and strips, a 2-out-of-3 coincidence for the triplet wire planes, and 1-out-of-2 possible hits for the triplet strip planes. The final trigger decision in the muon end-cap system is done by merging the results of the $R - \phi$ coincidence and the information from the EI/FI chambers in the inner station (see Section 6.7.1). As the $\eta - \phi$ coverage of the EI/FI chambers is limited, the coincidence requirements depend on the trigger region,

in order to keep a uniform efficiency in the end-cap region. Six sets of windows are constructed
7040 around the infinite-momentum path, corresponding to three different high- p_T and three different
low- p_T thresholds. Trigger signals from both doublets and the triplet are involved in identifying
the high- p_T candidates, while in case of the low- p_T candidates the triplet station may be omitted
to retain high efficiency, given the geometry and magnetic field configuration of a specific region.

System implementation The trigger scheme outlined above is implemented in purpose-built
7045 electronics, partly mounted on and near the TGC chambers, and partly located in the USA15 count-
ing room. A schema of the trigger signal and readout chain is shown in Fig. 169. The wire and
strip signals emerging from the TGC's are fed into ASD boards physically attached to the edge of
a TGC and enclosed inside the TGC electrical shielding. Each ASD board handles 16 channels.
From the ASD boards signals are routed to the so-called PS-boards, which integrate several func-
7050 tions in one unit. Each PS-board receives signals from up to 20 ASD's. First the signals are routed
to a patch-panel section, which also receives timing signals from the TTC system. Signal align-
ment and bunch-crossing identification (BCID) is performed at this stage, and physical overlaps of
TGC chambers are handled. In addition, detector control system and other control and monitor-
ing signals are routed to the other parts of the electronics mounted on the chambers. The aligned
7055 signals are passed to the so-called slave section, where the coincidence conditions are applied and
readout functions are performed. The PS-boards are placed on the accessible outer surfaces of the
TGC wheels: the electronics for the two doublets are mounted on the outside of the outer doublet
wheel M3 and those for the triplets on the inner surface of the triplet wheel M1. The EI/FI PS-
boards are installed in racks located near the EI/FI chambers. Signals from the doublet and triplet
7060 slave boards are combined to identify high- p_T track candidates in coincidence boards combining
all three trigger planes (M1, M2, M3), so-called high- p_T boards, located in dedicated mini-racks
around the outer rim of the triplet wheel. Wire (R -coordinate) and strip (ϕ -coordinate) information
is still treated separately at this point. Signals from high- p_T boards are sent to sector logic boards
containing an $R - \phi$ coincidence unit and a track selector to select the highest- p_T coincidences.
7065 The sector logic also receives directly the signals from the EI/FI slave boards and can incorporate
them into the trigger logic. The sector logic boards are located in the USA15 counting room. The
resulting trigger information for 72 separate trigger sectors per side is sent to the MUCTPI.

Full-information data sets are read out through the data acquisition system in parallel with the
primary trigger logic. For readout purposes the slave boards of one or more trigger sectors are
7070 grouped into local data acquisition blocks. Each slave board is connected to a so-called star switch,
which manages the data collection for a local data acquisition block. From the star switch, the data
are passed on to the ROD's located in the USA15 counting room, and from there to ROB's.

System segmentation and latency The trigger-sector segmentation of one pivot-plane oc-
tant is shown in Fig. 168 (right). The pivot plane is divided into two regions, end-cap ($|\eta| < 1.92$)
7075 and forward ($|\eta| > 1.92$). Each octant of the end-cap region is divided into six trigger sectors in
 ϕ , where a trigger sector is a logical unit which is treated independently in the trigger. Trigger
sectors are constructed to be projective with respect to the interaction point, and therefore may
cross chamber boundaries (see Fig. 168, left). Each octant of the forward region is divided into
three trigger sectors. There are hence 48 end-cap trigger sectors and 24 forward trigger sectors per




Figure 169. Schema of the trigger signal and readout chain of the L1 muon end-cap trigger. See text for details.

7080 end-cap of TGC detectors. Each trigger sector consists of independent sub-sectors corresponding to eight channels of wire groups and eight channels of readout strips, 148 for each end-cap trigger sector and 64 for each forward trigger sector. The trigger sub-sectors correspond to the RoI's sent to the L2 trigger for events accepted by the L1 trigger.

The latency of the muon end-cap trigger is about $2.1 \mu s$, well within the allowed envelope.

7085 8.2.2.3 Muon to central trigger processor interface

Functional overview The results from the muon barrel and end-cap trigger processors which form the input to the Muon to Central Trigger Processor Interface (MUCTPI) provide information on up to two muon-track candidates per muon trigger sector. The information includes the position and p_T threshold passed by the track candidates. The MUCTPI combines the information from all the sectors and calculates total multiplicity values for each of the six p_T thresholds. These multiplicity values are sent to the CTP for each bunch-crossing. For each sector either all muon candidates may be taken into account, or only the candidate with the highest p_T per sector. In forming the multiplicity sums, care has to be taken to avoid double-counting of muon candidates in regions where trigger chambers overlap. As described above, many cases of overlaps are resolved within the barrel and end-cap muon trigger processors. The remaining overlaps to be treated by

the MUCTPI are those in ϕ direction between neighbouring barrel trigger sectors, and between barrel and end-cap trigger sectors. The maximum overall multiplicity is seven candidates. Larger multiplicities will appear as a multiplicity of seven.

7100 Additional functions of the MUCTPI are to provide data to the L2 trigger and to the data acquisition system for events selected at L1. The L2 trigger is sent a subset of all muon candidate information which form the RoI's for the L2 processing. The muon RoI's sent to L2 are ordered according to decreasing p_T . The data acquisition receives a more complete set of information, including in addition the computed multiplicity values.

System implementation and latency The MUCTPI is divided into a number of building 7105 blocks which are housed in one 9U VMEbus crate. The different functionalities of the MUCTPI are implemented in three types of VMEbus modules which are connected to each other via an active back-plane, and controlled by a commercial CPU unit acting as VMEbus master.

7110 A total of 16 octant input boards each receive data corresponding to an octant in the azimuthal direction and half the detector in the η direction. They form muon-candidate multiplicities for this region, correctly taking into account the overlap zones between barrel sectors, and between barrel and end-cap sectors. There is no overlap between muon trigger sectors associated with different octant boards. The interface board to the CTP collects the multiplicity sums for the six p_T thresholds over the custom back-plane described below. The sums are transmitted to the CTP for each bunch-crossing. The interface board is also responsible for distributing time-critical 7115 control signals to the rest of the MUCTPI system. The readout driver of the system sends candidate information to the data acquisition and the L2 trigger for each accepted event. All modules are connected via a custom-built back-plane. It contains two components: an active part forms the total candidate multiplicities by adding the multiplicities of the input boards and a passive part contains a bus system to transfer data to the readout driver of the system on receipt of a L1A.

7120 The latency of MUCTPI is included in the latency numbers for the barrel and end-cap muon trigger systems quoted above, to which it contributes $0.2 \mu\text{s}$.

8.2.3 Central trigger processor

8.2.3.1 Functional overview The Central Trigger Processor (CTP) [239] receives trigger infor- 7125 mation from the calorimeter and muon trigger processors, which consists of multiplicities for electrons/photons, τ -leptons, jets, and muons, and of flags indicating which thresholds were passed for total and missing transverse energy, and for total jet transverse energy. Additional inputs are provided for special triggers such as a filled-bunch trigger based on beam-pickup monitors, and a minimum-bias trigger based on scintillation counters. Up to 372 signals can be connected to the input boards of the CTP; however, only up to 160 can be transmitted internally. The selection of the 7130 signals used from all signals available at the input boards is programmable. The currently foreseen input signals listed in Table 48 sum up to 150 bits and will therefore all be available in parallel.

In the next step the CTP uses look-up tables to form trigger conditions from the input signals. Such a condition could be, for example, that the multiplicity of a particular muon threshold has exceeded one, i.e. at least two muons in this event have passed this threshold. For such an event, 7135 this trigger condition would be set to *true*. Further trigger conditions are derived from internally generated trigger signals: two random triggers, two pre-scaled clocks, and eight triggers for pro-

Cable origin	Number of bits	Trigger information
Muon processor	6 thresholds \times 3 bits	muon multiplicities
Cluster processor 1	8 thresholds \times 3 bits	electron/photon multiplicities
Cluster processor 2	8 thresholds \times 3 bits	electron/photon or τ multiplicities
Jet/energy processor 1	8 thresholds \times 3 bits 4 bits	jet multiplicities total jet transverse energy
Jet/energy processor 2	2 \times 4 thresholds \times 2 bits	forward-jet multiplicities for each side,
Jet/energy processor 3	4 \times 1 bit 8 \times 1 bit	total transverse-energy sum, missing transverse-energy sum
CTP calibration	28 bits	up to 28 input bits for additional trigger inputs from beam pick-ups (see Section 9.10), beam condition monitors (see Section 3.4.1), luminosity detectors (see Sections 7.1 and 7.2), zero-degree calorimeters (see Section 7.3), and others.

Table 48. Trigger inputs to the CTP of the L1 trigger. The number of bits implies the maximum multiplicity which can be encoded, i.e. up to seven for three bits. Multiplicities larger than this value will be set to the possible maximum, in this case seven.

grammable groups of bunch-crossings. The maximum number of trigger conditions at any one time is 256.

The trigger conditions are combined to form up to 256 trigger items, where every trigger condition may contribute to every trigger item. An example for a trigger item would be that the following conditions have been fulfilled: at least two muons have passed a particular threshold, and at least one jet has passed a particular threshold. Furthermore each trigger item has a mask, a priority (for the dead-time generated by the CTP), and a pre-scaling factor (between 1 and 2^{24}). The L1A signal generated by the CTP is the logical OR of all trigger items.

The CTP provides an eight-bit trigger-type word with each L1A signal. This indicates the type of trigger, and can be used to select options in the event data processing in the front-end electronics and readout chain. The CTP sends, upon reception of each L1A signal, information about the trigger decision for all trigger items to the L2 trigger (RoI builder) and the data acquisition (ROS). Part of the readout data of the CTP is the number of the current luminosity block. A luminosity block is the shortest time interval for which the integrated luminosity, corrected for dead-time and pre-scale effects, can be determined. In case of detector failures, data can be rejected from the boundary of the last luminosity block known to be unaffected, and the interval should therefore be as small as possible to avoid unnecessary data loss. On the other hand, each luminosity block should contain enough data such that the uncertainty of the luminosity determination is limited by




Figure 170. Layout of the VMEbus crate for the central trigger processor of the L1 trigger. The calibration module has the further function of receiving additional trigger signals, which are transmitted to one of the input module connectors via a front panel cable. The LTP links are the connection to the individual sub-detector systems.

7155 systematic effects, not by the available statistics in the interval. For ATLAS this interval will be on
the order of minutes. A luminosity block transition is initiated by the CTP, which will momentarily
pause the generation of triggers, increment the luminosity block number in a register located on
the CTP decision module, and release the trigger again. From this location the number is included
in the readout data for each event. At each transition a set of scalers is read out from the CTP and
7160 stored, marked with the luminosity block number, in a database. These scalers keep track of the
number of triggers generated by the trigger logic, the number of triggers surviving the pre-scale
veto, and the number of triggers surviving the dead-time veto. The values of these counters are
needed to later derive the corresponding corrections of the luminosity value associated with each
luminosity block. For monitoring purposes the CTP provides bunch-by-bunch scalers of inputs and,
7165 integrated over all bunches, scalers of trigger inputs and trigger items before and after pre-scaling.

In addition to its function in the selection chain, the CTP is also the timing master of the detector. The clock signal synchronised to the LHC beams arrives at the CTP, and is distributed from here together with the L1A and other timing signals to all other sub-systems.

8.2.3.2 System implementation and latency The CTP consists of six types of modules which
7170 are housed in a single 9U VMEbus crate, as shown in Fig. 170. Internal communication between
the controller CPU and the modules, and between the modules proceeds by bus systems imple-
mented on the back-planes of the crate. In addition to VMEbus, the CTP modules use custom
busses for the synchronised and aligned trigger inputs (PITbus, where PIT = pattern in time), for
the common timing and trigger signals (COMbus), and for the sub-detector calibration requests
7175 (CALbus). These extra busses are implemented on a custom-built backplane installed in the CTP

crate.

Six different module types are employed in the CTP system. The timing signals from the LHC are received by the machine interface module (designated LHC inputs in Fig. 170), which can also generate these signals internally for stand-alone running. This board also controls and monitors the internal and external busy signals, as for example the busy signal transmitted from a sub-detector in case of overload on its data acquisition system. The module sends the timing signals to the COMbus, thereby making them available to all of the other modules in the CTP.

The trigger input modules receive trigger inputs from the muon and calorimeter trigger processors and other sources. The input boards select and route the trigger inputs to the PITbus, after synchronising them to the clock signal and aligning them with respect to the bunch-crossing. Three boards with four connectors of 31 trigger input signals each allow for a total of 372 input signals to be connected, of which up to 160 can be made available on the PITbus at any given time.

The trigger decision module receives the trigger inputs from the PITbus. It combines them and additional internal triggers using several look-up tables to form up to 256 trigger conditions. A typical example is a 3-to-4 look-up table, with the three input bits encoding a threshold multiplicity between 0 and 7, and the four output bits enabled if the multiplicity exceeds 0, 1, 4, 6, respectively. The first bit in this list may then be used as the trigger condition that one or more objects fulfilled the corresponding trigger threshold. In a further step the trigger conditions are combined using content-addressable memories to form up to 256 trigger items. Any of the up to 256 trigger conditions may participate in any of the up to 256 trigger items. The trigger masks, pre-scales, and dead-time generation following the forming of the trigger items are implemented in this module. The trigger results are transmitted to the COMbus. The trigger decision module also acts as the readout driver of the system, sending information to the L2 trigger and the data acquisition for each accepted event.

The monitoring module receives the 160 trigger inputs from the PITbus and monitors their behaviour on a bunch-by-bunch basis. The frequency of signals on each input line can be monitored and histogrammed, and can be retrieved via VMEbus. Trigger signals encoding multiplicities are decoded before they are monitored.

The output module (labelled LTP links in Fig. 170) receives the timing and trigger signals from the COMbus and fans them out to the sub-detectors. The module receives back from the sub-systems the busy signals, which are sent to the COMbus, and 3-bit calibration trigger requests, which are routed to the CALbus.

The calibration module time-multiplexes the calibration requests on the CALbus and sends them via a front-panel cable to one of the input modules. The calibration module also has front-panel inputs for beam pick-up monitors, minimum-bias scintillators, and test triggers.

The latency of the CTP is contained in the latency numbers for the barrel and end-cap muon trigger systems and the calorimeter trigger system quoted above, to which it contributes 100 ns.

On the sub-detector side, the timing signals are received by the Local Trigger Processor (LTP) [240], which acts as an interface between the CTP and the timing distribution system of each sub-detector. During stand-alone data taking of a sub-detector the LTP can generate all timing signals locally and also provides inputs for locally generated triggers. The LTP is fully programmable and can therefore act as a switch between the global and locally generated signals without the need for re-cabling. The timing distribution system of a sub-detector may be partitioned into several

parts, each with its own LTP. In this case LTP's can be daisy-chained in order to save output ports
7220 on the CTP and the associated cabling. The LTP is complemented by an interface module which
provides an additional input and output port, such that interconnections of sub-detectors are possi-
ble without removing the link to the CTP. Both the LTP and the interface module are implemented
as 6U VMEbus boards.

From the LTP the timing signals are distributed to the detector front-end electronics using the
7225 Timing, Trigger and Control system (TTC). The ATLAS TTC system is based on the optical fan-out
system developed within the framework of RD12 [241]. Clock and orbit signals synchronous to the
LHC beams arrive at the machine interface module of the CTP after passing through the RF2TTC
interface module shown in Fig. 171. In the RF2TTC, the signals are cleaned and delays may be
applied to account for any drift in the signal phase. From the CTP the signals are transmitted to the
7230 LTP's together with detector-specific timing and control signals like the L1A or the event counter
reset signal. From the LTP onwards TTC components are available again to serialise the signals
(TTCvi) and transmit them (TTCex) via optical fibres to the detector front-end electronics, where
TTC receiver chips (TTCrx) decode the transmitted information and make it available as electrical
signals for further use. The implementation and use of the TTC system is sub-system specific. As
7235 an example the muon trigger systems use TTC standard components to transmit the timing signals
all the way to the electronics mounted on the chambers, while in case of the inner tracking detector
a custom-built distribution system is used to transmit the signals from the counting rooms to the
cavern.

8.3 Data acquisition system and high-level trigger

7240 8.3.1 Overview

As explained in Section 8.1, the main components of the data acquisition system/High-Level Trig-
ger (DAQ/HLT) are: readout, L2 trigger, event-building, event filter, configuration, control and
monitoring. An overview of the event selection performed by the HLT is given in Section 10.9.
Here the movement of data from the detectors to the HLT and subsequently to mass storage is
7245 described. The main features of each component are described below.

A block diagram of the DAQ/HLT is shown in Fig. 160. The movement of events from the
detector to mass storage commences with the selection of events by the L1 trigger. During the
latency of the L1 trigger selection, up to $2.5 \mu\text{s}$, the event data are buffered in memories located
within the detector-specific front-end electronics. On selection by the L1 trigger the event data
7250 is transferred to the DAQ/HLT system over 1574 Readout Links (ROL's), having first transited
through the detector-specific ROD's. The 1574 event fragments are received into the 1574 Readout
Buffers (ROB's) contained in the Readout System (ROS) units where they are temporarily stored
and provided, on request, to the subsequent stages of the DAQ/HLT system.

For every selected event, the L1 trigger sub-systems (calorimeter, muon, and CTP) also pro-
7255 vide the RoI information on eight ROL's, a dedicated data path, to the RoI builder where it is
assembled into a single data structure and forwarded to one of the L2 supervisor (L2SV). As its
name suggests, the L2SV marshals the events within the L2 trigger. It receives the RoI's, assigns
each event to one of the L2 trigger's processing units (L2PU's) for analysis, and receives the result
of the L2PU's analysis.




Figure 171. Schema of the distribution of timing signals from the LHC radio-frequency system to ATLAS and within the experiment. Here ROD (Readout Driver) more generally denotes the readout electronics in the counting rooms which receive the timing signals, while front-end denotes electronics mounted on the detector components in the main cavern.

7260 Using the RoI information, requests for event data are made to the appropriate ROS's. The sequence of data requests is determined by the type of RoI identified by the L1 trigger and the configuration of the L2 trigger processing, i.e. the order of items in the trigger menu and the order of the algorithms per trigger item. The result, accept or reject, of the analysis is returned to the L2SV which subsequently forwards it to the DataFlow Manager (DFM). In addition to sending
 7265 the result of its analysis to the L2SV, an L2PU also sends a summary of the analysis which it has performed to a L2 trigger-specific ROS.

The DFM marshals the events during the event-building. For those events which were found not to fulfil any of the L2 selection criteria, the DFM informs all the ROS's to expunge the associated event data from their respective ROB's. Each event which has been selected by the L2 trigger
 7270 is assigned by the DFM to an event-building node (called SFI). The SFI collects the event data from the ROS's and builds a single event-data structure, the event. An SFI can build more than one event at a time and the requests to the ROS's for their data are dispatched according to an algorithm which ensures the quantity of data being received by the SFI does not exceed its available input bandwidth. The full event structure is sent to the event filter for further analysis. On completing
 7275 the building of an event an SFI notifies the DFM, which subsequently informs all the ROS's to expunge the associated event data from their respective ROB's.

The event filter, in addition to the selection, classifies the selected events according to a pre-determined set of event streams and the result of this classification is added to the event structure. Selected events are subsequently sent to the output nodes (SFO's) of the DAQ/HLT system. Conversely, those events not fulfilling any of the event filter selection criteria are expunged from the system. The events received by an SFO are stored in its local file system according to the classification performed by the event filter. The event files are subsequently transferred to CERN's central data-recording facility.

8.3.2 Control

The overall control of the experiment covers the control and monitoring of the operational parameters of the detectors and experiment infrastructure, as well as the coordination of all detector, trigger and data acquisition software and hardware associated with data-taking. This functionality is provided by two independent, complementary and interacting systems: the data acquisition control system, and the Detector Control System (DCS). The former is charged with controlling the hardware and software elements of the detectors and the DAQ/HLT needed for data-taking, while the DCS handles the control of the detector equipment and related infrastructure. The DCS is described in Section 8.5.

The DAQ/HLT system and detector systems are composed of a large number of distributed hardware and software components which in a coordinated manner provide for the data-taking functionality of the overall system. Likewise, their control and configuration is based on a distributed control system. The control system has two basic components: the process manager and the run control.

On each computer a process management daemon waits for commands to launch or interrupt processes. On the reception of such commands it interrogates the access manager and the resource manager to ascertain whether the requested operation is permitted. It is a task of the access manager to indicate whether the requester is authorised to perform the operation, while it is a task of the resource manager to check that the resources are available to perform the operation.

A hierarchical tree of run controllers, which follows the functional de-composition into systems and sub-systems of the detector, steers the data acquisition by starting and stopping processes and by carrying all data-taking elements through a finite state machine, which ensures that all parts of ATLAS are in a coherent state. As with the handling of commands to the process manager, run control commands also have to be authorised by the access manager. In addition to implementing a global finite state machine and managing the lifetime of processes, the run controllers are further customised according to the sub-system for which they are in charge. One example of a customised controller is the root controller, the starting point of the run control tree, which retrieves the run number from the run number service before starting any new run and drives luminosity block changes during the data-taking.

Another fundamental aspect of the control is the diagnostic and error recovery system. Several aspects of it are integrated into the run control processes. Errors raised by any data-taking node enter the error reporting system and can elicit an appropriate reaction. The diagnostic system can launch a set of tests to understand the origin of the reported problem and the recovery system can then take corrective actions. These aspects of the control have been implemented using an expert system. In order to allow for analysis of errors a posteriori, all error and information messages are

7320 archived in a database via the log service. A user interface is provided for efficient searching of messages.

Computer management and monitoring Access to the experiment's local area network, hence all computers located at the experimental site, is gained via an application gateway. User accounts and passwords are stored in a central directory, which is used to authenticate all the users for logging into the computers associated with the experiment. In addition, the same directory holds
7325 the configuration of various system services, such as the servers of the auto-mounted directories, and the user roles and policies required by the role-based access control scheme adopted.

All PC's and single board computers boot over the network using the pre-boot execution environment. The kernel and boot image files for the nodes are provided by a system of Local File Servers (LFS's) each of which serve approximately thirty clients. The initial boot image is a reduced version of Scientific Linux CERN (SLC), which provides the minimum set of binaries and
7330 libraries to operate a node. The remaining non-essential parts of the OS, e.g. the X-window environment, is then loaded via NFS from the LFS's. The LFS also provide, again via NFS, the ATLAS software and disk space to the nodes they serve.

A Central File Server (CFS) holds the master copy of the ATLAS software, which is distributed to other CFS's and to the LFS on a daily basis or on request. The unique installation of the
7335 operating system and the ATLAS software in conjunction with booting over the network, ensures the uniformity of the software throughout the computer cluster.

The disks served by the LFS are primarily used by the disk-less nodes as scratch space and for storage of detector specific software. Another use of the LFS is for running various central
7340 services related to the configuration and control of the trigger and data acquisition systems, thus providing a hierarchical structure for these services to provide the required scaling. Examples are the information server and remote database server (see Section 8.3.3).

All nodes are monitored by a customised version of Nagios, an application which monitors the health of the cluster by performing checks of various node services, e.g. disk utilisation, at
7345 regular intervals. The results of these checks are displayed by Nagios on a web page and recorded in a database for subsequent retrieval and analysis. Based on the collected data, Nagios allows the graphical tracking and analysis of the health of the cluster or an individual node. It is also configured to set alarms, send notifications via emails and SMS (Short Message Service) messages and, for some services, perform recovery operations, for example the restarting of a network interface.

7350 **8.3.3 Configuration**

The description of the hardware and software required for data-taking are maintained in configuration databases. A configuration is organised as a tree of linked segments according to the hierarchy describing the DAQ/HLT and detector systems. A segment defines a well defined sub-set of the
7355 hardware, software, and their associated parameters. The organisation of the data is described by common object-oriented database schemas which may be extended to describe the properties of specific hardware and software.

To support concurrent access to the configuration data by thousands of applications, and to notify control applications of changes to the configuration data during run time, remote database servers are used to ensure that access times to the configuration data do not scale with the number

7360 of deployed applications. Additional servers are also deployed to cache the results of queries to the relational databases, e.g. the conditions database.

Trigger configuration At any point in time the complete trigger chain needs to be consistently configured. For L1, i.e. the CTP, a trigger menu comprising up to 256 trigger items which should cause an event to be selected, see Section 8.2.3, is defined. Moreover the calorimeter and muon trigger systems have to be configured such that they deliver the information required by the trigger menu (multiplicities for trigger thresholds).

To ensure a coherent and consistent configuration of the L1, L2, and event filter, all components of the trigger are configured using an integrated system. It makes the configuration parameters available to all systems participating in the L1 trigger decision, and to all nodes forming the HLT farms.

The trigger configuration system itself contains a central relational database which stores the configuration data, tools to populate the database and ensure its consistency, the functionality to archive the configuration data and interfaces to extract the data in preparation for data taking or other purposes, e.g. simulation and/or data analysis. A detailed description of the system is given in [242].

Partitioning Partitioning refers to the ability to operate subsets of the ATLAS detector in parallel and disjointly, thus facilitating the concurrent commissioning and operation of subsets of the detector. Once two or more partitions have been commissioned they may then be operated together as a single partition. Partitions are in this way combined into a fully integrated and operational detector.

A partition maps to a TTC partition, therefore defining the subset of detector components within a partition (see Section 8.2.3.2 and also Table 51). In addition, the static point-to-point connections between the detector ROD's and the ROS's uniquely associates a set of ROS's to a partition. Other components of the DAQ/HLT (i.e. event-building nodes, event filter nodes and SFO's) are connected by multi-layered networks and can therefore be assigned to a partition as required by the operations to be performed. The management of resources, e.g. event-building nodes, between partitions is achieved by the resource manager.

The RoI builder drives the input to the L2 trigger and, from an operational perspective, can only be operated as a single unit. Thus the L2 trigger cannot be partitioned. Analogously, the CTP can only be operated as a single unit, therefore the complete L1 trigger may not be operated in more than a single partition. However, to facilitate calibration and checking of L1Calo input signals, it is possible to operate a partition which consists of the L1Calo and the LAr and/or tile calorimeters but independent of the CTP. Similarly, the L1 muon trigger can operate, for example the RPC's with the MDT's, as a separate partition without the CTP.

7395 **8.3.4 Monitoring and information distribution**

The monitoring component provides the framework for the routing of operational data and their analysis. Operational data ranges from physics event data, to histograms and the values of parameters. The routing of operational data is performed by the information, on-line histogramming and event monitoring services. The information service provides the distribution of the values of

7400 simple variables. Consumers of the information are able to subscribe to notifications of changes
to one or more information items. It also provides a means for any application to send commands
to any of the information providers, specifically for the control of information flow, e.g. an appli-
cation may ask a particular provider of information to increase the frequency at which it publishes
a particular piece of information. Complementing the exchange of the values of simple variables,
7405 the message reporting system transports messages among trigger and data acquisition applications.
Messages may be used to report debug information, warnings or error conditions. The message
reporting system allows association of qualifiers and parameters to a message. Moreover, receivers
of messages are able to subscribe to the service to be notified about incoming messages and apply
filtering criteria.

7410 The On-line Histogramming Service (OHS) extends the functionality of the information ser-
vice to histograms, in particular raw and ROOT histograms. Within the DAQ/HLT there are many
instances of the same application, e.g. L2PU's, active at any one time producing histograms. Via
the OHS, a gatherer application sums histograms of the same type and in turn publishes, via the
OHS, the resulting histograms. The visual presentation of histograms is based on ROOT and Qt,
7415 and allows for the presentation of reference histograms, fitting, zooming and the sending of com-
mands to histogram providers.

The event filter processing application is based on the off-line computing framework. The sub-
stitution of the selection algorithms with a monitoring or calibration algorithm allows for monitor-
ing and/or calibration tasks based on the offline computing framework to operate on-line, receiving
7420 events from the SFI's. It is also possible to configure these applications to receive events from the
event monitoring service. The latter provides a framework to enable the sampling and distribution
of event data as they flow through the DAQ/HLT system. Monitoring applications are able to re-
quest event fragments according to the values of elements in the event fragment, e.g. trigger and/or
sub-detector types, from a specific sampling point, e.g. a particular ROS (part of an event) or SFI
7425 (a complete event). Examples of monitoring applications using this service are the event dump and
event display.

To complement the viewing and analysis of histograms by an operator, a data quality monitor-
ing framework provides the automatic comparison of recently acquired data to reference data (e.g.
reference histograms), statistical checks and alarm generation. More specifically, user-supplied al-
7430 gorithms and/or reference data are used to automatically analyse the large quantities of monitoring
data, and generate alarms when deviations from the specified criteria occur.

8.3.5 Readout system

As described in Section 8.3.1, the Readout System (ROS) receives event data from the detector
ROD's via 1574 ROL's. All of the ROL's have the same design and implementation, based on the
7435 S-link interface. It allows for the transmission of 32-bit data at 40.08 MHz, i.e. up to 160 Mbyte/s,
and implements flow control and error detection [243]. ROB's are the buffers located at the receiv-
ing end of the ROL's, there being one ROL associated to one ROB. Three ROB's are physically
implemented on a module called a ROBIN, and up to six ROBIN's can be located in a ROS, which
is implemented on a server-class PC. The ROS provides the multiplexing of up to 18 ROL's to the
7440 subsequent components of the DAQ/HLT, i.e. L2 trigger and event-building, reducing the number
of connections by approximately an order of magnitude.




Figure 172. Expected average RoI request rate per ROS for a luminosity of $10^{33} \text{ cm}^{-2} \text{ s}^{-1}$.

A request by an L2PU for data involves, on average, one or two ROB's per ROS, whereas the requests for data from the event-building nodes concern the event data from all the ROB's of a ROS. In either case, the ROS replies to the requester with a single data structure. At the L1 trigger rate of 75 kHz, and an average of 1 kbyte received per RoI, the ROS is able to concurrently service up to approximately 20 kHz of data requests from the L2 trigger, up to 3.5 kHz of requests from event-building nodes, and expunge events on request from the DFM. The rate of data requests received by a specific ROS depends on the $\eta - \phi$ region of the data it receives over the RoI's and from which detector it receives data. For example, a ROS which receives data from the liquid-argon calorimeter barrel region is solicited for data more frequently than a ROS associated with the barrel MDT's. The expected average rate of RoI requests as a function of 135 ROS's, which participate to the L2 trigger is shown in Fig. 172 (see Section 10.9.3 for examples of initial trigger menus). Figure 173 shows the expected maximum L1 trigger accept rate sustainable by the ROS which is most solicited for RoI data by the L2 trigger, for different values of the L2 trigger's acceptance, i.e. the event-building rate. Also shown is the expected operating point at high luminosity.

ROBIN The ROBIN component provides the temporary buffering of the individual event fragments produced by the ROD's for the duration of the L2 trigger decision and, for approximately 3% of the events, the duration of the event-building process. In addition, it services requests for data at up to rates of approximately 20 kHz. As a consequence of the rates which have to be supported, the ROBIN is a custom-designed and built PCI-X mezzanine [244]. All functions related to the receiving and buffering of event fragments are realised in an FPGA. A PowerPC is used to implement the functions of memory management, servicing of data requests, control and operational monitoring.

8.3.6 L2 trigger

7465 The L2 trigger is achieved by
the combined functionality of the
RoI builder, L2SV, L2PU and
L2 trigger-specific ROS (pseudo-
ROS). The RoI builder receives the
7470 RoI information from the different
sources within the L1 trigger on
eight input RoI's and merges them
into a single data structure. It is
thus at the boundary between the
7475 L1 and L2 trigger systems and op-
erates at the L1 trigger rate (see
next Section). The single data
structure containing the RoI data
is transmitted by the RoI builder
7480 over one of the output RoI's to
the L2SV's. As described in Sec-
tion 8.3.1, L2SV's marshal the events through the L2 trigger.

The principal component of the L2 trigger is the L2 processing farm, where the event selection is executed. The system is designed to provide an event rejection factor of about 30, with an average
7485 throughput per farm node of about 200 Hz, using (but not exclusively, see Section 10.9.4.5) only the
data located in the RoI's, i.e. 1–2% of the full data of an event. The number of L2PU applications
performing the physics selection per node is configurable. On the hardware currently deployed
(see Section 8.4) there are eight L2PU's per node, and one L2PU per processing core of the node,
thus the average event processing time per L2PU should be less than 40 ms.

7490 The transmission of a summary of the L2 trigger's selection is achieved by the deployment of a
pseudo-ROS. At the end of its event analysis the L2PU sends to the pseudo-ROS information which
summarises the results of its analysis. Subsequently, the pseudo-ROS participates in event-building
like any other ROS within the system, its event data being the L2 trigger's summary analysis. In
this way, the results of the L2 trigger's analysis are built into the final event and subsequently used
7495 by the event filter to seed its selection.

The failure of one or more L2PU's during run time does not incur system down time. The system continues to operate at a reduced rate while the failed application, the L2PU, can be restarted under the supervision of the run control.

Steering of the event selection The HLT starts from the RoI's delivered by the L1 trigger
7500 and applies trigger decisions in a series of steps, each refining existing information by acquiring
additional data from increasingly more detectors. A list of physics signatures (trigger chains),
implemented event reconstruction (feature extraction) and selection algorithms are used to build
signature and sequence tables for all HLT steps. Feature extraction algorithms typically request
7505 detector data from within the RoI and attempt to identify features, e.g. a track or a calorimeter
cluster. Subsequently, a hypothesis algorithm determines whether the identified feature meets the




Figure 173. The maximum sustainable L1 trigger accept rate as a function of the L2 trigger acceptance for the ROS which is most solicited for RoI data by the L2 trigger. Also shown is the expected operating point at high luminosity.

criteria (such as a shower shape, track-cluster match or E_T threshold) necessary to continue. Each signature is tested in this way. The decision to reject the event or continue is based on the validity of signatures, taking into account pre-scale and pass-through factors. Thus events can be rejected early after an intermediate step if no signatures remain viable. In this manner the full data set associated with the RoI is transferred only for those events which fulfil the complete L2 trigger selection criteria, i.e. the amount of data transferred between the ROS's and the L2 trigger is minimised for those events which are rejected. The stepwise and seeded processing of events in the HLT is controlled by the steering.

The steering runs within the L2 and event filter processing tasks. It implements two of the key architectural and event-selection strategies of the trigger and data acquisition systems: RoI-based reconstruction and step-wise selection. Both are designed to reduce processing time and L2 network bandwidth. The steering takes the static configuration described in Section 8.3.3 and applies it to the dynamic event conditions (which RoI's are active and the status of their reconstruction) in order to determine which algorithms should be run on which RoI's and in which order, and ultimately to decide whether the event has fulfilled the criteria for acceptance.

Counters, maintained by the steering, for each trigger chain and step are made available by the online monitoring system, see Section 8.3.4, enabling the real time monitoring of the trigger rates. The steering also provides a mechanism for the selection algorithms to publish parameters necessary for monitoring the quality of the event selection.

Region-of-interest builder The RoI builder [245] is one of only three custom-built components within the DAQ/HLT system. It is a 9U VMEbus system composed of an input stage, an assembly stage, and a single-board computer for the purpose of configuration, control and operational monitoring. The input stage consists of three input cards which each receive and buffer, over three ROL's, the RoI data from three of the eight L1 trigger sources, namely: four CP ROD's, two JEP ROD's, the MUCTPI, and the CTP.

The eight RoI fragments are subsequently routed over a custom back-plane to builder cards in the assembly stage, where they are assembled into a single data structure (RoI record). The assignment of each event to a specific builder card for assembly is based on a token-passing mechanism between the builder cards. Each builder card has four output ROL's which are used to transfer the assembled RoI records to up to four L2SV's according to a round-robin algorithm.

Detector calibration using RoI data The calibration of the muon MDT chambers requires large data samples within a well-defined time window to establish the relationship between the drift path and measured time as a function of time. This measurement has to be made from the data of the MDT's alone using candidate tracks, and is based on an iterative procedure starting from a preliminary set of constants.

The L1 single-muon trigger rate at a luminosity of $10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ for a threshold of 6 GeV is approximately 25 kHz. For these candidate events the first step of the L2 trigger selection is the reconstruction of tracks in the muon system. To facilitate the calibration of the MDT's, each L2PU can be configured to additionally write, to a pre-defined buffer, the data of the candidate tracks and the results of its analysis, i.e. the RoI information and the results of track fits. Subsequently the

data are transferred from this buffer to a L2-wide calibration server which stores the data to disk prior to sending it to a remote calibration farm for processing.

8.3.7 Event-building

7550 The event-building functionality is provided by the DFM, ROS's and SFI's [246]. The SFI is the application which collects the event data from the ROS's and assembles the event as a single formatted data structure. An SFI is configured with a randomised list of the ROS's within the system, which is used to define the order in which data requests are sent to the ROS's. This results in the randomisation of the traffic pattern in the underlying network and hence improved network performance. To meet the rate requirements a number of SFI's work in parallel, each instance
7555 building a number of events concurrently. Each SFI informs the DFM of its readiness to receive events, and the DFM allocates events to the SFI's so as to ensure that the load is balanced across all available SFI's.

The default behaviour of the SFI is to collect all the event data associated with a given event into a single formatted data structure. However, a subset of the events accepted by the L2 trigger
7560 are for the purposes of detector calibration and do not necessitate the collection of all the event data. For this type of event, the SFI is capable of collecting a subset of the available event data. The subset is identified by the L2 trigger and communicated to the SFI via the DFM. The subset can range from a few ROB's to a whole sub-detector.

If a requested ROS data fragment is not received within a configurable time budget, the out-
7565 standing data fragment can be requested again. Only if several consecutive requests are not fulfilled does the SFI abandon the inclusion of the missing data and assemble an incomplete event. After an event has been moved to the event filter the SFI marks its buffers for re-use. If, for whatever reason, the buffers of the SFI become full, the SFI informs the DFM, i.e. exerts back-pressure, which subsequently suspends the allocation of events to the specific SFI until the SFI indicates it is
7570 again available.

The event-building system is designed to function even in case of failure of one or more SFI nodes. In this situation, the DFM ceases to assign events to the failed SFI's. Once the failed nodes become available again, they can be re-integrated into the event-building system without the system incurring down time.

7575 8.3.8 Event filter

The event filter is a processing farm; on each processing node a configurable number of independent processing tasks receive and process events. Unlike the L2 trigger, these tasks are based on standard ATLAS event reconstruction and analysis applications. The steering of the event selection is the same as L2, as described in 8.3.6. For those events passing the selection criteria, a subset
7580 of the data generated during the event analysis is appended to the event data structure, enabling subsequent offline analysis to be seeded by the results from the event filter. An integral part of the selection process is the classification of the events according to the ATLAS physics streams, see Section 8.3.9. To this end, for those events which fulfil the selection criteria, a tag is added to the event data structure identifying into which physics stream the event has been classified.

7585 The failure of one or more event-filter processing tasks or of a complete node during run-time does not provoke any system down-time. The system continues to operate at a reduced rate while

Stream	e	μ	Jet	γ	E_T^{miss} & τ	B -physics
e	31 ± 7.9	0.0056 ± 0.00058	$0.00053 \pm 6.2 \times 10^{-5}$	1.2 ± 0.4	1.4 ± 0.035	$(1.3 \pm 1.3) \times 10^{-5}$
μ	–	34 ± 8.7	0.021 ± 0.015	0.0028 ± 0.002	0.22 ± 0.022	0.076 ± 0.0043
Jet	–	–	38 ± 5.9	0.48 ± 0.4	0.71 ± 0.4	0 ± 0
γ	–	–	–	22 ± 5.7	0.22 ± 0.073	0 ± 0
E_T^{miss} & τ	–	–	–	–	32 ± 7.9	$(15 \pm 6.4) \times 10^{-6}$
B -physics	–	–	–	–	–	9.5 ± 5.5

Table 49. Overlap (Hz) between the data streams at a luminosity of $10^{33} \text{ cm}^{-2} \text{ s}^{-1}$.

the failed application, or node, can be restarted under the supervision of the run control. To ensure that no events are lost during such failures, each event on arrival in the event filter is written to a memory mapped file. On the restart of the failed application or of the node itself, an attempt can be made to re-analyse the event or accept the event without analysis.

8.3.9 Event output

The main functionality of the event-filter output nodes (SFO's) is to receive events which have passed the event filter selection criteria, interface the DAQ/HLT to CERN's central data-recording facility, and de-couple the data-taking process from possible variations in the central data-recording service.

The SFO maintains, locally, a set of files into which it records events at a peak event rate of up to 400 Hz. In the eventuality of a prolonged failure in the transmission of data to CERN's central data recording service, there is sufficient local storage capacity to buffer all events locally for up to 24 hours. Under normal operating conditions, this storage capacity is only partially used. The set of files maps to the ATLAS-defined data streams: electrons, muons, jets, photons, E_T^{miss} and τ -leptons, and B -physics. Each event is recorded in one or more files according to the stream classification made by the event-filter processing task. Table 49 shows the rates for each of the data streams and in the off-diagonal elements, the rates of the overlaps between them.

In addition to the data streams mentioned above, a subset of the events is also written to calibration streams and an express stream. The express stream is a subset of the events selected by the event filter and fulfil additional criteria which select the events as being useful for monitoring the quality of the data and the detector. The calibration stream provides the minimum amount of information needed for detector calibration, possibly at a rate higher than the data streams provide. These events will only contain a subset of the event data.

8.4 Implementation and capabilities of the DAQ/HLT

Most of the DAQ/HLT functionality is implemented on commodity, rack-mountable, server-class PC's. The PC's run Scientific Linux CERN and are interconnected by multi-layer gigabit-Ethernet networks, one for control functionality and another for data movement. The majority of PC's have similar specifications (e.g. two CPU sockets, two gigabit-Ethernet connections, support for IPMIv2.0), and differ only by the number and type of CPU's implemented and the amount of memory. The main features per component and the number of nodes deployed for initial operations

Component	Number of nodes	Number of racks	Number of CPU's/node	Memory (Gbyte)	Type of CPU
ROS	145	16	1	0.512	3.4 GHz Irwindale
SFI	48	3	2	2	2.6 GHz Opteron 252
DFM	12	1			
L2SV	10				
HLT	1116	36	8	8	Xeon E5320 1.86 GHz
SFO	6	2	2	4	Xeon E5130 2.0 GHz
Monitoring	32	4	4	8	Xeon E5160 3.0 GHz
Operations	20		2	4	Xeon E5130 2.0 GHz

Table 50. The main data-acquisition system components to be deployed for initial operation: the readout system (ROS), the event-building node (SFI), the dataflow manager (DFM), the L2 supervisor (L2SV), the high-level trigger (HLT) and the event filter output nodes (SFO).

in 2008 are given in Table 50. A few components, the RoI builder, ROL and ROBIN, are, however, implemented in custom hardware.

The ROS PC's are installed in standard ATLAS 52U racks, while all other PC's are installed in standard 47U or 52U server racks. The number of racks (for initial operation) for each component type is given in Table 50. In addition to the PC's, each rack also contains a local file server and two gigabit-Ethernet switches. The latter form part of the multi-layered gigabit-Ethernet network which implements the control and data networks. Each rack is also equipped with a water-cooled heat-exchanger, designed for the horizontal airflow within a rack, which provides up to 9.5 kW of cooling power. The number of 1U PC's per rack is typically just over thirty, constrained by cooling power, power distribution (particularly in-rush current) and weight limits.

For initial operations, the DAQ/HLT system will be fully configured in the area of configuration, control and monitoring functionality. The operations PC's are used to provide the various central services for configuring and controlling the trigger and data acquisition systems (e.g. run control, error logging). The monitoring PC's are used to monitor the system and sampled event data.

The initial system will also support full detector readout, over the 1574 point-to-point ROL's, into the ROS's at L1 trigger rates up to 75 kHz. The number of ROD's, ROL's and ROS's per detector TTC partition are given in Table 51. Also given in this table is the expected size of event data per L1 trigger for each part of the detector for a luminosity of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$.

As described in Section 8.3.7, the event-building functionality is performed by a set of SFI's and scales linearly with the number of SFI's, each SFI contributing 60 Hz and approximately 90 Mbyte/s to the total event-building rate and aggregate bandwidth. For initial operations, forty-eight SFI's are deployed allowing a sustained event-building rate of approximately 2.0 kHz, for an

TTC Partition			Number of ROD's	Number of ROL's	Number of ROS's	Data per L1A signal (kbyte)
Inner detector	Pixel	Layer 0	44	44	4	60
		Disks	24	24	2	
		Layers 1-2	64	64	6	
	SCT	End-cap A	24	24	2	110
		End-cap C	24	24	2	
		Barrel A	22	22	2	
		Barrel C	22	22	2	
	TRT	End-cap A	64	64	6	307
		End-cap C	64	64	6	
		Barrel A	32	32	3	
		Barrel C	32	32	3	
	Calorimetry	Tile	Barrel A	8	16	2
Barrel C			8	16	2	
Extended barrel A			8	16	2	
Extended barrel C			8	16	2	
LAr		EM barrel A	56	224	20	576
		EM barrel C	56	224	20	
		EM end-cap A	35	138	12	
		EM end-cap C	35	138	12	
		HEC	6	24	2	
		FCal	4	14	2	
Muon spectrometer	MDT	Barrel A	50	50	4	154
		Barrel C	50	50	4	
		End-cap A	52	52	4	
		End-cap C	52	52	4	
	CSC	End-cap A	8	8	1	10
		End-cap C	8	8	1	
L1	Calorimeter	CP	4	8	1	28 (can be varied)
		JEP	2	8	1	
		PP	8	32	3	
	Muon RPC	Barrel A	16	16	2	12
		Barrel C	16	16	2	
	Muon TGC	End-cap A	12	12	1	6
		End-cap C	12	12	1	
	MUCTPI		1	1	1	0.1
	CTP		1	1	1	0.2
Total			932	1574	145	1311

Table 51. Numbers of readout drivers (ROD's), readout links (ROL's) and readout systems (ROS's) per detector TTC partition, as well as expected data size per L1A signal for a luminosity of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$.

7640 average event size of approximately 1.3 Mbyte.

In addition to the features given in Table 50, the PC's for the SFO functionality are each equipped with three RAID controllers each managing eight 500 Gbyte SATA II disks. The three sets of disks are operated as a circular buffer: while events are being written to the event streams of one set of disks, a second set of disks is used to send data to CERN's central data recording
7645 service. In this configuration, a single RAID controller does not perform both writing and reading operations simultaneously, thus maximising the throughput of an SFO. The deployed set of SFO's fulfil the final design specifications: a sustained output bandwidth of 300 Mbyte/s and a peak rate of 600 Mbyte/s. Thus for an average event size of 1.3 Mbyte, this gives a sustained event rate of 200 Hz.

7650 Of the 36 HLT racks available for initial operations, eight racks (248 nodes) are dedicated to the event filter selection while the remaining twenty-eight racks (868 nodes) can be configured to perform either the L2 trigger or event filter selection, i.e. the amount of computing power apportioned to the L2 trigger and or the event filter will be adjusted according to the data-taking conditions. The baseline apportioning of these nodes envisages nine racks (279 nodes) for the L2
7655 trigger and twenty-seven racks (837 nodes) for the event filter. In the DAQ/HLT Technical Design Report [247], the algorithm processing times and rejection rates were based on single-core processors with an expected clock speed of about 8 GHz, giving processing times per event of order 10 ms at L2 and 1 s at the event filter. These figures correspond to approximately 40 ms and 4 s respectively per core on today's quad-core processors operating at a clock speed of 2 GHz.
7660 Measurements with simulated raw data show that the processing times per event at L2 and event filter are consistent with the available computing resources for acceptable trigger rates and with a representative mixture of simulated events passing the L1 trigger. Therefore the system deployed for initial operations should be able to handle an initial L1 trigger rate of approximately 40 kHz, about half of the final design specification.

7665 **8.5 Detector control system**

In order to enable coherent and safe operation of the ATLAS detector, a Detector Control System (DCS) has been defined and implemented [248]. The DCS puts the detector hardware into selected operational conditions, continuously monitors and archives its run-time parameters, and performs automatically corrective actions if necessary. Furthermore, DCS provides a human interface for
7670 the full control of ATLAS and its sub-detectors. Figure 174 shows the general system architecture consisting of two parts: the front-end systems and the back-end control.

The front-end consists of the hardware components used to control the detector, ranging from simple sensors to complex devices such as software controlled power supplies. In order to minimise the effort of integration of devices into the DCS and to achieve a homogeneous system, a small set
7675 of commercial devices, such as crates or power supplies, has been selected as standard. The readout of these devices is normally done using the industry-standard protocol OPC.

Due to the special conditions in the experiment cavern, strong magnetic field and ionising radiation, a general-purpose I/O concentrator, the Embedded Local Monitor Board (ELMB) [249], has been developed. An ELMB comprises a multiplexed ADC (64 channels with 16 bit resolution),
7680 24 digital I/O lines and a serial bus SPI to drive external devices. The ELMB can be configured for




Figure 174. Architecture of the DCS.

various types of sensors. A micro-controller pre-processes the data (e.g. calibration, threshold detection) before they are transferred via CANbus to the back-end. The ELMB is designed and tested to be radiation tolerant to a level of about 1 Gy/y and can hence also be placed inside the detector at places shielded by the calorimeter. The ELMB is used in two ways: either directly embedded in the detector electronics, or attached to a general-purpose motherboard to which sensors can be connected. In total about 5000 ELMB's are installed and they are controlled by an OPC server using the CANopen protocol.

The back-end is organised in 3 layers: the Local Control Stations (LCS) for process control of subsystems, the Sub-detector Control Stations (SCS) for high-level control of a sub-detector allowing stand-alone operation, and the Global Control Stations (GCS) with human interfaces in the ATLAS control room for the overall operation. Each station is a PC running the commercial controls software PVSS-II [250], which provides the necessary supervisory functions such as data analysis with the possibility of triggering the execution of pre-defined procedures, raising alarms, data visualisation, and archiving. This supervisory control and data acquisition software package (SCADA) has been chosen for all LHC experiments in the frame of the Joint Controls Project [251], which provides a set of tools on top of PVSS-II and also software components for the standardised devices. In total, the back-end consists of more than 150 PC's, connected as a distributed system for which PVSS-II handles inter-process communication via the local area network. The full back-end hierarchy down to the level of individual devices is represented by a distributed finite state machine

7700 allowing for the standardised operation and error handling in each functional layer.

The LCS are connected to the front-end of a specific sub-system and read, process, and archive the respective data. They execute the commands issued from the SCS and GCS layers and additionally allow the implementation of closed-loop control.

7705 The SCS enable full stand-alone operation of a sub-detector by means of the finite state machine and provide a user interface to control the different subsystems. The individual controls functions of the sub-detectors are described in previous chapters of this paper. The SCS also handles the synchronisation with the data acquisition system.

The GCS in the top layer provide all functions needed in the ATLAS control room to operate the complete detector. The operator interface shows the detector status and allows to navigate in the tree-like finite state machine hierarchy and to execute high level commands which are propagated to the layers below. A data viewer provides selection and plotting of all data available in the DCS. The alarm system collects and displays all parameters which are outside of pre-defined ranges, classified in three severity levels: Warning, Error, and Fatal. An information server (DCS IS) provides an interface to external control systems such as the sub-detector gas systems, the liquid-argon and helium cryogenics systems, the ATLAS magnets, and the electricity distribution for the detector. The data collected from those systems is made available inside the distributed DCS to the individual sub-detectors. Summarised status information is openly available via the World Wide Web.

7720 In addition to the individual sub-detector control stations, one SCS is dedicated to Common Infrastructure Control (CIC) to supervise the common environment and services of the detector. Each of the five geographical zones defined for the CIC (three electronics rooms underground, the cavern of the experiment, and the trigger and data acquisition systems computer rooms) is controlled by an LCS. In each zone, the environmental parameters are monitored, the operational parameters of the electronics racks are supervised, and the electricity distribution is controlled. 7725 Furthermore, the CIC includes a large network of I/O points in the experimental cavern, consisting of about 100 ELMB's. It reads the data of radiation monitors, the movement sensors of the safety system to look for personnel inside the ATLAS area (see Section 9.9), and some 200 temperature sensors positioned at the support structure of the barrel toroid. Additional readout capacity for future system extensions is available.

7730 The DCS comprises a set of common software tools and packages, used by sub-detector controls and the CIC. A configuration database stores the settings needed for the different operational modes of the detector. All status information and measured data can be transferred to the ATLAS-wide conditions database COOL. Another package allows the synchronisation and information exchange between the DCS and the data acquisition system.

7735 **9. Integration and installation**

9.1 Introduction

The ATLAS detector has been installed underground in a huge cavern, situated in Switzerland at point 1 on the LHC ring, directly opposite the main entrance to the CERN site. The installation task has involved a large team from across the whole collaboration (see Annex C).

7740 The soil conditions at point 1 were favourable for the large amount of excavation work required. Once the experiment was approved, preparations for the infrastructure began: civil engineering, electrical power distribution, cooling water, ventilation, etc. In parallel to the work on the infrastructure, many studies on the assembly, integration, layout, support structures, services (pipes and cables), safety issues and access requirements/means for the experiment were carried out.

7745 The construction of the component parts for ATLAS was distributed over many institutions around the world. The components then had to be brought to CERN in a timely manner, a considerable challenge in itself in terms of their size, complexity and fragile nature. In most cases, final assembly and testing were done at CERN on the surface, prior to installation underground.

7750 The main cavern is 92.5 metres underground and the detector is almost as large as the cavern in which it is housed. Supervising the construction of such a complex project meant that rather formal and uniform management and engineering tools had to be used to monitor and document the progress of the project and ensure that items arrived on time and satisfied the requirements. Ensuring that all the pieces of the puzzle fitted together turned out to be a particularly difficult challenge, since the physics goals and the geometry of the detector require minimal clearances
7755 between neighbouring parts. One of the most stringent requirements of the ATLAS detector is to ensure hermetic coverage over most of the solid angle: installation of the detector had therefore to be performed to great accuracy, in order to guarantee optimal coverage.

7760 After a brief description of the processes and tools used to fulfill the infrastructure, integration and installation tasks (Section 9.2), this chapter briefly describes the mechanical integration (Section 9.3), the overall infrastructure and services at point 1 (Section 9.4), the support and access structures (Section 9.5), the detector installation process (Section 9.6), the detector opening and access scenarios (Section 9.7), the beam-pipe (Section 9.8), safety issues (Section 9.9), and the interfaces to the LHC machine (Section 9.10).

9.2 Organisational issues

7765 The ATLAS project involves many people who are spread around the world. It has also generated a huge, complex and multi-disciplinary volume of data which needs to be organised and shared in an easy and transparent way. The management of the integration and installation work was the responsibility of the technical coordination team. In order to help manage the design, production and installation phases of the ATLAS project, various organisational processes and computing tools
7770 were developed [252]. The design phase of the project required the production of drawings, schedules and specifications for procurement. A number of review processes were included during this phase of the project. The goal of these reviews was to evaluate the feasibility and technical validity of the proposed designs. In addition to these reviews, and before launching the production of major items, internal design reviews and production-readiness reviews were organised. Then, during the

Tool	Function
Engineering data management system (EDMS)	Structured storage and retrieval of engineering data
CERN drawing directory (CDD)	Processing of technical drawings
Project progress tracking (PPT)	Regular Web-based notification and reporting system
Equipment management database (EMD) and manufacturing and test folder (MTF) database	Traceability of all equipment installed in the cavern
Rack wizard and ATLASeditor3D	Configuration of electronics rack connectivity (from detector to counting room)
ATLASsurvey3D	Monitoring of ATLAS sub-system displacements (also uses survey data)
Cable database	Assist cable installation team (labels, routing, connector specifications, etc)
The Glance Project	Interacts with all equipment data residing in distinct and geographically spread repositories

Table 52. List and function of the various software tools used by ATLAS technical coordination during the installation of the detector.

7775 production phase of the sub-systems, production-advancement reviews were implemented to check progress and compare it with the production milestones in the schedule. In case of specific and major technical issues, experts were called in from within and outside the collaboration to solve such problems in a timely fashion.

7780 The ATLAS technical management board meets on a monthly basis and provides a forum for regular reporting of the status and problems in all areas relevant to the work at point 1 to the collaboration scientists and engineers. During these meetings, the installation schedule is discussed and proposed future strategies are agreed upon. During the installation of ATLAS, the progress and status of the work were monitored on a weekly basis in dedicated meetings with each main detector system. The Web-based tools shown in Table 52 were used to assist in the communication and organisation process.

9.3 Mechanical integration

7790 The mechanical integration process had to address both static issues related to installation and survey of major detector components and dynamic issues related to detector placement, movements of parts during installation, and to access and maintenance (see Section 9.7). This section deals with the mechanical integration aspects related to installation.

7795 The mechanical integration process defined the overall experimental layout, where each nested sub-system has its well-defined shape and position and has no overlaps with any other sub-system. This integration process started from an initial input for the positioning of the sub-systems and for the space needed for access and services. It then defined mechanical envelopes and the overall three-dimensional layout of the ATLAS detector, using most of the tools listed in Table 52.

9.3.1 Envelopes (individual, global, dynamic)

Envelopes define the space allocated to each part of each sub-system. Three types of envelopes have been created, individual, global and dynamic, and they are defined as follows:

- The individual envelope is the space allocated for the manufactured object, including some space added to the nominal design drawing envelope, in order to take into account fabrication and assembly tolerances.
- The global envelope includes, in addition to the individual envelope, some space dedicated to the inaccuracy of the positioning inside the detector and the deformations applied during installation and operation.
- The dynamic envelope includes, in addition to the global envelope, space for deviations and deformations during displacements (e.g. during access) of the object inside the detector.

After the manufacture and installation of each sub-system, the envelopes were checked and compared to the measurements performed by the survey team. Envelopes have been created as 3D objects with the help of various CAD systems. All this work on modelling and conflict-checking has been most important in order to facilitate the installation process and avoid cost and schedule problems between conflicting objects during installation.

9.3.2 Survey and placement strategy

ATLAS is being assembled in a cavern which is not much larger than the detector itself. Thus any available space had to be optimised once an installation was complete. As soon as the cavern was delivered to ATLAS by the civil engineers, and before any infrastructure was installed, an exhaustive scan was carried out in order to check the as-built work. The task of surveying for the ATLAS detector has been a very challenging one due to the size, nature, complexity and global scale of the work.

9.3.2.1 Survey reference grid in the cavern The nominal beam-line was defined and used during the installation and positioning of the detectors in the cavern. It is defined by the best-fit alignment line of the low- β quadrupole magnets, located at a distance of 30 m from either side of the interaction point. This reference line can deviate from the real beam-line by as much as 2-3 mm. It is determined by the reference sockets in the tunnel, from which the machine elements are installed. The final control is carried out on the machine elements themselves (inner focusing quadrupole triplets included) relative to each other. Spatial uncertainties from 0.5 mm to 1.2 mm at the one sigma level was estimated for any fiducial mark with respect to the nominal beam-line depending on the location of the given target.

The datum (interaction point, radial orientation of the colliding beams and reference plane) is given by the initial geometry in the tunnel and the final positioning of the low- β quadrupole magnets. The survey grid reference in the cavern is linked to the machine geometry by standard geometrical measurements and permanent monitoring systems. These include hydrostatic and wire positioning capacitive sensors, implemented in the survey galleries and joining the low- β quadrupoles via the cavern and radial tubes [253]. The reference grid will thus be monitored throughout the lifetime of the detector.

7835 **9.3.2.2 Stability measurements of the floor and the bed-plates** Civil-engineering calculations indicated possible vertical floor movements of up to 6 mm settlement due to the loading of the detector and a 1 mm per year lift due to excavation heave. ATLAS has a very limited adjustment capability once the detector elements have been placed in-situ. A placement strategy was therefore developed to position all elements within the best achievable mechanical tolerance, relatively to the interaction point and the nominal beam-line [254].

7840 To monitor these predicted movements, periodical measurements have been carried out on about 20 reference marks embedded in the cavern floor. The measurements are referenced to the machine levelling and deep references in the tunnel.

7845 In addition to these measurements, a permanent hydrostatic system has been implemented in the ATLAS support feet bed-plates. It consists of six capacitive sensing stations monitoring the water plane in two 25 m long tubes, 55 mm in diameter, parallel to the beam and linked by a transversal tube. Two additional stations have been installed in the extreme trenches, recognised as stable zones, and linked to the bed-plate system. Altogether, this is equivalent to a reference water plane of 75 m length, inspected by eight sensors attached to the structure (bed-plate and stable floor) within an accuracy of better than 20 microns.

7850 The results from the measurements on the floor, as displayed in Fig. 175, show that, after an initial stable period, during which presumably the heave of the floor was balanced by the loading, a global heave of the floor of up to 1.2 mm happened in the central part of the cavern between March 2004 and March 2006. By August 2006, approximately 85% of the total load had been installed and, over the past 12 months, the cavern floor seems to have settled more or less and the upward heave is no longer visible.

7855 The hydrostatic system in the bed-plates gives immediate movements of the supporting structure with an accuracy of a few microns and has been used to monitor local movements when inserting the calorimeters. It will now be operated continuously to monitor movements in real time during data-taking.

7875 **9.3.2.3 Placement of ATLAS sub-systems** The ATLAS detector had to be assembled in the main cavern mostly because of the nature of the barrel toroid magnet structure, which is 26 m long




Figure 175. Vertical displacements of the ATLAS cavern floor, measured as a function of time in various reference points since August 2003. The beam interaction point is at the origin ($x = 0, z = 0$). The labels used for the reference points are the following: C23-7 represents a point on side C at a distance of 23 m longitudinally from the interaction point and of 7 m radially opposite to the centre of the LHC ring. In contrast, A23+10 represents a point on side A at the same longitudinal distance from the interaction point, but at a distance of 10 m radially towards the centre of the LHC ring. The points labelled B are located right in the middle of the cavern, nominally at $z = 0$.

and 20 m in diameter. As discussed above, once assembled and cabled, the detector and its main components cannot in practice be adjusted any more relative to the nominal beam-line. Careful consideration had therefore to be given to the optimal placement of sub-systems during the assembly process. Many sub-systems were prepared on the surface, while others were prepared in the main cavern, prior to their installation in their final position. The aim was to place all the detectors such that they will all be at their appropriate positions relative to the nominal beam-line, once the installation is completed. In the process of defining the initial placements, the following input had to be taken into account:

- 7885 • floor movement in the ATLAS main cavern (see above);
- deflections of the barrel toroid structures as they were loaded with the muon chambers and the services passing throughout ATLAS;
- deviations of as-built dimensions with respect to the nominal ones for neighbouring sub-systems;
- 7890 • deviations from nominal of the relative placements of different components, which were installed as one assembly. For example, the solenoid and the LAr barrel electromagnetic calorimeter were assembled on the surface inside their common vacuum vessel with a non-negligible relative error with respect to their theoretical placement. When placing the whole assembly in the cavern, the priority was therefore given to the solenoid regarding the relative importance of nominal placement for ATLAS operation;
- 7895 • uncertainties on the position of the interaction point from the uncertainties on the closed-orbit calculations and on the placement of the machine components.

Once an assembly of different sub-systems was completed, the relative position of its components was fixed. As the most prominent example, the barrel calorimeter assembly consists of: the barrel tile calorimeter, itself assembled from 64 individual tile modules, the barrel LAr cryostat, which includes the solenoid, the barrel LAr electromagnetic calorimeter and the support rails for the inner detector, located on the inner bore of the cryostat.

The strategy for the placement of a given assembly was the following:

1. determine the best position for the assembly, taking into account the input described above;
- 7905 2. after the initial placement was completed, survey the assembly as installed and perform adjustments wherever possible to come as close as possible to the optimal position;
- 7910 3. once a major component has been placed, update the envelope drawings to take advantage of the as-installed envelopes of the various components. In most cases, the components respected the assigned envelopes, and it was possible to recuperate some space, which was used either to increase the stay-clear area between assemblies or to optimise further the position of subsequent assemblies. In rarer cases, such as the barrel and end-cap inner detector assemblies, the distance between the assemblies had to be increased by 5 mm with respect to nominal. These deviations are accounted for in the final detector description of the as-installed detector.

Assembly	Component	Placement accuracy (mm)	Average position in y (mm)
Barrel toroid		10 (all coils)	Magnetic axis at -8
Barrel calorimetry	Solenoid	2	Magnetic axis at -2.3
Barrel calorimetry	EM calorimeter (side A)	3	-3.1
Barrel calorimetry	EM calorimeter (side C)	3	-2.3
Barrel calorimetry	Cryostat	2	-0.1
Barrel calorimetry	Tile calorimeter	3	Axis at -0.5
ID pixel		0.2	-1.2
ID barrel	SCT	0.5	-1.3
ID barrel	TRT	0.5	-1.6
ID end-cap (side A)	SCT	0.5	-1.9
ID end-cap (side A)	TRT	0.5	-1.6
ID end-cap (side C)	SCT	0.5	-2.4
ID end-cap (side C)	TRT	0.5	-1.6

Table 53. Placement accuracy and current average position along the vertical y direction of different components of the ATLAS barrel system. The average positions in y are given as an illustration of the priorities set in the placement strategy between conflicting requirements from different components of the barrel system. A more global picture of these current positions in the $x - z$ and $y - z$ planes is shown in Figs. 176 and 177.

7915 Table 53 shows the placement accuracies achieved along the vertical axis for the main components of the ATLAS barrel system relative to the nominal beam-line. Table 53 also shows the measured offsets in the vertical direction, illustrating the priority given to the placement of the solenoid as close as possible to the nominal beam-line and to the subsequent placement of the inner-detector components as close as possible to the nominal solenoid axis.

7920 The placement accuracy along the beam-line is also in the few mm range, but there are exceptions in some cases due to conflicts in envelopes between neighbouring sub-systems. The components of a given assembly are grouped together under the same heading to indicate clearly that their relative position with respect to each other was determined prior to final installation. These relative positions are in many cases more precise than the overall placement accuracy of the assembly, as for example in the case of the SCT and TRT, or even as in the case of the pixels with respect to the ID barrel, for which special care was taken to adjust the fixation points between the two sub-systems on the surface after assembly of the complete ID barrel in order to ensure that the geometrical axis of the pixel sub-system will be within 0.2 mm of the nominal beam-line. It is also important to note that the placement of the ID components with respect to the inner warm vessel of the barrel cryostat has been surveyed with an accuracy of approximately 0.3 mm, which is very similar to that achieved for the survey of the mapping machine used for the mapping of the solenoidal field (see Section 2.2.4).

9.3.2.4 Monitoring of the placement of the detector In a similar way, albeit less complete because of installation work still ongoing in the end-cap regions of the detector, Table 54 summarises

Assembly	Component	Placement accuracy (mm)	Average position in y (mm)
End-cap toroid (side A)		3	Align magnetic axis with barrel toroid
End-cap toroid (side C)		3	Align magnetic axis with barrel toroid
End-cap calorimeter (side A)	FCal	1	Align FCal symmetrically around beam-pipe
End-cap calorimeter (side A)	EMEC	3	-0.8
End-cap calorimeter (side A)	HEC1	3	-0.6
End-cap calorimeter (side A)	HEC2	3	0.5
End-cap calorimeter (side A)	Tile	3	-1.9
End-cap calorimeter (side C)	FCal	1	Align FCal symmetrically around beam-pipe
End-cap calorimeter (side C)	EMEC	3	-0.9
End-cap calorimeter (side C)	HEC1	3	-0.1
End-cap calorimeter (side C)	HEC2	3	-0.3
End-cap calorimeter (side C)	Tile	3	-0.4
Small wheels	MDT/CSC/TGC	4	Align on nominal beam-line
Big wheels	MDT/TGC	4	Align on nominal beam-line
End-wall chambers	MDT	4	Align on nominal beam-line

Table 54. Placement accuracy and current average position along the vertical y direction of different components of the ATLAS end-cap systems.

7935 the positioning status of the end-cap calorimeters, the small wheels, the big wheels, the end-cap
toroid magnets and the end-wall chambers. All the end-cap detectors can in principle be adjusted
relative to the nominal beam-line, when the ATLAS detector is in the open position. In order to
monitor possible small changes of detector positions after opening and closing the apparatus, an
active optical system has been installed, which will provide a precise monitoring ($20 \mu\text{m}$ in the
7940 transverse plane and $100 \mu\text{m}$ along the beam) of the relative position of these components with
respect to the barrel system.

The placement strategy explained above has been quite successful and all major ATLAS com-
ponents are located well within the initial target of being aligned to within a few mm from the
nominal beam-line at the start of data-taking.

7945 More detailed information on the location of individual assemblies is kept in the various
databases and will be updated as movements are monitored over time. As an example, Figs. 176
and 177 show the current deviations of the geometrical axes of the main components of the ATLAS
barrel from the nominal beam-line, respectively in the $x-z$ and $y-z$ planes. Since the full load-
ing of the detector in the barrel region, a larger than expected compression of the support feet has
7950 been observed. This has led to a 1.5 mm downward shift of the whole barrel, as can be seen more
precisely from the average y-positions of the axes of the inner-detector sub-systems in Fig. 177.
All these initial positions are recorded in the overall database and they will be monitored with time
as soon as new survey numbers become available or as the access and maintenance scenarios will




Figure 176. Deviations of geometrical axes of main components of the ATLAS barrel from nominal. Shown are the deviations in mm in the $x-z$ plane.




Figure 177. Deviations of geometrical axes of main components of the ATLAS barrel from nominal. Shown are the deviations in mm in the $y-z$ plane.




Figure 178. Layout of surface buildings and of access shafts to the ATLAS cavern at point 1. The main areas of underground activity are the main cavern (UX15) and the main counting room and service cavern (USA15). The main control room is in building SCX1 on the surface.

require.

7955 It will also be necessary to monitor the global movements of the detector to understand possible future deviations in position between the detector and the actual beam. In the case of such deviations at the level of 1 mm, the beams can be adjusted and steered, using the field in the triplet magnets and/or adjusting the position of the last magnet. In the event of larger movements, it will be necessary to adjust the magnets over the last 300 m on either side of the interaction point. It has
7960 been expected that such an exercise can be done once every three to four years.

9.4 Infrastructure and detector services

This section describes the infrastructure for the surface and underground areas and the services, fixed and mobile, which are connected to the ATLAS detector [255].

9.4.1 Civil engineering

7965 The civil engineering work for the ATLAS experimental area started in November 1997 whilst the previous accelerator (LEP) was still in operation and the situation remained so until the end of 1999. The work included the excavation and concreting of two new shafts, two new large caverns along with the linking galleries, and the erection at the surface of six new buildings, as shown in Fig. 178.

The underground work included the excavation and concreting of the following:

- 7970 • The PX14 and PX16 shafts, respectively 18 m and 12.6m in diameter, both 60 m deep.
- The PX15 shaft was an existing shaft that required concreting only.
- The main cavern UX15 (50 m long, 30 m wide and 35 m high), which houses the ATLAS detector. Due to the continued operation of LEP in 1999, the vault had to be concreted before the support walls were put in place, leading to a non-standard anchoring technique of the 7000 tonne roof of the UX15 cavern. UX15 was delivered to CERN for the installation of the infrastructure in June 2003.
- 7975 • The counting room and service cavern USA15, which houses the electronics racks and services that need to be close to the detector and is accessible during machine runs.
- The seven linking galleries for personnel access and the distribution of services between the UX15 and USA15 caverns.
- 7980

The surface buildings include the following:

- SX1, a steel frame building located on top of the main shafts and housing the travelling cranes, maximum 280 t capacity. Used as short term storage of detector components, prior to their lowering into the main cavern. It also acts as a shelter from the external environment and allows better control of the conditions in the cavern. This building was the first to be structurally achieved allowing the infrastructure works to take place in dry conditions.
- 7985 • SDX1, located on top of the personnel access shaft and used for the personnel and material access control to the underground areas. It also contains the uninterruptible power supplies for the detector services, an electrical sub-station, and the DAQ room.
- 7990 • SUX1, the ventilation building, containing the water chillers and air-conditioning units for the underground areas and the new surface buildings.
- SCX1, the main ATLAS control building, used during the detector assembly phase as an engineering design office.
- 7995 • SH1, a concrete building, containing the noisy cryogenic compressors and other cryogenic equipment.
- SF1, the new cooling towers for the final dissipation of heat recovered from all the underground areas and surface buildings via the water cooling circuits.

9.4.2 Electrical power distribution

8000 The total electrical power required in the ATLAS site at point 1 is about 13 MW, which has resulted in the installation of a new 66 kV sub-station, the reorganisation of the existing 18 kV sub-station, and a new 3.3 kV sub-station. Some 21 new transformers with a total power of 42 MW have



Figure 179. Air ducts installed for ventilation in the shaft and main cavern.

been installed both on the surface and underground to bring the 3 kV or 400 V/230 V power to all ATLAS systems.

There are different power networks to deal with the appropriate levels of continuous and safe operation:

8005

- Assured power: the normal power is backed up by two diesel generators of 1 MW, each located on the surface;
- Secured power: the technical solution is as for assured power, but it is not interrupted when emergency stop buttons are activated;
- Un-interruptible power: these very critical systems are fed by battery back-ups with a total power of approximately 700 kW.

8010

9.4.3 Air-conditioning and cooling systems

The global environmental requirements in the bulk of the ATLAS cavern are driven by general considerations and by more specific ones from the muon spectrometer chambers. The temperature should remain stable at $25\pm 3^\circ\text{C}$ and the humidity should be between 25% and 60%. The cooling and ventilation installation was one of the first infrastructure items to be installed, once the surface buildings had been handed over to CERN, in order to rapidly provide ventilation to the underground areas. This installation included new cooling towers in SF1, the liquid chillers in SUX1, 16 air-handling units in SUX1, SDX1, SX1, and USA15 and kilometres of water pipes and air-ducts. The system is able to cope with emergency cases in fast-extraction mode (see Section 9.9).

8015

8020

One of the most challenging tasks was the installation of air supply and extraction ducts in the PX14 shaft, as shown in Fig. 179, and on the vault of the UX15 cavern, this system having the

System	Medium	Capacity (kW)	Channel count	Operating temperature (°C)
Tile calorimeter	Water	55	24 cooling loops supplying 256 tile fingers	17 to 22
LAr calorimeter	Water	250	24 cooling loops supplying 60 electronic boxes	17 to 22
Diffusion pumps	Water	50	12 cooling loops supplying 26 diffusion pumps	14 to 19
Muon spectrometer and general-purpose stations	Water	300	26 cooling loops supplying CSC's plus racks for small and big wheels	17 to 22
ID evaporative (pixel and SCT)	C ₃ F ₈	60	4 distribution areas supplying 204 cooling channels	-30 to 10
ID mono-phase (TRT)	C ₆ F ₁₄	70	4 distribution areas supplying 176 cooling channels	14 to 22
ID mono-phase (cables)	C ₆ F ₁₄	80	32 distribution manifolds placed all over the detector	14 to 22

Table 55. Overview of main characteristics of the major cooling systems operating in the ATLAS cavern.

task of extracting the 200 kW of heat released into the air of the cavern by the ATLAS detector. The air-conditioning system has been up and running since summer of 2004. Provision has been made for thermal screens to be incorporated inside the detector to prevent a too large temperature gradient across the inner layer of muon chambers.

Table 55 summarises the characteristics of the main cooling systems used for extracting the heat from the detector itself (silicon-sensor leakage currents and ionisation in the TRT gas, as summarised in Table 19), from the on-detector and off-detector electronics, and from dissipation in the power cables. Most cooling systems are leakless by design and the inner detector has chosen fluorinert systems to minimise risks to the detector in case of leaks. In the case of the pixel and SCT sub-systems, a novel and complex evaporative system has been designed and brought into operation to minimise the amount of material devoted to cooling pipes, fluids and connectors on the detector itself. The detector-specific cooling systems are installed in the USA15 side cavern and have been brought into operation in 2005 and 2006.

9.4.4 Gas distribution

The ATLAS detector requires a variety of gases for its normal operation. Table 56 summarises the characteristics of the main gas systems used for normal operation.

The surface gas building SGX1 was designed for the storage, distribution, and mixing of inert and flammable gases in accordance with the CERN regulations. The building was completed in 1997. Large quantities of liquid gases are stored, both inside and outside this building, in particular N₂, CO₂, Ar, Xe, He, C₂H₂F₄, CH₄, C₄H₁₀, and n-C₅H₁₂. The building is also fitted with

System	Gas mixture	Pressure (bar)	Volume (m ³)	Channel count	Impurity limits
TRT active gas (normal operation)	Xe/CO ₂ /O ₂ 70/27/3	1.005	2.5 (detector) 2.0 (gas system)	48 for barrel 14 per end-cap	< 100 ppm CF ₄
TRT active gas (cleaning mode)	Ar/CO ₂ /CF ₄ 70/26/4				< 100 ppm H ₂ O
TRT ventilation (flushed to atm.)	CO ₂ (barrel)	1±0.001	2.9	48	< 100 ppm H ₂ O
TRT cooling (closed loop)	CO ₂ (end-cap)	1±0.001	6.0	4	< 1% N ₂ < 100 ppm H ₂ O
SCT ventilation (flushed to atm.)	N ₂	1.004		4	< 350 ppm H ₂ O
Pixel ventilation (flushed to atm.)	N ₂	1.004		1	< 350 ppm H ₂ O
ID ventilation (flushed to atm.)	CO ₂	1.0005		13	< 1% N ₂ in ID
MDT active gas	Ar/CO ₂ /H ₂ O 93/7/(≤ 1000 ppm)	3 (abs.)	710	112 for barrel 81 per end-cap	< 100 ppm O ₂
CSC active gas	Ar/CO ₂ 80/20	1±0.001	0.5	32	< 100 ppm O ₂
RPC active gas	C ₂ H ₂ F ₄ /Iso-C ₄ H ₁₀ /SF ₆ 94.7/5/0.3	1±0.001	14	128	< 1% O ₂ < 100 ppm H ₂ O
TGC active gas	CO ₂ /n-pentane 55/45	1±0.001	16	128 for active 56 for CO ₂ purge	< 5 ppm at entrance of n-pentane liquefier

Table 56. Overview of main characteristics of the gas systems operating in the ATLAS cavern.

an anti-deflagration roof, mechanical ventilation, permanent gas extraction and gas detection, and an alarm system. The mixing room contains the mixing systems for the different sub-detectors. The TRT, MDT's, CSC's and RPC's use non-flammable gases. Only the muon TGC's will use flammable gas (CO₂/n-pentane).

The underground installation consists of a large network of stainless-steel pipes, which convey the gases to the gas room located in the USA15 cavern and then to the gas racks in the main cavern. From the gas racks, many kilometres of pipes have been installed. They connect to all the different types of muon chambers (see Section 6.8.1) and to the inner detector (see Section 4.2.2).

9.4.5 Cryogenic systems

The ATLAS detector includes two independent systems requiring cryogenic technologies: the superconducting magnets and the liquid argon calorimeters. The cryogenic systems for the magnets and the LAr detectors have each been divided into three parts:

1. External cryogenics, which comprise all the equipment needed to provide the required cooling capacity at given temperature levels, including refrigeration plants and infrastructure. This equipment is located on the surface in the SH1 building and underground in the USA15




Figure 180. Layout of the underground external and proximity cryogenics lines for the LAr calorimeters.

cavern. Also, six large helium gas storage tanks (3 m diameter, 21 m long) have been installed on the surface behind the SX1 building.

- 8060
2. Proximity cryogenics, which comprise all the equipment linking the internal cryogenics to the external cryogenics. This equipment is located in the main cavern, on the steel structure (HS) that surrounds the detector (see Section 9.5).
 3. Internal cryogenics, which comprise all the devices located inside the system concerned (magnets or liquid argon calorimeter).

8065 The cryogenics systems for the magnets are described in some detail in Section 2.1.4.2. This section is devoted to a brief description of the cryogenics systems for the LAr calorimeters.

8070 The primary cooling source for the liquid argon calorimeter installation is a 20 kW nitrogen refrigerator, which operates at 80 K. The compressor station is placed in the SH1 surface building and the cold box in the USA15 side cavern. The high- and low-pressure gas lines connecting these two items pass through the PX-15 shaft. The cold box delivers its cooling power to a 15,000 litre phase-separator dewar placed in the main cavern. Figure 180 shows the underground layout of the proximity and external cryogenics for the LAr calorimeters (shown as if installed alone for

convenience). One can clearly see the fixed cryogenic lines supplying the barrel calorimeter at the top and also the cryogenics lines in the flexible chains, which supply the two end-cap calorimeters and which follow them whenever they have to move for access and maintenance of the detector.

Two 50,000 litre liquid nitrogen storage tanks placed on the surface will supply liquid nitrogen to the phase-separator dewar via a 283-metre long transfer line, in case of problems with the nitrogen refrigerator system. A cryogenic centrifugal pump circulates the liquid nitrogen from the phase-separator dewar through the thirteen heat exchangers placed in the liquid argon cryostats. Each cryostat has been equipped with a valve box, which regulates the mass flow and pressure of the liquid passing through each of the individual heat exchangers. These valve boxes are placed on the HS surrounding structure.

The gaseous nitrogen coming from the heat exchangers is returned to the phase-separator dewar and from there returned to the nitrogen refrigerator system or, in case it is not operational, vented to the surface through a 120-metre long gas line. The three cryostats, placed at the heart of the detector, are linked by large-diameter argon lines to their individual expansion vessel placed on the HS structure. The liquid/gaseous argon boundary of each of the cryostats is located in these expansion vessels.

The need to move the calorimeter end-cap cryostats over a 12 m distance required the implementation of a movement system for the argon and nitrogen lines connecting these cryostats with their expansion vessels and the nitrogen regulation valve boxes. These movement systems are located on the HS structure and are described in Section 9.7.

The 84 m³ of liquid argon present in the cryostats can, in the event of problems, be emptied by gravity into two 50 m³ argon storage tanks placed at the lowest point of the main cavern. A DN500 safety valve line collects any gas coming from the pressure safety valves placed on the cryostats, or storage tank volumes venting it to the surface.

9.4.6 Racks and cables

Prior to the start of the civil engineering work, the detailed cabling and rack needs were not known. The consolidated data from the various sub-systems was provided only much later and the correct provisions for the distribution of cables from the detector to the electronics racks in the counting room and service caverns were therefore made just before the start of the installation work. These included:

- locations for 100 racks in the main cavern, supported from the HS steel structure which surrounds the detector;
- the arrangement of the service cavern and counting room (USA15) with provision for 250 racks on two floors, equipped with a 2.5 MW water-cooling system;
- two main cable distribution galleries, which connect the main cavern to the counting room with provision for 44 cable trays, each 600 × 100 mm² in size;
- the distribution of these 44 cable trays on the HS supporting structure around the perimeter of the detector and to the HO supporting structure on the end-walls of the cavern (see Section 9.5).

- locations for 70 racks in another service cavern (US15) together with about 20 holes, each 300 mm in diameter, through the 2 m-thick connecting wall for the passage of cables, with the associated cable-tray distribution system in the existing false floor and with a dedicated cooling system of 500 kW capacity;

8115

- locations for 100 racks in a self-contained data-acquisition room in the SDX1 surface building with a 500 kW total water-cooling capacity and a dedicated 100 kW air-conditioning system.

The various electronics units for the detector are thus installed in racks, implanted in USA15, UX15 and US15 in the underground areas, and in the data-acquisition room on the surface. The total number of racks (electronics, gas and water) for ATLAS amounts to approximately 500. The Rack Wizard tool mentioned in Section 9.2 is essential to monitor the evolution of the racks with time. The power requirements, specific contents and the connections are constantly updated as they evolve: this is required not only for the maintenance of the detector and the understanding of its evolution on the long term, but also to meet (Installation Nucléaire de Base) regulations (see Section 9.9).

8120

8125

8130




Figure 181. Quantities of cable and flexible-pipe bundles installed by the cabling team.

It was necessary to design new types of racks with respect to what existed in the previous projects, due to the increased power consumption, which therefore required more cooling capacity. The rack cooling is provided by turbine units located at the top of the racks, which push the air down in ducts at the sides of the racks. A deflector at the bottom directs the air upwards through the equipment to be cooled, which is interspersed with heat exchangers. The turbine units comprise an ELMB-based monitoring system (temperatures, humidity, air flow, etc.), which is supervised by the common infrastructure control of the DCS (see Section 8.5). It was also found useful to increase the width of some racks. All these requirements have led in the end to the production of four types of racks, from 46 to 57 units in height, from 900 to 1000 mm in depth, and 600 mm in width.

8135

8140

Most of the installed cable trays are made of stainless steel to minimise the perturbations to the magnetic field (see Section 2.2.2). A total of about 50,000 cable bundles, 3000 flexible pipes and 3500 metallic pipes (see Fig. 181) were installed over a period of two years (May 2005 to May 2007). Many additional proximity cables were installed by the individual sub-systems. Space had to be found to route the large quantity of cables and pipes of the inner detector and barrel calorimeter systems through the muon spectrometer; this was accomplished by routing most of them radially outwards at $z = 0$ and at fixed azimuthal locations, as illustrated in Fig. 182. All the relevant cable and pipe data are stored in the cable database and constantly updated for the same reasons as those described for the racks above.

8145

8150




Figure 182. Detailed three-dimensional layout and routing of cables and services for the ATLAS barrel system. The three flexible chains for the end-cap calorimeters can be seen in the horizontal plane (right) on the side where the end-cap calorimeter trigger cables can reach the main service cavern (USA15) along the shortest possible path, and at 45° below the horizontal plane. One also sees the cryogenic lines for liquid argon at the top and bottom of the drawing. The inner barrel muon chambers in the central region are shown. One clearly sees the holes in the acceptance caused by the considerable volume of services exiting the detector at $z \sim 0$.

9.4.7 Drag-chains and mobile services

Many of the ATLAS sub-systems must move away from the run position to allow access into the detector. As well, the end-cap calorimeters need to remain in a cold bath of liquid argon for the duration of the experiment. The end-cap toroids are cooled with liquid helium and these are also to remain cold during the movement to avoid the lengthy cool-down and warm-up periods (20 to 40 days).

In order to satisfy the above requirements it was decided to use so-called drag-chains, which allow for the services to be supported in a flexible structure. They are used as follows:

- For each end-cap calorimeter there are three chains, as shown in Figs. 182 and 183, non-standard commercial products, each around 30 m-long with parts in stainless steel, specifically developed for this application. They also have a force-assist system that enables the chain to be pulled back into its stored position when the calorimeter is being closed. Particularly challenging was the construction of two of the three chains are at 45° downwards with respect to the beam axis.

- For each muon inner layer (small wheel) there are four chains, each about 3 m in length inside the ATLAS detector.
- For each end-cap toroid there are two chains. These are in aluminium and have been custom made. They run unsupported over a 9 m length at 24 m from the floor when the toroids are moved to their fully open position off the beam axis.

8170

9.4.8 Grounding and electromagnetic compatibility

The ATLAS detector consists of many complex components, installed and operating in contiguous volumes, resulting in a large amount of installed equipment with multiple interconnections and shared services. For these reasons, the performance of the detector components could be heavily affected by electromagnetic interference or induced electronic noise, if these issues are not properly taken care of already at the design stage and systematically followed up on during installation. This is particularly important for detector systems with a large dynamic range and with analogue front-ends and/or readout.

8175

8180

8185

8190

8195

An ATLAS policy was developed and adopted already in the design phase [128, 256] to minimise possible electromagnetic interference effects. Apart from safety considerations, one of the main concerns has been the prevention of ground-loop currents which could couple to the signals of the detector systems. The proposed implementation of the ATLAS policy on grounding and electromagnetic compatibility has been put in place during the construction and installation phases [129]. The main guidelines can be summarised in three rules:

all detector systems are electrically isolated, there are no connections to ground other than through the safety network, and there are no connections between different detector systems.

Even though these rules have been implemented and validated early on in the design phase, it has been a challenge to preserve the electrical isolation of certain large systems during installation. For these reasons, several alarm systems [129] have been installed and operated to ensure that the various detector components remain electrically isolated and grounded to single points. This effort was especially relevant for large structures with many connections like the LAr cryostats, which act as a support for other detector elements and are themselves supported by other parts of ATLAS (see Section 5.5.1.1).

8200



Figure 183. The end-cap calorimeter on side A in its fully open position with all three drag-chains and the flexible LAr fill-line connected.




Figure 184. The ATLAS feet and rail system after installation and prior to the installation of the first barrel toroid coil. Also shown are the blue steel surrounding structures (HS and HO), and, in the background, one of the orange HF trucks.

8205 9.5 Support and access structures

9.5.1 Feet and rail system

The feet and rail system is shown in Fig. 184 shortly after their installation in the pit and before the lowering of the first barrel toroid coil. This system is the main support, the back-bone, of the ATLAS detector. It is made of nine pairs of feet, bound by girders that altogether support the two
 8210 bottom coils of the barrel toroid magnet. On top of these feet are two rails, and their supports, on which the central part of ATLAS can slide. The total load that the feet and rail system has to cope with is about 6000 tonnes (of which approximately 1000 tonnes correspond to the barrel toroid, which is only supported by the feet). The feet and rail systems are mounted on bed-plates, which give the detector its 1.24% incline with respect to the cavern floor, an angle which matches the
 8215 inclination of the LHC accelerator tunnel.

Since the two bottom toroid coils are placed inside the feet, there was a strong requirement for the material to be non-magnetic. In addition, the total deformation was to be kept to a minimum,




Figure 185. The blue support structures (HS on the sides and HO at the ends of the main cavern) at the beginning of ATLAS installation. The arches which now connect the two sides of HS at the top of the main cavern were left out at the time for the installation of the barrel toroid. A barrel toroid coil is in the process of being lowered onto its temporary supports (see Section 9.6).

and stresses well below the elastic limit. Low-carbon austenitic stainless steel was chosen for its good mechanical properties and very low magnetic permeability. One of the main technical issues
 8220 has been to produce non-magnetic welds for such a huge number of welded joints (up to 15 tonnes of filler metal in total).

In order to obtain a precise and reproducible geometrical path of the loads during the move-
 ment of the sub-detectors on their air-pad movement systems (see Section 9.7.2), and also to pre-
 serve the integrity of the beam vacuum system (see Section 9.8), the flatness requirement on the
 8225 rails was one millimetre over their total length (more than 25 m), and 0.2 mm over any length of
 one metre. The maximum deflection of the system remains below 1 mm during the movement of
 the loads.

The requirement to preserve maximum acceptance for the muon spectrometer resulted in
 special chambers in the region of the feet and additional chambers alongside the rails (see Sec-
 tion 6.3.2). Numerous improvements in the feet design were introduced to cope with constraints
 8230 from the muon alignment system, with various designs of the muon support rails, as well as with
 the barrel toroid magnet instrumentation and contact surfaces.

9.5.2 Trucks

8235 The so-called HF trucks are normal steel structures, which are placed directly below the two shafts
of the main cavern. They allow for the main components of the detector to be lowered underground
using the surface cranes and remain there temporarily before moving into their final position, either
by using the cavern travelling cranes (as for the barrel toroid magnet coils) or by using the air-pad
movement systems (as for the calorimeters).

8240 These structures also have the role of supporting the end-cap toroid magnets, as well as the
end-cap calorimeters, in the opening sequences of the ATLAS detector (see Section 9.7). During
installation, they also support the forward shielding, whereas they will only support part of it when
it is inserted in its final location. They are therefore able to cope with the 1000 tonnes of maximum
static load from the barrel calorimeter and they have to allow for the translation of the end-cap
toroid magnets and of the forward shielding away from the beam (total weight of 400 tonnes). For
8245 such movements, air-pads will be attached to the base of these structures to allow them to slide on
the cavern floor.

9.5.3 Surrounding structures (HS and HO)

The blue HS structures, which surround the ATLAS detector, as shown in Fig. 185, have dual roles
of providing personnel access to the periphery of the detector and to support all the equipment that
8250 has to be located close to the detector: proximity cryogenics, electronics racks, gas-distribution
racks, electrical switchboards, and services distribution lines (gas, water, coolants, power). These
structures were the most tricky to assemble, since they are very close to the detector in certain
places and had to be assembled in two stages: the large pillars and gangways up to a height of 20 m
were installed at the same time as the HO structures, but the tops of the HS arches were installed
8255 only after the completion of the installation of the barrel toroid magnet and with very little margin
left. The two structures, which span distances of more than 20 m in three dimensions had to match
each other to within 2 cm. The last arch was finally and successfully installed at the end of 2005.

The main role of the HO structures, which are to be found at the two ends of the cavern, are
to support the end-wall muon chamber stations. They also serve as a useful means of assembling
8260 the sectors of the muon big wheels, before they are hung from the rail system. These structures
also serve as viewing platforms for the thousands of visitors to the ATLAS cavern . Approximately
1000 tonnes of normal steel have been used to build these 13-storey-high structures.

9.5.4 Muon barrel access structures

The aluminium access platforms inside the barrel muon spectrometer have several functions:

- 8265 ● permanent access inside the barrel toroid, so that the muon chambers and their service con-
nections can be accessed in a very short intervention (for example, to disconnect the gas
supply to a specific chamber;
- permanent access to the patch-panels of the inner detector (PP2), so that they can also be
accessed during a very short intervention;
- 8270 ● an emergency exit (through sector 1) in case of access to the barrel calorimeter;




Figure 186. One of the assembled TGC big wheels in the ATLAS cavern. The chambers are fixed to an aluminium structure, which was pre-assembled into sectors on the surface and then assembled as a complete wheel in the cavern itself.

- access for the installation of the barrel muon chambers;
- access to the vacuum pumps of the barrel toroid.

9.5.5 Big wheels

The muon spectrometer (see Section 6.1) includes four big moving wheels at each end, each wheel measuring 23 metres in diameter (see Fig. 186). Of the eight wheels in total, six are composed of thin-gap chambers (TGC's) for the muon trigger system and the other two consist of monitored drift-tube chambers (MDT's) to measure precisely the position of the muons. The so-called big wheels comprise aluminium structures which support the muon end-cap chambers. These big wheels resemble bicycle wheels and are made of sectors, which had been pre-assembled on the surface prior to their transport to the cavern, where they were assembled on the end-wall HO structures. Once one of the wheels is completed on the HO structure, it is lifted onto the traction system, which allows it to move longitudinally towards the barrel toroid magnet and reach its final position in the closed configuration of the detector. It is important to note that the big wheels in their final position need to be inclined with a slope of 1.24% with respect to the vertical to account for the angle between the horizontal cavern floor and the inclination of the machine tunnel.

9.6 Detector Installation

The installation of the detector can be sub-divided into six main phases, which are briefly described below. The barrel toroid magnet, once installed, occupies the central region of the cavern leaving the two sides, A and C, for the lowering and assembly of the remaining large detectors and magnets.

8290 9.6.1 Phase 1: infrastructure in the main cavern, feet and rails

The main cavern was handed over to ATLAS in May 2003. The first operation was to install the general infrastructure (metallic structures around the cavern walls, temporary electricity and lighting, ventilation ducts, and the overhead travelling cranes).

8295 With the steel structures installed, the first elements of the ATLAS detector to be brought down were the bed-plates, which were bolted to the concrete cavern floor. After the bed-plates, the stainless steel support feet, 18 in total, were lowered one by one and installed. The main rails were installed and surveyed once positioned on the feet. The feet provide the mechanical support for most of the ATLAS sub-systems, namely the barrel toroid magnet, the calorimeters, the barrel muon chambers, the end-cap toroid magnets, the services and the access structures, amounting to
8300 about 6000 tonnes.

9.6.2 Phase 2: barrel calorimetry and barrel toroid

Side A: barrel toroid The first barrel toroid coil was delivered to point 1 in October 2004. The coil with its weight of 100 tonnes and total length of 25 m, was lifted by the surface crane, tilted with hydraulic winches, lowered, in an inclined orientation, through the 18 m diameter shaft down
8305 into the cavern. It was then turned back to the horizontal orientation, before being lowered onto the temporary supports (see Fig. 185). From there, it was picked up by the two 65 tonnes underground travelling cranes and put into its final position inside the ATLAS feet. Once the coils were in position, the aluminium struts and girders were installed so that the next coil could be attached to them. This process was repeated until the assembly was completed. In parallel with the barrel-
8310 toroid assembly, the first 100 muon barrel chambers were installed in between the struts/girders and the ATLAS feet.

Side C: barrel calorimeter The lower part of the tile calorimeter was lowered in March 2004. Individual tile calorimeter modules were then assembled together, one by one, until 32 of the 64 modules were completed. The LAr barrel calorimeter cryostat was then lowered into this
8315 half-cradle in October 2004, as shown in Fig. 187. The tile module assembly was then continued until the mechanical assembly of the full barrel calorimeter was completed.

Barrel: completion of barrel toroid and calorimeter installation The last aluminium girder was put in place in September 2005, completing the mechanical assembly of the barrel
8320 toroid structure. Then the hydraulic jacks, which were supporting the complete structure during the assembly, were released. At this moment the load was transferred from the external temporary supporting structure (used during the magnet assembly) to the support feet at the bottom. The temporary support structure was then cut and removed to give space for the barrel calorimeter, which was moved inside the barrel toroid in October 2005.




Figure 187. Lowering of the barrel LAr calorimeter down to the cavern in October 2004. The first barrel toroid coil can also be seen on a temporary support platform before it is installed in the cradles of the feet.

9.6.3 Phase 3: end-cap calorimeters and muon barrel chambers

8325 **Side C: end-cap calorimeter** With the barrel calorimeter installed inside the bore of the magnet, the space on side C was now vacant for the assembly of the first end-cap calorimeter. This assembly was very similar to that of the barrel and was finished in January 2006. It was then moved inside the barrel toroid in February 2006, once the installation of the services (pipes, cables, etc.) was completed.

8330 **Side A: barrel muon chambers and end-cap calorimeter** On side A, the first of the 656 barrel muon chambers was installed in February 2006. When the assembly of the second end-cap calorimeter started in March 2006, it had to be carried out in parallel with the muon-chamber installation. The second end-cap calorimeter was mechanically completed in May 2006. It stayed outside the barrel toroid for a further two months for the installation of the services while
8335 the end-cap calorimeter on side C was moved to its nominal position and the magnetic field of the solenoid was switched on and measured in June 2006 (see Section 2.2.4).

9.6.4 Phase 4: muon big wheels, inner detector and completion of muon barrel

Side C: big wheels In April 2006, work started on the first end-cap muon middle station (often referred to as a big wheel) with the mounting of the tooling on the end-wall structure (HO).




Figure 188. View of installed barrel muon spectrometer stations and end-cap calorimeter on side A.




Figure 189. View of barrel calorimeter and inner-detector end-flange after installation of the first inner-detector end-cap in early June 2007 (left). This was followed shortly thereafter by the installation of the second inner-detector end-cap and of the pixel detector with the central VI section of the vacuum pipe (right).

8340 The first sector was installed in July 2006. Work progressed with an average rate of two sectors per week and this first wheel was mechanically completed in September 2006. After installing the services, the wheel was released from the end-wall structures and moved against the barrel magnet in November 2006.

In March 2007, the second of four big
8345 wheels was completed and the first wheel was
then opened for the lowering of the inner de-
tector end-cap C. Also all the remaining barrel
muon chambers were installed before closing
8350 the end of the barrel on side C with the com-
pleted big wheels.

Side A: big wheels After finishing the
solenoid field mapping, the barrel section of
the inner detector was lowered and installed
8355 inside the bore of the barrel cryostat in Au-
gust 2006. While the work on the connections of
the inner detector services continued, the end-
cap calorimeter was moved partially inside to
allow space for the completion of the muon bar-
8360 rel chambers. By the end of December 2006,
600 chambers or 90% of the total, had been in-
stalled (see Fig. 188).

In January 2007, the preparations for the
muon big-wheel assembly started. The first sec-
8365 tor was installed in March 2007, because the
end-cap calorimeter needed to be moved to the
open position to allow the lowering of the first
inner detector end-cap.

8370 **Barrel** The installation of the barrel
muon chambers continued in parallel with the assembly of the first muon big wheels. The in-
stallation of the services for the inner detector, calorimeters and muon chambers also continued.
In May and June 2007, the two inner detector end-caps and the pixel detector together with the
central VI section of the beam-pipe (see Section 9.8) were lowered into the pit and installed, as
8375 shown in Fig. 189.

9.6.5 Phase 5: end-cap toroid magnets and muon small wheels

The two end-cap toroids were lowered onto the trucks in June-July 2007, as illustrated in Fig. 190.
The muon end-cap small wheels are in the process of being assembled to the shielding disks on the
surface (see Chapter 11).

8380 9.6.6 Phase 6: beam-pipe and forward shielding

The last elements to be installed will be the beam-pipe and the forward shielding, and this will
require that all the sub-systems are progressively moved into their closed positions along the beam
axis. There will be no access to regions in between the sub-systems.



Figure 190. Lowering of the first end-cap toroid magnet onto the truck on side A in June 2007. One of the TGC big wheels can be seen on the right of the picture.

9.7 Access and detector opening

8385 9.7.1 Access scenarios

Three access scenarios have been defined, depending on the duration of the shutdown period and the degree of dismantling of the detector. These can be characterised as follows:

- 8390 1. Very short accesses are typically of the order of a few hours. Such accesses can be provided immediately after the machine shut-down. They can happen on a daily basis, but are not scheduled. As a consequence, no detector components are moved and the access shaft to the surface is not opened (there is therefore no crane access through the shaft). All magnetic fields stay on.
- 8395 2. Short accesses have a duration from a few weeks to five months. The shorter ones will be based on the needs of the ATLAS sub-systems. In agreement with the other sub-systems, the other LHC experiments and the LHC machine, such accesses can be provided for a short period. Short access is also considered as the standard configuration during the annual LHC shut-down for a period of approximately five months.

8400 During such accesses, the cavern shaft is opened so that crane access to the surface is possible. The removable elements of the forward shielding (see Section 3.2) are brought up to the surface, while the muon big wheels, the end-cap toroids, the small wheels and the end-cap calorimeters can be moved along the beam axis. The beam-pipe is left in place, but at atmospheric pressure, and flushed with very pure neon gas (see Section 9.8). All magnetic fields are turned off. A maximum of ten persons are allowed inside at each end of the detector.

- 8405 3. Long accesses are dedicated to the inner detector and small-wheel removal and installation. Such accesses are also for non-standard interventions, which require a break of the beam-pipe. Their duration is the same as that of the LHC annual shut-down (of the order of five months), but their frequency is expected to be much lower and will be related to requests of the experiment for a detector upgrade or for a major maintenance operation. In contrast to short accesses, the beam-pipe is dismantled and one of the end-cap toroids is moved sideways. A second truck is installed along the axis of the detector in order to move back the corresponding small wheel and lift it to the surface. The corresponding end-cap calorimeter is moved back so that sufficient access is possible to the inner detector. All magnetic fields are turned off. The number of people allowed access is defined according to the evacuation plan of the cavern and the detailed operations which need to be performed.

8415 Given the high levels of induced radioactivity expected in the regions of the detector closest to the beams, as discussed in Section 3.5, strict access control and compliance with regulations as laid out in Section 9.9 will be of paramount importance during access to any part of the main cavern.

9.7.2 Movement system

8420 During access, a number of sub-systems move into their position on air-pads: the end-cap toroid magnets, the shielding disks (small wheels) and the end-cap calorimeters. The equipment for each detector movement system is basically the same: in the closed configuration, the detectors rest

on hydraulic cylinders called blocking jacks. They are equipped with nuts so that the load can be transferred to solid feet, without the need for oil pressure. During movement, the load is transferred from the blocking jacks to the air-pads, which consist of two main components: a rubber air-skirt, 8425 which allows the lifting of the detector on a thin film of air, and a hydraulic jack, which allows for the height to be adjusted to a set limit during the movement. Thus, the detector can slide on its rails using the air-pad system with a low friction factor of 0.01. The number of air-pads underneath a sub-system will depend on its weight. They are grouped so that the load is supported by three iso-static points. The movement itself is provided by two hydraulic cylinders, parallel to the rails, 8430 and the detectors are moved step by step according to the stroke of the cylinders.

Because of the sensitivity of the detectors to vibrations, shocks, or tilt, the movement must be smooth and well controlled. Moreover, the clearance between detectors and the beam-pipe is only about 15 mm, a distance of similar size to that of the air-pad lift. Therefore a compensation of the pneumatic action has been implemented, so that the sub-system under air-lift is not raised by 8435 more than 5 mm. Four height sensors, located on each mobile sub-system, provide feedback to the controller, which drives the hydraulic valves of the air-pads.

The movement of the sub-systems is further complicated because of the services connected to them through the drag-chains, as described in Section 9.4. Some of these chains are equipped with their own movement system, therefore it is necessary to monitor these movements with respect to 8440 those of the main movement system.

9.8 Beam-pipe

The beam vacuum system represents the main interface between the experiment and the LHC machine. It must therefore fulfil a dual set of requirements:

- the ATLAS requirements, particularly excellent transparency to particles, limited beam-gas 8445 backgrounds and conformity with environmental constraints, in terms of radiation, electro-magnetic noise and thermal behaviour;
- the accelerator requirements, namely safe operation of the machine, adequate beam aperture and static and dynamic vacuum conditions compatible with the ultimate LHC performance.

The ATLAS beam vacuum system consists of seven beam-pipes of 38 m total length, spanning 8450 the distance between the two TAS collimators located at each end of the cavern. They are bolted together with flanges to form an ultra-high vacuum system, which can be fully baked out in situ. The central chamber, called vacuum inner detector (VI), is centred about the interaction point. It has an inner diameter of 58 mm and is constructed of beryllium metal with a thickness of 0.8 mm (see Fig. 61). The remaining six chambers are installed symmetrically on both sides 8455 of the interaction point and named after the detector, which supports them: VA (vacuum argon end-cap), VT (vacuum toroid end-cap) and VJ (vacuum forward shielding). They are constructed of thin-walled stainless steel tubes with diameters increasing progressively from 60 mm to 80 mm and finally to 120 mm. Chambers inside different detectors are mechanically decoupled by vacuum bellows, which also serve to absorb thermal expansion during bakeout.

8460 The VI chamber was integrated into the pixel detector on the surface, and installed as part of the pixel package (see Section 4.8.1). It is aligned on the beam axis using a system of laser and

CCD cameras, which measure the chamber deformation. The VA chambers are centred inside the warm bore of the LAr end-cap cryostats by sliding supports, which allow the detector to move longitudinally along the beam-pipe. Special minimised ultra-high-vacuum flanges, with only 35%
8465 of the volume of a standard flange, have been developed to pass through the bore. The VT chambers are held by retractable jack supports on rails in the forward shielding. These can be adjusted from the back-face of the end-cap toroid or fully retracted to allow the end-cap toroids to move longitudinally along the beam-pipe. The VJ chambers are cantilevered from the forward shielding located on the cavern wall, inside a conical support designed to fit inside the opened end-cap
8470 toroid. The flanges between the VJ chambers and the TAS collimators are remotely actuated from the outside of the forward shielding, because of the high activation expected in this region at design luminosity.

This supporting system is conceived to allow ATLAS to rapidly move to a short access without the need to open the beam vacuum to air and hence re-activate the Non-Evaporable Getter (NEG) system (see below). However, the chambers are not able to support the stresses induced by offsets expected during opening whilst under vacuum. The chambers will therefore be vented to neon gas at atmospheric pressure, purified to the ppb level by a specially developed gas-purifying system mounted on side A of the HO structure. Neon is not pumped by the NEG system, so the beam vacuum system can be rapidly made operational at the end of a short intervention by simply
8480 re-pumping the neon gas.

The main pump used to eliminate desorbed gasses in the system is a non-evaporable getter (NEG) film sputtered onto the whole of the inner surface of the beam-pipe. After activation by heating the beam-pipe to $\sim 200^\circ\text{C}$, this NEG film gives a very high distributed pumping speed for chemically active gasses. Chemically-inert gasses not pumped by the NEG system are removed
8485 by two minimised sputter-ion pumps [257] at ± 3.8 m and by larger pumps at ± 19 m from the interaction point.

The whole length of the vacuum system is permanently equipped with a mass-minimised system of heaters, thermocouples and insulation which allow the NEG system to be re-activated annually. This bakeout system consists of polyimide foil heaters wrapped with silica aerogel, polyimide
8490 tape and foil for the chambers. Flexible bellows, pumps and transitions are equipped with semi-permanent flexible heating jackets.

Significant optimisation of the forward beam-pipe chambers is planned for the LHC machine upgrade, as discussed in Section 3.5. Stainless steel will be replaced by aluminium or other low-Z materials wherever possible to minimise both the background radiation in the muon chambers and
8495 access problems due to beam-pipe activation.

9.9 Safety in ATLAS

The safety responsibilities for ATLAS include the safety of the personnel as well as the protection of the environment, equipment and infrastructure during the installation and the various phases of operation of the detector (data-taking, access and maintenance).

8500 The main risks are located in the underground experimental area, especially in the main cavern and the adjacent technical-service caverns. A risk assessment of these areas has been performed prior to the beginning of the installation. This was continuously revised and updated during installation and commissioning. The main risks are human operational errors, fire, cryogenic-fluid leaks,

and radiation during beam operation. There are also dangers linked to the presence of magnetic fields, electrical hazards, laser beams, flammable gases and CO₂ gas. Other risks are related to the mechanical integrity of the detector components, in the case of major incidents or even of seismic events.

Potential risks, pertaining to the installation process of the various components, as well as to all operations of opening and closing of the detector during the shut-down periods, have received special attention. These risks are associated with the difficulties related to working at heights, to multiple parallel activities carried out by various working teams, which have to share the same working space, and to the manipulation tools for heavy objects. In order to minimise such risks, actions are taken at various levels:

- a safety organisation has been established in the experimental area and is enforced with an effective in situ presence;
- all activities are managed via the concept of work packages. Each activity is prepared, described and analysed before work can commence. All safety issues are discussed, and tasks optimised as appropriate to minimise risks;
- access to the underground areas is restricted to specialised and trained personnel;
- safety aspects are considered from the early design phase of the equipment and infrastructure, through all the installation and commissioning phases. For example, the barrel toroid coils have been equipped with surface-mounted heaters to warm the eight magnet cryostats and thus prevent condensation and ice formation in the event of a vacuum loss of the magnet system;
- safety systems have been designed and implemented to detect at a very early stage any possible sources of danger and to activate alarms and trigger the required safety actions;
- all alarm informations concerning underground safety and access are collected and managed in the ATLAS control room by the Shift Leader In Matters Of Safety (SLIMOS). This person acts in real time, a necessary condition to guarantee the highest level of safety for all personnel and equipment.
- specialised safety courses are required for all personnel working underground;
- dedicated courses for people doing specialised work such as electrical power, etc.

9.9.1 Organisation of safety

The ATLAS safety organisation is led by the GLIMOS (Group Leader In Matters Of Safety). The GLIMOS supervises the various activities, the specialised safety officers and the territorial safety officers, who are responsible for the safe operation of the underground areas and surface buildings. The specialised safety personnel includes officers for radio-protection, cryogenics, lasers, flammable gases, and electrical hazards. The territorial safety officers are responsible for the safety of the buildings and underground areas around the ATLAS site. Their duty, in particular for the underground area, is to ensure daily safety controls and visual inspections and to take

appropriate actions where required. For the main cavern, given the size and the complexity of the work during installation or access, the territorial safety officer leads a team of technicians.

8545 There is also an external safety coordinator, who leads a small independent team to verify the safety-condition levels inside the experimental area. This team has been active during the construction phase and will be kept operational during the access and maintenance periods. This group is reinforced by a team of engineers, who are in charge of supervising the installation, the commissioning and the maintenance of the various safety systems (see Section 9.9.3).

8550 From the beginning of the LHC operation, an additional safety organisation will be put in place around the SLIMOS in the ATLAS control room. The SLIMOS will be continuously on duty, as described in Section 9.9.5.

8555 The work packages for the underground activities are agreed upon and are integrated into the general planning to minimise overlap of work and resolve potential conflicts. These work packages cover all activities, from infrastructure or detector installation, to commissioning or maintenance work, and to the movements of heavy objects. A work package is declared active only when all crucial technical and safety issues have been reviewed and agreed upon.

9.9.2 Access control

8560 **9.9.2.1 General aspects** Access to the underground areas is restricted to persons who participate in an ongoing declared activity (work package), are authorised and have completed specific safety-training courses. These cover, in addition to the standard general safety training, specific training associated with the hazards which may be encountered in ATLAS: evacuation of the underground areas, cryogenic risks, hazards associated with static magnetic fields, radio-protection, electrical hazards, and handling and removal of equipment inside the caverns.

8565 The control of the access authorisation and the verification of the training and personal biometrical parameters are performed by the LHC access control and safety systems. Personnel and material access control devices are implemented at the top of the lifts and at the entry points of the ATLAS main cavern. In addition to these checks, the access system of the main cavern (UX15) will deliver to each person a safety token during controlled accesses.

8570 **9.9.2.2 INB regulations** By a convention signed in 2000, CERN and the French nuclear authorities have agreed to apply the INB (Installation Nucléaire de Base) rules and regulations to the LHC machine and experiments. These rules and regulations govern and impose stringent limitations on the operation, maintenance and future dismantling and disposal of the ATLAS detector. They are written down in two documents, the Règles Provisoires de Sûreté and the Règles Générales d'Exploitation. In particular, they define yearly integrated dose-rate limits and assign specific labels to different regions of the detector depending on the induced activation.

8575 For what concerns the long term and in particular the final disposal of the ATLAS detector, the regions of ATLAS closest to the beams have already been classified as radioactive, whereas regions further away from the beams will remain classified as conventional. This is based on calculations using as input a scenario corresponding to ten years of operation and two years of cool-down.

8580 Detailed rules of operation are therefore required, in particular for managing the flow and traceability of equipment and materials to and from the experiment. A first document [258] on radiological controls of material from the main cavern describes the control procedures which

ATLAS will put in place. All equipment leaving the cavern will be measured for radioactivity and tracked.

9.9.3 Safety systems

8585 Following the various risk assessments related to the underground work environment and especially the ones concerning fire and cryogenic leaks, a number of dedicated safety systems have been implemented under the direct supervision of the ATLAS GLIMOS and of the CERN Safety Commission. These safety systems have been designed and implemented so as to detect at a very early stage any event which might endanger the safety of personnel, environment or ATLAS equipment. The readout of most of these systems uses the standard DCS tools described in Section 8.5 (ELMB, CANbus and basic communication software).

9.9.3.1 Hazard-detection systems The main and service caverns are equipped with standard detectors, which detect the presence of smoke inside the infrastructure and service areas of the caverns. The electronics racks have been equipped with smoke-detection points and some of them 8595 with an associated CO₂ gas extinguisher system.

Due to the large quantity of liquid argon (84 m³) in the three LAr cryostats, which might fill a large part of the main cavern in only a few minutes with an asphyxiating gas in case of a catastrophic failure, three large trenches have been built in the floor of the cavern. In case of a major leak, the cryogenic liquids and the cold and heavy gases would be contained in these trenches. 8600 Access is restricted to these areas and there is an oxygen-deficiency detection system installed. In normal conditions, air is permanently extracted from the lowest point of these trenches. If a leak is detected, the gas extraction can be increased to a massive rate of 32,000 m³/h.

The TGC's in the small and big wheels are filled with a flammable gas mixture (see Section 6.7 and Table 56). Their distribution racks have therefore been equipped with flammable-gas 8605 detection heads. The internal areas of the ATLAS detector are equipped with air-sampling tubes or sniffers, which may detect the presence of smoke, CO₂ or flammable gases, and hence a subsequent deficiency of oxygen. These tubes run on the inside of the various sub-systems and along the detector platforms. They serve as a protection of equipment and personnel working inside the ATLAS detector.

8610 The barrel toroid warm structure, which supports the eight barrel toroid coils, is made of aluminium. In case of a major fire inside the ATLAS detector, the aluminium will begin to lose part of its structural properties at a rather low temperature of approximately 200°C. In order to minimise the risk of any mechanical-instability problem of the toroid warm structure, temperature sensors are fixed on these aluminium parts. These send the temperature information to the ATLAS 8615 SLIMOS desk in the control room.

The various safety systems, fixed detection systems and sniffers, generate alarms. Two different alarm-threshold values are defined for each type of detection. The first threshold generates a warning and triggers preventive actions on the ATLAS detectors via the Detector Safety System (DSS). The second threshold indicates that there is a serious danger to the personnel or the 8620 environment, which requires the immediate intervention of the fire brigade. In addition, this second threshold also triggers the evacuation from the ATLAS underground areas and immediate actions on

the detector via DSS and on the infrastructure (for example, modification of the cavern-ventilation configuration as described above in the case of a cryogenic leak).

9.9.3.2 Foam extinguishing system In addition to standard fire-fighting means, such as portable
8625 fire extinguishers and hose reels, a foam extinguishing system has been implemented in the vault
of the cavern. This foam system may be used in the extreme case, to protect the detector and
the CERN firemen in the event of a fire getting out of control. The system consists of 12 large
blowers installed in the vault of the UX15 cavern which are fed by a mixture of water and detergent
and can fill-up the cavern in less than 15 minutes, suffocating any fire. Since this foam has only a
8630 1/1000 water content, personnel trapped in the foam would survive without problems until the foam
settles (approximately one hour). Tests have also demonstrated that the foam does not penetrate
into electronics racks.

9.9.3.3 Finding people inside ATLAS The FPIAA system (Finding People Inside ATLAS Ar-
eas) detects the presence of persons in all areas of the ATLAS main cavern, including those inside
8635 the detector itself. This system does not require any special device to be worn by the personnel.
It is based on the use of approximately 500 passive infra-red sensors, appropriately modified to
be radiation-tolerant and to operate in a magnetic field. Each one of the 500 small volumes in
the cavern and inside ATLAS is continuously monitored: if a person were to disappear without
reappearing in the adjacent volume, an alarm would be generated.

8640 **9.9.4 Detector safety system**

The Detector Safety System, DSS [259], is the central tool to bring (parts of) the ATLAS detector in
a safe state in cases when an abnormal operational situation arises or a safety hazard is detected. Its
main task is to protect the detector equipment. DSS works ATLAS-wide, i.e. across sub-detector
boundaries and including all common infrastructure components of the ATLAS detector. It has its
8645 own sensors to detect potentially dangerous situations (e.g. over-temperatures) and receives input
from the hazard detection systems described in Section 9.9.3. This information is collected by
DSS stations distributed over the different counting rooms and is analysed centrally by a redundant
system based on programmable logic controllers. In a matrix-like fashion, all input signals can be
combined by logic operations to trigger the appropriate action, usually a shut-down procedure of
8650 the relevant equipment. This process is fully automatic: operator intervention is only needed to
analyse and correct the fault and to bring the detector back into operation. Care has been taken
when implementing the DSS system not to rely on external services, such as computer networks or
normal electricity supplies. The principle of positive safety has been used throughout, i.e. in case of
missing sensor information or possible internal system problems of DSS, all relevant safety actions
8655 are executed. A dedicated operator interface in the ATLAS control room provides the SLIMOS
with the detailed status of the DSS at all times.

9.9.5 Safety during operation

As described above, the safety organisation and access control will be coordinated during operation
around the SLIMOS desk in the ATLAS control room. Responsibility for access control will nor-
8660 mally be transferred to the SLIMOS from the central LHC control room (CCC). The SLIMOS will

Machine operation configuration	Nominal	75 ns	Roman pots
Number of bunches	2808	936	43
Number of protons per bunch (10^{11})	1.15	0.9	0.1
Bunch spacing (ns)	25	75	2025
β function at the interaction point (m)	0.55	1–11	2625
Crossing angle (μrad)	285	250	0
Peak luminosity ($\text{cm}^{-2}\text{s}^{-1}$)	10^{34}	10^{33}	10^{28}

Table 57. Main operational parameters of the LHC machine for a few configurations: the nominal one (left), the initial one with a bunch spacing of 75 ns (centre), and the specialised one for the measurement of the total cross-section (right).

be in charge of controlling in real time the safety conditions inside the cavern via the various safety systems described above. The SLIMOS will also be responsible for providing information to the fire-brigade on the status of the main cavern and the detector, including: beam status, configuration of the detector, detailed instructions for accessing the region of intervention, number of people in the underground areas, radiation levels and environmental conditions, relevant information concerning the status of the ATLAS detector and the infrastructure (cooling, cryogenics, magnets), and status of all possible safety alarms.

9.10 Interface to the LHC machine

For safe and optimal operation of both the LHC machine and the ATLAS detector, the two parties will continuously exchange information about their overall status as well as about the status of relevant individual sub-systems. This data exchange will be used to synchronise actions during the different states of operation, to provide online feedback on tuning operations, to rapidly react to errors, and to understand quickly and efficiently their causes.

ATLAS and the LHC machine exchange most data over the network through the DCS information server (see Section 8.5). In addition, dedicated hardware links are used for critical signals that have to be transmitted on time and in a reliable fashion, such as the beam permission signals and timing signals.

The LHC communicates to ATLAS the total beam and individual bunch intensities, the average 2-dimensional beam size, the average bunch length, the luminosity at the four interaction points, the average beam loss, and the average horizontal and vertical beam positions. Table 57 lists basic beam properties for some of the interesting configurations envisaged for machine operation.

ATLAS reports to the LHC information that allows the machine to optimise and monitor the conditions of the beams, in particular the quality of collisions and machine-induced backgrounds. Experience from previous colliders shows that the machine-induced background in the detectors is very hard to predict. A number of different factors intervene in a complex manner:

- the local vacuum pressure as well as the vacuum at more distant places, such as the arcs, affects the halo entering the detector;
- inefficiencies of both the betatron and momentum-cleaning systems and the detailed settings of the collimators will also heavily influence the observed background levels;
- 8690 • other factors, which have a direct impact on the beam halo, are of course the total beam current, the beam tune shift and the orbit positions.

It is therefore of prime importance to the experiment to define reliable background indicators and to communicate them to the LHC control room. These background indicators must be continuously available to the operating crew for monitoring, in particular before stable beam conditions have been reached during the setting-up phase of the machine. They must therefore be available in the experiment independently of the main data acquisition. The ATLAS beam conditions monitor (or BCM as described in Section 3.4.1) meets these requirements and will be used in this context.

Among the parameters that ATLAS sends to the LHC are: the total luminosity, the luminosity per bunch, indicators for quality of collisions and amount of machine-induced backgrounds, counting rates for individual bunches, and the position and size of the luminous region.

The 40 MHz bunch clock of the LHC and a pulse per revolution is transmitted from the LHC radio-frequency system at point 4 to ATLAS over a total length of 14 km of optical fibre. Once received in the ATLAS counting room, these signals are fine-adjusted in phase and distributed via the L1 central trigger processor to all ATLAS sub-systems (see Section 8.2.3).

ATLAS receives for each beam one signal from a beam position monitor, which is located 175 m upstream of ATLAS. These signals provide a precise timing reference in order to monitor the phase of the LHC clock with respect to the bunches. In addition, they serve as inputs to the L1 trigger, for which they provide a filled bunch trigger signal for each beam and a time reference with respect to the abort gap of the LHC bunch train.

When in operation, the LHC machine undergoes a sequence of operational modes such as filling, ramping, adjust, stable beams, and unstable beams. The current machine operational mode is received by ATLAS via software, which is appropriate in most cases for synchronising ATLAS operation with LHC operation. Before a state transition, a hand-shake protocol between the LHC and ATLAS is used: the LHC operators request from ATLAS confirmation before going into e.g. the state (adjust mode), where the low- β squeeze and other adjustments take place. A similar protocol is used before a scheduled beam dump by the LHC operator.

A fail-safe and reliable beam interlock system is installed around the LHC ring, with several systems giving permission for beams. The absence of a beam permission signal leads to an immediate beam dump: the safe extraction of the beam from the LHC in less than 300 μ s. The ATLAS beam interlock system (BIS) consists of three parts, each of which gives beam permission: the detector BIS, the spectrometer magnet BIS and the Roman-pot position BIS. The detector BIS takes inputs from the BCM and possibly other detectors and gives beam permission only when background conditions allow safe operation of the detector.

Additional flags related to the machine modes are transmitted from the LHC to ATLAS through a fast, safe and reliable hardware link, as they are used in the context of the movement control of

the Roman pots and for the Roman-pot position BIS (see Section 7.2). The ATLAS BIS is complemented by additional interlocks, for instance to inhibit injection into the LHC and to apply more sophisticated logic. The system is flexible enough so that it can evolve with the experience obtained in the operation of the LHC and ATLAS.

8730 **10. Expected performance of the ATLAS detector**

10.1 Introduction

Since the publication of the ATLAS Detector and Physics Performance Technical Design Report [260] in 1999, all the detector components of the experiment have been constructed and integrated and most of them have been installed (see Section 9.6). A detailed understanding of their features (geometry, amount of material and placement accuracy) has been achieved over this
8735 period, as described elsewhere in this article.

The purpose of this chapter is to present an overview of the main performance features of the ATLAS experiment, as expected today from the latest round of simulations and the current version of the reconstruction software and as validated wherever possible using test-beam measurements. It is therefore a snapshot of the present understanding of the performance of the detector. Somewhat
8740 in contrast to earlier documents, in particular the Detector and Physics Performance TDR [260], the performance results presented here will focus on the initial low-luminosity scenario for the beginning of data-taking at the LHC. Since the luminosity is expected to rise over the first year or so from $10^{31} \text{ cm}^{-2} \text{ s}^{-1}$ to $10^{33} \text{ cm}^{-2} \text{ s}^{-1}$, most of the results presented below will correspond
8745 to simulated events without pile-up nor cavern background, except in certain explicit cases where their contributions at luminosities of $\approx 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ have been considered.

The first two sections are devoted to the expected tracking performance in ATLAS and its powerful but complex magnet system (see Chapter 2). The overall expected performance of the inner detector is described in Section 10.2, while that of the muon spectrometer, both stand-alone and
8750 combined with the inner detector, is presented in Section 10.3. Sections 10.4 (electrons and photons), 10.5 (hadronic jets), 10.6 (missing transverse energy), 10.7 (hadronic τ -decays), 10.8 (tagging of heavy flavours) and 10.9 (trigger) describe the expected performance of the overall ATLAS detector to trigger, reconstruct, identify and measure the major final-state objects over the required range of energies for most of the physics channels of interest at the LHC.

8755 **10.1.1 Realistic data challenge**

Over the past seven years, a large and modular suite of software tools for simulation and reconstruction has been developed and integrated into the ATLAS computing model and first full-scale exercises of the operation of this computing model have begun. A large number of high-statistics samples of Monte-Carlo events have been run through the complete ATLAS simulation, reconstruction and analysis chain to assess the readiness of the overall system to cope with the initial
8760 data. Results from this data challenge, in particular from its calibration and alignment component, will be presented wherever relevant in this chapter.

As part of the preparations for initial data-taking, the simulation software has been adapted to describe and simulate, in addition to the ideal ATLAS detector description most commonly
8765 used, an ATLAS experimental set-up with alignment and placement shifts which are similar in size to those anticipated in the real detector. These have been included from macro-assembly to individual module level, as for example in the inner detector. In addition, magnetic field and material distortions have been included wherever relevant, as well as calibration distortions of the electromagnetic and hadronic calorimeters for certain specific studies. The results presented here
8770 are based on many tens of millions of events, originating from a variety of physics processes and

event generators, and with the detector response simulated using GEANT 4 (version G4.7.1.p01 and QGSP GN physics list) [228, 261].

8775 The results published more than eight years ago in Ref. [260] correspond to a detector description which is quite different from the current one. Several real changes happened to the layout of the ATLAS detector:

- the η -coverage of the TRT has been decreased from $|\eta| < 2.5$ to $|\eta| < 2.0$, resulting in a significant loss of performance in that region (momentum resolution, tracking performance and electron identification);
- 8780 • the end-cap cryostats and the extended barrel tile calorimeters have been recessed by 40 mm in z to make room for inner-detector services;
- certain end-cap muon chambers dedicated to momentum measurements in the transition region between the barrel and end-cap toroids have been deferred in terms of construction and installation in ATLAS, resulting in a significant loss in stand-alone performance (efficiency and momentum resolution).

8785 In addition, the description of the installed detector has improved considerably:

- the amount of material in the inner detector and just in front of the electromagnetic calorimeter has increased substantially;
- the amount of material in the muon spectrometer has increased substantially in several areas.

8790 For these reasons, the expected performance is somewhat worse than that published in Ref. [260]. Only the most striking examples can be given in this article:

- 8795 • the 25% degradation in the expected resolution for the invariant mass of four muons from Higgs-boson decay for $m_H = 130$ GeV reconstructed in stand-alone mode (see Fig. 230). This degradation is due in equal proportions to the missing chambers in the transition region between the barrel and end-cap toroids and to the increase of the material in the description of the muon spectrometer;
- the 12% degradation in the expected resolution for the reconstructed invariant mass of four electrons from Higgs-boson decay for $m_H = 130$ GeV (see Fig. 250) and of two photons from Higgs-boson decay for $m_H = 120$ GeV (see Fig. 251);
- the expected degradations in performance are smaller for other channels such as $Z \rightarrow \tau\tau$.

8800 The model of the set-up used for the results presented here differs nevertheless from the reality in the ATLAS cavern, as it has been described in the inner-detector, calorimeter and muon-spectrometer chapters, in several important respects since it had to be frozen for large-scale simulation:

- 8805 • the amount of material in the inner detector has increased in some services regions of the active volume by a few percent of a radiation length, X_0 (at maximum 7% X_0);




Figure 191. Sketch of the ATLAS combined test-beam set-up.

- the amount of material in the inner detector outside the active volume and therefore near to the barrel and end-cap cryostats of the LAr calorimeter has increased by substantial amounts in certain regions: by 3.5% X_0 for $|\eta| < 0.7$, by 35–40% X_0 for $|\eta|$ increasing from 0.8 to 1.1, by 50–80% X_0 for $1.1 < |\eta| < 1.8$ and by 15% X_0 for $1.8 < |\eta| < 2.2$;
- 8810 • the amount of material in the muon spectrometer is larger in certain areas than what has been included in the detector description for the results presented here. The largest missing items are the support structures for the small and big wheels (a few tens of tonnes), the saddle support structures for the barrel calorimetry, the inner-detector PP2 patch-panels, and more generally specific mechanical supports and services throughout the muon spectrometer. The
8815 uncertainties on the knowledge of this extra material will remain large until the installation of the last few macro-components in the pit has been completed.

10.1.2 Combined test-beam

The understanding of the detector components has improved considerably over the many years of construction, especially with extensive measurements in test-beams of the stand-alone and combined performance of the various calorimeters in the H6 and H8 test-beam facilities at CERN. The
8820 main results of these measurements are summarised in Section 5.6.

A dedicated effort to understand the combined performance of a complete slice of the ATLAS detector, from the pixel detectors to the outermost stations of the muon chambers, took place in 2004 with the large-scale combined test-beam (CTB) exercise. Figure 191 shows a sketch of the



Figure 192. Picture of the combined test-beam set-up for the inner detector and the calorimeters. The beam particles come from the left of the picture, traverse the magnet and then hit the calorimeters on the right side of the picture. On the left, just behind the pole tips of the magnet in which the pixel and SCT modules were installed, are the barrel TRT modules. On the yellow rotating support table is the cryostat housing the LAr electromagnetic calorimeter modules and behind it (right side of the picture) are the tile calorimeter barrel (not visible) and extended barrel modules.



Figure 193. Picture of the combined test-beam set-up for the end-cap muon chamber system. The beam particles come from the right side of the picture, traverse the barrel muon chamber set-up, which is mostly hidden by the concrete blocks, and then go through three stations of end-cap MDT and TGC chambers.

8825 layout of the CTB set-up, and Figs. 192 and 193 show respectively pictures of some of the actual
tracking and calorimeter components and of some of the muon chamber components, as they were
operated in 2004.

This effort has led to an improved detector description, and also to first sets of detailed calibra-
tion and alignment procedures, essential to the initial understanding of the detector performance
8830 and to the extraction of the first physics results. The main results obtained from these measure-
ments and from their comparison to the detailed simulation of the detector (used both for the CTB
and for ATLAS itself) are presented in this article:

- the inner-detector alignment results are presented in Section 10.2.2;
- the TRT electron identification results are presented in Section 10.2.5;
- 8835 • the muon-chamber alignment results are presented in Section 10.3.2;
- the electromagnetic calorimeter energy measurement results are presented in Section 5.6
together with all the other test-beam results related to stand-alone and combined calorimeter
performance;

A general consequence of the various combined calorimeter test-beam efforts and of the CTB
8840 data analysis in particular is that the detector description of the barrel electromagnetic calorimeter
and the calibration software of the various calorimeters have been considerably refined to reach
agreement between test-beam data and simulation. These refinements are fully integrated into
the ATLAS software framework for the experiment itself, a necessary condition to the desired
tight coupling between test-beam simulation and data analysis and the actual simulation of physics
8845 collisions in ATLAS.

10.2 Reconstruction and identification of charged particles in the inner detector

Charged particle tracks with transverse momentum $p_T > 0.5$ GeV and $|\eta| < 2.5$ are reconstructed
and measured in the inner detector and the solenoid field. The efficiency at low momentum is,
however, limited because of the large amount of material in the inner detector (see Section 4.10
8850 and Fig. 72). The intrinsic measurement performance expected for each of the inner-detector
sub-systems is described in Section 4.1. This performance has been studied extensively over the
years [61], both before and after irradiation of production modules, and also, more recently, during
the combined test beam (CTB) runs in 2004 as described in Section 10.1, and in a series of cosmic-
ray tests in 2006 as described in Section 4.9. The results have been used to update and validate
8855 the modelling of the detector response in the Monte-Carlo simulation. This section describes the
expected performance of the inner detector in terms of alignment, tracking, vertexing and particle
identification.

10.2.1 Track reconstruction

The inner-detector track reconstruction software [262] follows a modular and flexible software
8860 design, which includes features covering the requirements of both the inner-detector and muon-
spectrometer reconstruction (see Section 10.3 for a description of the strategies used for muon

reconstruction). These features comprise a common event data model [263] and detector description [264], which allow for standardised interfaces to all reconstruction tools, such as e.g. track extrapolation, track fitting including material corrections, and vertex fitting. The extrapolation package combines propagation tools with an accurate and optimised description of the active and passive material of the full detector [265] to allow for material corrections in the reconstruction process. The suite of track-fitting tools includes global- χ^2 and Kalman-filter techniques, and also more specialised fitters, such as dynamic noise adjustment [266], Gaussian-sum filters [267] and deterministic annealing filters [268]. Other common tracking tools are provided, e.g. to apply calibration corrections at later stages of the pattern recognition, to correct for module deformations or to resolve hit-association ambiguities.

Track reconstruction in the inner detector is logically sub-divided into three stages:

1. A pre-processing stage, in which the raw data from the pixel and SCT detectors are converted into clusters and the TRT raw timing information is turned into calibrated drift circles. The SCT clusters are transformed into space-points, using a combination of the cluster information from opposite sides of a SCT module.
2. A track-finding stage, in which different tracking strategies [262, 269], optimised to cover different applications, are implemented. The default tracking exploits the high granularity of the pixel and SCT detectors to find prompt tracks originating from the vicinity of the interaction region. First, track seeds are formed from a combination of space-points in the three pixel layers and the first SCT layer. These seeds are then extended throughout the SCT to form track candidates. Next, these candidates are fitted, outlier clusters are removed, ambiguities in the cluster-to-track association are resolved, and fake tracks are rejected. This is achieved by applying quality cuts, e.g. on the number of associated clusters, with explicit limits set on the number of clusters shared between several tracks and the number of holes per track (a hole is defined as a silicon sensor crossed by a track without generating any associated cluster). The selected tracks are then extended into the TRT to associate drift-circle information in a road around the extrapolation and to resolve the left-right ambiguities. Finally, the extended tracks are refitted with the full information of all three detectors and the quality of the refitted tracks is compared to the silicon-only track candidates and hits on track extensions resulting in bad fits are labelled as outliers (they are kept as part of the track but are not included in the fit).

A complementary track-finding strategy, called back-tracking, searches for unused track segments in the TRT. Such segments are extended into the SCT and pixel detectors to improve the tracking efficiency for secondary tracks from conversions or decays of long-lived particles.

3. A post-processing stage, in which a dedicated vertex finder is used to reconstruct primary vertices. This is followed by algorithms dedicated to the reconstruction of photon conversions and of secondary vertices.

8900 **10.2.2 Alignment of the inner detector**

The alignment of the inner detector is a crucial component in reaching the required tracking performance. The alignment procedure must determine accurately the actual positions in space of the silicon modules (pixel and SCT) as well as of the straws (or groups of straws) in the TRT modules. The task therefore corresponds to the determination of six degrees of freedom for each module, if it is treated as a rigid body. It will also be necessary to correct for imperfections within the modules, due to temperature gradients, module bows or other distortions. To ensure that the misalignment of silicon modules does not inflate the track parameter uncertainties by more than 20% above the intrinsic resolution at high- p_T , the module positions need to be determined with a precision of approximately $10\ \mu\text{m}$ or better in the bending plane [61]. For a precision measurement of the mass of the W -boson, an understanding of the module positions at the level of $1\ \mu\text{m}$ or better is required. The expected as-built and survey precisions of the inner-detector components before data-taking are described in Section 4.3, and their overall placement accuracy inside the inner bore of the barrel LAr cryostat is summarised in Table 20.

Alignment constants for the inner detector will be derived from a dedicated stream of tracks selected at a rate of $\sim 10\ \text{Hz}$, and will be updated if required every 24 hours. To reach a precision of $10\ \mu\text{m}$ on the silicon-module positions, approximately one million good tracks with various topologies will be selected within this 24 hour period and written out to the calibration and alignment stream at the time of the final high-level trigger decision.

Several different track-based alignment techniques have been applied to CTB data, to cosmic-ray data and to Monte-Carlo simulations of a misaligned inner detector. All the approaches are based on the minimisation of hit residuals from high-momentum tracks, which are preferred because of their lower multiple-scattering distortions. The minimisation of track residuals is a necessary but not sufficient requirement for the alignment of the inner detector. Certain global distortions of the geometry may not be or may only weakly be constrained by such tracks (these are called "weak modes"). While preserving the helical trajectory of the track, these modes would, in general, lead to biases on the measured track parameters. Sagitta distortions, which arise from systematic biases in the measurement of q/p_T , where q is the charge of the track, are one of the prominent examples.

To constrain and eliminate these weak modes, it is important to use tracks with different topologies:

- tracks from the interaction point, which will always constitute the bulk of the sample of tracks used for alignment. Using the primary vertex as an additional constraint will help to eliminate certain weak modes;
- cosmic-ray tracks, which have the advantage of providing a continuous helical trajectory across the whole inner detector, thereby mimicking a pair of opposite-sign equal-momentum and back-to-back tracks, when they pass close to the interaction point. In addition, a large fraction of the cosmic-ray tracks will cross the inner detector far from the beam axis, thereby providing additional constraints to eliminate certain weak modes;
- tracks from beam halo will help to constrain the initial alignment of the end-cap regions;




Figure 194. Distributions of pixel (left) and SCT (right) residuals for the most accurate measurement coordinate, as obtained for charged pions with an energy of 100 GeV in the combined test-beam data. The results are shown for tracks reconstructed in the pixel and SCT detectors before (dashed histogram) and after (full histogram) alignment. The curves represent Gaussian fits to the residuals after alignment. Because of the large misalignments of certain modules, most of the entries before alignment lie outside the boundaries of the plots.

- 8940 • tracks passing through the overlap regions of adjacent modules. These constrain the circumference of cylindrical geometries and thus improve the determination of the average radial position of the modules;
- track pairs from Z and J/ψ decays. Fitting these tracks to a common decay vertex and to a known invariant mass will provide sensitivity to systematic correlations between different
- 8945 detector elements;
- finally, additional constraints are provided by the information from survey measurements, which are, however, limited in practice to the relationships between nearby detector elements connected by rigid support structures.

With the unprecedented number of detector modules in the inner detector, the alignment task is

8950 immense in its scope and complexity. With the aim of simplifying it, the alignment procedure can be broken down into several steps. As a first step, the large detector structures (the barrel and the end-caps of each of the three sub-systems) are aligned with respect to each other. By treating these large-scale structures as rigid bodies with only a few degrees of freedom, the procedure converges on a global alignment with only limited statistics of reconstructed tracks. To achieve this goal,

8955 it is planned that sufficient cosmic-ray data be taken before LHC turn-on. In a second step, the individual barrel layers and end-cap disks can be aligned with respect to each other, leading to a system with several hundreds of degrees of freedom. In a third step, the complete alignment of all the detector modules implies resolving a system with almost 36,000 degrees of freedom (1744 pixel modules, 4088 SCT modules and 136 TRT modules) and therefore requires the large

8960 samples of tracks mentioned above. The last step in the whole process requires the study of possible residual biases, using resonances decaying to muons, E/p measurements combining inner detector

and electromagnetic calorimetry, and combined muon measurements with the muon spectrometer (see Section 10.3.2).

10.2.2.1 Alignment in the combined

test-beam The alignment procedure has been applied to CTB data [270] using charged-hadron beams with energies between 5 and 180 GeV. The results obtained are shown in Fig. 194, in the case of one beam energy of 100 GeV, for the pixel and SCT residuals for the most accurate measurement coordinate. The striking improvement observed in the residual distributions after alignment arises from alignment constants changing by typically 100–200 μm for some of the pixel and SCT modules. The measured resolutions after alignment are in agreement with those expected from Monte-Carlo simulation of the CTB set-up with a perfect alignment.




Figure 195. Fractional momentum resolution for pions as a function of pion momentum. The results are compared between data after alignment (see text) and simulation.

Figure 195 compares the measured momentum resolution for pions after alignment with that expected from Monte-Carlo simulation as a function of the pion momentum which ranges from 5 to 100 GeV. The mean of the beam momentum as measured using the pixels and the SCT is correct to $\sim 1\%$ at the highest energy of 100 GeV, indicating that residual misalignments are small. The disagreement observed at the lower end of the momentum spectrum is most likely related to the quality of the data taken with low-energy beams (a mix of electrons and pions taken early in the run when the operation of the pixel and silicon detectors was less stable than for the higher-energy runs).

10.2.2.2 Misaligned inner detector in ATLAS simulation Within the context of the realistic data challenge described in Section 10.1, the inner-detector alignment algorithms are undergoing stringent tests, based on events simulated with a misaligned inner-detector geometry and reconstructed with the nominal geometry. The main focus of these tests is to study the various alignment approaches [271, 272, 273] within a realistic and full-scale scenario typical of what can be expected with initial data. One issue of particular interest is the implementation and validation of robust methods to determine and eliminate the weak modes, especially the sagitta distortions. The misalignments introduced for the realistic data challenge do not, however, cover all possible misalignment scenarios. In particular, systematic deformations of large scale structures like end-cap disks or barrel layers are not simulated. Twists or radial deformations of the barrel layers are known to correspond to weak modes of the alignment.

Nevertheless, a number of systematic displacements and rotations of large and smaller-scale structures were introduced, in addition to the smaller and random misalignments introduced at the module level. This resulted in initially low efficiency for reconstructing certain tracks and in track-parameter distortions of large magnitude (the expected mass peak for $Z \rightarrow \mu\mu$ decays was




Figure 196. Difference between the reconstructed and true mass of dimuon pairs from $Z \rightarrow \mu\mu$ decay. The results are shown in the case of an ideal (perfectly aligned) inner detector (open circles) and for the inner detector after a first-pass alignment (full circles), based on high- p_T muon tracks and cosmic rays (see text).




Figure 197. Asymmetry between negative and positive muons as a function of p_T , as obtained for the sample of $Z \rightarrow \mu\mu$ decays reconstructed in the inner detector and described in the text. The results are shown in the case of an ideal (perfectly aligned) inner detector (open circles) and for the inner detector after a first-pass alignment (full circles).

not initially visible). To converge on the first-pass alignment results presented here, reconstructed tracks were constrained to the beam-line and tracks from simulated cosmic-ray events were also
9005 used to provide additional constraints.

The impact of global sagitta distortions on the reconstructed invariant masses of neutral resonances decaying into oppositely charged particles is in principle only a small effect, since the momenta of the positively and negatively charged daughters are shifted in opposite directions. However, ϕ -dependent sagitta distortions may give rise to larger effects, which can become very
9010 significant at relatively high momentum (depending on the size and systematic nature of the residual misalignments). Figure 196 shows the effect of these residual misalignments on reconstructed $Z \rightarrow \mu\mu$ decays after applying the corrections obtained from a first-pass alignment of the inner detector based on high- p_T muons and cosmic rays. The monitoring of the evolution of the alignment constants during the various stages of this first-pass alignment has shown that residual
9015 distortions on the track parameters remain, even after using cosmic rays to remove some of the weak modes to which tracks originating from the interaction point are not very sensitive. The residual distortions are observed to be much smaller in the barrel than in the end-caps, for which the constraints provided by cosmic rays are much weaker. The fitted Gaussian widths of the reconstructed $Z \rightarrow \mu\mu$ peaks in Fig. 196 are 2.6 GeV for the ideal (or perfectly aligned) inner detector
9020 and 3.9 GeV for the inner detector after completing the first-pass alignment.

A measure of these residual distortions can be extracted, as one would do with real data, by searching for possible asymmetries between the p_T -spectra of negative and positive muons from $Z \rightarrow \mu\mu$ decays. This is illustrated in Fig. 197 which clearly demonstrates a significant residual asymmetry between the two spectra after the first-pass alignment. This large asymmetry is clearly
9025 related to the large residual contribution of 2.9 GeV to the resolution on the reconstructed dimuon

Track parameter	0.25 < $ \eta $ < 0.50		1.50 < $ \eta $ < 1.75	
	$\sigma_X(\infty)$	p_X (GeV)	$\sigma_X(\infty)$	p_X (GeV)
Inverse transverse momentum ($1/p_T$)	0.34 TeV ⁻¹	44	0.41 TeV ⁻¹	80
Azimuthal angle (ϕ)	70 μ rad	39	92 μ rad	49
Polar angle ($\cot \theta$)	0.7×10^{-3}	5.0	1.2×10^{-3}	10
Transverse impact parameter (d_0)	10 μ m	14	12 μ m	20
Longitudinal impact parameter ($z_0 \times \sin \theta$)	91 μ m	2.3	71 μ m	3.7

Table 58. Expected track-parameter resolutions (RMS) at infinite transverse momentum, $\sigma_X(\infty)$, and transverse momentum, p_X , at which the multiple-scattering contribution equals that from the detector resolution. The momentum and angular resolutions are shown for muons, whereas the impact-parameter resolutions are shown for pions (see text). The values are shown for two η -regions, one in the barrel inner detector where the amount of material is close to its minimum and one in the end-cap where the amount of material is close to its maximum.




Figure 198. Relative transverse momentum resolution (left) as a function of $|\eta|$ for muons with $p_T = 1$ GeV (open circles), 5 GeV (full triangles) and 100 GeV (full squares). Transverse momentum, at which the multiple-scattering contribution equals the intrinsic resolution, as a function of $|\eta|$ (right).

mass after the first-pass alignment. If this residual width were for example ten times smaller, then a few days of data-taking at a luminosity of $10^{31} \text{ cm}^{-2} \text{ s}^{-1}$ would be required to actually detect a significant effect using $Z \rightarrow \mu\mu$ decays.

10.2.3 Tracking performance for single particles and particles in jets

9030 The expected performance of the tracking system for reconstructing single particles and particles
in jets is determined using a precise modelling of the individual detector response, geometry and
passive material in the simulation. In this section, a consistent set of selection cuts for reconstructed
tracks has been used throughout. Only prompt particles with $p_T > 1$ GeV and $|\eta| < 2.5$ are
considered. Standard quality cuts require reconstructed tracks to have at least seven precision hits
9035 (pixels and SCT); in addition, the transverse and longitudinal impact parameters at the perigee must
fulfill respectively $|d_0| < 2$ mm and $|z_0 - z_v| \times \sin \theta < 10$ mm, where z_v is the position of the
primary vertex along the beam and θ is the polar angle of the track. Stricter selection cuts, called




Figure 199. Transverse impact parameter, d_0 , resolution (left) as a function of $|\eta|$ for pions with $p_T = 1$ GeV (open circles), 5 GeV (full triangles) and 100 GeV (full squares). Transverse momentum, at which the multiple-scattering contribution equals the intrinsic resolution, as a function of $|\eta|$ (right).




Figure 200. Modified longitudinal impact parameter, $z_0 \times \sin \theta$, resolution (left) as a function of $|\eta|$ for pions with $p_T = 1$ GeV (open circles), 5 GeV (full triangles) and 100 GeV (full squares). Transverse momentum, at which the multiple-scattering contribution equals the intrinsic resolution, as a function of $|\eta|$ (right).

b -tagging cuts, are defined by: at least two hits in the pixels and one in the vertexing layer, as well as $|d_0| < 1$ mm and $|z_0 - z_v| \times \sin \theta < 1.5$ mm. A reconstructed track is matched to a Monte-Carlo particle if at least 80% of its hits were created by that particle. The efficiency is defined as the fraction of particles which are matched to reconstructed tracks passing the quality cuts, and the fake rate is defined as the fraction of reconstructed tracks passing the cuts which are not matched to a particle.

The resolution of a track parameter X can be expressed as a function of p_T as:

$$\sigma_X = \sigma_X(\infty)(1 \oplus p_X/p_T),$$

where $\sigma_X(\infty)$ is the asymptotic resolution expected at infinite momentum and p_X is a constant representing the value of p_T , for which the intrinsic and multiple-scattering terms in the equation are equal for the parameter X under consideration. This expression is approximate, working well at high p_T (where the resolution is dominated by the intrinsic detector resolution) and at low p_T (where the resolution is dominated by multiple scattering). Figures 198, 199 and 200 show the




Figure 201. Charge misidentification probability for high energy muons and electrons as a function of p_T for particles with $|\eta| \leq 2.5$ (left) and as a function of $|\eta|$ for $p_T = 2$ TeV (right).




Figure 202. Track reconstruction efficiencies as a function of $|\eta|$ for muons, pions and electrons with $p_T = 5$ GeV. The inefficiencies for pions and electrons reflect the shape of the amount of material in the inner detector as a function of $|\eta|$.




Figure 203. Track reconstruction efficiencies as a function of $|\eta|$ for pions with $p_T = 1, 5$ and 100 GeV.

9050 momentum resolution for muons and the transverse and longitudinal impact parameter resolutions
 9055 for pions, all without a beam constraint and assuming the effects of misalignment to be negligible.
 Table 58 shows the values of $\sigma_X(\infty)$ and p_X for tracks in two η -regions, corresponding to the barrel
 and end-caps. The TRT measurements are included in the track fits for tracks with $|\eta| < 2.0$, be-
 yond which there are no further TRT measurements. The impact parameter resolutions are quoted
 only for tracks with a hit in the vertexing layer (this requirement has a very high efficiency, as
 9055 illustrated below).

The determination of the lepton charge at high p_T is particularly important for measuring
 charge asymmetries arising from the decays of possible heavy gauge bosons (W' and Z'). Typi-
 cally, such measurements require that the charge of the particle be determined to better than 3σ .
 Whereas the muon charge can be reliably measured at the highest momenta in the muon system,
 9060 only the inner detector can measure the charge of electrons. The fraction of electrons for which




Figure 204. Track reconstruction efficiencies and fake rates as a function of $|\eta|$, for charged pions in jets in $t\bar{t}$ events and for different quality cuts (as described in the text). "Reconstruction" refers to the basic reconstruction before additional quality cuts.




Figure 205. Track reconstruction efficiencies and fake rates as a function of the distance ΔR (defined as $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}$) of the track to the jet axis, using the standard quality cuts and integrated over $|\eta| < 2.5$.

the sign of the charge is wrongly determined is shown in Fig. 201, together with the same fraction for muons, included as a reference (perfect alignment has been assumed). For the muons, the fraction is well described by the nominal (Gaussian) resolution, whereas electrons are more complicated since they are subject to bremsstrahlung. This should help for charge determination since the momentum is reduced, but sometimes the electrons overlap with subsequent conversion electrons from the bremsstrahlung photons, which may cause pattern-recognition problems because of extra hits and overlaps. For $p_T < 1$ TeV, the sign of the curvature of a track is sufficiently well determined that the benefit from bremsstrahlung is small and the overlap problem dominates the electron reconstruction, causing the electron charge determination to be of worse quality than for muons. However at 2 TeV, the poor intrinsic resolution is the dominant factor, and the effect of bremsstrahlung compensates for the pattern-recognition problems.

Figure 202 shows the efficiencies for reconstructing muons, pions and electrons with a transverse momentum of 5 GeV, whereas Fig. 203 shows the efficiencies for reconstructing pions with p_T varying from 1 to 100 GeV. In addition to multiple-scattering, pions are affected by hadronic interactions in the inner-detector material, while electrons are subject to even larger reconstruction inefficiencies because of bremsstrahlung. As a result, the efficiency curves as a function of $|\eta|$ for pions and electrons reflect the shape of the amount of material in the inner detector (see Figs. 72 and 73). As expected, the efficiency becomes larger and more uniform as a function of $|\eta|$ at higher energies.

Figure 204 shows the track reconstruction efficiency for prompt pions (produced before the vertexing-layer) and the fake rate for tracks in jets in $t\bar{t}$ events as a function of $|\eta|$. For these events, the mean jet p_T is 55 GeV, and the mean p_T of the accepted tracks which they contain

Event type	x-y resolution (μm)	z resolution (μm)	Reconstruction efficiency (%)	Selection efficiency (%)
$t\bar{t}$ (without beam constraint)	18	41	100	99
$t\bar{t}$ (with beam constraint)	11	40	100	99
$H \rightarrow \gamma\gamma$ (without beam constraint)	36	72	96	79
$H \rightarrow \gamma\gamma$ (with beam constraint)	14	66	96	79

Table 59. Primary vertex resolutions (RMS), without and with a beam constraint in the transverse plane, for $t\bar{t}$ events and $H \rightarrow \gamma\gamma$ events with $m_H = 110$ GeV in the absence of pile-up. Also shown, in the presence of pile-up at a luminosity of $10^{33} \text{ cm}^{-2} \text{ s}^{-1}$, are the efficiencies to reconstruct and then select the hard-scattering vertex within $\pm 300 \mu\text{m}$ of the true vertex position in z . The hard-scattering vertex is selected as the primary vertex with the largest Σp_T^2 , summed over all its constituent tracks.

is 4 GeV. The loss of efficiency at $|\eta| = 0$ with the b -tagging criteria arises from inefficiencies in the pixel vertexing-layer, which are assumed here to be 1%; this improves at higher $|\eta|$, owing to the presence of larger clusters when the track incidence angle decreases. Beyond $|\eta| \sim 1$, the tracking performance deteriorates, mostly because of increased material. As shown in Fig. 205, the fake rate increases near the core of the jet, where the track density is the highest and induces pattern-recognition problems. This effect increases as the jet p_T increases. A few percent efficiency can be gained at the cost of doubling the fake rate in the jet core.

10.2.4 Vertexing performance

Vertexing tools constitute an important component of the higher-level tracking algorithms. The residuals of the primary vertex reconstruction are shown in Fig. 206, as obtained without using any beam constraint, for $t\bar{t}$ events and $H \rightarrow \gamma\gamma$ events with $m_H = 110$ GeV. The results shown here for $H \rightarrow \gamma\gamma$ events are based on tracks reconstructed from the underlying event and do not make use of the measurement of the photon direction in the electromagnetic calorimeter, which is discussed in Section 10.4. The primary vertex in $t\bar{t}$ events has always a rather large multiplicity and includes a number of high- p_T tracks, resulting in a narrower and more Gaussian distribution than for $H \rightarrow \gamma\gamma$ events. Table 59 shows the resolutions of the primary vertex reconstruction in these $t\bar{t}$ and $H \rightarrow \gamma\gamma$ events, without and with a beam constraint in the transverse plane, as well as the efficiencies to reconstruct and select correctly these primary vertices in the presence of pile-up at a luminosity of $10^{33} \text{ cm}^{-2} \text{ s}^{-1}$.

The resolutions for the reconstruction of the radial position of secondary vertices for three-prong hadronic τ -decays in $Z \rightarrow \tau\tau$ events, with a mean p_T of 36 GeV for the τ -lepton, and for $J/\psi \rightarrow \mu\mu$ decays in events containing B -hadron decays, with a mean p_T of 15 GeV for the J/ψ , are shown respectively in Figs. 207 and 208. In the first more challenging example, the vertex resolutions are Gaussian in the central region, but have long tails as can be seen from the points showing 95% coverage in Fig. 207. Finally, Fig. 209 shows the resolution as a function of decay radius for the reconstruction of the radial position of secondary vertices for K_s^0 decays with mean p_T of 6 GeV in events containing B -hadron decays. The reconstruction is performed in three dimensions and hence requires at least two silicon hits. Consequently, the efficiency falls rapidly for decay radii larger than 30 cm. The effect of crossing the three successive pixel layers is clearly




Figure 206. Primary vertex residual along x , in the transverse plane (left), and along z , parallel to the beam (right), for events containing top-quark pairs and $H \rightarrow \gamma\gamma$ decays with $m_H = 110$ GeV. The results are shown without pile-up and without any beam constraint.




Figure 207. Resolution for the reconstruction of the radial position of the secondary vertex for three-prong hadronic τ -decays in $Z \rightarrow \tau\tau$ events, as a function of the pseudorapidity of the τ . The τ -leptons have an average transverse momentum of 36 GeV.




Figure 208. Resolution for the reconstruction of the radial position of the secondary vertex for $J/\psi \rightarrow \mu\mu$ decays in events containing B -hadron decays, as a function of the pseudorapidity of the J/ψ . The J/ψ have an average transverse momentum of 15 GeV.

visible as well as the degraded resolution for decays beyond the last pixel layer. Figure 210 shows the resolution as a function of decay radius for the reconstruction of the invariant mass of the charged-pion pair for the same $K_s^0 \rightarrow \pi^+\pi^-$ decays.

9115 10.2.5 Particle identification, reconstruction of electrons and photon conversions

The reconstruction of electrons and of photon conversions is a particular challenge for the inner detector, since electrons have lost on average between 20 and 50% of their energy (depending on $|\eta|$) when they leave the SCT, as illustrated in Fig. 211. In the same region, between 10% and 50% of photons convert into an electron-positron pair, as illustrated in Fig. 212.




Figure 209. Resolution for reconstruction of radial position of secondary vertex for $K_s^0 \rightarrow \pi^+\pi^-$ decays in events containing B -hadron decays, as a function of the K_s^0 decay radius.




Figure 210. Resolution for reconstruction of the invariant mass of the charged-pion pair for $K_s^0 \rightarrow \pi^+\pi^-$ decays in events containing B -hadron decays, as a function of the K_s^0 decay radius.




Figure 211. Fraction of energy lost on average by electrons with $p_T = 25$ GeV as a function of $|\eta|$, when exiting the pixel, the SCT and the inner-detector tracking volumes. The fraction of energy lost is not a strong function of the electron energy. For $|\eta| > 2.2$, there is no TRT material, hence the SCT and TRT lines merge.




Figure 212. Probability for a photon to have converted as a function of radius for different values of $|\eta|$, shown for photons with $p_T > 1$ GeV in minimum-bias events. The probability is not a strong function of the photon energy.

9120

The TRT plays a central role in electron identification, cross-checking and complementing the calorimeter, especially at energies below 25 GeV. In addition, the TRT contributes to the reconstruction and identification of electron track segments from photon conversions down to 1 GeV and of electrons which have radiated a large fraction of their energy in the silicon layers.

9125

By fitting electron tracks in such a way as to allow for bremsstrahlung, it is possible to improve the reconstructed track parameters, as shown for $|\eta| > 1.5$ in Fig. 213 for two examples of bremsstrahlung recovery algorithms. These algorithms rely exclusively on the inner-detector information and therefore provide significant improvements only for electron energies below ~ 25 GeV (see Section 10.4.2 for a discussion of bremsstrahlung recovery using also the position information of the electromagnetic calorimeter). The dynamic-noise-adjustment (DNA) method extrapolates




Figure 213. Probability distribution for the ratio of the true to reconstructed momentum for electrons with $p_T = 25$ GeV and $|\eta| > 1.5$. The results are shown as probabilities per bin for the default Kalman fitter and for two bremsstrahlung recovery algorithms (see text).




Figure 214. Probability for reconstructed invariant mass of electron pairs from $J/\psi \rightarrow ee$ decays in events with $B_d^0 \rightarrow J/\psi(ee)K_s^0$. The results are shown for the default Kalman fitter and for two bremsstrahlung recovery algorithms (see text). The true J/ψ mass is shown by the dotted line.




Figure 215. Average probability of a high-threshold hit in the barrel TRT as a function of the Lorentz γ -factor for electrons (open squares), muons (full triangles) and pions (open circles) in the energy range 2–350 GeV, as measured in the combined test-beam.




Figure 216. Pion efficiency shown as a function of the pion energy for 90% electron efficiency, using high-threshold hits (open circles), time-over-threshold (open triangles) and their combination (full squares), as measured in the combined test-beam.

9130 track segments to the next silicon detector layer. If there is a significant χ^2 contribution, com-
 9135 patible with a hard bremsstrahlung, the energy loss is estimated and an additional noise term is
 included in the Kalman filter [266]. The Gaussian-sum filter (GSF) is a non-linear generalisation
 of the Kalman filter, which takes into account non-Gaussian noise by modelling it as a weighted
 sum of Gaussian components and therefore acts as a weighted sum of Kalman filters operating in
 parallel [267]. Figure 214 shows the improvements from bremsstrahlung recovery for the recon-
 structed $J/\psi \rightarrow ee$ mass. Without any bremsstrahlung recovery, only 50% of events are recon-
 structed within ± 500 MeV of the nominal J/ψ mass, whereas with the use of the bremsstrahlung
 recovery, this fraction increases to approximately 60% for both algorithms.

Using pion, electron and muon samples
 9140 in the energy range between 2 and 350 GeV,
 the barrel TRT response has been measured
 in the CTB in terms of the high-threshold
 hit probability, as shown in Fig. 215. The
 transition-radiation X-rays contribute signifi-
 9145 cantly to the high-threshold hits for elec-
 tron energies above 2 GeV and saturation sets
 in for electron energies above 10 GeV. Fig-
 ure 216 shows the resulting pion identification
 efficiency for an electron efficiency of 90%,
 9150 achieved by performing a likelihood evalua-
 tion based on the high-threshold probability
 for electrons and pions for each straw. Fig-
 ure 216 also shows the effect of including
 time-over-threshold information, which improves the pion rejection by about a factor of two when
 9155 combined with the high-threshold hit information. At low energies, the pion rejection (the inverse
 of the pion efficiency plotted in Fig. 216) improves with energy as the electrons emit more transition
 radiation. The performance is optimal at energies of ~ 5 GeV, and pion-rejection factors above 50
 are achieved in the energy range of 2–20 GeV. At very high energies, the pions become relativistic
 and therefore produce more δ -rays and eventually emit transition radiation, which explains why
 9160 the rejection slowly decreases for energies above 10 GeV.




Figure 217. Expected pion efficiency as a function of $|\eta|$ for an efficiency of 90% for electrons with $p_T = 25$ GeV.

The electron-identification performance expected for the TRT in ATLAS, including the time-over-threshold information, is shown as a function of $|\eta|$ in Fig. 217 as the pion identification efficiency expected for an electron efficiency of 90%. The shape observed is closely correlated to the number of TRT straws crossed by the track, which decreases from approximately 35 to a minimum of 20 in the transition region between the barrel and end-cap TRT, $0.8 < |\eta| < 1.1$, and which also decreases rapidly at the edge of the TRT fiducial acceptance for $|\eta| > 1.8$. Because of its more efficient and regular foil radiator, the performance in the end-cap TRT is better than in the barrel TRT (see Section 4.3.3).

Figure 218 shows the efficiency for reconstructing conversions of photons with $p_T = 20$ GeV
 9170 and $|\eta| < 2.1$ as a function of the conversion radius, using the standard tracking algorithm com-
 bined with the back-tracking algorithm described in Section 10.2.1. At radii above 50 cm, the
 efficiency for reconstructing single tracks drops and that for reconstructing the pair drops even
 faster because the two tracks are merged. If both tracks from the photon conversion are recon-
 9175 structed successfully, vertexing tools can be used to reconstruct the photon conversion with high
 efficiency up to radii of 50 cm. The overall conversion-finding efficiency can be greatly increased
 at large radii by defining single tracks as photon conversions under certain conditions. Only tracks
 which have no hits in the vertexing layer, are not associated to any fitted primary or secondary
 vertex, and pass a loose electron identification cut requiring more than 9% high-threshold hits on
 the TRT segment of the track, are retained. The resulting overall efficiency for finding photon
 9180 conversions is almost uniform over all radii below 80 cm, as shown in Fig. 219.




Figure 218. Efficiency to reconstruct conversions of photons with $p_T = 20$ GeV and $|\eta| < 2.1$, as a function of the conversion radius. Shown are the efficiencies to reconstruct single tracks from conversions, the pair of tracks from the conversion and the conversion vertex. The errors are statistical.




Figure 219. Efficiency to identify conversions of photons with $p_T = 20$ GeV and $|\eta| < 2.1$, as a function of the conversion radius. The overall efficiency is a combination of the efficiency to reconstruct the conversion vertex, as shown also in Fig. 218, and of that to identify single-track conversions (see text). The errors are statistical.




Figure 220. Sagitta measurement in the muon combined test-beam barrel sector set-up as a function of the value of systematic displacements of the middle barrel chamber in the direction indicated by the sketch.




Figure 221. Sagitta measurement in the muon combined test-beam barrel sector set-up as a function of the value of systematic rotations of the inner barrel chamber around the axis indicated by the sketch (x -axis parallel to the drift tubes).

10.3 Muon reconstruction and identification

10.3.1 Introduction

The collisions at the LHC will produce a broad spectrum of final-state muons, ranging from low-momentum non-isolated muons in b -jets to high-momentum isolated muons from W/Z -boson decays or from possible new physics. The experiment will detect and measure muons in the muon spectrometer and will also exploit the measurements in the inner detector and the calorimeters to improve the muon identification efficiency and momentum resolution. Muon measurements are a combination of accurate measurements in the muon spectrometer and in the inner detector.




Figure 222. For muons with $p_T = 100$ GeV, expected fractional momentum resolution as a function of $|\eta|$ for stand-alone and combined reconstruction. The degradation in the region with $1.1 < |\eta| < 1.7$ is due to the absence of the middle muon stations in the barrel/end-cap transition region for the initial data-taking, to the low bending power of the magnetic field in the transition region between the barrel and end-cap toroids and to the material of the coils of the end-cap toroids.




Figure 223. For muons with $p_T = 100$ GeV, expected fractional momentum resolution as a function of ϕ for stand-alone and combined reconstruction. The resolution is degraded at $\phi = 240$ and 300 degrees, due to the additional material introduced by the feet which support the barrel part of the detector.

9190 The muon spectrometer also efficiently triggers on muons over a wide range of energies and over
 9195 $|\eta| < 2.4$. as described in detail in Section 6.5 for the detectors and in Section 10.9 for the actual trigger performance. The inner detector provides the best measurement at low to intermediate momenta, whereas the muon spectrometer takes over above 30 GeV. The toroidal field guarantees excellent momentum resolution even at the highest values of η . A complex grid of Hall probes is required to reconstruct the actual toroid coil positions with respect to the muon chambers (see Section 2.2).

This section describes the alignment results obtained in the combined test-beam (CTB), which have validated the overall alignment strategy for both the barrel and end-cap muon-chamber systems, and the expected muon reconstruction performance in terms of momentum resolution, track-finding efficiency and mass resolution for selected channels.

9200 10.3.2 Calibration and alignment

In order to achieve the required performance for combined muon reconstruction, the inner detector and the muon spectrometer must be calibrated and aligned internally and with respect to each other. The alignment of the inner detector is described in Section 10.2.2.

9205 In the muon spectrometer, movements of most of the precision chambers (MDT and CSC) are monitored by a system of optical sensors with an accuracy of a few μm (see Section 6.3.4). In principle, the optical system alone should provide the chamber positions with an accuracy such that the alignment contribution to the error on the sagitta measurement does not exceed $40 \mu\text{m}$. Muon tracks, however, are required to align the chambers with no (or poor) optical connection, to align




Figure 224. For muons with $p_T = 100$ GeV, expected fractional stand-alone momentum resolution as a function of ϕ and $|\eta|$. The results are based on a parametrisation using the material distribution in the muon spectrometer shown in Fig. 126, the magnetic field configuration in the muon spectrometer, and the spatial resolution of the muon chambers. No momentum measurement is possible at $|\eta| < 0.1$ over most of the azimuth, nor at $|\eta| = 1.3$ because of holes in the acceptance of the muon spectrometer (see text).

the end-caps with respect to the barrel, and to align the muon spectrometer with respect to the inner
 9210 detector with an accuracy of approximately $200 \mu\text{m}$ in z and 1 mm in $R\phi$.

10.3.2.1 Performance of optical alignment system in test-beam The optical alignment concept
 for the muon spectrometer underwent a final round of testing and validation with one full barrel
 sector and one full end-cap sector in the H8 muon beam line at CERN in 2002-2004 (see Figs. 191
 and 193). Figures 220 and 221 show as examples the measured track sagittas, after applying the
 9215 corrections obtained from the optical alignment system, for a specific displacement of the middle
 chamber of the barrel sector and for a specific rotation of the inner chamber of the barrel sector,

respectively. Alignment accuracies of approximately $20\ \mu\text{m}$ have been achieved in these tests, well within the design specifications of the alignment system (see Section 6.3.4) [202, 203, 274].

10.3.2.2 Alignment of the muon spectrometer with tracks In the muon spectrometer (see Section 6.3.2 and Table 34), some chambers are not optically linked (BIS.8, BEE), or the optical connection does not have the required precision for the sagitta measurement (barrel chambers of the small sectors). During normal data-taking, these chambers can be aligned precisely using muon tracks passing through overlap regions with the optically aligned neighbouring chambers. Similarly, the alignment of the two end-caps with respect to the barrel will use tracks fully reconstructed in the barrel and passing through one end-cap chamber: one example of such an overlap is that between BIS-EIL-BML-BOL.

As an additional independent test of the achieved alignment accuracy, it is foreseen to run for some short periods without magnetic field in the toroids, while the solenoid is at full field. This will yield straight tracks in the muon spectrometer, which can be selected to have e.g. $p_T > 10\ \text{GeV}$, using the matching track reconstructed in the inner detector to limit the impact of multiple scattering. If the chamber alignment were perfect, the measured sagittas would be centred around zero with a variance determined by multiple scattering and the position resolution of the chambers. Significant deviations from zero in certain η - ϕ regions would point to errors of chamber positioning, as obtained from the optical alignment. A statistical accuracy of $30\ \mu\text{m}$ on the average sagitta can be obtained with 5000 tracks with $p_T > 10\ \text{GeV}$ per chamber triplet. This corresponds to a less than one day of data-taking at a luminosity of $10^{33}\ \text{cm}^{-2}\ \text{s}^{-1}$. A similar procedure can be used during cosmic-ray data-taking to align parts of the spectrometer independently of LHC operation.

10.3.2.3 Overall calibration and alignment strategy The drift-time measurements of the MDT's are synchronised with an accuracy of 200 ps by measuring the minimum drift time from the raw drift-time spectra of the individual tubes. The space-to-drift-time relationships, R - t , are iteratively determined from the residuals of reconstructed muon track segments in the chambers. The required R - t accuracy of $20\ \mu\text{m}$ can be achieved with 2000 track segments per chamber.

Both the alignment constants obtained from tracks and the MDT calibrations will be produced on a daily basis and will have to be ready within 24 hours to be used in the reconstruction. In order to collect enough statistics for these tasks, a dedicated stream of high- p_T single muons will be provided at a rate of 1 kHz as a direct output of the L2 muon trigger [176].

10.3.3 Reconstruction strategies

Muons with momenta ranging from approximately 3 GeV to 3 TeV are identified and measured with optimal acceptance and efficiency through the use of a combination of three track-reconstruction strategies (see Section 10.2.1 for a brief description of the tracking software common to inner-detector and muon-spectrometer reconstruction):

- Stand-alone: muon track reconstruction based solely on the muon spectrometer data over the range $|\eta| < 2.7$ (defined by the spectrometer acceptance).
- Combined: combination of a muon-spectrometer track with an inner-detector track over the range $|\eta| < 2.5$ (defined by the inner-detector acceptance).




Figure 225. Expected stand-alone and combined fractional momentum resolution as a function of p_T for single muons with $|\eta| < 1.1$.




Figure 226. Expected stand-alone and combined fractional momentum resolution as a function of p_T for single muons with $|\eta| > 1.7$.




Figure 227. Efficiency for reconstructing muons with $p_T = 100$ GeV as a function of $|\eta|$. The results are shown for stand-alone reconstruction, combined reconstruction and for the combination of these with the segment tags discussed in the text.




Figure 228. Efficiency for reconstructing muons as a function of p_T . The results are shown for stand-alone reconstruction, combined reconstruction and for the combination of these with the segment tags discussed in the text.

- Segment tag: combination of an inner-detector track with a muon-spectrometer segment, i.e. a straight-line track, in an inner muon station.

Track reconstruction in the muon spectrometer is logically sub-divided into the following stages: pre-processing of raw data to form drift-circles in the MDT's or clusters in the CSC's and the trigger chambers (RPC's and TGC's), pattern-finding and segment-making, segment-combining, and finally track-fitting. Track segments are defined as straight lines in a single MDT or CSC station. The search for segments is seeded by a reconstructed pattern of drift-circles or clusters or by drift-circles or clusters lying in a region of activity, which is defined by the trigger chambers and has a size of the order of 0.4×0.4 in $\eta - \phi$ space.

Full-fledged track candidates are built from segments, starting from the outer and middle stations and extrapolating back through the magnetic field to the segments reconstructed in the other stations. Each time a reasonable match is found, the segment is added to the track candidate. The final track-fitting procedure takes into account, in full detail, the geometrical description of the

traversed material and the magnetic field inhomogeneities along the muon trajectory.

9270 The muon-spectrometer track parameters are determined at the inner stations, which yield
the first set of measurements in the muon spectrometer. The track is then propagated back to the
interaction point and the momentum is corrected for the energy loss in the calorimeters (and in
the inner detector). The energy lost by dE/dX in the calorimeters is estimated by an algorithm,
which uses either the parametrised expected energy loss or the measured calorimeter energy. The
9275 measured energy is used only if it exceeds significantly the most probable energy loss and if the
muon track is isolated.

The combination of the stand-alone tracks reconstructed in the muon spectrometer with tracks
reconstructed in the inner detector is performed in the region $|\eta| < 2.5$, which corresponds to
the geometrical acceptance of the inner detector. This combination will considerably improve the
9280 momentum resolution for tracks with momenta below 100 GeV, but will also suppress to a certain
extent backgrounds from pion punch-through and from pion or kaon decays in flight.

In the case of segment tags, inner-detector tracks are extrapolated to the inner muon stations
and either associated directly to reconstructed muon segments or used to select muon drift-circles
and clusters in a cone with typically a size of 100 mrad, from which track segments are then
9285 reconstructed. The muons reconstructed through this procedure provide an important improvement
to the stand-alone muon reconstruction for three main reasons:

- at momenta below typically 6 GeV, muon tracks do not always reach the middle and outer
muon stations;
- in the barrel/end-cap transition region with $1.1 < |\eta| < 1.7$, the middle stations are missing
9290 for the initial data-taking (EES and EEL chambers in Table 35) and the stand-alone recon-
struction efficiency is reduced in this region;
- in the difficult regions at $\eta \approx 0$ and in the feet, the geometrical acceptance of the muon
stations is considerably reduced.

10.3.4 Muon reconstruction performance for single muons

9295 Three main quantities can be used to summarise the performance of the muon reconstruction and
identification algorithms: the momentum resolution, the efficiency and the misidentification or
fake rate. This section presents the expected performance of the three first strategies described
above for single muons. Both the stand-alone and combined results shown here have been obtained
using as an example the algorithms described in Ref. [275]. Except where directly relevant to the
9300 performance (e.g. for estimates of the fake rates), the results presented here do not include any
effects arising from cavern background or pile-up.

Figure 222 shows the expected fractional momentum resolution, averaged over ϕ , for single
muons with $p_T = 100$ GeV, as obtained for stand-alone and combined muon tracks. Over a large
fraction of the acceptance, the stand-alone resolution is close to 3%, as shown in more detail in
9305 Fig. 223, which shows its variation as a function of ϕ in the region $0.3 < |\eta| < 0.65$. One clearly
sees the degradation in resolution due to the feet which support the experiment and are situated
close to $\phi = 240$ and 300 degrees. In the region $1.1 < |\eta| < 1.7$, the large degradation of the stand-
alone momentum resolution is due to several effects. In the region $1.1 < |\eta| < 1.3$, the degradation

9310 is due to the absence of the middle muon stations in the barrel/end-cap transition region for the
 initial data-taking, which results in a large degradation of the resolution since the measurement
 is limited to an angle-angle measurement between the inner and outer stations. At larger values
 of $|\eta|$, the degradation is due to the combination of the low bending power of the magnetic field
 in the transition region between the barrel and end-cap toroids and of the large amount of material
 in the coils of the end-cap toroid in limited regions in ϕ . The contribution of the inner detector
 9315 to the combined resolution is therefore more important in this η -region. In the barrel region, the
 contribution of the inner detector remains significant, whereas it basically vanishes for $|\eta| > 2.0$.
 This is due to the intrinsically worse momentum resolution in the inner detector because of the
 absence of any TRT measurements in this η -region, of the solenoidal field non-uniformity, and of
 the shorter length of the tracks in the inner-detector magnetic volume.

9320 The stand-alone momentum resolution of muons with $p_T = 100$ GeV can be calculated based
 on the spatial resolution of the chambers, the material distribution, and the magnetic-field config-
 uration in the muon spectrometer [276]. The result of this calculation is shown as a function of ϕ
 and $|\eta|$ in Fig. 224. No momentum measurement is possible at $|\eta| < 0.1$ and $|\eta| = 1.3$ because
 of holes in the acceptance of the muon spectrometer. The expected stand-alone momentum reso-
 9325 lution is approximately 3% over most of the $\eta - \phi$ plane. It is degraded to 5% at $|\eta| = 0.2, 0.3$
 and 0.7, due to support structures of the barrel toroid magnet coils. The degradation in the regions
 corresponding to $1.2 < |\eta| < 1.7$ and to ϕ -values which are multiples of 22.5° is caused by the
 small bending power of the magnetic field in these regions. The resolution expectations from this
 analytical model are in good agreement with the results shown in Figs. 222 and 223, which are
 9330 based on full simulation and reconstruction.

Figures 225 and 226 show the expected stand-alone and combined momentum resolutions as
 a function of p_T , excluding the η -region $1.1 < |\eta| < 1.7$, respectively for the barrel and end-cap
 muon spectrometer. The stand-alone resolution displays its characteristic behaviour with optimal
 resolution achieved at ~ 100 GeV. At lower transverse momenta, the stand-alone resolution is dom-
 9335 inated by fluctuations in the energy loss in the calorimeters, whereas at higher transverse momenta,
 it is dominated by the intrinsic MDT tube accuracy, assumed to be $80 \mu\text{m}$ in the case of a cali-
 brated and aligned detector. At low transverse momenta, the combined resolution reflects directly
 the dominant performance of the inner detector, which is itself limited by multiple scattering for
 transverse momenta below ~ 10 GeV (see Section 10.2.3).

9340 In Figs. 227 and 228, the single muon reconstruction efficiency is shown, respectively as a
 function of $|\eta|$ for muons with $p_T = 100$ GeV and as a function of p_T . The efficiency is defined
 as the fraction of simulated muons which are reconstructed within a cone of size $\Delta R = 0.2$ of the
 initial muon. The results are shown for stand-alone reconstruction, for combined reconstruction,
 and for the overall combination of these with the segment tags discussed above. The efficiency
 9345 for stand-alone tracks drops to very low values in the region with $\eta \sim 0$ because of the large
 gap for services, in which there are very few muon stations. The stand-alone efficiency also drops
 substantially close to $\eta = 1.2$, which corresponds to a region in the barrel/end-cap transition region
 where several stations are missing. The efficiency for combining stand-alone muon tracks with the
 inner detector is very high in the central region, starts to drop for $|\eta| > 2.0$ and decreases rapidly
 9350 to 0 for $|\eta| > 2.4$. The segment tags contribute only to a limited extent to the overall efficiency for
 $1.4 < |\eta| < 2.0$ for muons with high p_T , but Fig. 228 shows that, as expected, their contribution is

substantial for lower p_T values.

The efficiencies presented above must be compared to the expected fake rates, especially in the presence of cavern background, which permeates the whole muon spectrometer, and of pile-up, which affects mostly the high- η region. Electromagnetic showers triggered by energetic muons traversing the calorimeters and support structures lead to low-momentum electron and positron tracks, which accompany the muons in the muon spectrometer. These low-momentum tracks are an irreducible source of fake stand-alone muons. Most of them can be rejected by a cut on their transverse momentum. For example, a cut requiring $p_T > 5$ GeV reduces the fake rate to a few percent per triggered event. Such fakes can be almost entirely rejected by requiring a match of the muon-spectrometer track with an inner-detector track.

The second source of fake stand-alone muons is the background of thermal neutrons and low-energy γ -rays in the muon spectrometer (the so-called "cavern background"). Most of these fakes also have transverse momenta smaller than 5 GeV. The expected fake rate with $p_T > 5$ GeV from cavern background at $10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ is below 2% per triggered event. This rate is proportional to the background counting rate and can be reduced by almost an order of magnitude by requiring a match of the muon-spectrometer track with an inner-detector track.

10.3.5 Reconstruction of $Z \rightarrow \mu\mu$ and $H \rightarrow \mu\mu\mu\mu$ decays

The large expected rates of $Z \rightarrow \mu\mu$ decays provide an excellent tool to untangle various effects which might lead to distortions of the measured dimuon invariant mass spectrum. One example is shown in Fig. 229 for stand-alone muon measurements, where the performance obtained with a misaligned layout is compared to that expected from a perfectly aligned layout. The misalignments introduced for this study were random displacements of typically 1 mm and random rotations of typically 1 mrad. These lead to a distribution of the difference between the dimuon reconstructed invariant mass and the true dimuon mass with a fitted Gaussian resolution of approximately 8 GeV. The fitted Gaussian resolution obtained for the same distribution in the case of the perfectly aligned layout is 2.5 GeV.

Figures 225 and 226 show the expected stand-alone and combined momentum resolutions as a function of p_T , excluding the η -region $1.1 < |\eta| < 1.7$, respectively for the barrel and end-cap muon spectrometer. The stand-alone resolution displays its characteristic behaviour with optimal resolution achieved at ~ 100 GeV. At lower transverse momenta, the stand-alone resolution is dominated by fluctuations in the energy loss in the calorimeters, whereas at higher transverse momenta, it is dominated by the intrinsic MDT tube accuracy, assumed to be $80 \mu\text{m}$ in the case of a calibrated and aligned detector. At low transverse momenta, the combined resolution reflects directly




Figure 229. For stand-alone muon reconstruction, reconstructed invariant mass distribution of dimuons from $Z \rightarrow \mu\mu$ decays for an aligned layout of the chambers and for a misaligned layout, where all chambers are displaced and rotated randomly by typically 1 mm and 1 mrad.




Figure 230. For $H \rightarrow \mu\mu\mu\mu$ decays with $m_H = 130$ GeV, reconstructed mass of the four muons using stand-alone reconstruction. The results do not include a Z -mass constraint.




Figure 231. For $H \rightarrow \mu\mu\mu\mu$ decays with $m_H = 130$ GeV, reconstructed mass of the four muons using combined reconstruction. The results do not include a Z -mass constraint.

the dominant performance of the inner detector, which is itself limited by multiple scattering for transverse momenta below ~ 10 GeV (see Section 10.2.3).

The muon reconstruction and identification efficiency will also be measured from data using $Z \rightarrow \mu\mu$ decays and the tag-and-probe method described in Section 10.9.7 with similar results in terms of accuracy of the measurement. These in situ measurements will be extended to lower-mass resonances, using J/ψ and Υ decays at lower initial luminosities.

Finally, Figs. 230 and 231 show the four-muon invariant mass distributions from respectively stand-alone and combined reconstruction without using any Z -mass constraint for $H \rightarrow \mu\mu\mu\mu$ decays in the case of a Higgs-boson mass of 130 GeV. The stand-alone resolution is 3.3 GeV, whereas the combined resolution is 2.1 GeV. The non-Gaussian tails in the distribution amount to 29% (resp. 18%) of events which lie further than 2σ away from the peak for the stand-alone (resp. combined) reconstruction. They are partially due to radiative decays, but mostly to muons poorly measured in certain regions of the muon spectrometer, especially in the case of the stand-alone measurements.

10.4 Electrons and photons

Efficient and accurate reconstruction and identification of electrons and photons will be a task of unprecedented difficulty at the LHC, where the ratios of inclusive electrons and photons to jets from QCD processes are expected to be between one and two orders of magnitude worse than at the Tevatron (as an example, the electron-to-jet ratio is expected to be $\sim 10^{-5}$ at $p_T = 40$ GeV). In addition, the large amount of material in front of the electromagnetic calorimeters and the harsh operating conditions at the LHC design luminosity provide a difficult challenge in terms of preserving most of the electrons and photons with their energies and directions measured as well as would be expected from the intrinsic performance of the electromagnetic calorimeters measured in test-beams. This section is devoted to a summary of the calibration and expected performance of the electromagnetic calorimeter, of electron and photon identification in the energy range of interest for initial physics, and of the strategies under evaluation for the validation and certification of the performance in situ.




Figure 232. Average energy loss in GeV as a function of $|\eta|$ for electrons with an energy of 100 GeV. The results are shown before the presampler (open circles) and the strip layer (crosses).




Figure 233. Fraction of photons converting at a radius of below 80 cm (115 cm) in open (full) circles as a function of $|\eta|$.

10.4.1 Calibration and performance of the electromagnetic calorimeter

The results presented in this section are based on detailed simulation studies, validated by extensive test-beam studies over the past years (see Section 5.6) and using reconstruction procedures developed for test-beam data analysis. Compared to Ref. [260], the material budget in front of the calorimeter has increased substantially. The large amount of material in front of the presampler and the electromagnetic calorimeter leads to substantial energy losses for electrons, as shown in Fig. 232 (see also Fig. 211 for more details on electron energy loss in the inner-detector material itself) and to a large fraction of photons converting, as shown in Fig. 233 (see also Fig. 212 for details on the photon conversion probability in the inner-detector material).

Electron and photon reconstruction is seeded using a sliding-window algorithm with a window size corresponding to 5×5 cells in the middle layer of the electromagnetic calorimeter (see Table 3 for a detailed description of the granularity and η -coverage of the electromagnetic calorimeter). A cluster of fixed size is then reconstructed around this seed. For electrons, the energy in the barrel electromagnetic calorimeter is collected over an area corresponding to 3×7 cells in the middle layer or 0.075×0.175 in $\Delta\eta \times \Delta\phi$. This choice optimises the balance between the conflicting requirements of collecting all the energy even in the case of hard bremsstrahlung and of preserving the energy resolution by minimising the contributions from noise and pile-up. For unconverted photons, adequate performance is obtained by limiting the area to 3×5 cells in the middle layer, whereas converted photons are treated like electrons. Finally, for the end-cap electromagnetic calorimeters, an optimal area of 5×5 cells in layer 2 has been chosen for both electrons and photons.

Position corrections are applied as a first step in the precise reconstruction of the electromagnetic cluster. Corrections for modulations of the local energy response as a function of the extrapolated impact point of the electron in both η and ϕ are shown in Figs. 234 and 235, respectively. These corrections do not modify the global energy scale and are rather small in terms of the relative response: typically, the η -variation is, minimum to maximum, around 1%, whereas the ϕ -modulation correction due to the accordion structure of the absorbers is, minimum to maximum,




Figure 234. Electron energy response modulation as a function of the η offset within the cell. The curve represents a fit to the points used to parametrise the correction.




Figure 235. Electron energy response modulation as a function of the ϕ offset from the absorber. The curve represents a fit to the points used to parametrise the correction.




Figure 236. Difference between measured and true energy normalised to true energy for electrons with an energy of 100 GeV at $\eta = 0.325$.




Figure 237. Difference between measured and true energy normalised to true energy for electrons with an energy of 100 GeV at $\eta = 1.075$.

around 0.4%. The parabolic component of this latter correction is smaller than the one in η because of the energy sharing between adjacent cells in ϕ .

The most important corrections to optimise at the same time the energy resolution and the linearity of the response are incorporated using η -dependent longitudinal weights, similarly to what is described for the electromagnetic calorimeter test-beam results in Section 5.6.1:

$$E = s(\eta)[c(\eta) + w_0(\eta) \cdot E_{\text{PS}} + E_{\text{strips}} + E_{\text{middle}} + w_3(\eta) \cdot E_{\text{back}}], \quad (10.1)$$

9450 where s is an overall scale factor, c is an offset, w_0 corrects for energy losses upstream of the pre-
 9455 sampler, and w_3 corrects for longitudinal leakage, while E_{PS} , E_{strips} , E_{middle} and E_{back} represent
 the energies measured in the successive compartments of the electromagnetic calorimeter (pre-sam-
 pler, strips, middle and back). The weights are determined as functions of $|\eta|$, using simulated
 single-particle events (electrons and photons) with energies from 5 GeV to 200 GeV. The weights
 are calculated separately for electrons (matched track required) and photons (no matched track
 required) and applied to the corresponding cluster energies. In the future, this method will be re-




Figure 238. Difference between measured and true energy normalised to true energy for all photons with an energy of 100 GeV at $\eta = 1.075$.




Figure 239. Difference between measured and true energy normalised to true energy for unconverted photons with an energy of 100 GeV at $\eta = 1.075$.




Figure 240. Expected relative energy resolution as a function of energy for electrons at $|\eta| = 0.3, 1.1,$ and 2.0 . The curves represent fits to the points at the same $|\eta|$ by a function containing a stochastic term, a constant term and a noise term.




Figure 241. Expected relative energy resolution as a function of energy for photons at $|\eta| = 0.3, 1.1,$ and 2.0 . The curves represent fits to the points at the same η by a function containing a stochastic term, a constant term and a noise term.

placed by a more complex algorithm, which corrects the different types of true energy loss one by one, by correlating each of them with measured observables.

In Figs. 236 and 237, the energy response, plotted as the difference between measured and true energy divided by the true energy, is shown for electrons with an energy of 100 GeV and for two illustrative η -positions in the barrel electromagnetic calorimeter. The central value of the energy is reconstructed with excellent precision ($\sim 3 \times 10^{-4}$) if one assumes perfect knowledge of the material in front of the calorimeter. Both the Gaussian core and the non-Gaussian component of the tail of the energy distribution are significantly worse at the point with the larger η due to the larger amount of material in front of the calorimeter (see Fig. 73). As shown in Figs. 238 and 239, the resolution and non-Gaussian tails are better for photons than for electrons, but are somewhat worse for all photons than for unconverted photons, i.e. photons not converting before leaving the volume of the inner detector.




Figure 242. Expected relative energy resolution as a function of $|\eta|$ for electrons with an energy of 100 GeV.




Figure 243. Expected relative energy resolution as a function of $|\eta|$ for photons with an energy of 100 GeV.




Figure 244. Expected η -position resolution as a function of $|\eta|$ for photon showers with an energy of 100 GeV and for the two main layers of the barrel and end-cap electromagnetic calorimeters.




Figure 245. Expected precision on the polar angle θ of photons from $H \rightarrow \gamma\gamma$ decays as a function of $|\eta|$, expressed in units of $\text{mrad} \cdot \sqrt{E}$, where E is the measured energy of the photon shower in GeV.

The energy resolution as a function of energy is shown in Figs. 240 and 241, respectively for
 9470 electrons and photons and for three illustrative values of $|\eta|$. The results shown here include the
 expected electronic noise contributions at 100 GeV of 190, 190 and 230 MeV (respectively 180, 180
 and 230 MeV) for the three η -values for electrons (respectively photons).

As expected in the case of the points at the larger η -values, the resolution is degraded with
 respect to the one at the more central value of η . Fits to these results similar to those described
 9475 in Section 5.6.1 and expressed in Eq. 5.2 yield stochastic terms of respectively 10.0%, 15.1%
 and 14.5% for the electrons at the three η -values shown. The corresponding terms for photons
 are found to be 10.2%, 12.4% and 12.1%, once again showing that photons are less sensitive than
 electrons to the material in front of the calorimeter. This can also be clearly seen when comparing
 Figs. 242 and 243, which show for electrons and photons the expected relative energy resolution as
 9480 a function of $|\eta|$ for a fixed energy of 100 GeV. The η -region between 1.37 and 1.52 corresponds to
 the difficult transition region between the barrel and end-cap cryostats, where the energy resolution
 degrades significantly despite the presence of scintillators in the crack between the barrel and end-

cap cryostats to correct for the energy lost in the barrel cryostat flange (see Section 5.4). This crack region is not used for photon identification nor for precision measurements with electrons.

9485 In Fig. 244, the expected η -position resolution is shown for the two main layers (strips and middle layer) of the barrel and end-cap calorimeters. The resolution is fairly uniform as function of $|\eta|$ and is $2.5 - 3.5 \times 10^{-4}$ for the strips (which have a size of 0.003 in η in the barrel electromagnetic calorimeter) and $5 - 6 \times 10^{-4}$ for the middle-layer cells (which have a size of 0.025 in η). The regions with worse resolution correspond to the barrel/end-cap transition region and, for the
9490 strips, to the region with $|\eta| > 2$, where the strip granularity of the end-cap calorimeter becomes progressively much coarser (see Table 3). The results shown in Section 5.6.1 are somewhat better because they correspond to a higher electron energy of 245 GeV.

Because of the fine lateral and longitudinal granularity of the electromagnetic calorimeter, these η -position measurements can be used to determine the direction of the axis of the shower
9495 development in the η -direction (or polar angle θ). To achieve the best performance, one requires an accurate parametrisation of the shower depth (R -coordinate in the barrel and z -coordinate in the end-caps), as determined by Monte-Carlo simulations for both layers. The resulting resolution on the polar angle of photon showers is shown in Fig. 245 for a representative sample of photons from $H \rightarrow \gamma\gamma$ decays. A resolution of 50–75 mrad $/\sqrt{E \text{ (GeV)}}$ is obtained, which should be sufficient
9500 to e.g. measure accurately the invariant mass of photon pairs without using any primary vertex information.

In addition to the calorimeter-seeded electron and photon reconstruction, a second electron reconstruction and identification algorithm uses good-quality tracks as a seed and constructs a cluster around the extrapolated impact point in the calorimeter [277]. This algorithm relies more
9505 on the electron identification capabilities of the inner detector and has been developed to improve the efficiency for low- p_T electrons (see Section 10.4.3) as well as for electrons close to jets (see Section 10.8.5). The algorithm matches good-quality inner-detector tracks to small clusters of electromagnetic energy. For a given track, only the energy contained in a small window along the track extrapolation is used and the contribution of neighbouring hadronic showers is therefore reduced.
9510 The identification procedure takes full advantage of the tracking and electron-identification capabilities of the TRT in the inner detector (over $|\eta| < 2.0$, as described in Section 10.2.5), as well as of the granularity of the electromagnetic calorimeter. A likelihood ratio combines inner-detector information (measured track momentum and transition-radiation hits) with shower-shape variables from the calorimeter.

9515 In the following, unless specified otherwise (as in Section 10.4.3), only the results of the calorimeter-seeded algorithm will be discussed.

10.4.2 Electron and photon reconstruction and identification

For the standard reconstruction of electrons and photons, a seed cluster is taken from the electromagnetic calorimeter and a loosely matching track is searched for among all reconstructed tracks.
9520 Additionally, the candidate is flagged if it matches a photon conversion reconstructed in the inner detector. Electron and photon candidates are thus separated reasonably cleanly, by requiring the electrons to have an associated track but no associated conversion. In contrast, the photons are defined as having no matched track, or as having been matched to a reconstructed conversion.

Cuts	$E_T > 17 \text{ GeV}$			$E_T > 8 \text{ GeV}$		
	Efficiency (%)		Jet rejection	Efficiency (%)		Jet rejection
	$Z \rightarrow ee$	$b, c \rightarrow e$		Single electrons ($E_T=10 \text{ GeV}$)	$b, c \rightarrow e$	
Loose	87.9 ± 0.5	38 ± 1	570 ± 10	75.7 ± 0.5	46 ± 1	510 ± 10
Medium	76.7 ± 0.5	27 ± 1	2200 ± 20	64.8 ± 0.5	36 ± 1	1280 ± 10
Tight (TRT)	61.3 ± 0.5	20 ± 1	$(8 \pm 1)10^4$	46.4 ± 0.5	25 ± 1	$(4.7 \pm 0.5)10^4$
Tight (isol.)	63.6 ± 0.5	16 ± 1	$(9 \pm 1)10^4$	48.7 ± 0.5	24 ± 1	$(4.3 \pm 0.5)10^4$
	Relative populations of surviving candidates (%)			Relative populations of surviving candidates (%)		
	Isolated	Non-isolated	Jets	Isolated	Non-isolated	Jets
Medium	0.9	6.4	92.6 (1.5 + 91.1)	–	7.7	92.3 (2.2 + 90.9)
Tight (TRT)	10.5	56.1	33.4 (4.3 + 29.0)	–	63.2	36.8 (4.0 + 32.8)
Tight (isol.)	13.0	53.4	33.6 (4.6 + 29.0)	–	62.8	37.2 (4.4 + 30.3)

Table 60. Expected efficiencies for isolated and non-isolated electrons and corresponding jet background rejections for the three standard levels of cuts used for electron identification. The results are shown for simulated inclusive jet samples corresponding to E_T -thresholds of the electron candidates of 17 GeV (left) and 8 GeV (right). The three bottom rows show, for each of the inclusive jet samples, the fractions of all surviving candidates which originate from the different categories for the medium cuts and the two sets of tight cuts. The isolated electrons are prompt electrons from W , Z and top-quark decay and the non-isolated electrons are from b , c decay. The residual jet background is split into its two dominant components, electrons from photon conversions and Dalitz decays (first term in brackets) and charged hadrons (second term in brackets). The quoted errors include part of the systematics, but do not include the larger systematic uncertainties from the physics input and detector simulation.

For all electron and photon candidates, shower-shape variables (lateral and longitudinal shower profiles, etc.) are calculated using the fine granularity of the electromagnetic calorimeter, and typically more than 50 calorimeter cells are summed to collect the full cluster energy. Additionally, combined reconstruction properties, such as the ratio of energy (calorimeter) to momentum (inner detector), the difference between the coordinates η and ϕ reconstructed by the cluster and the track extrapolated into the calorimeter, and the ratio of high-threshold transition radiation hits to low-threshold hits on the track, are used to identify electrons.

The energy of high- p_T electrons is obtained from the energy measured in the calorimeter (the inner-detector momentum measurement is not expected to improve the accuracy of the calorimeter energy measurement significantly for energies above 20–30 GeV). The η and ϕ directions are, however, more precisely determined using the associated track. For photons, everything is derived from the calorimeter information, the energy, the ϕ -direction using the precisely known average transverse position of the primary vertex, and the η -direction as described above.

10.4.2.1 Electrons The standard identification for isolated high- p_T electrons is based on cuts on the shower shapes, on information from the reconstructed track and on the combined reconstruction. Jet rejections are computed with respect to truth-particle jets reconstructed using particle four-momenta within a cone of size $\Delta R = 0.4$. Three sets of cuts have been studied depending on




Figure 246. Overall reconstruction and identification efficiency of various levels of electron cuts: loose, medium, and tight isol. as a function of E_T for single electrons (open symbols) and for isolated electrons in a sample of physics events with a busy environment (full symbols).




Figure 247. Jet rejection as a function of overall reconstruction and identification efficiency for electrons, as obtained using a likelihood method (full circles). The results obtained with the standard cut-based method are also shown in the case of tight TRT (open triangle) and tight isol. (open square) cuts.

the signal efficiency and jet rejection requirements of the physics samples under study:

- "loose cuts" consisting of simple shower-shape cuts (longitudinal leakage, shower shape in the middle layer of the electromagnetic calorimeter) and very loose matching cuts between reconstructed track and calorimeter cluster;
- 9545 • "medium cuts", which add shower-shape cuts using the important information contained in the first layer of the electromagnetic calorimeter and track-quality cuts similar to the standard reconstruction cuts quoted in Section 10.2.3;
- 9550 • "tight cuts", which tighten the track-matching criteria and the cut on the energy-to-momentum ratio. These cuts also explicitly require the presence of a vertexing-layer hit on the track (to further reject photon conversions) and a high ratio between high-threshold and low-threshold hits in the TRT detector (to further reject the background from charged hadrons), as shown in Section 10.2.5. Additionally, further isolation of the electron may be required by using calorimeter energy isolation beyond the cluster itself. Two sets of tight selection cuts are used in this section to illustrate the overall performance of the electron identification. They are labelled as "tight (TRT)", in the case where a TRT cut with approximately 90% efficiency for electrons is applied, and as "tight (isol.)", in the case where a TRT cut with approximately 95% efficiency is applied in combination with a calorimeter isolation cut.
- 9555

The performance of the cut-based analysis is summarised in Table 60 and in Fig. 246 for 9560 electrons. As can be seen from Table 60, the signal from prompt electrons is dominated by initially non-isolated electrons from heavy flavours, which explains the much lower efficiency observed for these electrons. Dedicated algorithms might improve this efficiency somewhat, but these electrons will nevertheless provide the most abundant initial source of isolated electrons and will be used

for alignment of the electromagnetic calorimeters and the inner detector, for E/p calibrations, and more generally to improve the understanding of the material of the inner detector. For tight cuts and an electron p_T of ~ 20 GeV, the isolated electrons from W , Z and top-quark decays represent less than 20% of the total prompt electron signal and are only at the level of ~ 30 –40% of the residual jet background. For the lower E_T -threshold of 8 GeV, the expected signal from isolated electrons is negligibly small. Not surprisingly, the tight TRT cuts are more efficient to select non-isolated electrons from heavy-flavour decay, while the tight isol. cuts are more efficient at selecting isolated electrons. After tight cuts, the signal-to-background ratio is close to 2:1, and depends only weakly on the E_T -threshold. The residual background is dominated by charged hadrons. Further rejection could be possible at the expense of loss of efficiency by stronger cuts (TRT and/or isolation) and by improving the photon conversion reconstruction (see Section 10.2.5).

Figure 246 shows in more detail the overall reconstruction and identification efficiencies for the three sets of electron cuts discussed above: the E_T dependence of the efficiencies is shown for single electrons of fixed E_T as well as for physics processes containing isolated electrons from cascade decays of supersymmetric particles to illustrate the rather stable behaviour of the cuts when moving from the ideal case of single particles to a busy environment with many additional jets in the event. The somewhat worse efficiency observed in complex events is attributed to the fraction of cases when the electron candidate is close to or even within a high- p_T jet. The overall efficiency of the cuts remains stable for even higher electron energies (the efficiency of the tight isol. cuts is 68% for electrons of $E_T = 500$ GeV).

In addition to the traditional cut-based analysis, multivariate techniques have been developed, based on similar variables, and the performance of a likelihood technique is shown as an example in Fig. 247. Compared to the tight cuts described above, a gain of 4–8% in efficiency for the same fixed rejection against jets or of 40–60% in rejection for the same fixed efficiency can be obtained, using this likelihood method for isolated electrons with energies typical of those expected from $Z \rightarrow ee$ decays.

As discussed already to some extent in Section 10.2.5, certain dedicated tracking algorithms improve the momentum reconstruction for electrons with transverse momenta up to 10 GeV. However, as shown in Fig. 248 for electrons with $p_T = 25$ GeV, a significant reduction of the tails due to bremsstrahlung can only be achieved at higher energies by combining the inner-detector measurements with the accurate measurement of the ϕ -position of the electromagnetic shower. This latter constraint, when combined with the extrapolated track impact in the calorimeter, provides enough information to estimate with reasonable accuracy the origin and energy of a hard bremsstrahlung




Figure 248. For electrons with $p_T = 25$ GeV and $|\eta| > 1.5$, integral probability for ratio of true to reconstructed transverse momentum to exceed a given value. The various symbols represent different track-fitting algorithms (see Section 10.2.5) and the bremsstrahlung recovery algorithm, which uses the accurate measurement of the shower position in ϕ in the electromagnetic calorimeter (see text).

Selection cuts	$E_T > 25$ GeV		$E_T > 40$ GeV	
	Quark jets	Gluon jets	Quark jets	Gluon jets
Before isolation	1770 ± 50	15000 ± 700	1610 ± 100	15000 ± 1600
After isolation	2760 ± 100	27500 ± 2000	2900 ± 240	28000 ± 4000

Table 61. Jet rejections obtained before and after applying track-isolation cuts for photon candidates with $E_T > 25$ GeV and $E_T > 40$ GeV and for a photon efficiency of approximately 84%. The rejection values are shown with their statistical errors separately for quark and gluon jets.




Figure 249. For reconstructed photon candidates with $E_T > 25$ GeV (left) and with $E_T > 40$ GeV (right), jet rejection as a function of photon efficiency, as obtained using a likelihood method. The results obtained with the standard cut-based method are also shown for reference.

photon. As shown in Fig. 248, this combined bremsstrahlung recovery procedure will reduce considerably the tails in the E/p distribution, which will be an important tool for studying the uniformity of calibration of the electromagnetic calorimeter, as well as material and alignment effects.

9610 **10.4.2.2 Photons** Photons are much harder to extract as a signal from the jet background than
certain specific isolated electron signals, such as those expected from $Z \rightarrow ee$ or $W \rightarrow ev$ decays.
A single set of photon identification cuts, equivalent to the "tight cuts" defined for electrons, has
been optimised based on the shower shapes in the calorimeter with special emphasis on separating
single π^0 's from photons using the very fine granularity in η of the strip layer. In addition, a
9615 simple track-isolation criterion has been added to further improve the rejection while preserving
the vast majority of converted photons. Using these criteria, an efficiency of 84% has been obtained
for photons with an energy spectrum as expected from $H \rightarrow \gamma\gamma$ decay with $m_H = 120$ GeV. This
efficiency is quite uniform over the whole η -range except for the crack between the barrel and
end-cap calorimeters mentioned above. For this value of the photon efficiency, a jet rejection
9620 of ~ 5000 (without track isolation) to 9000 (with track isolation) has been achieved, averaged over
the parton flavours corresponding to the inclusive di-jet background sample used. The expected
jet rejections are shown in Table 61 separately for quarks and gluons and for two relevant values
of the E_T -threshold applied to the photon candidates. The larger rejection expected against gluon
jets is due to the softer fragmentation and therefore broader lateral extent of gluon jets compared to




Figure 250. Expected distribution for the invariant mass of the four electrons from Higgs-boson decays with $m_H = 130$ GeV. The energies of the electrons are determined only from the electromagnetic calorimeter measurements. The results do not include a Z -mass constraint.




Figure 251. Expected distribution for the invariant mass of the two photons from Higgs-boson decays with $m_H = 120$ GeV. The shaded plot corresponds to events in which at least one of the two photons converted at a radius below 80 cm.

9625 light jets which are dominant in the quark-jet sample. The residual background from jets is mostly
 composed of isolated π^0 's, so the fine-grained strip compartment of the electromagnetic calorimeter
 is an important element to achieve such rejections. As for the electrons, the jet rejections are
 computed with respect to truth-particle jets reconstructed using particle four-momenta within a
 cone of size $\Delta R = 0.4$.

9630 Multivariate methods have also been developed for the more difficult case of photon identi-
 fication. These can be seen in Fig. 249, which shows as an example the expected performance
 for a likelihood technique compared to the standard cut-based analysis. For photon candidates
 with $E_T > 25$ GeV and a fixed efficiency of 84%, the rejection with respect to the cut-based selec-
 tion is improved by 6% for the likelihood method.

9635 **10.4.2.3 Reconstruction of $H \rightarrow eee$ and $H \rightarrow \gamma\gamma$ final states** The performance of the recon-
 struction, including calibration, with the identification criteria discussed above is shown in Fig. 250
 for decays of a Higgs boson with a mass of 130 GeV to four electrons (loose electron cuts applied)
 and in Fig. 251 for decays of a Higgs boson with a mass of 120 GeV to two photons (tight pho-
 ton cuts applied and barrel/end-cap transition region excluded). A global constant term of 0.7%
 9640 has been included in the electromagnetic calorimeter resolution for these plots. In the case of
 $H \rightarrow \gamma\gamma$ decays, the photon directions are derived from a combination of the direction measure-
 ment in the electromagnetic calorimeter described above (see Fig. 245) with the primary vertex
 information from the inner detector (see Table 59).

In the case of the Higgs-boson decay to four electrons, the central value of the reconstructed
 9645 invariant mass is correct to ~ 1 GeV, corresponding to a precision of 0.7%, and the expected
 Gaussian resolution is $\sim 1.5\%$. The non-Gaussian tails in the distribution amount to 20% of events
 which lie further than 2σ away from the peak. They are mostly due to bremsstrahlung, particularly
 in the innermost layers of the inner detector, but also to radiative decays and to electrons poorly
 measured in the barrel/end-cap transition region of the electromagnetic calorimeter.




Figure 252. Distribution of the transverse energy accumulated in $\Delta\eta \times \Delta\phi = 0.1 \times 0.025$ middle-layer regions with a few hours of minimum bias events. The full histogram corresponds to the hemisphere with a nominal amount of inner-detector material in the simulation for $1.8 < \eta < 1.9$, whereas the dotted histogram corresponds to the hemisphere with a 25% increase in the amount of material in the same η -region.




Figure 253. Distribution of the fraction of energy deposited in the strip compartment by electrons from W/Z decays corresponding to the statistics expected for an integrated luminosity of 50 pb^{-1} . The full histogram corresponds to the hemisphere with a nominal amount of inner-detector material in the simulation for $1.8 < \eta < 1.9$, whereas the dotted histogram corresponds to the hemisphere with a 25% increase in the amount of material in the same η -region.




Figure 254. Statistical accuracy expected from inter-calibration of the electromagnetic calorimeter as a function of the number of reconstructed $Z \rightarrow ee$ decays or of the integrated luminosity (see text). These results assume a perfect knowledge of the material in front of the electromagnetic calorimeter.




Figure 255. Expected differential cross-section for low-mass electron pairs using the 2e5 trigger menu item discussed in Section 10.9.3. Shown is the invariant di-electron mass distribution reconstructed using tracks for $J/\psi \rightarrow ee$ decays (dotted histogram), $\Upsilon \rightarrow ee$ decays (dashed histogram) and Drell-Yan production (full histogram). Also shown is the expected background after the offline selection described in the text (full circles).

9650

In the case of the Higgs-boson decay to two photons, the central value of the reconstructed invariant mass is correct to $\sim 0.2 \text{ GeV}$, corresponding to a precision of 0.3%, and the expected resolution is $\sim 1.2\%$. Figure 251 also clearly shows that most of the non-Gaussian tails at low values of the reconstructed mass of the photon pair are due to photons which converted in the inner detector.

9655 **10.4.3 Assessment of performance in situ with initial data**

One important ingredient in the calibration strategy for the electromagnetic calorimeter is the use of large-statistics samples of $Z \rightarrow ee$ decays to perform an accurate inter-calibration of regions with a fixed size of $\Delta\eta \times \Delta\phi = 0.2 \times 0.4$ [278]. It is expected that such a scheme will decrease the initial spread from region to region, conservatively assumed to be approximately 1.5–2%, to values comparable to the expected constant term of $\sim 0.5\%$ in each region. This however assumes an excellent knowledge of the material in front of the electromagnetic calorimeter. The material in the inner detector should be eventually mapped out very accurately using e.g. photon conversions, but other less sensitive but more robust methods will also be used, exploiting the high granularity of the electromagnetic calorimeter. The energy flow measured in the second layer of the electromagnetic calorimeter, for example in minimum-bias events, provides such a tool, as illustrated in Fig. 252. Only energy deposits more than 5σ above the electronic noise level are considered for these measurements. With approximately two million minimum-bias events, corresponding to roughly one day of data-taking, additional material inside the inner detector amounting to 20% X_0 would be identified in any region of size $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ [279]. Another more sensitive possibility is the study of the ϕ -symmetry of the fraction of energy deposited in the first layer of the electromagnetic calorimeter by isolated electrons, as shown in Fig. 253. Combining this information with that from the other layers in the calorimeter and with the momentum measurement of the electrons will provide higher sensitivity (for example in η) than the minimum-bias results.

Figure 254 shows the result of such an inter-calibration procedure applied to simulated $Z \rightarrow ee$ decays with an initial 2% spread from region to region. Once the material in front of the electromagnetic calorimeter is sufficiently well understood, an inter-calibration accuracy of 0.7% could be achieved for a total of approximately 50,000 $Z \rightarrow ee$ decays, reconstructed with the medium set of identification cuts described above, and corresponding to an integrated luminosity of $\sim 150 \text{ pb}^{-1}$.

As described in Section 10.9.3 for initial luminosities of $10^{31} \text{ cm}^{-2} \text{ s}^{-1}$, a trigger on low-mass di-electron pairs (the 2e5 signature in Table 64) should provide good statistics of $J/\psi \rightarrow ee$ and $\Upsilon \rightarrow ee$ decays. An example of the signal and background samples which will be provided by the low-mass pair di-electron trigger in early data is shown in Fig. 255. For this study, the track-seeded algorithm introduced in Section 10.4.1 has been used with tight electron cuts as described above. The signal-to-background ratio obtained is larger than one at the J/ψ and Υ peaks, but the extraction of electron pairs from Drell-Yan will require further studies (tighter identification or kinematic cuts). With an integrated luminosity of 100 pb^{-1} and an efficient identification and reconstruction of these low-mass pairs, approximately 100,000 J/ψ decays and 30,000 Υ decays could be isolated for detailed studies of the electron identification and reconstruction performance, in particular in terms of matching energy and momentum measurements at a scale quite different from that of the more commonly used $Z \rightarrow ee$ decays.

9680 **10.5 Jet reconstruction**

The ATLAS calorimeters have very high lateral granularity and several samplings in depth over $|\eta| < 3.2$ (see Table 3 for an overview of the properties of the various ATLAS calorimeters). The forward calorimeters, which cover the region $3.2 < |\eta| < 4.9$, also provide sufficient granularity to reconstruct jets with small polar angles with reasonable accuracy and efficiency. For the re-

construction of jets in the wide variety of physics processes of interest at the LHC, specific care has therefore been taken to devise a modular and generic design of the corresponding software. The implementation allows for the use of a variety of jet clustering algorithms using as input any reconstruction object having a four-momentum representation. These inputs can vary from calorimeter cells, or charged tracks, to Monte-Carlo truth objects, such as stable particles or final-state partons from the generator. It also supports easy implementation of jet-clustering algorithms different from the ones most commonly used, and has followed the guidelines collected for Run II at the Tevatron [280].

10.5.1 Jet clustering algorithms

The two default jet-clustering algorithms in ATLAS are a seeded fixed-cone algorithm and a successive recombination algorithm. Both algorithms are used in two different configurations, one producing narrow jets for e.g. W -mass spectroscopy in $t\bar{t}$ events or events containing large multiplicities of jets as in supersymmetric models, and the other producing wider jets for e.g. QCD studies of di-jet and multi-jet final states at luminosities below $10^{33} \text{ cm}^{-2} \text{ s}^{-1}$.

The seeded cone algorithm uses two parameters, the transverse energy threshold for a seed, $E_T = 1 \text{ GeV}$ for all cone jets, and the cone size, $R_{\text{cone}} = \sqrt{\Delta\eta^2 + \Delta\phi^2}$, with $R_{\text{cone}} = 0.4$ for narrow jets and $R_{\text{cone}} = 0.7$ for wide jets. In all cases, a split-and-merge step follows the actual cone building, with an overlap fraction threshold of 50%. The cone algorithm in this particular implementation is fast and therefore also used in the high-level trigger (see Section 10.9).

The k_{\perp} algorithm in ATLAS is implemented following the suggestions in [281], which makes it efficient even for a rather large number of input objects and avoids the usual pre-clustering step. The distance parameter $R = \sqrt{\Delta\eta^2 + \Delta\phi^2}$ is adjusted for narrow jets to $R = 0.4$ and for wide jets to $R = 0.6$. The physics performance is very similar to the one of the corresponding cone configurations. In all cases the full four-momentum recombination is used to calculate the jet kinematics after each clustering step.

10.5.2 Input to jet reconstruction

Typical inputs for jet-finding in ATLAS are final-state particles for truth-particle jets, and calorimeter signals for reconstructed or calorimeter jets. Naturally, truth-particle jets are only available in simulated data. They are formed by applying a jet algorithm to all stable neutral and charged particles in the final state within $|\eta| < 5$. These particles can emerge from the hadronisation of the hard-scattered parton, from initial- and final-state radiation, and from the underlying multiple interactions in the event. The kinematic properties of these particles are taken at their generation vertex, before any interaction with the detector and its magnetic field.

Figure 256 presents an overview of the reconstruction flow for calorimeter jets. Calorimeter jets are reconstructed by applying a jet-clustering algorithm to calorimeter signals, typically followed by a calibration step. Two different signals from the calorimeter are used for jet-finding, towers and topological clusters. Towers are formed by collecting cells into bins of a regular $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ grid, depending on their location, and summing up their signals, or a fraction of their signal corresponding to the overlap area fraction between the tower bin and the cell in $\Delta\eta$ and $\Delta\phi$. This summing stage is non-discriminatory, meaning all calorimeter cells are used in the towers. Towers with negative signals are dominated by noise, and cannot be used in jet-finding.

They are recombined with nearby positive signal towers until the net signal is positive, i.e. the resulting towers have a valid physical four-vector and can directly be used by the jet finders. This approach can be understood as an overall noise cancellation rather than suppression, since the noisy cells still contribute to the jets at initial luminosities of $10^{31} \text{ cm}^{-2} \text{ s}^{-1}$ to $10^{33} \text{ cm}^{-2} \text{ s}^{-1}$.

Topological cell clusters represent an attempt to reconstruct three-dimensional energy depositions in the calorimeter [282, 154]. First, nearest neighbours are collected around seed cells with a significant absolute signal above the major seed threshold, i.e. $|E_{cell}| > 4\sigma_{cell}$ of the total noise (electronics plus pile-up). Energy equivalents of the σ of the electronic noise alone in the various calorimeter cells are shown in Fig. 100, while Fig. 257 shows estimates for the total σ when fluctuations from pile-up at a luminosity of $2 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ are included at the cell level. These neighbouring cells are collected independently of the magnitude of their own signal. If the absolute value of their signal significance is above a secondary seed threshold, typically such that $|E_{cell}| > 2\sigma_{cell}$, they are considered secondary seeds and their direct neighbours are also collected. Finally, all surrounding cells above a very low threshold (typically set to 0σ) are added if no more secondary seeds are among the direct neighbours. A final analysis of the resulting cluster looks for multiple local signal maxima. In case of more than one maximum in a given cluster, it is split into smaller clusters, again in three dimensions, along the signal valleys between the maxima.




Figure 256. Jet reconstruction flow for calorimeter jets from towers or clusters.

Contrary to the signal tower formation, topological cell clustering includes actual noise suppression, meaning that cells with no signal at all are most likely already not included in the cluster. This results in substantially less noise, as shown in Fig. 258, and less cells, as shown in Fig. 259, in these cluster jets than in tower jets. Topological cell clusters are under study for use as the basis for the local hadronic energy calibration, which attempts to correct for detector effects, such as calorimeter responses with $e/h > 1$ and dead-material energy losses, outside of the jet context itself. Although very promising in terms of noise suppression, the topological cell clustering will require careful validation with real data, in particular in terms of the possible impact of long-range noise correlations and of detailed studies of pile-up effects as the luminosity increases.




Figure 257. Expected noise from the electronics and pile-up at $2 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ in individual cells of the various compartments of the calorimeters as a function of $|\eta|$. See Fig. 100 for the pure electronic noise expected from the various compartments of the calorimeters. Note that the presampler noise is corrected for by the appropriate sampling fractions as discussed in Section 5.5.2.1.




Figure 258. Average electronic noise contribution to cone jets with $R_{cone} = 0.7$ in QCD di-jet events, reconstructed from towers (open circles) and topological cell clusters (full circles), as a function of $|\eta|$.




Figure 259. Average total number of cells contributing to cone jets with $R_{cone} = 0.7$ in QCD di-jet events, reconstructed from towers (open circles) and topological cell clusters (full circles), as a function of the jet energy.

10.5.3 Jet calibration

The strategy currently adopted for calorimeter jet calibration in ATLAS is the application of cell signal weighting similar to the original approach developed for the H1 calorimeter [283]: all

calorimeter cells with four-momenta (E_i, \vec{p}_i) , where $E_i = |\vec{p}_i|$, in tower or cluster jets are considered and re-summed with weighting functions, w , such that the resulting new jet four-momentum is:

$$(E_{rec}, \vec{p}_{rec}^{jet}) = \left(\sum_i^{N_{cells}} w(\rho_i, \vec{X}_i) E_i, \sum_i^{N_{cells}} w(\rho_i, \vec{X}_i) \vec{p}_i \right). \quad (10.2)$$

The weighting functions w depend on the cell signal density, $\rho_i = E_i/V_i$, and on the cell location in the calorimeter, \vec{X}_i , consisting basically of module and compartment identifiers. They are fitted using simulated QCD di-jet events, covering the whole kinematic range expected at the LHC, and matching calorimeter cone-tower jets, with $R_{cone} = 0.7$, with nearby truth-particle cone jets of the same size and with energy E_{truth} , and then constraining E_{rec} in Eq. 10.2 to E_{truth} by:

$$\frac{\partial \chi^2}{\partial w(\rho_i, \vec{X}_i)} = \frac{\partial}{\partial w(\rho_i, \vec{X}_i)} \left[\sum_{\text{matched jets}} \frac{((E_{rec} + E_{DM}) - E_{truth})^2}{E_{truth}} \right] = 0. \quad (10.3)$$

The weighting functions determined in this way absorb all detector effects, including missing signals from charged truth particles with less than ~ 400 MeV transverse momentum, which are bent away from the calorimeter by the solenoid magnetic field in the inner detector cavity. Implicitly included also are corrections for energy loss in inactive materials, except for losses between the electromagnetic barrel and tile barrel calorimeters, which are parametrised in Eq. 10.3 as:

$$E_{DM} = \alpha \sqrt{E_{EMB3} E_{TILE0}}, \quad (10.4)$$

where E_{EMB3} is the sum of the energies of the cells in the last compartment of the barrel electromagnetic calorimeter belonging to the jet and E_{TILE0} is the corresponding sum in the first compartment of the hadronic tile calorimeter. Both quantities are reconstructed at the electromagnetic energy scale. The parameter α was assumed to be independent of energy and of η and was determined together with $w(\rho_i, \vec{X}_i)$ in a combined fit according to eq.(10.3).

Naturally, the calibration applied in this way only corrects to the level of the truth-particle jet. The extracted weighting functions were obtained for cone-tower jets with $R_{cone} = 0.7$ and are not universal, since they depend on the choice of calorimeter signals used, on the jet algorithm chosen and on its specific configuration, and on the choice of (simulated) physics calibration samples used to extract them. Residual mis-calibrations for all cluster cone and cluster k_{\perp} jets are corrected for by functions depending on $|\eta|$ and p_T of each measured jet. Similar corrections are applied to tower cone jets with $R_{cone} = 0.4$ and to the tower k_{\perp} jets. These corrections have been derived by comparing the calorimeter jets after applying the cell signal weights and dead-material corrections with the matching truth-particle jet in the simulated QCD di-jet events.

10.5.4 Jet signal characteristics

All signal features discussed in the following are extracted from simulations including a model for the electronic noise in each calorimeter cell, tuned with parameters extracted from various test-beam measurements. The results shown here are, unless stated otherwise, based on the jet-calibration procedure described above, called from now on global calibration. Pile-up fluctuations are not included.




Figure 260. Signal linearity for cone-tower jets with $R_{cone} = 0.7$, as expressed by the ratio of reconstructed tower jet energy to the matching truth-jet energy E_{rec}/E_{truth} , in two different regions of $|\eta|$ and as a function of E_{truth} . Jet signals calibrated at the electromagnetic energy scale are compared to the fully calibrated jets.




Figure 261. Fractional energy resolution for calibrated cone-tower jets reconstructed with $R_{cone} = 0.7$ and $R_{cone} = 0.4$ in two different regions of $|\eta|$ and as a function of E_{truth} .




Figure 262. For cone-tower jets reconstructed with $R_{cone} = 0.7$, distribution of ΔR between reconstructed and matched particle jet axes for two different transverse energy and η -ranges.




Figure 263. For cone-cluster jets reconstructed with $R_{cone} = 0.7$, distribution of ΔR between reconstructed and matched particle jet axes for two different transverse energy and η -ranges.

The most important requirements for the jet signal after global calibration are a linear response across all jet energies, a uniform response as independent as possible from the jet direction, and a fractional energy resolution within the specifications laid out in Table 1.

9800 **10.5.4.1 Jet signal linearity and energy resolution** The signal linearity for calorimeter jets in ATLAS is expressed by the ratio of the reconstructed jet energy and the matched truth-jet energy, E_{rec}/E_{truth} , in simulated QCD di-jet events.

9805 Figure 260 shows, for two different regions in $|\eta|$, that the the signal linearity for cone jets made from towers with $R_{cone} = 0.7$ is reasonable over the whole energy range after the global calibration is applied. Figure 260 also shows the deviations from signal linearity expected for jets reconstructed at the electromagnetic energy scale, i.e. without any hadronic calibration applied. In this case, the reconstructed jet signals correspond to only $\sim 50\%$ (at the lowest energies) to $\sim 80\%$ (at the highest energies) of the true jet energy.




Figure 264. Signal uniformity for QCD di-jets in two different E_T ranges, as a function of $|\eta|$ of the matched truth-particle jet. The results are shown for cone-tower jets with $R_{cone} = 0.7$ and $R_{cone} = 0.4$.




Figure 265. Jet energy resolution for QCD di-jets in two different E_T ranges, as a function of $|\eta|$ of the matched truth-particle jet. The results are shown for cone-tower jets with $R_{cone} = 0.7$ and $R_{cone} = 0.4$.

The fractional energy resolution for the same jets, again after global calibration, is shown as a function of E_{truth} and for two different η -regions in Fig. 261. In addition, the resolution for a smaller cone size $R_{cone} = 0.4$ is shown. The curves show the results of a three-parameter fit to the energy resolution function:

$$\frac{\sigma}{E} = \sqrt{\frac{a^2}{E} + \frac{b^2}{E^2} + c^2} \quad (10.5)$$

For central jets in the region $0.2 < |\eta| < 0.4$, the stochastic term is $\approx 60\% \sqrt{\text{GeV}}$, while the high-energy limit of the resolution, expressed by the constant term c , is $\approx 3\%$ with the current global calibration. One important contribution to the η -dependence of the jet energy resolution is the noise, which varies quite rapidly due to the increasing readout-cell size and the change in calorimeter technology in the hadronic calorimeters from the low-noise tile calorimeter to the (higher-noise) LAr calorimeter with increasing η . The noise term b in the energy resolution function is found to increase from 0.5 GeV to 1.5 GeV when going from the barrel to the end-cap η -ranges shown in Fig. 261.

10.5.4.2 Jet direction measurement The highly granular ATLAS calorimeters provide a precise measurement not only of the jet energy, but also of the jet direction together with the knowledge of the primary vertex position from the inner detector (see Section 10.2.4). Figures 262 and 263 show very similar distributions of the distance ΔR between reconstructed and matched truth-particle jet directions for tower and cluster cone jets with $R_{cone} = 0.7$ in two different pseudo-rapidity regions and two different transverse energy ranges. The choice of calorimeter signal obviously does not significantly affect the direction reconstruction of the jet. The general conclusion is that for both tower and cluster jets with transverse energies above 100 GeV, basically all reconstructed jets fall within the default matching cuts, $\Delta R < 0.2$. At lower transverse energies, however, it is clear that the precision with which the jet axis is reconstructed is degraded and a non-negligible fraction of reconstructed jets will fall outside the default matching cut. This issue is rediscussed below in Section 10.5.5 with wider matching cuts to assess the efficiency and purity of reconstruction of low- p_T jets.




Figure 266. Efficiency of jet reconstruction in VBF-produced Higgs-boson events as a function of p_T of the truth-particle jet for cone-tower and cone-cluster jets with $R_{cone} = 0.7$.




Figure 267. Purity of jet reconstruction in VBF-produced Higgs-boson events as a function of p_T of the reconstructed jet for cone-tower and cone-cluster jets with $R_{cone} = 0.7$.

9830 **10.5.4.3 Jet signal uniformity** The variation of the jet energy response as a function of the jet
direction is a measure of the uniformity of the jet signal across the full rapidity coverage of the
calorimeters. Figure 264 shows for tower jets the ratio of reconstructed to matching truth-particle
energy as a function of $|\eta|$ for jets in two different bins of E_T . The dips in response, corresponding
to the two transition regions, $1.2 < |\eta| < 2.0$ and $2.8 < |\eta| < 3.4$, are much more apparent at low
9835 transverse energies. The dip in response in the last η -bin is a reflection of the limited fiducial cov-
erage of jet reconstruction for $|\eta| > 4.4$. The η -dependence of the corresponding fractional energy
resolution in Fig. 265 can be understood: the energy E_{jet} of jets with $30 < p_T < 40$ GeV increases
from $E_{jet} = 30$ GeV at $|\eta| = 0$ to $E_{jet} \approx 1.8$ TeV at $|\eta| = 4.5$. Following the parametrisation
in Eq.(10.5), the fractional energy resolution, σ/E , improves dramatically over this energy range.
9840 The energy-dependent stochastic, a/\sqrt{E} , and noise, b/E , terms dominate over a large part of the
kinematic regime. For jets with $480 < p_T < 640$ GeV, the jet energy range is $480 \leq E_{jet} < 7000$
GeV from $|\eta| = 0$ up to $|\eta| \approx 3.1$, which is the kinematic limit at the LHC. In this region, σ/E is
essentially independent of E_{jet} , i.e. dominated by the constant term, $c \gg a/\sqrt{E} \gg b/E$.

10.5.5 Jet reconstruction performance

9845 The evaluation of the jet reconstruction performance includes not only the required signal features
discussed above, but also parameters which are more oriented towards physics analysis, such as
jet-finding efficiency and purity, jet vetoing, and jet tagging.

The jet reconstruction efficiency is defined as:

$$\varepsilon(R_m) = \frac{\# \text{ matches of truth particle jets with reconstructed jets}}{\# \text{ truth particle jets}} = \frac{N_m^{jets}(R_m)}{N_{truth}^{jets}}, \quad (10.6)$$

where $R_m = \sqrt{(\eta_{reco} - \eta_{truth})^2 + (\phi_{reco} - \phi_{truth})^2}$ is the chosen matching radius (typically $R_m = 0.2$). The purity π of the jet reconstruction can be expressed as:

$$\pi(R_m) = \frac{\# \text{ matches of truth particle jets with reconstructed jets}}{\# \text{ reconstructed jets}} = \frac{N_m^{jets}(R_m)}{N_{reco}^{jets}}. \quad (10.7)$$




Figure 268. Efficiency of jet reconstruction in VBF-produced Higgs-boson events as a function of the rapidity y of the truth-particle jet for $p_T^{jet} > 10$ GeV and $p_T^{jet} > 20$ GeV and for cone-tower and cone-cluster jets with $R_{cone} = 0.7$.




Figure 269. Purity of jet reconstruction in VBF-produced Higgs-boson events as a function of the rapidity y of the reconstructed jet for $p_T^{jet} > 10$ GeV and $p_T^{jet} > 20$ GeV and for cone-tower and cone-cluster jets with $R_{cone} = 0.7$.

The fake jet reconstruction rate f is then simply $f = 1 - \pi(R_m)$. In all cases, only one match is allowed for each reference jet. In case of two or more nearby jets, the one closest to the chosen reference is taken.

The two different calorimeter signal definitions used for jet reconstruction (towers and clusters) are expected to produce different efficiencies and purities. This is particularly important for searches for specific exclusive final states, where the requirement that no additional jet be present in the event is often used as a powerful tool to reject certain backgrounds. For example, one of the interesting production channels for the Higgs boson is vector-boson fusion (VBF), which has a very characteristic final state with two forward-going quark jets (often called tag jets) and, for non-hadronic Higgs-boson decay modes, no jets from the hard-scattering process itself in the central region of the detector. In this specific case, the efficiency of the jet-finding in the forward region, as defined in Eq. 10.6, is a measure of the jet-tagging probability. The purity of the jet reconstruction in the central region then measures the efficiency for vetoing low- p_T jets.

The resulting efficiencies and purities are shown for cone-tower and cone-cluster jets with $R_{cone} = 0.7$, respectively, as a function of p_T and y in Figs. 266, 267, 268 and 269 for the specific case of VBF produced $H \rightarrow \tau\tau$ decays with $m_H = 120$ GeV and for a looser matching radius $R_m = 0.5$ (see Eqs. 10.6 and 10.7). These results show that for $p_T > 40$ GeV, the performances of the tower and cluster jets are very similar. For lower values of p_T , however, the cluster jets are found with both higher efficiency and purity than tower jets.

For jets reconstructed with $p_T > 10$ GeV, the fake rates in the central region are quite high, ranging from 30% for cluster jets to 45% for tower jets. In the forward regions, the jet-tagging efficiencies are close to 90% for cluster jets while they are only around 50% for tower jets with, however, significantly higher fake rates of $\sim 10\%$ for the cluster jets. These results are clearly also quite sensitive to pile-up, so it is important to stress here that the numbers above apply only for initial data-taking at luminosities between $10^{31} \text{ cm}^{-2} \text{ s}^{-1}$ and $10^{33} \text{ cm}^{-2} \text{ s}^{-1}$.




Figure 270. Jet response for seeded cone-tower jets ($R_{cone} = 0.4$) in $\gamma + \text{jet}$ events, averaged over η and calculated by the missing transverse momentum fraction method, as a function of the jet energy. The calorimeter signals are reconstructed at the electromagnetic energy scale.




Figure 271. Jet response for seeded cone-tower jets ($R_{cone} = 0.4$) from the same events and again calculated by the missing transverse momentum fraction method, as a function of the jet direction, $|\eta_{jet}|$. The degraded response in the calorimeter crack regions is clearly visible.

10.5.6 Validation of jet calibration with in-situ measurements

9875 There are several final states at the LHC which provide signals for validation of the jet energy calibration, and, in some cases, even the extraction of further corrections. In general, final states with a well measured electromagnetic object balancing one or more jets in transverse momentum, such as in $\gamma + \text{jet(s)}$ and $Z + \text{jet(s)}$ events, are good choices for this task. The $\gamma + \text{jet(s)}$ process provides high statistics in the transverse momentum range from 40 to 400 GeV, but lower purity than the $Z + \text{jet(s)}$ process, which should, however, cover precisely the lower edge of the transverse momentum range, up to 100 GeV.

9885 As an example, one approach to measure the jet response using $\gamma + \text{jet(s)}$, which has been developed at the Tevatron, is the missing transverse momentum projection fraction. The basic idea of this method is to project the hadronic transverse-momentum vectors onto the transverse-momentum vector of the photon and to measure the apparent E_T^{miss} fraction. In events where the photon is back-to-back with the jet (to better than approximately ten degrees in ϕ , the jet response R_{jet} can then be determined by

$$R_{jet} = - \frac{\sum_{\text{signals}} \vec{p}_{T,had} \cdot \hat{\eta}_\gamma}{p_{T,\gamma}}. \quad (10.8)$$




Figure 272. Ratio of the reconstructed di-jet mass from $W \rightarrow jj$ decays in $t\bar{t}$ events to the nominal mass as a function of the transverse momentum of the W -boson, p_T^W , for globally calibrated cone-tower jets with $R_{cone} = 0.7$. Shown are the results for the nominal jet-selection cuts, $p_T > 40$ GeV (open circles), for jets reconstructed with $p_T > 10$ GeV (open squares) and for jets re-scaled to obtain a more uniform response as a function of $|\eta|$ (full triangles).

Here $\hat{\eta}_\gamma = \vec{p}_{T,\gamma}/p_{T,\gamma}$ is the direction of the photon in the transverse event plane. The hadronic transverse momentum can be calculated using the reconstructed jet(s) ($\vec{p}_{T,had} = \vec{p}_{T,jet}$), or just using the sum of cluster signals without the jet context ($\vec{p}_{T,had} = \vec{p}_{T,calo}$). Figure 270 shows the jet response for cone-tower jets with $R_{cone} = 0.4$ at the electromagnetic energy scale, as a function of the jet energy, for simulated $\gamma + \text{jet}$ events. This variable can be measured directly and can thus become the basis for a global jet energy scale calibration derived from collision data. The η -dependence of the jet response for the same jets and events is shown in Fig.271. The shape of the response clearly indicates the effect of the crack regions of the ATLAS calorimeter system on the jet energy measurement.

The other important final state for jet calibration are hadronically decaying W bosons ($W \rightarrow qq$), which in ATLAS can only be used with high purity in $t\bar{t}$ production. Here, m_W constrains the energy scale of the two quark jets. Figure 272 shows the ratio between the reconstructed di-jet mass from $W \rightarrow jj$ decays and the nominal W -boson mass as a function of the true transverse momentum, p_T^W , of the W -boson. For the nominal selection cuts used to reconstruct $t\bar{t}$ events, this ratio departs significantly from unity at low values of p_T^W because of the high p_T -threshold of 40 GeV applied to the jets, as illustrated in Fig. 272. With further in-situ corrections aimed at re-scaling jet energies as a function of $|\eta|$ to obtain a uniform response, e.g. as shown in Fig. 271, a linearity of better than 2% can be achieved up to values of p_T^W as high as 200 GeV.

10.6 Missing transverse energy

A very good measurement of the missing transverse energy, E_T^{miss} , is a critical requirement for the study of many physics channels in ATLAS, in particular in the search for signals from new physics such as supersymmetry or extra dimensions. A good E_T^{miss} measurement in terms of linearity and accuracy is also important for the reconstruction of the top-quark mass from $t\bar{t}$ events with one top quark decaying semi-leptonically. It is crucial for the efficient and accurate reconstruction of the Higgs-boson mass when the Higgs boson decays to a pair of τ -leptons. Another important requirement on the measurement of E_T^{miss} is to minimise the impact of tails induced by imperfections in the detector coverage or detector response. The η -coverage of the forward calorimeters minimises by design any tails from particles escaping at very large η , but there are several transition regions in the calorimetry, which will lead to incorrect measurements of E_T^{miss} in a certain fraction of the cases. This could significantly enhance for example the backgrounds from QCD multi-jet events to a possible signal from supersymmetry or the backgrounds from $Z \rightarrow ll$ decays accompanied by high- p_T jets to a possible signal from Higgs-boson decay into two leptons and two neutrinos. This section describes briefly the reconstruction and calibration of E_T^{miss} in ATLAS, illustrates the expected performance with a few examples, and finally concludes with a discussion of the possible sources of fake E_T^{miss} .

10.6.1 Reconstruction and calibration of E_T^{miss}

The E_T^{miss} reconstruction in ATLAS is based in a first step on the calibrated calorimeter cell energies (following the global calibration scheme described in Section 10.5.3) and on the reconstructed muons. The E_T^{miss} muon term is calculated from the momenta of the muons measured using the stand-alone muon-spectrometer reconstruction (see Section 10.3). Energy lost by muons in the calorimeter is thus not double-counted, since it is already taken into account in the calorimeter




Figure 273. Linearity of response for reconstructed E_T^{miss} as a function of the average true E_T^{miss} for different physics processes covering a wide range of true E_T^{miss} and for the different steps of E_T^{miss} reconstruction (see text). The points at average true E_T^{miss} of 20 GeV are from $Z \rightarrow \tau\tau$ events, those at 35 GeV are from $W \rightarrow e\nu$ and $W \rightarrow \mu\nu$ events, those at 68 GeV are from semi-leptonic $t\bar{t}$ events, those at 124 GeV are from $A \rightarrow \tau\tau$ events with $m_A = 800$ GeV, and those at 280 GeV are from events containing supersymmetric particles at a mass scale of 1 TeV (left). Linearity of response for reconstructed E_T^{miss} as a function of the true E_T^{miss} for $A \rightarrow \tau\tau$ events with $m_A = 800$ GeV (right).




Figure 274. Resolution σ of the two components of the E_T^{miss} vector after refined calibration as a function of the total transverse energy, ΣE_T , measured in the calorimeters for different physics processes corresponding to low to medium values of ΣE_T (left) and to higher values of ΣE_T (right). The curves correspond respectively to the best fit, $\sigma = 0.53\sqrt{\Sigma E_T}$, through the points from $Z \rightarrow \tau\tau$ events (left) and to the best fit, $\sigma = 0.57\sqrt{\Sigma E_T}$, through the points from $A \rightarrow \tau\tau$ events (right). The points from $A \rightarrow \tau\tau$ events are for masses m_A ranging from 150 to 800 GeV and the points from QCD jets correspond to di-jet events with $560 < p_T < 1120$ GeV.

9930 term. Only good-quality muons with a matched track in the inner detector are considered, which reduces considerably possible contributions from fake muons, sometimes created from high hit multiplicities in the muon spectrometer in events with very energetic jets.

9935 In a second step, the E_T^{miss} reconstruction accounts for the so-called cryostat term, which corrects for the energy lost in the cryostat between the barrel LAr electromagnetic and tile calorimeters. This correction is applied following the recipe described in Section 10.5.3 and Eq. 10.4 and is found to be non-negligible for high- p_T jets: it represents a 5% contribution per jet with p_T

above 500 GeV.

In a final step, a refined calibration of E_T^{miss} is performed through the association of each high- p_T object in the event to its globally calibrated cells. Starting from the reconstructed identified objects in a carefully chosen order, namely electrons, photons, hadronically decaying τ -leptons, b -jets, light jets and muons, each calorimeter cell is associated to its parent high- p_T object. The refined calibration of E_T^{miss} then replaces the initial contribution from globally calibrated cells by the contribution from the corresponding calibrated high- p_T objects themselves. The cells which survive a noise cut optimised in terms of E_T^{miss} measurements and which do not contribute to any reconstructed object are also calibrated using the global calibration scheme and accounted for in the E_T^{miss} calculation.

10.6.2 Evaluation of E_T^{miss} performance

The E_T^{miss} performance is evaluated by comparing the final reconstructed and calibrated value of E_T^{miss} with the true E_T^{miss} , calculated using all stable and non-interacting particles in the final state, for a number of physics processes of interest, involving a variety of topologies and final states over a wide range of energies. Although this evaluation focuses primarily on the linearity of response and on resolution, other features, such as the direction of the E_T^{miss} vector (in the transverse plane) and tails in the measurement of E_T^{miss} have also been carefully studied.

The expected performance in terms of E_T^{miss} linearity of response as a function of true E_T^{miss} is shown for a number of physics processes of interest in Fig. 273. The evolution of the linearity of response is illustrated for each of the major steps in the E_T^{miss} reconstruction described above:

- the uncalibrated E_T^{miss} corresponds to the use of cell energies at the electromagnetic scale, which therefore creates a large systematic bias of 10–30% in the response (the bias is smaller for events containing little hadronic activity on average, such as $W \rightarrow e\nu$ and $W \rightarrow \mu\nu$ decays);
- the reconstructed E_T^{miss} based on globally calibrated cell energies and reconstructed muons provides a correct response to within 5%;
- the reconstructed E_T^{miss} including in addition the cryostat correction provides excellent linearity of response for all processes except $W \rightarrow e\nu$;
- the refined E_T^{miss} calibration in the specific case of $W \rightarrow e\nu$ events amounts to correcting the globally calibrated cells of the electron shower back to the electromagnetic scale and the linearity of response is then also restored in this case.

Figure 274 shows that the E_T^{miss} resolution σ follows an approximate stochastic behaviour over a wide range of values of the total transverse energy deposited in the calorimeters. A simple fit to a function $\sigma = a \cdot \sqrt{\Sigma E_T}$ yields values between 0.53 and 0.57 for the parameter a , for ΣE_T values between 20 and 2000 GeV. The refined E_T^{miss} calibration yields somewhat better results for the E_T^{miss} resolution for e.g. $W \rightarrow e\nu$ decays. Departures from this simple behaviour are expected and observed for low values of ΣE_T where noise plays an important contribution and for very high values of ΣE_T where the constant term in the jet energy resolution dominates.

10.6.3 Measurement of E_T^{miss} direction

9975 Figure 275 shows the E_T^{miss} azimuthal angular resolution as a function of the true E_T^{miss} for three different physics processes. The measurement of the E_T^{miss} azimuth is clearly more accurate for $W \rightarrow e\nu$ events, which contain in general one high- p_T electron and moderate hadronic activity in addition, than for $t\bar{t}$ events, which contain much more hadronic activity. Figure 275 also shows that, for values of the true E_T^{miss} below 40 GeV, the accuracy on the measurement of the direction of a E_T^{miss} vector with small modulus degrades rapidly. In contrast, for high values of the true E_T^{miss} , azimuthal accuracies below 100 mrad can be achieved.

9985 As discussed in Section 10.6.5, large fluctuations in the jet energy measurements, due in particular to cracks in the fiducial acceptance of the calorimeters, may lead to fake E_T^{miss} with a vector of large modulus pointing in the same direction as the mis-measured jet in the azimuthal plane. A good accuracy on the measurement of the E_T^{miss} azimuth will therefore be needed to apply a cut, requiring that the measured E_T^{miss} vector be isolated from all high- p_T jets in the event, with a high efficiency for signal events with large true E_T^{miss} .




Figure 275. Accuracy of the measurement of the azimuth of the E_T^{miss} vector as a function of the true E_T^{miss} for three different physics processes: semi-leptonic $t\bar{t}$ events, $Z \rightarrow \tau\tau$ and $W \rightarrow e\nu$ events.

10.6.4 Use of E_T^{miss} for mass reconstruction

10000 The reconstructed E_T^{miss} vector can be used to improve the overall reconstruction of final-state topologies with only one neutrino in the final state (e.g. in $t\bar{t}$ events with one hadronic and one semi-leptonic top-quark decay). But, under certain simplifying assumptions and only for pairs which are not back-to-back [284, 285], one can even use the reconstructed E_T^{miss} vector in $Z \rightarrow \tau\tau$ and $A \rightarrow \tau\tau$ decays, despite the presence of several neutrinos in the final state, to reconstruct the invariant mass of the $\tau\tau$ pair. The results of such a procedure are shown in Fig. 276 for the reconstruction of $Z \rightarrow \tau\tau$ and $A \rightarrow \tau\tau$ decays with $m_A = 450$ GeV, where A is a supersymmetric Higgs boson. The reconstructed masses are correct to $\sim 2\%$ and the mass resolution is approximately 11%. Nevertheless, significant tails remain in the distributions because of the highly non-Gaussian effects induced by mis-measurements of E_T^{miss} and by the approximations used.

10.6.5 Fake E_T^{miss}

10010 Fake E_T^{miss} , defined simply as the difference between reconstructed and true E_T^{miss} , can arise at a significant level from a number of different sources: beam-gas scattering and other machine backgrounds, displaced interaction vertices, hot/dead/noisy cells (or regions) in the calorimeters, and mis-measurements in the detector itself, due to high- p_T muons escaping outside the fiducial acceptance of the detector (see also Section 10.3) and to large losses of deposited energy in cracks or inactive materials (see also Section 10.5.4.3). These latter two effects might effectively limit




Figure 276. Expected distributions for the reconstructed invariant mass of τ -lepton pairs, with one τ -lepton decaying to a lepton and the other one decaying to hadrons. The results are shown for $Z \rightarrow \tau\tau$ decays (left) and for $A \rightarrow \tau\tau$ decays with $m_A = 450$ GeV (right).




Figure 277. For QCD di-jet events containing at least one jet with $560 < E_T < 1120$ GeV, distribution of fake E_T^{miss} (circles), calculated as the difference between reconstructed and true E_T^{miss} , compared to the true E_T^{miss} (triangles) expected in this di-jet sample (left). Also shown is the distribution of fake E_T^{miss} due to muons (squares), calculated as the difference between the fake E_T^{miss} and the residual fake E_T^{miss} obtained using only the true muons in the event. The fake and true E_T^{miss} distributions are shown (right) after applying an isolation cut on the azimuth of the reconstructed E_T^{miss} vector. This cut requires that the distance in azimuth between the reconstructed E_T^{miss} vector and the direction of any high- p_T jet reconstructed in the event be larger than 17° .

the performance of the E_T^{miss} reconstruction in the longer term and have therefore been studied in detail.

10015

Figure 277 shows the distributions of fake and true E_T^{miss} for QCD di-jet events containing at least one jet with $560 < E_T < 1120$ GeV and for two cases: in the first case, all events are considered and Fig. 277 shows that fake E_T^{miss} dominates the spectrum up to true E_T^{miss} values of 200 GeV for this particular sample. In the second case, events are considered only if the reconstructed E_T^{miss} vector is isolated in azimuth from all reconstructed jets in the event. The isolation cut requires that the distance in azimuth between the reconstructed E_T^{miss} vector and the direction of any high- p_T

10020

jet reconstructed in the event be larger than 17° . In this case, true E_T^{miss} dominates the spectrum even for true E_T^{miss} values well below 100 GeV. These results confirm that the main source of fake E_T^{miss} in these events arises from mis-measurements of jets in certain regions of the calorimeter.

For those events with a fake E_T^{miss} larger than 100 GeV, Fig. 278 shows the distribution of $|\eta|$ of the mis-measured jet, defined as the jet which deviates the most from its matching truth jet in terms of its energy measurement. Clear excesses of mis-measured jets are observed around $\eta = 1.5$ (crack region) and around $\eta = 0$ (sensitive region for calibration of the overall calorimeter response). As already shown in Fig. 277, these excesses can be significantly reduced with simple topological cuts, but other tools can also reduce them further if required, such as the use of jets reconstructed from tracks to further improve the isolation in azimuth of the reconstructed E_T^{miss} .




Figure 278. For QCD di-jet events containing at least one jet with $560 < E_T < 1120$ GeV and with a fake E_T^{miss} larger than 100 GeV, distribution of $|\eta|$ of the mis-measured jet.

10.7 Hadronic τ -decays

Hadronic decays of τ -leptons will play an important role at the LHC, especially as probes for new phenomena spanning a wide range of theoretical models. Based on this motivation, two complementary approaches, one track-based and the other calorimeter-based (see Section 10.7.3), have been developed to efficiently reconstruct and identify these decays, whilst providing the required large rejection against the otherwise overwhelming backgrounds from hadronic jets. The equally difficult task of triggering on these decays as inclusively as possible is addressed in Section 10.9.

In general, hadronically decaying τ -leptons are reconstructed by matching narrow calorimeter clusters with a small number of tracks. Specific analyses may require exactly one or three tracks with total charge consistent with the charge of a τ -lepton, and, if more than one, the tracks may be required to be quite collimated and to be consistent with originating from a common secondary vertex. The visible reconstructed energy of the hadronically decaying τ -lepton is concentrated in a narrow cone around the leading (highest- p_T) track (typically a cone of half-angle $\Delta R = 0.2$ is sufficient to collect this energy). It can be estimated using only the calorimeter information or using a more refined scheme (often called energy flow), which combines the reconstructed track momenta with the energy of localised electromagnetic clusters within the chosen narrow cone.

Several key variables, which are characteristic of the properties of hadronic τ -decays, are used for the purpose of identification: the profile of the shower in the electromagnetic calorimeter, the isolation of the narrow calorimeter cluster used to identify the τ -candidate, the number and energy-weighted width of strips, the ratio between the transverse energy deposited in the calorimeter and the transverse momentum of the leading track, the number of associated tracks (passing some quality criteria), the momentum-weighted width and invariant mass of the track system and the signed impact parameter significance. Both traditional cut-based selections and multi-variate discrimina-

tion techniques (likelihood, neural networks, etc.) have been applied to this set of identification variables (see Section 10.7.3).

10065 Two specific performance aspects of particular interest for the reconstruction of hadronic τ -decays are first discussed in this section and are followed by the more general discussion of the overall performance in terms of reconstruction and identification efficiency versus rejection of the large backgrounds from QCD jets expected at the LHC.

10.7.1 Track reconstruction in hadronic τ -decays

10070 The efficiency and quality of the track reconstruction in the inner detector are discussed in some detail in Section 10.2. For hadronic τ -decays from a representative sample of $W \rightarrow \tau\nu$ and $Z \rightarrow \tau\tau$ decays studied with the track-based algorithm, particular attention has been given to minimise the amount of charge misidentification and of migration between the single- and three-prong categories in the reconstruction. In the low- p_T range, the performance is degraded due to hadronic interactions in the inner-detector material (see for example Fig. 202). For hadronic 10075 τ -decays with high energy, the performance for three-prong decays will be degraded due to the strong collimation of the tracks. Figures 279 and 280 show the efficiency for reconstructing tracks from single-prong and three-prong τ -decays for τ -leptons from W/Z -boson decays as a function of the track transverse momentum and pseudorapidity. The quality criteria used are the standard ones discussed in Section 10.2.3 and the results shown in Fig. 280 are in agreement with those shown for 10080 single particles in Fig. 203, except for three-prong τ -decays at high energy for which a degradation in efficiency is observed.

The charge of the identified hadronic τ -decay is determined as the sum of the reconstructed track charges. For the leading track, which is required e.g. by the track-based algorithm to have 10085 a transverse momentum larger than 9 GeV, charge misidentification is limited to $\sim 0.2\%$ with the standard quality cuts. The overall charge misidentification probability for the τ -lepton is, however, dominated by combinatorial effects: single-prong decays may migrate to the three-prong category due to photon conversions or the presence of additional tracks from the underlying event, or a three-prong decay may be reconstructed as a single-prong decay due to inefficiencies of the track 10090 reconstruction and selection. This overall misidentification is estimated to be below $\sim 3\%$ without requiring further quality cuts.

The rejection of leptonic τ -decays misidentified as single-prong hadronic τ -candidates is based on dedicated algorithms optimised to veto electrons and muons in the kinematic configurations of interest here. The rejection obtained against electron tracks from $W \rightarrow e\nu$ decays is approxi- 10095 mately 50 for a τ -efficiency of 95%. Using only information from the calorimeter combined with the inner detector, the rejection obtained against muons from $W \rightarrow \mu\nu$ decays is sufficient, reaching a value of approximately 30 for a τ -efficiency of 99%.

10.7.2 Electromagnetic clusters in single-prong decays

10100 Because of the very fine granularity of the electromagnetic calorimeter, electromagnetic clusters created by showers from photons from π^0 decays can be identified and measured with reasonable efficiency and accuracy within the narrow cone used to reconstruct hadronic τ -decays. The results reported in this section have been obtained using the three-dimensional topological clustering de-




Figure 279. Reconstruction efficiency for charged-pion tracks as a function of the pion transverse momentum for single- and three-prong hadronic τ -decays from $W \rightarrow \tau\nu$ and $Z \rightarrow \tau\tau$ signal samples.




Figure 280. Reconstruction efficiency for the charged-pion track as a function of $|\eta|$ for three different ranges of pion p_T , for single-prong hadronic τ -decays from $W \rightarrow \tau\nu$ and $Z \rightarrow \tau\tau$ signal samples.




Figure 281. Fractional energy response obtained for the visible transverse energy (see text), for single-prong hadronic τ -decays from a $W \rightarrow \tau\nu$ signal sample with one reconstructed electromagnetic cluster.




Figure 282. Distribution of reconstructed invariant mass of visible decay products (see text), for single-prong hadronic τ -decays from a $W \rightarrow \tau\nu$ signal sample with at least one reconstructed electromagnetic cluster.

scribed in more detail in Section 10.5 applied only to the first two layers of the electromagnetic calorimeter.

10105

As an example, in the case of single-prong decays, the reconstructed charged track in the inner detector and the reconstructed isolated clusters in the electromagnetic calorimeter may be used to obtain the energy and invariant mass of the visible products of the hadronic τ -decay. The resulting performance has been evaluated for $W \rightarrow \tau\nu$ decays and is shown in Table 62 for inclusive single-prong decays and also for exclusive decays containing a ρ or a_1 meson compared to decays containing only one single charged pion. Figure 281 shows the response and resolution obtained by this algorithm for reconstructing the visible transverse energy from the τ -decay, in the cases where one such isolated electromagnetic cluster is identified: the response is correct to $\sim 2.5\%$ and the fractional energy resolution is $\sim 5\%$, i.e. far better than that obtained for normal hadronic jets in the same energy range of 20–50 GeV. In the cases where several such clusters are identified,

10110

Decay mode	No cluster	One cluster	More than one cluster
All single-prong τ -decays	32%	35%	33%
$\tau \rightarrow \pi^\pm \nu$	65%	20%	15%
$\tau \rightarrow \rho^\pm (\rightarrow \pi^0 \pi^\pm) \nu$	15%	50%	35%
$\tau \rightarrow a_1^\pm (\rightarrow 2\pi^0 \pi^\pm) \nu$	9%	34%	57%

Table 62. Expected probabilities for observing a specific multiplicity of localised clusters in the electromagnetic calorimeter within the narrow cone ($\Delta R = 0.2$) used to identify the τ -candidate, for inclusive and exclusive single-prong hadronic τ -decays from $W \rightarrow \tau \nu$ and $Z \rightarrow \tau \tau$ signal samples,

10115 their energy-weighted barycentre is calculated and the fractional energy resolution is somewhat degraded to $\sim 7\%$. Finally, in the cases where at least one such cluster is identified, Fig. 282 shows the reconstructed invariant mass of the system for three single-prong final states. The use of certain specific final states in hadronic τ -decays will be of great interest in polarisation and spin analyses in searches for new particles decaying into τ -leptons.

10120 10.7.3 Identification of hadronic τ -decays and rejection of QCD jets

Two complementary algorithms for τ -identification and reconstruction have been studied, as outlined above:

- 10125 • a track-based algorithm [286], which relies on tracks reconstructed in the inner detector and adopts an energy-flow approach. This algorithm has been optimised for visible transverse energies in the 10–80 GeV range, which corresponds to hadronic τ -decays from $W \rightarrow \tau \nu$ and $Z \rightarrow \tau \tau$ processes;
- a calorimeter-based algorithm [287], which relies on clusters reconstructed in the calorimeter and has been optimised for visible transverse energies above 30 GeV, which corresponds to hadronic τ -decays from heavy Higgs-boson production and decay.

10130 Figures 283 and 284 show the expected performance of the two algorithms, expressed as curves describing jet rejection versus efficiency for single- and three-prong hadronic τ -decays separately and for different ranges of the visible transverse energy. The jet rejections are computed with respect to truth jets reconstructed using particle four-momenta within a cone of size $\Delta R = 0.4$. The behaviour of the respective rejection versus efficiency curves reflects the different optimisations performed for the two algorithms. Whereas the track-based algorithm has been tuned to preserve similar performance for single- and three-prong decays, the calorimeter-based algorithm has been tuned to provide the best possible rejection at medium-to-high energies and is therefore more performant for single-prong decays than the track-based algorithm. For an overall efficiency of 30% for single-prong decays, the rejection against jets is typically between 700 and 6000, as is
10135
10140 illustrated more quantitatively and as a function of the visible transverse energy in Table 63.

The track-based algorithm requires a good-quality track system, in which the leading track has transverse momentum above 9 GeV, as a seed for building a hadronic τ -candidate. This pro-




Figure 283. Expected rejection against hadronic jets as a function of the efficiency for hadronic τ -decays for the track-based algorithm using a neural-network selection. The results are shown separately for single- and three-prong decays and for two ranges of visible transverse energy.




Figure 284. Expected rejection against hadronic jets as a function of the efficiency for hadronic τ -decays for the calorimeter-based algorithm using a likelihood selection. The results are shown separately for single- and three-prong decays and for two ranges of visible transverse energy.




Figure 285. Track multiplicity distributions obtained for hadronic τ -decays with visible transverse energy above 20 GeV using the track-based τ -identification algorithm. The distributions are shown after reconstruction, after cut-based identification and finally after applying the neural network (NN) discrimination technique for an efficiency of 30% for the signal.




Figure 286. Track multiplicity distributions obtained for the background from QCD jets with visible transverse energy above 20 GeV using the track-based τ -identification algorithm. The distributions are shown after reconstruction, after cut-based identification and finally after applying the neural network (NN) discrimination technique for an efficiency of 30% for the signal.

10145 provides already after reconstruction considerable rejection against QCD jets with high track multiplicities. This is illustrated in Figs. 285 and 286, which show respectively the normalised track-multiplicity spectra for hadronic τ -candidates, with visible transverse energy above 20 GeV, from $Z \rightarrow \tau\tau$ decays and from QCD jets. The distributions are shown after the reconstruction step, after a cut-based identification algorithm and finally after applying a multi-variate discrimination technique using a neural network. The track multiplicity in the QCD jet sample is quite different from that in the signal sample, for any of the cuts applied. At the same time, Fig. 285 shows that

Algorithm		$E_T = 10\text{-}30$ GeV	$E_T = 30\text{-}60$ GeV	$E_T = 60\text{-}100$ GeV	$E_T > 100$ GeV
Track-based (neural network)	1-prong	740 ± 70	1030 ± 160		
	3-prong	590 ± 50	590 ± 70		
Calorimeter-based (likelihood)	1-prong		1130 ± 50	2240 ± 140	4370 ± 280
	3-prong		187 ± 3	310 ± 7	423 ± 8

Table 63. Rejection of track-based and calorimeter-based τ -identification algorithms over the range of visible transverse energy in which they have been optimised. The values are given separately for single- and three-prong decays and for an efficiency of 30%. The quoted errors are statistical.

10150 the fractions of single-prong and three-prong decays in the signal sample approach those expected
from an ideal signal sample: for single-prong (respectively three-prong) candidates, the fractions
of correctly assigned decays improve from 87% (respectively 74%) after reconstruction to 91%
(respectively 86%) after cut-based identification and to 92% (respectively 93%) after applying the
neural-network discrimination technique.

10155 Figures 285 and 286 also show that the candidates with track multiplicity above three may
be used to normalise the QCD background. This would allow a reasonably precise calibration of
the performance of the τ -identification algorithms using real data, provided the rejection against
QCD jets is proven to be sufficient to extract a clean signal in the single-prong and three-prong
categories. The sensitivity of such a method can be enhanced by also studying the track multi-
10160 plicity outside the narrow cone used for τ -identification and combining this information with that
presented in Fig. 286.

10.8 Flavour tagging

The ability to tag hadronic jets arising from heavy flavours is an important asset for many physics
analyses, such as precision measurements in the top-quark sector and searches for Higgs bosons
10165 or other new physics signatures. This section describes the b -tagging performance which can be
achieved using different methods [288]. In the results presented in this section, the impact of
possible residual misalignments on the b -tagging performance has not been taken into account.

10.8.1 Ingredients of b -tagging algorithms

Except when explicitly stated otherwise, the results presented in this section are based on simula-
10170 tions without pile-up and with a perfect alignment of the inner detector. Jets are reconstructed in
the calorimeters using standard algorithms (see Section 10.5.1) and the jets with $p_T > 15$ GeV and
 $|\eta| < 2.5$ are considered for b -tagging. Only reconstructed tracks within a distance $\Delta R < 0.4$ from
the jet axis are used for b -tagging.

To assess quantitatively the b -tagging performance, the Monte-Carlo truth is used to determine
10175 the type of parton from which a jet originates. This labelling procedure is somewhat ambiguous.
For the results presented here, a quark-based labelling has been used: a jet is labelled as a b -jet if
a b -quark with $p_T > 5$ GeV is found in a cone of size $\Delta R = 0.3$ around the jet direction. A jet is
labelled as a c -jet (or τ -jet) if a c -quark (or τ -lepton) with $p_T > 5$ GeV is found in the cone instead
of a b -quark. When no heavy quark nor τ -lepton satisfies these requirements, the jet is labelled as




Figure 287. Signed transverse impact parameter, d_0 , distribution for b -jets, c -jets and light jets.




Figure 288. Signed transverse impact parameter significance, d_0/σ_{d_0} , distribution for b -jets, c -jets and light jets.

10180 a light jet. No attempt is made to distinguish between u -, d -, s -quarks and gluons. It is important to note that this labelling procedure defines as b -jets most gluon jets splitting to a $b\bar{b}$ pair in the parton-shower process.

10.8.1.1 Track selection and track impact parameters

The track selection for b -tagging is designed to select well-measured tracks and to reject fake tracks in jets and secondary tracks from K_s^0 , Λ and hyperon decays, as well as electrons from photon conversions. Only tracks with transverse momentum above 1 GeV are considered. At least seven precision hits (pixels and strips) are required, of which at least two must be in the pixel detector and one in the pixel vertexing layer. The transverse (d_0) and longitudinal (z_0) impact parameters at the perigee must fulfil respectively, $|d_0| < 1$ mm and $|z_0 - z_v| \sin \theta < 1.5$ mm, where z_v is the reconstructed primary vertex position in z and θ is the measured polar angle of the track. The efficiency of these cuts and the resulting fake-track rate in jets are discussed in Section 10.2.3 (see in particular Fig. 204).

For the b -tagging algorithms, the impact parameters of tracks are computed with respect to the primary vertex (see Section 10.2.4). The transverse impact parameter is signed using the jet direction as measured by the calorimeters: tracks crossing the jet axis behind the primary vertex have a negative impact parameter. The distribution of the signed transverse impact parameter, d_0 , is shown in Fig. 287 for tracks reconstructed in b -jets, c -jets and light jets. Figure 288 shows the corresponding significance distribution, d_0/σ_{d_0} , which gives more weight to precisely measured tracks.

10.8.1.2 Secondary vertices

To further increase the discrimination between b -jets and light jets, the inclusive vertex formed by the decay products of the B -hadron, including the products of the subsequent charm hadron decay, can be reconstructed. The search starts by combining all track pairs which form a good vertex, using only tracks with a high impact-parameter significance in order to remove the tracks which are compatible with the primary vertex. The invariant mass of the




Figure 289. Properties of secondary vertices reconstructed in b -jets and light jets: invariant mass of all tracks originating from the vertex (left), the ratio of the sum of the energies of the tracks originating from the vertex to the sum of the energies of all tracks in the jet (middle) and number of two-track vertices (right).

particles originating from the secondary vertex candidate and the location of this vertex candidate
10205 are used to reject vertices which are likely to come from K_S^0/Λ decays and photon conversions or
from secondary interactions in material such as the beam-pipe or the vertexing layer. All tracks
from the remaining two-track vertices are combined into a single vertex and three of its properties
are exploited: the invariant mass of all the tracks originating from the vertex, the ratio of the sum
of the energies of the tracks originating from the vertex to the sum of the energies of all tracks in
10210 the jet, and the number of two-track vertices. These properties are illustrated in Fig. 289 for b -jets
and light jets. The secondary-vertex reconstruction efficiency depends quite strongly on the event
topology and the typical efficiencies achieved are higher than 60% for the $t\bar{t}$ and WH events studied
here.

10.8.2 Likelihood-ratio tagging algorithms

10215 For both the impact-parameter tagging and the secondary-vertex tagging, a likelihood-ratio method
is used: the discriminating variables are compared to pre-defined smoothed and normalised distri-
butions for both the b - (signal) and light- (background) jet hypotheses. Multi-dimensional probabili-
ty density functions are used as well for some b -tagging algorithms. The ratio of the probabilities
defines the track or vertex weight, which can be combined in a jet weight as the sum of the loga-
10220 rithms of the individual weights. The distribution of such a weight is shown in Figs. 290 and 291
for b -, c - and light jets for two different b -tagging algorithms: the first one combines only the
transverse impact parameter significance of tracks, while the second one combines in two dimen-
sions the transverse and longitudinal impact parameter significances of tracks as well as the three
variables from the secondary vertex search discussed above. The former algorithm is simpler and
10225 more robust than the latter which will require more time to commission. Currently, no use is made
of probability density functions for c -jets, and these are not considered when creating the reference
distributions for the signal and background hypotheses.

10.8.3 Jet activity and jet purification

A difficulty arises as soon as the jet multiplicity is high and various jet flavours are present in the
10230 same event: a light jet close in ΔR to a b -jet will sometimes be labelled as a light jet, even though




Figure 290. Jet b -tagging weight distribution for b -jets, c -jets and purified light jets (see Section 10.8.3). The b -tagging algorithm is based on the transverse impact parameter significance of tracks.




Figure 291. Jet b -tagging weight distribution for b -jets, c -jets and purified light jets (see Section 10.8.3). The b -tagging algorithm uses the transverse and longitudinal impact parameter significances of tracks as well as the properties of the secondary vertex found in the jet.

tracks from B -hadron decay with high lifetime content may be associated with it. This leads to an artificial degradation of the estimated performance, which is not related to the b -tagging algorithm itself but to the labelling procedure which strongly depends on the activity in the event. In order to obtain a more reliable estimation of the b -tagging performance, a purification procedure has been devised: light jets for which a b -quark, a c -quark or a τ -lepton are found within a cone of size $\Delta R = 0.8$ around the jet direction are not used to assess the b -tagging performance.

The performance estimated after purification represents the intrinsic power of the b -tagging algorithms and should be similar for different kinds of physics events; in contrast, results obtained using all the light jets, regardless of their environment, are more dependent on the underlying activity in the event. These latter results are, however, more representative of the actual b -tagging performance to be expected for a given physics analysis. This is illustrated in Figs. 292 and 293 for two types of physics processes. The WH events correspond to events in which the W decays leptonically and the Higgs boson decays to a $b\bar{b}$ pair (signal case) or is forced to decay to a $u\bar{u}$ or $c\bar{c}$ pair (background case). Such events therefore usually have only two high- p_T and well-separated jets and the light-jet rejection obtained is similar with and without jet purification, as shown in Fig. 292. For semi-leptonic $t\bar{t}$ events on the contrary the jet activity is quite high and therefore the two performance curves with and without purification shown in Fig. 293 differ in the region of b -jet efficiencies below 80%, where the lifetime content dominates over resolution effects. It is also important to note that the purification procedure discards jets coming from gluon-splitting to heavy quarks.




Figure 292. Rejection of light jets and c -jets with and without purification versus b -jet efficiency for WH events with $m_H = 120$ GeV, using the b -tagging algorithm based on the 3D impact parameter and secondary vertices.




Figure 293. Rejection of light jets and c -jets with and without purification versus b -jet efficiency for $t\bar{t}$ events, using the b -tagging algorithm based on the 3D impact parameter and secondary vertices.

10.8.4 Expected b -tagging performance

As shown in Figs. 292 and 293, a light-jet rejection higher than 100 can be achieved for a b -jet efficiency of 60%. The performance depends strongly on the jet momentum and pseudorapidity. This is illustrated in Figs. 294 and 295 for the two b -tagging algorithms described above. At low p_T and/or high $|\eta|$, the performance is degraded mostly because of the increase of multiple scattering and secondary interactions. At high p_T , some dilution arises because the fraction of fragmentation tracks in the fixed-size cone increases, and more B -hadrons decay outside the vertexing layer: some gain should therefore be achieved by changing the track selection. At very high p_T , the performance degradation arises from pattern-recognition deficiencies in the core of very dense jets.

10.8.5 Soft-lepton tagging

Soft-lepton tagging relies on the semi-leptonic decays of bottom and charm hadrons. It is therefore intrinsically limited by the branching ratios to leptons: at most 21% of b -jets will contain a soft lepton of a given flavour, including cascade decays of bottom to charm hadrons. However, when a soft lepton is present, b -tagging algorithms based on soft leptons can exhibit high purity. More importantly, they have only small correlations with the track-based b -tagging algorithms, which is very important for checking and cross-calibrating performance with data.

Soft muons are reconstructed using two complementary reconstruction algorithms (see Section 10.3): combined muons, which correspond to a track fully reconstructed in the muon spectrometer and matched with a track in the inner detector, and muons with a low momentum, typically below ~ 5 GeV, which cannot reach the muon middle and outer stations and are identified by matching an inner-detector track with a segment in the muon spectrometer inner stations only. Muons reconstructed in this way and satisfying some basic selection criteria, $p_T > 4$ GeV and




Figure 294. Rejection of purified light jets as a function of the jet transverse momentum for two different b -tagging algorithms operating at a fixed b -tagging efficiency of 60% in each bin.




Figure 295. Rejection of purified light jets as a function of the jet pseudorapidity for two different b -tagging algorithms operating at a fixed b -tagging efficiency of 60% in each bin.




Figure 296. Rejection of light jets versus b -tagging efficiency in $t\bar{t}$ events (branching ratios to lepton and lepton identification efficiency included) for the soft-muon b -tagging algorithm. The results are shown without and with the pile-up and cavern background expected at a luminosity of $10^{33} \text{ cm}^{-2} \text{ s}^{-1}$.




Figure 297. Rejection of light jets versus b -tagging efficiency in WH events (branching ratios to lepton and lepton identification efficiency included) for the soft-electron b -tagging algorithm.

10275

$|d_0| < 4 \text{ mm}$, are associated to the closest jet provided their distance to the jet axis satisfies $\Delta R < 0.5$. Finally, the kinematic properties of the jet-muon system, such as the relative transverse momentum of the muon with respect to the jet axis, are used in order to reject the background caused by punch-through particles and decays in flight in light jets. As shown in Fig. 296 for $t\bar{t}$ events, the soft-muon b -tagging algorithm yields an efficiency of 10% (including branching ratios and identification efficiency) and a light-jet rejection of 200 for jets with $p_T > 15 \text{ GeV}$ and $|\eta| < 2.5$. The rejection against light jets decreases by approximately 30% when the expected

10280 contributions from pile-up and especially cavern background at $10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ are included.

Reconstructing soft electrons in jets in the electromagnetic calorimeter is more difficult because of the overlap of hadronic showers with the electron shower itself. This is achieved using the soft-electron algorithm [277] which matches an inner-detector track to an electromagnetic cluster, as described in Section 10.4.1. The performance of this algorithm is, however, highly dependent
10285 on the track density in the jets as well as on the amount of material in front of the electromagnetic calorimeter (photon conversions). As shown in Fig. 297, a light-jet rejection of 90 can be achieved for an efficiency of 7% in WH events. Currently, for a 7% (respectively 10%) b -tagging efficiency, about 75% (respectively 40%) of the surviving light jets are tagged by electrons originating from photon conversions: the performance would therefore substantially improve if these conversions
10290 could be rejected further.

10.9 Trigger performance

10.9.1 Overview

This section gives an overview of the performance achieved on simulated raw data using the online physics selection strategy of ATLAS. As already mentioned in Section 8.1, components of the
10295 reconstruction and analysis software, implemented mostly in the offline environment in previous experiments, have had to be embedded within the trigger system to achieve the required rejection power while retaining excellent sensitivity to the various physics signatures of interest. A great deal of flexibility is provided by the three-level trigger system to adapt to changes in the luminosity (from fill-to-fill and even during a single fill), to variations in the background conditions, and to new
10300 requirements which will undoubtedly arise as the understanding of the physics, trigger performance and detector develops.

The approach taken to guarantee good acceptance for as broad a spectrum of physics as possible is to use mainly inclusive criteria for the online selection, i.e. signatures mostly based on single-
10305 and di-object high- p_T triggers. The choice of the thresholds is made to have a good overlap with the reach of the Tevatron and other colliders, and to ensure good sensitivity to new particles over a wide range of masses and decay channels. This high- p_T inclusive selection is complemented where necessary with more focussed signatures, such as the presence of several different physics objects or the use of topological criteria.

10.9.2 Selection strategy

10310 The architecture of the trigger and data acquisition system is described in Section 8.3 (see in particular Fig. 160) and is based on a three-level trigger system, with a first level (L1) using hardware based on ASIC's and FPGA's, and the other two (L2 and EF or event filter, collectively also called high-level triggers or HLT) using software algorithms running on farms of commercial computers. At L2, the event selection is based on specialised algorithms, optimised for speed, whereas the
10315 EF uses more complex algorithms, basically identical to those used in the offline reconstruction software.

The L2 and EF algorithms are usually "seeded", meaning that reconstruction is normally guided by the previous trigger level to access and process data only in a "Region-of-Interest" (RoI)

containing particle candidates. This significantly reduces the processing time (and also data move-
10320 ment from the holding buffers to the L2 processors), without degrading the selection performance.
For the EF, and even for L2 where necessary, data can be accessed and processed from the full
detector, within the constraints of available data-movement and processing resources. This applies
for example to scans of the complete inner detector for low- p_T tracks for the B -physics selection or
to the processing of all calorimeter cells for an improved calculation of missing transverse energy.
10325 At L2, this can only be done in special cases and for a small fraction of the events due to bandwidth
limitations, whereas in the EF the full event data are available in memory. It is also possible to use
so-called secondary RoI's which did not contribute to the L1 selection, but provide the coordinates
of lower- p_T objects which can be included in the L2 selection.

In the HLT, "feature-extraction" algorithms are used to identify objects (such as electrons or
10330 jets) and determine their properties or to determine global characteristics of the event. The sequence
of execution of the algorithms (e.g. ordered according to complexity) is chosen to maximise the
physics potential and retain adequate flexibility within the available data-movement and processing
resources of the HLT. After each step in the sequence, hypothesis algorithms determine whether
a given signature is satisfied or not. The processing of any given RoI is stopped as soon as it is
10335 clear that it cannot contribute to the selection of the event. The event itself is rejected if none of the
signatures in the trigger menu is satisfied.

The initial implementation and capabilities of the DAQ/HLT system are described in Sec-
tion 8.4, where it is stated that the system should handle a L1 trigger rate of ~ 40 kHz, i.e. approx-
imately 50% of the design specification. Clearly, only the availability of real data will allow the
10340 whole strategy to be finalised. However, it is important to be able to face this initial phase with the
most complete set of tools possible and with a versatile selection architecture, in order to cope with
the surprises which are likely to appear at the time of LHC start-up.

10.9.3 Trigger menus

Trigger menus are tables which specify thresholds and selection criteria at each of the three trig-
10345 ger levels to address the physics-analysis requirements of ATLAS. The process of preparing the
menus takes into account an assessment of the rejection capabilities at each selection stage and for
each signature, and the rate capabilities of each level of the trigger and of the offline computing
system. This procedure is iterative and makes use of earlier studies of the L1 trigger and HLT, as
documented in [209, 247].

10350 Trigger items, defined as entries in the trigger menu corresponding to selected physics objects,
are identified using a notation where a symbol representing a particle type is preceded by a multi-
plicity value and followed by a E_T -threshold value, e.g. $2e5$ corresponds to a requirement of two or
more electrons, each with E_T above 5 GeV. The threshold value quoted for L1 is the raw E_T cut ap-
plied in the hardware, and high efficiency is only achieved for particles or jets of somewhat higher
10355 E_T ; this differs from the definition used in previous documents [247]. For inclusive selections, the
multiplicity requirement of one is implicit. An "i" following the threshold indicates that an isolation
requirement is made in addition. For example, $\tau 20i$ requires at least one hadronic τ candidate with
transverse energy above 20 GeV and with a specific calorimeter isolation requirement in addition.
The term "xE" is a short form for E_T^{miss} .

Signature	L1 rate (Hz)	HLT rate (Hz)	Comments
Minimum bias	Up to 10000	10	Pre-scaled trigger item
e10	5000	21	$b, c \rightarrow e, W, Z, \text{Drell-Yan}, t\bar{t}$
2e5	6500	6	Drell-Yan, $J/\psi, \Upsilon, Z$
γ 20	370	6	Direct photons, γ -jet balance
2 γ 15	100	< 1	Photon pairs
μ 10	360	19	$W, Z, t\bar{t}$
2 μ 4	70	3	B -physics, Drell-Yan, $J/\psi, \Upsilon, Z$
μ 4 + $J/\psi(\mu\mu)$	1800	< 1	B -physics
j120	9	9	QCD and other high- p_T jet final states
4j23	8	5	Multi-jet final states
τ 20i + xE30	5000 (see text)	10	$W, t\bar{t}$
τ 20i + e10	130	1	$Z \rightarrow \tau\tau$
τ 20i + μ 6	20	3	$Z \rightarrow \tau\tau$

Table 64. Subset of items from an illustrative trigger menu at $10^{31} \text{ cm}^{-2} \text{ s}^{-1}$.

10360 The steering and configuration of the trigger (see Section 8.3.6) support the description of both straightforward RoI-based triggers like single electrons, muons, τ -leptons and jets along with more complex triggers like E_T^{miss} and triggers for B -physics. For each trigger level, items in the menu can be pre-scaled to reduce their rates, or "pass-through" flags can be raised, where events are accepted irrespective of the HLT selection decision for the purpose of systematic studies.

10365 The initial start-up luminosity at the LHC is expected to be around $10^{31} \text{ cm}^{-2} \text{ s}^{-1}$. This provides convenient conditions for commissioning the trigger and the detector sub-systems, validating the trigger and offline software algorithms, and ensuring that basic Standard Model signatures can be observed. The trigger menu for this start-up scenario reflects these requirements and allows for low p_T -thresholds on final-state leptons and photons, without any pre-scaling at L1, and for higher
10370 p_T -thresholds, for which most of the HLT algorithms are executed in "pass-through" mode.

Table 64 presents an example of a sample of the triggers which will be used at start-up. The rates shown have been estimated using non-diffractive minimum-bias events with a total assumed cross-section of 70 mb. Triggering on single and di-leptons should be possible with quite low p_T -thresholds and without applying isolation or other complex criteria, which must be validated with
10375 real data at turn-on. With the exception of the minimum-bias selection, the items indicated are those which should be operable without pre-scaling at $10^{31} \text{ cm}^{-2} \text{ s}^{-1}$. The full menu contains a number of additional components, including many pre-scaled items with lower thresholds.

The rates for combined triggers which require two or more final-state leptons or photons are expected to be low in most instances, allowing them to be run without pre-scaling with very low
10380 thresholds. Significant bandwidth will be devoted to collecting large samples of minimum-bias data for use in physics analysis and for detector and trigger performance studies. Multi-jet triggers will be run at a comparatively high rate to test b-jet tagging in the HLT which is discussed in Section 10.9.6. A small amount of bandwidth is allocated for inclusive E_T^{miss} and scalar sum- E_T trig-

gers, as well as using the E_T^{miss} signature in combination with other criteria. Note that for the item
10385 $\tau 20i + xE30$ in Table 64, the E_T^{miss} selection is made only at the EF level, in case the corresponding
L1 selection takes time to commission. The rate of the $\tau 20i$ item at L1 is approximately 5 kHz.

The quoted trigger rates are subject to large uncertainties on the cross-sections for QCD pro-
cesses in proton-proton collisions at LHC energies, and on the modelling of the performance of
the detector. The rates indicated assume that the selection cuts will already have been reasonably
10390 well tuned to achieve high background rejection with good signal efficiency. There is still scope
to use tighter cuts and more delicate variables such as isolation after extensive optimisation and
thorough validation. Should the rates turn out to be higher than estimated, the inclusive thresholds
could be raised substantially without compromising much of the main initial physics programme,
but nevertheless to the detriment of an efficient collection of large data samples required for the
10395 initial understanding of the detector performance.

10.9.4 Examples of trigger performance

The expected trigger performance at an initial luminosity of $10^{31} \text{ cm}^{-2} \text{ s}^{-1}$ is illustrated in the fol-
lowing with representative examples from the menu discussed above. As documented in Ref. [247]
and discussed briefly in Section 10.9.6, the trigger also meets the physics requirements up to a
10400 luminosity of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$.

The performance results presented here were obtained using an exact simulation of the algo-
rithms which are implemented in the L1 hardware and using the same HLT algorithms as those
which are run online. The full HLT chain was used to obtain the performance results. As de-
scribed in Section 10.1, these studies have also been made for data simulated with a misaligned
10405 and mis-calibrated detector to verify, and improve if necessary, the robustness of the selection.

10.9.4.1 Electrons and photons The performance of the electron and photon triggers has been
evaluated for a luminosity of $10^{31} \text{ cm}^{-2} \text{ s}^{-1}$, using simulations of single particles and selected
physics channels. The trigger efficiencies are quoted with respect to events containing electrons
and photons identified with loose offline particle identification cuts (see Section 10.4). Inefficiency
10410 in the trigger selection arises mainly from tighter selection requirements needed to reduce the
background rate to an acceptable level. There are also small losses due to the coarser calorimeter
granularity used at L1 and the simpler (and faster) selection algorithms applied at L2 compared to
the offline reconstruction.

Figure 298 shows the L1, L2 and EF efficiencies as a function of E_T for the signature e10,
10415 the menu item selecting electrons with $E_T > 10 \text{ GeV}$, as estimated using simulations of single
electrons. The efficiency reaches a plateau value for E_T above $\sim 15 \text{ GeV}$ and is quite uniform as a
function of $|\eta|$, except for a 10–20% dip in the transition region between the barrel and the end-cap
calorimeters. Figure 299 shows the normalised relative rates expected from QCD jets satisfying
the e10 signature as a function of $|\eta|$ for the successive trigger levels. These relative rates are
10420 normalised for each trigger level to the total number of events selected and then the rate in each
bin is rescaled to that expected for a bin of fixed size $\Delta\eta = 0.5$. The rates are quite sensitive to the
result of the trigger efficiency optimisation and their uniformity reflects the lower efficiency in the
regions where the electromagnetic calorimeter performance is not optimal, as in the barrel/end-cap




Figure 298. Trigger efficiencies at L1, L2 and EF as a function of the true electron E_T for the e10 menu item. The efficiencies are obtained for single electrons and are normalised with respect to the medium set of offline electron cuts discussed in Section 10.4.




Figure 299. Relative rates versus $|\eta|$ for jets passing the L1, L2 and EF trigger selections for the e10 menu item. The relative rates are shown for each of the seven η -ranges used to optimise the offline selection of isolated electrons and are normalised as described in the text.




Figure 300. Trigger efficiencies at L1, L2 and EF as a function of the true photon E_T for the $\gamma 20i$ menu item. The efficiencies are obtained for single photons and normalised with respect to loose offline photon identification cuts.




Figure 301. Normalised relative rates versus $|\eta|$ for jets passing the L1, L2 and EF trigger selections for the $\gamma 20i$ menu item. The relative rates are shown for each of the six η -ranges used to optimise the offline selection of isolated photons and are normalised as described in the text. The bin corresponding to the barrel/end-cap transition region is not shown because the offline selection excludes it.

10425 transition region with $1.37 < |\eta| < 1.52$. Similar results are obtained for photons and shown for the signature $\gamma 20i$, the menu item selecting isolated photons with $E_T > 20$ GeV, in Figs. 300 and 301.

10.9.4.2 τ -leptons and E_T^{miss} The performance of the trigger for selecting high- p_T τ -leptons is illustrated with the $\tau 20i$ signature selecting hadronic τ -decays with true visible E_T of the hadronic τ -decay (defined as the summed transverse energy of all the decay products which are not neutrinos) larger than 20 GeV. Figure 302 shows the trigger efficiency after each trigger level, nor-




Figure 302. Trigger efficiencies at L1, L2 and EF as a function of the true visible E_T of the hadronic τ -decays for the $\tau 20i$ menu item.




Figure 303. Efficiency after EF for τ -trigger items with different thresholds as a function of the true visible E_T of the hadronic τ -decay.

10430 malised to an offline selection with loose requirements (see Section 10.7), for hadronic τ -decays
 10435 from $W \rightarrow \tau\nu$ and $Z \rightarrow \tau\tau$ decays. The efficiency exhibits a drop of approximately 15% after L2,
 mostly because of the τ -identification cuts applied. The efficiency turn-on rises more slowly than
 for the electron and photon triggers, especially at L1, reflecting the poorer resolution obtained for
 hadronic showers. Figure 303 shows the efficiency turn-on curves for various τ -trigger thresholds
 as a function of the true visible E_T of the hadronic τ -decay. The overall efficiency with respect to
 the offline selection is typically 85% on the plateau.

A somewhat special case is that of E_T^{miss} triggers which can be used either inclusively or in
 combination with other objects, in particular with jets or hadronic τ -triggers. Because E_T^{miss} is a
 global property of the event, the RoI-driven L2 trigger is not capable of substantially improving
 the L1 trigger. However, the E_T^{miss} algorithm in the EF improves substantially on L1 by accessing
 10440 the precision readout of the entire calorimeter and performing a simplified version of the offline
 algorithm.

A challenging goal of the τ -selection during the low-luminosity period is to collect a large
 sample of $W \rightarrow \tau\nu$ decays. This can be achieved using a τ -trigger in combination with a require-
 10445 ment (potentially only at the level of the EF) of substantial E_T^{miss} (see Table 64). Such events are
 obviously interesting for physics analyses, but are also needed to monitor the hadronic energy scale
 using single charged pions, and for other performance studies. An additional goal is to provide
 triggers with low p_T -thresholds and loose trigger requirements, in addition to the single high- p_T
 electron and muon triggers, for collecting efficiently Z -bosons decaying into two τ -leptons, where
 10450 one τ -lepton decays to an electron or muon and the other to hadrons. The background rates for
 these $e/\mu+\tau$ triggers are estimated to be in the range of one Hz or less at the initial luminosity
 of $10^{31} \text{ cm}^{-2} \text{ s}^{-1}$, with rather loose HLT cuts applied to the trigger objects.

10.9.4.3 Jets The inclusive jet trigger j120 presented in Table 64 is complemented by a series of
 pre-scaled items chosen to give an approximately uniform rate across the jet E_T -spectrum. Col-
 10455 lecting sufficient statistics over the entire jet E_T -spectrum is important for differential cross-section
 measurements and also for the measurement of detector, trigger and physics algorithm efficiencies.

Trigger item	j10	j18	j23	j35	j42	j70	j120	3j10	3j18	4j10	4j18	4j23
Pre-scale factor at L1	42000	6000	2000	500	100	15	1	150	1	30	1	1
L1 rate (Hz)	4	1	1	1	4	4	9	40	140	40	20	8
EF rate (Hz)	4	1	1	1	4	4	9	0.05	1	0.04	0.1	5

Table 65. Summary of L1 single-jet and multi-jet menu items, of L1 pre-scale factors and expected L1 and EF rates at a luminosity of $10^{31} \text{ cm}^{-2} \text{ s}^{-1}$.




Figure 304. Expected differential spectrum for single jets as a function of the reconstructed E_T of the leading jet. The solid line shows the distribution after applying the L1 trigger thresholds and pre-scale factors presented in Table 65, while the dashed line represents the distribution expected without any trigger requirements.




Figure 305. Efficiency as a function of the true jet E_T (as defined for a cone of size $\Delta R = 0.4$) for each of the single-jet L1 menu items shown in Table 65.

The set of threshold and pre-scale combinations is expected to be stable with rising luminosity for the first few years of data-taking. The strategy adopted to optimise the jet trigger menu for different luminosities is then primarily to modify the pre-scale factors associated with each jet-trigger threshold, rather than to change the set of thresholds on an ad-hoc basis.

Table 65 summarises a set of L1 jet-trigger items, L1 pre-scale factors and L1 and EF rates for a luminosity of $10^{31} \text{ cm}^{-2} \text{ s}^{-1}$. Since the jet rates cannot be reduced much by the HLT, the EF rates quoted in Table 65 are obtained through additional pre-scale factors applied wherever necessary. Figure 304 shows the corresponding reconstructed differential E_T spectrum of the leading jet after the L1 trigger accept. The differential distribution thus obtained is almost uniform over the range of L1 single-jet triggers run with different pre-scale factors, yielding about 10^8 leading jets with E_T in the range between 10 and 100 GeV for an integrated luminosity of 100 pb^{-1} . Figure 305 shows that over this range of jet E_T , the efficiency at threshold of the various pre-scaled jet trigger menu items turns on much more slowly than the corresponding curves for leptons because of the poorer resolution of the jet E_T reconstructed at L1.

10.9.4.4 Muons The geometrical coverage of the muon trigger detector system (see Section 6.5 for a detailed description) limits the overall acceptance for triggering on muons at L1, as illustrated




Figure 306. Acceptance map in η - ϕ space for the L1 muon trigger, which covers the η -range $|\eta| < 2.4$. The black points represent regions not instrumented with L1 trigger detectors because of the presence of various supports and services.




Figure 307. Estimated EF output rates for muons as a function of p_T -threshold at a luminosity of $10^{31} \text{ cm}^{-2} \text{ s}^{-1}$, integrated over the full η -range covered by the L1 trigger, $|\eta| < 2.4$.

in Fig. 306. The barrel trigger system covers approximately 80% of the η - ϕ plane (over $|\eta| < 1.0$), while the end-cap trigger extends over approximately 96% of the relevant η - ϕ space. The limitations of the barrel system can be seen in Fig. 306, and are dominated by the crack at $\eta < 0.1$ (largely to accommodate inner-detector and calorimeter services), by the regions occupied by the feet of the experiment and by the space taken by the barrel toroid ribs. The end-cap trigger coverage is limited only by the detector supports and by the holes needed for the optical alignment system. Within the fiducial acceptance of the trigger detectors, the L1 trigger efficiency for muons with p_T larger than the selection thresholds exceeds 99%. The L2 trigger then provides a first reduction of the L1 rates by confirming the muon candidates with a more precise measurement of their momentum and by matching them to inner-detector tracks.

The rates of muons at the output of the EF have been computed at a luminosity of $10^{31} \text{ cm}^{-2} \text{ s}^{-1}$, by summing the contributions from the barrel and end-cap regions of the muon spectrometer. As shown in Fig. 307, several physics processes contribute significantly to the rate. The rates given as a function of the p_T -threshold are for an inclusive muon selection, without applying an isolation requirement. The largest contributions to the total rate in the p_T -range from 4 to 6 GeV are from charm, beauty and in-flight decays of charged pions and kaons. Isolation, as well as refined matching requirements between the tracks in the inner detector and muon spectrometer, can be used to further reduce the rates.

10.9.4.5 B-physics The trigger for B -physics is initiated by a single- or di-muon selection at L1. At $10^{31} \text{ cm}^{-2} \text{ s}^{-1}$, a threshold $p_T > 4 \text{ GeV}$ will be used, rising to about 6 GeV at $10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ to match the rate capabilities of the HLT.

At the initial expected luminosity of $10^{31} \text{ cm}^{-2} \text{ s}^{-1}$, the dimuon final states are selected by the $2\mu 4$ trigger which is expected to have a rate of a few Hz. For single-muon triggers, searches can be made in the HLT for additional features using information from the inner detector and calorimeters, as well as from the muon spectrometer. Mass cuts and secondary-vertex reconstruction are used to select the B -decay channels of interest. Channels, such as $B_d \rightarrow J/\psi(\mu\mu)K_s^0$ and $B_{s,d} \rightarrow \mu\mu$,

are triggered by requiring two muons fulfilling J/ψ or $B_{s,d}$ invariant-mass cuts. Identification of the second muon can either originate from a separate L1 RoI, or from the HLT in an enlarged RoI around the first muon. For other channels containing muons, such as $B_d \rightarrow K^{*0}\mu\mu$ or $B_s \rightarrow \phi\mu\mu$, inner-detector tracks are combined to first reconstruct the K^{*0} or ϕ and then the muon tracks are added to reconstruct the $B_{s,d}$.

For hadronic final states like $B_s \rightarrow D_s^- \pi^+$ and $B_s \rightarrow D_s^- a_1^+$, inner-detector tracks are combined to reconstruct first the ϕ -meson from the D_s decay, then the D_s and finally the B_s . Two different strategies are used for finding the tracks, depending on luminosity. Full reconstruction over the whole inner detector can be performed at $10^{31} \text{ cm}^{-2} \text{ s}^{-1}$, since the L1 muon rate is comparatively modest, while at higher luminosities reconstruction will be limited to L1 jet RoI's with $E_T > 5 \text{ GeV}$. This latter approach has lower efficiency for selecting the signal, as shown in Fig. 308, but requires fewer HLT resources for a fixed L1 rate. If one combines triggers for hadronic final states and pre-scaled single muon triggers needed for trigger efficiency measurements, the overall rate for B-physics triggers is approximately 10 Hz at $10^{31} \text{ cm}^{-2} \text{ s}^{-1}$.

10.9.5 Trigger commissioning

A detailed strategy for commissioning the trigger during initial running with beam is being developed. It is assumed that the luminosity will be significantly less than $10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ during this period. A first step will be to establish a time reference for bunches of protons colliding at the interaction point in ATLAS. Signals from passive beam pick-ups will be used to form a filled-bunch trigger with known latency. This will be combined with the minimum-bias trigger, based on scintillation counters which are mounted in front of the end-cap cryostats (see Section 5.4), which will be used to signal inelastic proton-proton collisions. The resulting interaction trigger will be used in setting up the timing of the experiment for the detector readout and of the calorimeter and muon L1 trigger systems.

Once the timing-in of the detector is completed, the minimum-bias trigger will be used to collect data for initial physics studies, in parallel with continuing work on commissioning the rest of the trigger. Since the calorimeter and muon L1 trigger systems are digital, commissioning tasks such as calibration using real (minimum-bias) data can be done offline, comparing the results read out from the trigger systems with corresponding quantities from the detector readout. It is anticipated that the trigger system will be brought on-line progressively. A first step will be to use, in parallel with a pre-scaled minimum-bias selection, the L1 trigger with loose and simple selection criteria, with relaxed requirements in the muon trigger and not using calorimeter quantities such




Figure 308. Trigger efficiency for $B_s \rightarrow D_s^- a_1^+$ events passing the offline selection as a function of the p_T of the B -meson. The results are shown for two HLT scenarios, the most performant one, based on a full scan of the inner detector which can be used at low luminosity, and a RoI-based scan which can be used at higher luminosities (see text).

as isolation and global energy sums, which are sensitive to low-energy detector behaviour. Tighter and more complicated selections will be brought in progressively after thorough offline validation of their performance.

10545 Once the calorimeter and muon L1 triggers are operational, work will ramp up on commissioning the HLT. Many aspects can be addressed offline, using exactly the same algorithms as online, but running on data recorded previously. Then the L2 and EF algorithms will be used online in passive mode, while still recording all events selected by L1. The highest-priority physics channels will initially be covered by high-threshold L1 triggers which are passed through the HLT without further selection, while using the HLT actively elsewhere. Analysis of the recorded data will provide further optimisation of the algorithms and cross-checks on the efficiency of the HLT. As the luminosity increases towards $10^{33} \text{ cm}^{-2} \text{ s}^{-1}$, the full power of the trigger will be required to limit the event rate, while retaining high efficiency for the physics channels of interest.

10555 As an example, one can consider in more detail the commissioning of the L1 muon trigger. Given the low luminosity assumed for initial data-taking ($10^{31} \text{ cm}^{-2} \text{ s}^{-1}$), the related low cavern-background rates expected and the large bunch spacing foreseen (75 ns or more), the configuration parameters of the muon trigger system can initially be relaxed while maintaining acceptable rates. The data collected will be used to check and complete the commissioning of the muon trigger, which has already started using cosmic-ray data. In particular, large samples of muons will be needed to fine-tune the time calibration of the full system, with a required accuracy of about 3 ns in the barrel system. Initial coincidence roads will have been prepared based on simulation and will be available for several p_T -thresholds from about 4 GeV to 40 GeV. Once real data are available with large statistics, these roads will be checked and optimised with muons reconstructed over the full acceptance of the detector, using the p_T -measurement obtained with the inner detector. The commissioning of the muon trigger will use data collected with wide coincidence roads and also with other triggers (minimum bias, jets). The information recorded from the L1 muon trigger will be examined together with the results of the offline reconstruction, allowing measurements of the trigger efficiency for muons as a function of p_T .

10560 In a similar way, the start-up menu for the electron and photon selection must provide data samples needed to commission trigger and detectors, as well as for physics analyses. Relevant physics processes include $J/\psi \rightarrow ee$, $\Upsilon \rightarrow ee$, Drell-Yan, $Z \rightarrow ee$, $W \rightarrow e\nu$ and direct photon production. The menu discussed above selects such events with single electrons with E_T above ~ 10 GeV or single photons with E_T above ~ 20 GeV, in addition to the selection with double-object triggers at significantly lower thresholds.

10.9.6 Evolution to higher luminosities

10575 Building on the experience gained during the start-up phase, the trigger algorithms and parameters will be optimised to provide a trigger selection for use at higher luminosities. As the LHC luminosity ramps up towards its design value, tighter selections will be needed to control the rate. These will include using complex signatures involving multiple observables, higher p_T -thresholds, tighter selection criteria and requiring a more precise matching between different detector systems.

10580 The trigger reconstruction and selection software must be robust against higher detector occupancies, pile-up and cavern backgrounds, which may affect the performance significantly at luminosities above $10^{33} \text{ cm}^{-2} \text{ s}^{-1}$. Many studies have been made to assess the performance of

L1 signature	Rate (kHz)	HLT signature	Rate (Hz)
EM18I	12.0	e22i	40
2EM11I	4.0	2e12i	< 1
MU20	0.8	γ 55i	25
2MU6	0.2	2 γ 17i	2
J140	0.2	μ 20i	40
3J60	0.2	2 μ 10	10
4J40	0.2	j370	10
J36+XE60	0.4	4j90	10
TAU16I+XE30	2.0	j65+xE70	20
MU10+EM11I	0.1	τ 35i+xE45	5
Others	5.0	2 μ 6 for <i>B</i> -physics	10

Table 66. Subset of items from two illustrative trigger menus at L1 (left) and at the HLT (right) for a luminosity of $2 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$. The examples given are more to illustrate the evolution of the rates and thresholds as a function of luminosity, when comparing to Table 64, than to provide accurate predictions of the expected rates.

the trigger and data-acquisition system at high luminosities. Table 66 [247] shows an illustrative sample of L1 and HLT signatures, which could be used under stable operating conditions at luminosities around $2 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$.

The triggers should guarantee coverage of the full physics programme, including searches for new physics and precision measurements of Standard Model parameters. The signatures include single- and di-lepton, photon and jet triggers, similar to those used at $10^{31} \text{ cm}^{-2} \text{ s}^{-1}$, but with higher p_T -thresholds and tighter selection criteria. Requirements on lepton and photon isolation, large E_T^{miss} , and possibly other complex criteria such as flavour tagging, which will have been operated only in a passive or loose mode during the start-up phase, will surely play an important role to achieve a sufficient rate reduction.

As an example, one can consider the case of *b*-jet tagging at the HLT. The performance of the proposed HLT *b*-tagging algorithms is based on transverse and longitudinal impact parameters of charged tracks in jets. The L2 and EF *b*-tagging efficiencies are strongly correlated with the offline *b*-tagging efficiency. To preserve full acceptance for an offline analysis with its *b*-tagging selection criteria set for a given offline *b*-jet efficiency, the L2 and EF *b*-tagging algorithms must operate at an efficiency which is higher. Since most of the offline *b*-tagging results are obtained for *b*-jet efficiencies of $\sim 60\%$, the results quoted here for *b*-tagging in the HLT are given for *b*-jet efficiencies of approximately 80% for L2 and 70% for EF. Figure 309 shows that light-jet rejection factors larger than ten can be achieved, both at L2 and EF for a *b*-jet efficiency of 70% and *b*-jet




Figure 309. Light-jet rejection factor as a function of b -jet efficiency at L2 and EF.




Figure 310. Rate reduction after L2 and EF as a function of E_T , using b -jet signatures with an efficiency of 70% per b -jet.

tagging could thus allow a more flexible operation of the L1 multi-jet trigger menus. To illustrate this, the rate reduction which could be achieved at L2 or EF by requesting two or more b -jets, is shown as a function of E_T in Fig. 310.

10605 Far more accurate projections of the rates given in Table 66 will become possible once real data from the start-up phase have been accumulated and analysed. The total output rate of the trigger system at luminosities above $10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ should remain fixed at approximately 200 Hz, a rate defined by the capabilities of the offline computing system.

10.9.7 Measurements of trigger efficiency from data

10610 Since the trigger efficiency represents a basic element of any physics analysis, it is essential to have several independent methods for estimating it. It is important to depend as little as possible on Monte-Carlo models of LHC physics and on the detector operating conditions, particularly at the start-up of the LHC programme, given the large extrapolation from lower-energy measurements. Techniques under study include the "tag-and-probe" method, e.g. triggering events with the electron
 10615 in $Z \rightarrow ee$ decays and measuring the efficiency to trigger on the positron in addition, and the "bootstrap" method, e.g. using minimum-bias events to measure the efficiency to trigger on low- p_T jets, then triggering on low- p_T jets and using them to measure the efficiency to trigger on higher- p_T jets, etc. Redundant selections can also be used, in which one or more of the steps in the selection are skipped, thereby providing the possibility of determining the corresponding contributions to the
 10620 inefficiency.

As an example, studies have been made for the $Z \rightarrow ee$ tag-and-probe method, using events satisfying the e22i single-electron trigger selection, in which an opposite-charge electron pair has been identified by the offline reconstruction with an invariant mass near the Z peak. Using the second electron in these events as the probe which was not required to pass any trigger selection, the efficiency (relative to the offline selection) of a given trigger signature can be measured. Figure 311
 10625 shows the efficiency of the e22i trigger as a function of p_T of the electron, as measured without any reference to Monte-Carlo truth information in the simulated sample of $Z \rightarrow ee$ events. The shape of the trigger-threshold curves in Fig. 311, obtained using as a reference the Monte-Carlo truth information, are accurately reproduced by the tag-and-probe measurements, and the values agree




Figure 311. Trigger efficiencies as expected to be measured from data using the tag-and-probe method for electrons from approximately 25,000 $Z \rightarrow ee$ decays corresponding to an integrated luminosity of 100 pb^{-1} . The efficiencies are normalised with respect to a reference loose offline selection. The points with error bars show the measured efficiencies after L1 (full circles), L2 (open triangles) and the EF (full squares). Also shown as histograms are the corresponding distributions obtained using as a reference the Monte-Carlo truth information.




Figure 312. Difference between trigger efficiency as expected to be measured from data (using the tag-and-probe method for muons from $Z \rightarrow \mu\mu$ decays) and true efficiency (obtained using as a reference the Monte-Carlo truth information) normalised to true efficiency as a function of η . The efficiencies are normalised with respect to a reference loose offline selection. The results are shown after L1 (top), L2 (middle) and EF (bottom), and correspond to a sample of approximately 50,000 $Z \rightarrow \mu\mu$ decays for an integrated luminosity of 100 pb^{-1} .

10630 to better than 1% on the plateau for a sample of $Z \rightarrow ee$ decays corresponding to an integrated
 10635 $Z \rightarrow \mu\mu$ decays also corresponding to an integrated luminosity of 100 pb^{-1} . The results are plotted
 as the relative difference between the trigger efficiency measured using the tag-and-probe method
 and the true trigger efficiency as obtained from the Monte-Carlo truth information. The statistical
 accuracy achieved per bin is at the percent level.

10640 A similar method can be used to measure the efficiency for triggering on hadronic τ -decays,
 which can be measured using $Z \rightarrow \tau\tau$ samples collected with single electron and muon triggers (as
 shown in Table 64), where one of the τ -leptons decayed leptonically. In events where the second τ -
 lepton decays to hadrons, one can measure the fraction of τ -leptons reconstructed offline, which
 also pass the τ -trigger. This will be done by correlating the detailed information recorded from
 the trigger with the results of the offline reconstruction and will require more integrated luminosity
 10645 than in the case of the electron and muon triggers.

11. Conclusions and outlook

The broad range of physics opportunities and the demanding experimental environment of high-luminosity 14 TeV proton-proton collisions have led to unprecedented performance requirements and hence technological challenges for the general-purpose detectors at the LHC. The overall ATLAS detector design is the result of a complex optimisation process between conflicting requirements. These requirements can be expressed tersely as a set of four basic criteria over a large acceptance in pseudo-rapidity and basically full azimuthal coverage for all of the major detector systems (see Chapter 1 for details):

- very good electromagnetic calorimetry for electron and photon identification and measurements, complemented by full-coverage hadronic calorimetry for accurate jet and E_T^{miss} measurements;
- high-precision muon momentum measurements with the capability to guarantee accurate measurements at the highest luminosity using the muon spectrometer alone;
- efficient tracking at high luminosity for high- p_T lepton momentum measurements, electron and photon identification, τ -lepton and heavy-flavour identification, and full event reconstruction capability;
- efficient triggering with low p_T -thresholds on electrons, photons, muons and τ -leptons, thereby providing high data-taking efficiencies for most physics processes of interest at the LHC.

After approximately fifteen years of detector design, construction, integration and installation, the ATLAS detector is now completed and almost entirely installed in the cavern (see Chapter 9). All detector teams, together with the ATLAS performance and physics working groups, have developed detailed commissioning strategies using cosmic rays, single-beams, and initial data with colliding beams. As more and more detector components become operational, detector calibrations and extensive stand-alone and combined studies with cosmic-ray events are being carried out. These commissioning periods also exercise the full data acquisition chain, including the online and offline data-quality assessment tools and the streaming of events into several physics streams based on the trigger decision. During the spring of 2008, calibration tests and cosmic-ray data-taking will ramp up further, while the few remaining components of the detector are installed and commissioned. The ATLAS detector will be ready for the first LHC collisions in summer 2008.

11.1 Detector installation and hardware status

The status of the ATLAS detector systems at the time of submission of this paper in December 2007 is summarised below and in Table 67.

As described in Chapter 2, the superconducting magnet system comprises the central solenoid, the barrel toroid, two end-cap toroids, and their services. Both the central solenoid and the barrel toroid magnets have been successfully commissioned at full current, and their safety systems have been tested in situ. Their mechanical behaviour as well as the magnetic-field measurements have confirmed the design expectations (see Section 2.2). The magnetic field in the inner-detector cavity

Component	Operational readiness	Comments
Magnets	Residual RMS values between mapping and model	Solenoid and barrel toroid tested at nominal field. End-cap toroids tested at 50% field.
- Solenoid - Barrel toroid - End-cap toroids	~ 0.4 mT for all three field components. 1 – 6 mT depending on position. Analysis in progress.	Nominal field: 2T Preliminary analysis in one sector. Goal is 2 mT.
Inner detector	Fraction of fully functional channels	Mechanical installation complete. In situ cabling in progress.
- Pixel	99.7%	After final integration of system on surface.
- SCT	99.8%	After integration with TRT system on surface.
- TRT	98.4%	Measured in situ.
LAr calorimeters	Fraction of fully functional channels	Installed and operational. Electronics tuning ongoing.
- EM barrel/end-cap	99.98%	Tested cool on surface.
- HEC	99.91%	Tested cool on surface.
- FCAL	99.77%	Tested cool on surface.
Tile calorimeter	Fraction of fully functional channels	Installed and operational. Electronics/power supply tuning on-going.
- Barrel/extended barrel	99.2%	Measured in situ for part of detector.
Muon spectrometer	Fraction of fully functional channels	Installed except for small wheels and some end-wall chambers.
- MDT	99.9%	Tested on surface and partly in situ.
- RPC	99.5%	Tested on surface and partly in situ.
- TGC	99.9%	Tested on surface and partly in situ.
- CSC	99.9%	Tested on surface.
Trigger and data acquisition	System used for cosmic-ray tests and performance verified in stand-alone and commissioning tests.	Readout system installed and operational. Trigger processing-power limited to 40 kHz L1 rate.

Table 67. Hardware status summary of the major ATLAS detector systems. Depending on the installation status, the results are based on measurements on the surface prior to installation and/or measurements after installation in the main cavern.

has been carefully mapped and the residual fractional bending-power uncertainties are well within specifications, i.e. below 5×10^{-4} . In the muon spectrometer, a preliminary analysis of B-sensor readings in a quarter of the field volume yields systematic uncertainties a few times larger than the ultimate desired precision. The end-cap toroids have been tested successfully in stand-alone mode at 50% field. An extensive field-reconstruction campaign over the entire spectrometer volume will be carried out during the full ATLAS magnet-system test which is scheduled in spring 2008, a few months before the start of LHC operation.

As described in Chapter 4, the inner tracking detector combines three concentric detector systems, namely the pixel detector, the SCT and the TRT. Substantial parts of the integrated barrel and end-cap TRT and SCT systems as well as parts of the pixel detector have been successfully operated on the surface in cosmic-ray tests before these systems were installed in the cavern. The tests of the installed inner-detector components are ongoing in parallel with completing the connection of the inner-detector services (cables and pipes). The TRT system is already operational in cosmic-ray runs. The SCT and pixel service connections and subsequent stand-alone and cosmic-ray testing will follow next.

As described in Chapter 5, all three calorimeter cylinders, the barrel and the two end-caps, with the tile calorimeter surrounding the LAr cryostats, are installed in the cavern. The three cryostats are cold and filled with LAr. The main activities are currently focused on the installation of the remaining electronics and on the overall system commissioning. More than half of the calorimeter channels are already part of the regular readout chain.

As described in Chapter 6, the muon spectrometer is instrumented with precision chambers for momentum measurements (MDT's and CSC's) and with fast chambers for triggering (RPC's and TGC's). The construction of the various types of chambers has been completed for the initial detector configuration. The installation of the barrel stations and of the big end-cap wheels has also been completed (with the exception of less than 1% of the barrel stations, which have been left out temporarily for easier access). The installation of the end-wall chambers (MDT's) is ongoing and the pre-assembly of the completed small wheels (MDT's and CSC's) of the first station integrated with the shielding disk and a layer of TGC's is well advanced. In parallel with the completion of the installation of the end-wall chambers and of the small wheels over the next months, the commissioning with cosmic rays is ongoing for both the barrel and end-cap regions, gradually increasing the number of sectors involved in these tests.

As described in Chapter 8, the components of the L1 trigger, of the DAQ/HLT system, and of the detector control systems are in an advanced stage of installation. The L1 trigger system (with its calorimeter, muon and central trigger processor sub-systems) is in its final production and installation phase for both the hardware and the software. The calorimeter trigger installation is completed and the central trigger processor sub-system is in place and routinely used during detector commissioning runs. The readout system, the event builder, and the output to mass storage have been demonstrated in technical runs to deliver the required performance and data through-put rates. The HLT processing power, sufficient to handle a 40 kHz L1 acceptance rate, is planned to be installed in 2008. The HLT algorithms have been successfully tested with physics events pre-loaded in the readout system, and also with cosmic-ray muons. Besides their own commissioning, these systems are used extensively and routinely for the commissioning of specific detector systems and of the overall ATLAS experiment.

The ATLAS control room is fully operational and heavily used. It has become the centre of one of the most prominent activities in the collaboration over the past months, namely periods of global commissioning runs during which, in particular, cosmic-ray events are recorded with the components of the detector already installed and operational in the cavern.

11.2 Outlook on commissioning with data

Chapter 10 summarises the expected performance of the ATLAS experiment. Many of the results are supported by test-beam measurements, in particular for the single-particle response of the detector elements to electrons, photons, pions and muons at various benchmark energies. Other results on the expected performance rely solely on the simulation of the detector geometry, of the detector response, and of the underlying physics processes. These include jets, E_T^{miss} , hadronic τ -decays, b -tagging and trigger performance. Most of these results, particularly the expected trigger rates, are subject to large uncertainties because of the hitherto unexplored energy range for QCD processes at the LHC. At the LHC design luminosity, simulation uncertainties affecting the estimated detector performance also arise from pile-up of p - p interactions (mostly in the triggered bunch-crossing), and from background in the ATLAS cavern consisting predominantly of slow neutrons (see Chapter 3).

In all detector systems, calibration runs of various types are used to map noisy and dead channels. The tile calorimeter also performs dedicated laser and cesium-source calibration runs. These initial calibration data, combined with test-beam measurements performed over the past years, are critical to achieve a sufficient quality of the first collision data. The cosmic-ray data will provide important additional information for aligning the detectors relative to each other. As an example, these data will define an absolute geometry for most of the octants of the barrel muon spectrometer, and will be used as a reference for the alignment based on optical sensors. These data will also be used to define an initial alignment of the major components of the inner detector relative to each other. As shown in Chapter 10, cosmic-ray data are considered as an important ingredient in the overall alignment strategy of the inner detector.

The combination of the results of the detector-specific calibration and commissioning runs with those from the analysis of future large-scale cosmic-ray data will define to a large extent the expected calibration and alignment accuracies for the major ATLAS detector components at the LHC start-up. These ATLAS start-up goals and the ultimate design goals of the experiment, in terms of tracking and calorimeter performance, are summarised in Table 68.

At the start-up of the LHC, after timing-in the detector systems with the colliding LHC bunches and the trigger signals, minimum-bias triggers from scintillator counters will provide large event statistics for initial physics studies at luminosities of $10^{31} \text{ cm}^{-2} \text{ s}^{-1}$ or less. All the triggered events will be used to perform a thorough shake-down of the ATLAS detector systems, thereby refining and completing the dead, noisy and faulty channel maps. The large rates of rather high- p_T isolated tracks (leptons or pions) will be used to refine the inner-detector alignment. High- and low-threshold transition radiation hits from isolated electron and pion tracks will be compared to the expectations from simulation studies.

Minimum-bias events will help to monitor the azimuthal uniformity of the calorimeter response and, to a certain extent, the amount of material in the inner detector. In this initial phase, it will also be crucial to validate the ATLAS calorimeter simulation by comparing shower shapes

	Start-up of LHC	Ultimate goal	Physics goals
Electromagnetic energy uniformity	1–2%	0.5%	$H \rightarrow \gamma\gamma$
Electron energy scale	$\sim 2\%$	0.02%	W mass
Hadronic energy uniformity	2–3%	$< 1\%$	E_T^{miss}
Jet energy scale	$< 10\%$	1%	Top-quark mass
Inner-detector alignment	50–100 μm	$< 10 \mu\text{m}$	b -tagging
Muon-spectrometer alignment	$< 200 \mu\text{m}$	30 μm	$Z' \rightarrow \mu\mu$
Muon momentum scale	$\sim 1\%$	0.02%	W mass

Table 68. Expected calibration and alignment accuracies at the LHC start-up and the ultimate design goals. Examples of physics channels or measurements driving the requirements are indicated in the last column.

for isolated lepton and hadron tracks. The statistics corresponding to a few days of low-luminosity data-taking without toroid field should provide enough straight muon tracks to calibrate the muon optical alignment system to less than 100 μm . This will be improved to 30 μm at higher luminosity, which is required to take full benefit from the spatial resolution of 40 μm per muon chamber. These steps are all necessary to achieve the goal of measuring 1 TeV muon tracks with approximately 10% accuracy.

The commissioning of the overall trigger system will be a gradual process (see Section 10.9.5). Simple inclusive L1 calorimeter and muon triggers will be included first, followed by more complex L1 triggers, involving for example E_T^{miss} . At the same time, the HLT system will begin to operate, initially in pass-through mode in order to test the algorithms, and later using the full power of the HLT. The data collected with the complete low-luminosity trigger menu will contain copious quantities of low-energy leptons from heavy quark decays and also from direct J/ψ and Υ production. The data will contain approximately $5 \times 10^5 W \rightarrow \mu\nu$ and $5 \times 10^4 Z \rightarrow \mu\mu$ decays reconstructed per 100 pb^{-1} of integrated luminosity (the expected rates are somewhat lower for electrons). The low-luminosity trigger menu will also provide abundant samples of high- p_T jets, of prompt photons, mainly from γ -jet events, and of hadronic τ -decays.

All these events will be crucial for an initial validation of the ATLAS performance. More specifically, the inner-detector material can be mapped with photon conversions to an accuracy of 1% X_0 with the statistics available after several months of data-taking. Inclusive electrons can be used to test bremsstrahlung recovery in the inner detector. The inner-detector alignment is expected to converge to the required accuracy of approximately 10 μm soon after the full detector commissioning has started, allowing the constant term in the tracking resolution to be kept below 20% of the overall resolution. Residual inner-detector mis-alignments can be studied with the use of resonances of known mass and lifetime using their decays to lepton pairs, with E/p comparisons for well-measured electrons in the electromagnetic calorimeter, and with high- p_T muons in combined track fits with the muon spectrometer.

A preliminary electromagnetic inter-calibration can be obtained at low luminosity using the

azimuthal symmetry of inclusive isolated electrons from various sources. The next phase of the electromagnetic inter-calibration will use $Z \rightarrow ee$ events. If the inner-detector material is well understood at that point, data corresponding to an integrated luminosity of 100 pb^{-1} would be sufficient to significantly improve the expected initial uniformity of 1–2% to a statistical precision of approximately 0.7%. Further improvements will require the use of E/p distributions from inclusive electrons and/or $W \rightarrow e\nu$ decays.

Jet calibration will use E_T -balancing in di-jet, γ -jet and also Z -jet events. The latter two channels will be important to determine the global jet-energy scale with an expected precision of better than 5% after a few months of data taking. The expected number of ~ 500 fully reconstructed $t\bar{t}$ events for 100 pb^{-1} with one W decaying hadronically and the other one leptonically, will allow a calibration of the jet-energy scale using invariant mass fits to $W \rightarrow jj$ decays.

The most widely studied method to measure with data the performance of b -tagging algorithms at the LHC relies on the selection of $t\bar{t}$ events. However, recent developments show that the techniques extensively used by the Tevatron experiments, combining track-based and soft-muon b -tagging algorithms in di-jet events, could also be used at the LHC. Once large-statistics samples of $t\bar{t}$ events become available, b -jet samples with very high purity will be extracted and used to calibrate, for example, the b -tagging likelihoods directly, thereby reducing the reliance on Monte-Carlo simulation.

One of the most difficult detector observables to measure accurately is E_T^{miss} . Because it is sensitive to many new physics signatures, the tails of its distribution must be precisely calibrated with data before E_T^{miss} measurements can be used for discrimination and especially reconstruction purposes. A reliable measurement of E_T^{miss} requires the removal from the data sample of beam-halo muons, beam-gas collisions, cavern background and cosmic rays. Moreover, all calorimeter cells must be calibrated (for both electromagnetic and hadronic showers), and noise levels and deficient cells must be mapped and corrected for. Initial data-driven E_T^{miss} studies will use minimum-bias and di-jet events, analysing the missing E_T resolution as a function of the summed transverse energy. With larger statistics, the use of $W \rightarrow l\nu$ decays, of mass-constrained $t\bar{t}$ events and of $Z \rightarrow \tau\tau$ decays should lead to a calibration of the E_T^{miss} -scale to about 5%.

Initial physics measurements will primarily focus on Standard Model processes with high cross-sections. The most prominent among these will be the production of hadronic jets, of W and Z bosons, and also of $b\bar{b}$ and $t\bar{t}$ pairs. Analyses aiming at searches for new phenomena will first concentrate on the understanding of the detector performance and on these Standard Model processes. The ATLAS performance and physics working groups will exploit to the full the rich variety of known physics processes at the LHC to calibrate the analysis tools and thus to prepare for the exciting searches for new physics, which have been the driving motivation of large numbers of physicists during the many years of work which have brought the collaboration this far.

11.3 Future changes to the ATLAS detector system

As the luminosity of the LHC machine reaches its design value of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, the detector parts which have been staged due to budgetary constraints need to be completed. The main items falling into this category are a significant part of the HLT processing farm, some parts of the muon spectrometer including in particular the monitored drift-tube chambers in the transition region between the barrel and end-cap toroids, and also some of the shielding elements in the forward region.

During this phase, the performance of the ATLAS detector will be continuously evaluated and optimised, in particular as physics samples are used to further study and improve the calibration and alignment procedures. Pile-up effects will also need to be understood and dealt with as the luminosity increases.

After reaching design luminosity, the challenge will be to operate and optimise the ATLAS detector, its multi-faceted trigger system and the various physics analyses over several years of data-taking. The detector parts are generally designed for ten years of operation (conservatively estimated to correspond to an integrated luminosity of up to 700 fb^{-1}). The most critical element is the innermost layer of the pixel detector or vertexing layer, which is located at a radial distance of only 5 cm from the beam-pipe. This layer is designed to survive a 1 MeV neutron equivalent fluence of approximately 10^{15} cm^{-2} , which corresponds to less than half the integrated luminosity mentioned above. The pixel vertexing layer may need to be replaced earlier than other parts of the detector.

If the LHC luminosity were to be increased significantly beyond the current estimates, as suggested in some studies for the LHC machine upgrade on a time-scale not earlier than 2015, several detector components are likely to need substantial changes. In particular, the inner-detector system would need to be completely replaced, and certain calorimeter, muon and shielding elements in the forward directions would also require significant changes and improvements. Research and development work has started in earnest within the collaboration in several of the areas mentioned above. However, a decision about the necessity, scope and time-scale of such an upgrade can only be made after a few years of LHC and detector operation, considering both the physics results and the performance of the machine and the status of the experiments at that point.

A. Appendix: ATLAS acronyms

ADC	Analogue-to-Digital Converter
ALFA	Absolute Luminosity for ATLAS
ASD	Amplifier/Shaper/Discriminator
ASDBLR	Amplifier/Shaper/Discriminator/BaseLine Restoration
ASIC	Application-Specific Integrated Circuit
ASM	Amplification, Sampling, (digitization) and Multiplexing
ATLAS	A Toroidal LHC ApparatuS
BC	Bunch Crossing
BCID	Bunch-Crossing IDentification
BCM	Beam Conditions Monitor
BC-mux	Bunch-Crossing MultipleXing
BEE	Barrel End-cap Extra
BIL	Barrel Inner Large
BIR	Barrel Inner Rail
BIS	Barrel Inner Small
BM	Barrel Middle
BML	Barrel Middle Large
BMS	Barrel Middle Small
BO	Barrel Outer
BOC	Back Of Crate
BOL	Barrel Outer Large
BOS	Barrel Outer Small
BPM	BiPhase Mark
BT	Barrel Toroid
CANbus	Controller Area Network bus
CALbus	Calibration bus
CDD	CERN Drawing Directory
CERN	European Organization for Nuclear Research
CFS	Central File Server
CIC	Common Infrastructure Control
CIS	Charge Injection System
CMA	Coincidence Matrix chip
CM	Coincidence Matrix
CMM	Common Merger Module
CMOS	Complementary Metal-Oxide Semiconductor
COMbus	Common Timing and Trigger Bus
COOL	ATLAS-wide conditions database
CP	Cluster Processor
CPM	Cluster Processor Module
CRC	Cyclic Redundancy Check
CSC	Cathode Strip Chambers
CS	Central Solenoid
CSM	Chamber Service Module
CTP	Central Trigger Processor
DAC	Digital-to-Analogue Converter
DAQ	Data AcQuision system
DCS	Detector Control System
DFM	Data Flow Manager
DMILL	Durci Mixte sur Isolant Logico-Lineaire (a radiation-hard ASIC technology)
DORIC	Digital Opto-Receiver Integrated Circuit
DSP	Digital Signal Processors
DSS	Detector Safety System
EB	Extended Barrel
EC	End-Cap
ECT	End-Cap Toroid
EDMS	Engineering Data Management System
EEL	End-cap Extra Large
EES	End-cap Extra Small
EF	Event Filter

EI	End-cap Inner wheel
EIL	End-cap Inner Large
EIS	End-cap Inner Small
ELMB	Embedded Local Monitor Board
EMB	ElectroMagnetic Barrel calorimeter
EMD	Equipment Management Database
EMEC	ElectroMagnetic End-cap Calorimeter
EM	ElectroMagnetic
EM	End-cap Middle wheel
EML	End-cap Middle Large
EMS	End-cap Middle Small
ENC	Equivalent Noise Charge
EOL	End-cap Outer Large
EOS	End-cap Outer Small
EST	ElectroStatic Transformer
FADC	Flash ADC
FCal	Forward Calorimeter
FEB	Front-End Board
FECcont	Front-End Crate controller board
FEC	Front-End Crate
FE	Front-End
FIFO	First-In/First-Out
FI	Forward Inner wheel
flVPS	finger Low Voltage Power Supply
FPGA	Field-Programmable Gate Array
FPIAA	Find Persons Inside ATLAS Area
GSEL	Gain-SElector chip
GCS	Global Control Stations
HAD	HADron calorimeter
HEC	Hadronic End-cap Calorimeter
HF	Steel structures below access shafts
HLT	High-Level Trigger
HO	Blue support structure on ends of ATLAS cavern
HS	Blue support structure on sides of ATLAS cavern
HV	High Voltage
HVPS	High Voltage Power Supply
IC	Integrated Circuit
ID	Inner Detector
IS	Information Server
ITC	Inter TileCal scintillators
IWV	Inner Warm Vessel
JEM	Jet/Energy Module
JEP	Jet/Energy-sum Processor
JTAG	Joint Task Action Group
L1A	Level-1 Accept
L1Calo	Level-1 Calorimeter trigger
L1	Level-1 trigger
L2	Level-2 trigger
L2PU	Level-2 Processing Unit
L2SV	Level-2 SuperVisor
LAr	Liquid Argon
LB	Long barrel
LCS	Local Control Stations
LFS	Local File Server
LHC	Large Hadron Collider
LUCID	LUMinosity measurement using Cerenkov Integrating Detector
LVDS	Low-Voltage Differential Signalling
LV	Low Voltage
LVPS	Low Voltage Power Supply
MCC	Module Control Chip
MDT	Monitored Drift Tubes

MIP	Minimum Ionising Particle
MRS	Message Reporting Service
MTF	Manufacturing and Test Folder database
MUCTPI	MUon-to-Central-Trigger-Processor-Interface
NEF	Neutron Equivalent Fluence
NEG	Non-Evaporable Getter
NMR	Nuclear Magnetic Resonance
NRZ	Non Return to Zero
ODH	Oxygen Deficiency Hazard
OFC	Optimal Filtering Coefficients
OF	Optimal Filtering
OHS	Online Histogramming Service
OMB	Optical Multiplexer Board
OTx	Optical Transmitter
PEEK	PolyEther-Ether-Ketone
PITbus	Pattern In Time Bus
PL	Pad Logic
PMT	PhotoMultiplier Tube
PP	Patch Panel
PPM	Pre-Processor Module
PPr	Pre-Processor
PS	Presampler
PST	Pixel Support Tube
PU	Processor Unit
PVSS	Prozessvisualisierungs und Steuerungs System
QSP	Quarter Service Panel
QUAD	QUADrupole
ROBIN	ReadOut Buffer Module
ROB	ReadOut Buffer
ROC	ReadOut Crate
ROD	ReadOut Driver
RoIB	Region-of-Interest Builder
RoI	Region-of -Interest
ROL	ReadOut Link
ROS	ReadOut System
RPC	Resistive Plate Chambers
SCA	Storage Capacitor Array
SCA	Switched Capacitor Array
SCS	Sub-detector Control Stations
SCT	SemiConductor Tracker
SEU	Single Event Upset
SFI	Event building node
SFO	Event filter output node
SIC	System Interlock Card
SLB	SLave Board
SL	Sector Logic
SLC	Scientific Linux CERN
SPAC	Serial Protocol for Atlas Calorimeters
SSW	Star SWitch
TAN	Target Absorber Neutral
TAS	Target Absorber Secondaries
TBB	Tower Builder Board
TBM	Trigger and Busy Module
TC	Technical Coordination
TDAQ	Trigger and Data AcQuision
TDB	Tower Driver Board
TDC	Time-to-Digital Converter
TDR	Technical Design Report
TGC	Thin Gap Chambers
TileCal	Tile Calorimeter
TLA	Three Letter Acronym

TMB	Technical Management Board
ToT	Time over Threshold
TPG	Thermal Pyrolytic Graphite
TR	Transition Radiation
TRT	Transition Radiation Tracker
TTCR	Trigger, Timing, and Control Receiver
TTCrx	Trigger, Timing, and Control Receiver chip
TTC	Timing, Trigger, and Control
TTCvi	Timing, Trigger, and Control VME interface module
TVS	Transient Voltage Suppressor
USA	Underground Service Area
VA	Vacuum Argon end-cap
VCEL	Vertical Cavity Surface-Emitting Laser Diode
VDC	VCSEL Driver Chip
VI	Vacuum Inner detector
VJ	Vacuum forward shielding
VME	Versa Module Eurocard
VMEbus	Versa Module Euro bus
VT	Vacuum Toroid end-cap
WLS	WaveLength Shifting
ZDC	Zero Degree Calorimeter

B. Acknowledgements

We are greatly indebted to all CERN's departments and to the LHC project for their immense efforts not only in building the LHC, but also for their direct contributions to the construction and installation of the ATLAS detector and its infrastructure. We acknowledge equally warmly all our technical colleagues in the collaborating Institutions without whom the ATLAS detector could not have been built. Furthermore we are grateful to all the funding agencies which supported generously the construction and the commissioning of the ATLAS detector and also provided the computing infrastructure.

We acknowledge the support of ANPCyT, Argentina; Yerevan Physics Institute, Armenia; ARC and DEST, Australia; Bundesministerium für Wissenschaft und Forschung, Austria; National Academy of Sciences of Azerbaijan; State Committee on Science & Technologies of the Republic of Belarus; CNPq and FINEP, Brazil; NSERC, NRC, and CFI, Canada; CERN; NSFC, China; Ministry of Education, Youth and Sports of the Czech Republic, Ministry of Industry and Trade of the Czech Republic, and Committee for Collaboration of the Czech Republic with CERN; Danish Natural Science Research Council; IN2P3-CNRS and Dapnia-CEA, France; Georgian Academy of Sciences; BMBF, DESY, and MPG, Germany; GSRT, Greece; ISF, MINERVA, GIF, DIP, and Benozio Center, Israel; INFN, Italy; MEXT, Japan; CNRST, Morocco; FOM and NWO, Netherlands; The Research Council of Norway; Ministry of Science and Higher Education, Poland; GRICES and FCT, Portugal; Ministry of Education and Research, Romania; Ministry of Education and Science of the Russian Federation, Russian Federal Agency of Science and Innovations, and Russian Federal Agency of Atomic Energy; JINR; Ministry of Science, Serbia; Department of International Science and Technology Cooperation, Ministry of Education of the Slovak Republic; Slovenian Research Agency, Ministry of Higher Education, Science and Technology, Slovenia; Ministerio de Ciencia y Tecnología, Spain; The Swedish Research Council, The Knut and Alice Wallenberg Foundation, Sweden; State Secretariat for Education and Science, Swiss National Science Foundation, and Cantons of Bern and Geneva, Switzerland; National Science Council, Taiwan; TUBITAK, Turkey; The Science and Technology Facilities Council, United Kingdom; DOE and NSF, United States of America.

The ATLAS detector design and construction has taken about fifteen years, and our thoughts are with all our colleagues who sadly could not see its final realisation.

C. Annex

We are greatly indebted to all our colleagues who, together with the authors of this paper, have worked hard for years on the installation and we extend to all of them our heartfelt thanks for their invaluable contribution to this tremendous task.

10895 A. Abbasi, A. Abdallah, J. Adamek, E. Adli, G. Adroit, D. Afanasyev, V. Afonin, M. Afzal, S. Ageev, E. Ahmad, S. Ahmad, J. Ahmed, S. Ahmed, N. Akbar, V. Akhnazarov, G. Akhpertjanyan, K. Akhtar, N. Akhtar, S. Akhtar, F. Alberti, A. Alekhin, A. Ali, M. Alves, S. Ami, W. Andreazza, M. Andrianov, P. Anisimov, A. Anisovets, N. Anokhin, D. Antonov, R. Arsenikhina, W. Arshad, D. Arutinov, A. Asarsa, N. Ashfaq, A. Astakhov, G. Avoni, N. Awan Muhammad, 10900 A. Baig, F. Bais, S. Bakatura, S. Bakov, V. Balabanov, A. Balashov, C. Baldanza, F. Baltasar Dos Santos Pedrosa, H. Band, N. Baranenkov, V. Barozier, J-Ch. Barrière, D. Barzilay, C. Batailt, R. Baud, Ch. Bault, A. Bazanov, A. Bekele, N. Belikov, S. Belov, A. Belyakov, N. Belyakov, A. Belyushin, J. Bendotti, Ph. Benoit, E. Berbee, J. Bernaz, S. Berry, D. Bertet, E. Bertier, P. Bijoneau, V. Biryukov, H. Blampey, Z. Blaszczyk, M. Blumentritt, V. Bobrovnikov, Y. Bocharov, 10905 S. Bochkov, A. Bode, N. Bogdanov, A. Boguslavskiy, O. Bohner, M. Bonnet, F. Bonthond, O. Boriosov, Ch. Bourgeois, N. Bourgeois, G. Braga, H. Brettel, R. Brocard, J. Brown, B. Brunel, X. Brunel, S. Buda, F. Budano, R. Buis, V. Bukhterev, K. Bulatov, A. Bulychev, J. Buskop, A. Bychkov, A. Bykov, P. Bylinkin, M. Cadabeschi, D. Calabro', T. Calderon, A. Calderone, W. Cameron, E. Capitolo, C. Capocchia, M. Capponi, G. Capradossi, V. Carmona, P. Carrie, S. Cerioni, C. Chagas, Th. Chaminade, P. Charra, I. Chayukov, N. Cherkimov, S. Cherkasov, D. Chernichenko, 10910 N. Cherny, V. Cherpunov, A. Chiarini, A. Chichikin, G. Chiodi, A. Chomicova, J. Chorazak, R. Chritin, I. Chubrikov, A. Chuguev, J. Cid, X. Ciodaro, G. Cipolla, V. Cmirmov, L. Colly, R. Consentino, F. Corbaz, J-P. Corso, K. Crum, G. Cucurru, M. Cuvara, W-D. Cwienk, B. Dabrowski, D. Damianoglou, E. Danilevich, N. Danilov, V. Danilov, E. Dantoing, I. D'Antone, A. Daragan, S. Darf, A. d'Auria, S. David, L-P. De Menezes, D. De Paoli, N. Debiee, Ch. Dechelette, 10915 M. Decombaz, D. Demin, N. Denisov, V. Dergunov, L. Deront, D. Desforges, M. Despet, S. Di Marco, L. Di Stante, P. Dias Ferrao, M. Din, A. Ditta, N. Dixon, D. Diyakov, O. Dmitriev, I. Dmitrov, E. Dobrolyubov, N. Dobrynin, M. Dogaru, M. Dubiel, J. Dubois, B. DubouŐ, C. Duijsters, A. Dukin, V. Durastanti, M. Duret, B. Duthion, E. Dutkiewicz, V. Dyusov, F. Eduard, A. Efremov, A. Elahi, K. Elumakhov, A. Elvin, B. Erbel, A. Ershova, J. Escourrou, J. Etheridge, C. Evans, 10920 A. Evseev, E. Evseev, A. Evstifeev, F. Doherty, J. Farrel, G. Fausto, A. Fedorov, B. Fedoseev, P. Ferret, B. Ferreira, J. Ferrer Garcia, L. Filho, Y. Filippov, B. Finegan, S. Finelli, F. Fink Grael, M. Fink, G. Fiore, I. Firsae, A. Fischer, A. Florek, B. Florek, M. Fol, A. Folley, G. Fomichev, A. Foreste, D. Forshier, T. Franc, M. Frasc, A. Freddi, A. French, A. Froton, D. Gabriele, Y. Gal, 10925 F. Galleazzi, P. Gallno, P. Galukh, S. Galushchak, K. Galuszka, M. Gandel, V. Gapelchenkov, A. Garanin, S. Garanin, G. Gariano, J. Garwolinski, A. Gaweda, D. Geiss, E. Gennari, I. Georgadze, F. Gherarducci, H. Gjermansen, A. Gleyberman, F. Godin, H. Goldshtein, A. Golovin, A. Goncalves de Oliveria, A. Gongadze, Y. Gorbunov, P. Gorelikov, A. Gorokhov, V. Gorshkov, D. Grafov, Y. Grandjean, S. Granena, Ph. Gravelle, O. Gregoryew, O. Grigoriev, Y. Grinitzki, 10930 E. Grunthalova, N. Gruzdev, S. Gruzdev, V. Gruzdev, A. Gubanov, O. Gudkov, P. Guglielmini, A. Guiga, V. Guliaev, A. Gulin, D. Gvozdev, J. Habring, L. Haddad, F. Hamdi, E. Hanuska, A. Haq, J. Heintz, O. Heintz, A. Hoffman, J. Homma, C. Houd, J. Hrivna, J. Huber, M. Humeau, G. Hus-

sain, I. Hussain, M. Hussain, S. Hussain, T. Hussain, A. Iaciofano, A. Iacoangeli, B. Iankovski, V. Ignatyev, D. Ilin, I. Ilyashenko, H. Imran, A. Imtiaz, S. Infante, A. Innocente, A. Iqbal, Z. Iqbal, A. Isaev, G. Iuvino, N. Ivanchenko, V. Ivanov, Y. Ivanov, V. Izotov, O. Jamet, N. Javadov, L. Jeanty, M. Jeckel, D. John, D. Jourde, C. Jouve, J-N. Joux, G. Juban, J. Jusko, O. Jusko, B. Kachalkin, S. Kajfasz, M. Kalaev, M. Kalyamin, R. Kantor, J. Kantorski, A. Kapiehnikov, Y. Karalevich, G. Karam, V. Karnaukhov, V. Karpishin, V. Karpov, M. Kasay, I. Kashin, V. Katarushkin, O. Kehat, D. Khalin, A. Khan, F. Khan, M.I. Khan, M.R. Khan, A. Kharchenko, Y. Kharzheev, E. Khodoruk, A. Khokhlov, S. Khomayko, K. Khormovas, Y. Khromov, S. Khvorostinin, S. King, S. Kiraynov, V. Kir'ianov, V. Kirillov, E. Kiselev, V. Kladov, N. Kleinermann, A. Klementyev, P. Klemm, D. Kochev, H. Kok, A. Kolbasin, N. Kolesnikov, V. Kolgin, V. Kolomoets, V. Komissarchikov, N. Kopeykin, S. Koptelov, A. Korenchenko, V. Kormilitsyn, S. Kornishkin, A. Korobov, D. Korobov, A. Korotkov, A. Korporaal, I. Kostyukhina, Z. Kotek, S. Kotelnikov, Z. Kott, L-J. Kottelat, V. Koukol, A. Kovalev, B. Kovalevskiy, I. Kozhina, J. Kozien, O. Kozlova, M. Kral, S. Kramarenko, A. Kratochvil, M. Kraytor, O. Krivov, E. Kudinov, A. Kudryashov, A. Kudryavtsev, W. Kuilman, Y. Kukovenkov, A. Kulagin, S. Kulakov, A. Kuleshov, V. Kuleshov, A. Kulikov, M. Kulikov, V. Kulikov, P. Kulka, E. Kumakshin, A. Kursov, A. Kuznetsov, N. Kuznetsov, V. Kuznetsov, E. Kvashina, D. Kyslukha, M. Lages, P. Lancon, D. Laporte, L. Lars-Eric, H. Laskus, N. Latalin, G. Laurenti, A. Lavrentyev, D. Lawlor, I. Lax, D. Layton, A. Le Coguie, B. Lebegue, S. Leber, R. Leboeuf, Ch. Leboube, V. Ledl, R. Leguyt, A. Lekhtman, A. Leonhardt, A. Leonhardt, A. Leonov, A. Lepikov, K. Lepo, J. Lethinen, A. Levana, K. Levterov, Ch. Li, B. Liccardi, J. Ligocki, F. Limia- Conde, M. Lippert, M. Lloreda, M. Llyas, N. Lobanov, S. Lobanov, T. Locatelli, S. Loffredo, D. Lomidze, R. Lubov, V. Luchinin, R. Lunadei, F. Lupenko, M. Lyablin, O. Lykov, V. Lysikov, A. Lyubimtsev, T. Machaidze, J. Magalhaes, W. Mahboob, K. Mahmood, J. Maillefaud, J. Majka, E. Makarov, A. Maksimov, A. Malik, F. Malik, S. Malook, I. Malyukova, A. Mapelli, F. Mariani, I. Mar'in, V. Marousov, D. Martin-Gil Calderon, S. Martirosyan, A. Mashasha, A. Masson, M. Mathieu, M. Mathieu, Z. Matijas, M. Matin, M. Matiullah, E. Mattelon, A. Matyushin, A. Maximov, G.M.Mayers, L. Mazurin, F. McEwan, D. Mcfarlane, A. Medvedev, A. Mehmood, K. Mehmood, T. Mehmood, V. Melekhov, J.J. Melone, L. Mendel, S. Meneghini, D. Mergelkuhl, G. Merkulov, F. Merlet, J. Metselaar, A. Miccoli, R. Michinelli, V. Mikhaylov, A. Milenko, S. Mirek, Y. Mishagin, A. Mocova, A. Moita, J. Moita, J. Molina-Perez, G. Monshi, B. Monticelli, P. Mooshofer, J-C. Moreau, N. Morozov, R. Morton, N. Mosulishvili, G. Mouquin, S. Mousatov, A. Mozzhukhin, A. Muhammad, S. Muhammad, B. Munneke, J-F. Muraz, G. Murino, I. Muromskiy, A. Mysnik, S.ă Naydenov, S. Nazarov, V. Neapolitanski, R. Necca, S. Nechaev, E. Nedza Waclaw, A. Nefedyev, V. Nekrasov, T. Nephew, A. Neverov, I. Nikitin, L. Nikitina, V. Nikitina, I. Nikitine, M. Nila, T. Niwa, J. Noels, V. Novikov, N. Novozhenin, D. Nozdrin, L. Obartuch, K. Ohanyan, G. Olesen, A. Olyunin, R. Organezov, A. Orlov, N. Orlov, A. Orsini, W. Ostrowicz, Q. Ouyang, J. Palacky, E. Pankratov, M. Papiez, E. Paramoshkina, D. Parchet, G. Paredes, A. Park, B. Pasmantirer, E. Pastori, S. Pavis, P. Pavlicek, Y. Pavlov, D. Pchelov, G. Pellegrini, S. Pelletier, F. Pellrgrino, Y. Pénichot, V. Penna, J. Penwel, B. Peralva, G. Perez, E. Perez-Duenaz, D. Perminov, P. Pernaud, M. Perry, L. Pertov, V. Pertushin, P. Petit, V. Petrov, J. Pieron, M. Pikna, G. Pileggi, Ph. Pinget, B. Pirollet, O. Pisano, A. Pisarev, E. Pleshkov, M. Podmolik, N. Politsyn, I. Poluektov, P. Polyakov, E. Pominov, X. Pons, B. Ponzio, G. Popeneciu, A. Popov, S. Popov, I. Porizko, S. Porokhovoy, A. Potapov, N. Pravotorov, D. Prigent, S. Prodon, F. Profeti, A. Prudkoglyad, A. Prukglyad,

M. Pryanichnikov, S. Pshenichnyy, A. Pushkarev, S. Pyatt, J. Quetsch, M. Qureshi, M. Rachwalik, W. Rahman, L. Raines, N. Rama, M. Rashid, J. Rasmussen, S. Ravat, L. Reed, M.B.Reilly, Y. Reinert, S. Ren, S. Reshetnikov, A. Revin, M. Riaz, E. Richards, K. Richards, A. Rietmeijer, M. Rizwan, D. Rizzi, Y. Robert, R. Rocco, G. Roche, V. Rodiakino, V. Rodionchikov, J. Roevekamp, J. Rogowski, V. Romano, B. Romdhane, A. Romero, L. Rose-Dulcina, F. Rosset, A. Rovani, P. Rozanek, 10980 T. Rudenko, A. Ruggieri, D. Ruggieri, R. Ruggiero, V. Russo, I. Ryzhov, V. Ryzhov, S. Sablyukov, R. Sadilov, Ch. Sagianos, S. Salakhutdinov, A. Samoshkin, V. Samsonov, M. Sandomirskiy, Y. Sandomirskiy, L. Sargsyan, V. Satti Pervez, V. Saveliev, D. Savin, O. Savin, J. Saxon, S. Sbitnev, 10985 C. Scagliotti, Ch. Scalisi, J. Schmalzer, H. Schuijlenburg, J. Sedlak, K. Sedlak, R. Sedlmeyer, E. Seletskaya, A. Seletskiy, X. Serghi, E. Serochkin, M. Serochkin, V. Serukhin, A. Seweryn, E. Seweryn, K. Sexton, M. Shafiq, S. Shah, A. Shah Zaigham, M. Shahbaz, A. Shalimov, V. Sharapov, A. Sharmazanashvili, N. Sharmazanashvili, A. Shastov, A. Sheibin, N. Sherbakov, B. Shishigin, D. Shooltz, Y. Shovkun, O. Shporker, V. Shpyrkin, I. Shugaev, D. Sidorov, S. Sidorov, G. Silushin, A. Simonov, V. Sinkin, I. Siudek, E. Sizikov, F. Skzrecz, T. Slavicek, A. Smirnov, S. Smirnov, 10990 V. Smirnov, I. Smotrova, C. Snippe, R. Snopkov, A. Soares, A. Sobolev, M. Socha, S. Sokolov, N. Sokovinov, A. Solodkov, Y. Solomatin, A. Soloviev, J. Sorokina, A. Sorokovikov, M. Sowinski, M. Spairani, A. Starikov, J. Steffens, J. Stehlik, V. Stepanov, D. Steyaert, P. Stock, A. Stopa, A. Straczek, P. Straczek, S. Studenov, V. Sultanov, Y. Sun, A. Surenskiy, A. Surmava, V. Suslov, R. Swift, E. Syshchikov, A. Szczepanski, A. Szperlak, M. Tahir, H. Takia, O. Talamanov, N. Tamonov, M. Tanveer-Ahmed, A. Tazzioli, A. Tchuguev, J. Tempir, N. Terekhov, V. Terekhov, Y. Terentyev, V. Tereshin, D. TÓzier, G. Theodoropoulos, I. Thiel, Ch. Thomas, A. Tikhonov, G. Timokhin, I. Titenkov, E. Tkachenko, S. Tolaini, D. Topiline, E. Torki, G. Torromeo, M. Travaglini, N. Trubin, A. Trushchev, V. Tsishevskiy, V. Tsurkan, V. Tsvetkov, P. Turco, E. Tusi, A. Tveryakhin, V. Tyurin, V. Udovichenko, M. Umer, S. Unzhakov, A. Usachev, A. Ustinov, J. v. Loeben, A. Vakantov, A. Vallet, W. Van Emden, M. Van Overbeek, M. van Overbeek, J. Van Santen, A. Varga, 11000 O. Vasilevskiy, L. Vasiluev, M. Vasily, L. Vasilyev, A. Vassalo, T. Vassilieva, A. Vassine, A. Vassioukhine, F. Vercellati, V. Vetrov, J. Vettori, A. Vinogradov, E. Viscione, K. Vlasik, J. Vodolan, S. Vogt, R. Von Boehn-Buchholz, Ch. Vottis, A. Vyakhorev, A. Vyatkin, V. Vysokov, V. Vysokov, A. WanMer, A. Wasem, J. Weichert, H. Wetteskind, M. White, A. Wiart, A. Wimmer, S. Witt, 11005 T. Wojas, E. Wood, K. Wood, A. Yablokov, V. Yaichkov, A. Yakutin, A. Yandyan, I. Yanin, B. Yankovski, A. Yayechev, O. Yudin, S. Yukhimchuk, A. Yurchak, A. Zabaluev, S. Zaborshchikov, R. Zajac, A. Zaman, I. Zaytsev, Y. Zhadnov, V. Zhdanov, N. Zhukov, R. Zibziner, J. Zimmer, M. Znoy, D. Zoitoiu, Y. Zolnikov, P. Zotov, V. Zubrovskiy, M. Zuffa, O. Zvadyuk, S. Zvyagintsev

References

- 11010 [1] ATLAS Collaboration, Detector and Physics Performance Technical Design Report CERN/LHCC/99-14/15 (1999).
- [2] A. Yamamoto et al., The ATLAS Central Solenoid, Nucl. Inst. Meth. **A584** (2008) 53.
- [3] A. Dudarev et al., Construction of the ATLAS toroid magnet system, JINST **888** (2008) to be submitted to JINST.
- 11015 [4] ATLAS Collaboration, Magnet System Technical Design Report, CERN/LHCC/97-18 (1997).
- [5] ATLAS Collaboration, Barrel Toroid Technical Design Report, CERN/LHCC/97-19 (1997).
- [6] ATLAS Collaboration, End-cap Toroid Technical Design Report, CERN/LHCC/97-20 (1997).
- [7] ATLAS Collaboration, Central Solenoid Technical Design Report, CERN/LHCC/97-21 (1997).
- [8] H. H. J. ten Kate, Superconducting Magnet System for the ATLAS Detector at CERN, IEEE Tr. Appl. Superconductivity 9 (1999) 841.
- 11020 [9] H. H. J. ten Kate, The ATLAS Superconducting Magnet System: Status of Construction & Installation, IEEE Tr. Appl. Superconductivity 16 (2006) 499.
- [10] A. Yamamoto et al., Progress in the ATLAS Central Solenoid, IEEE Tr. Appl. Superconductivity 10 (2000) 353.
- 11025 [11] S. Mizumaki et al., Fabrication and Mechanical Performance of the ATLAS Central Solenoid, IEEE Tr. Appl. Superconductivity 12 (2002) 416.
- [12] R. Ruber et al., Quench Characteristics of the ATLAS Central Solenoid, IEEE Tr. Appl. Superconductivity 16 (2006) 533.
- [13] R. Ruber, Ultimate performance of the ATLAS Superconducting Solenoid, IEEE Tr. Appl. Superconductivity 17 (2007) 1201.
- 11030 [14] P. Védérine et al., Manufacturing and integration progress of the ATLAS barrel toroid magnet at CERN, IEEE Tr. Appl. Superconductivity 14(2) (2004) 491-494.
- [15] J.-M. Rey et al., Cold mass integration of the ATLAS Barrel Toroid Magnets at CERN, IEEE Tr. Appl. Superconductivity 16 (2006) 553.
- 11035 [16] P. Védérine et al., Completion of the Manufacturing of the ATLAS Barrel Toroid Magnet at CERN, IEEE Tr. Appl. Superconductivity 16 (2006) 504.
- [17] P. Miele et al., The ATLAS magnet test facility at CERN, IEEE Tr. Appl. Superconductivity 11 (2000) 1713.
- [18] A. Dudarev et al., First full-size ATLAS barrel toroid coil successfully tested up to 22 kA at 4 T, IEEE Tr. Appl. Superconductivity 15(2) (2005) 1271-1275.
- 11040 [19] A. Dudarev et al., On-Surface Test of the ATLAS Barrel Toroid Coils: Overview, IEEE Tr. Appl. Superconductivity 16 (2006) 508.
- [20] C. Beriaud et al., Quench Behavior of the ATLAS Barrel Toroid, Paper presented at IEEE International Magnet Technology Conference, Philadelphia (2007).
- 11045 [21] J.-J. Rabbers et al., Experimental and theoretical investigation of the ramp losses in the conductor and coil casing of the ATLAS Barrel Toroid Coils, IEEE Tr. Appl. Superconductivity 16 (2006) 549.

- [22] A. Foussat et al., Assembly Concept and Technology of the ATLAS Barrel Toroid, *IEEE Tr. Appl. Superconductivity* 16 (2006) 565.
- 11050 [23] D. E. Baynham et al., Engineering Status of the End Cap Toroid Magnets for the ATLAS Experiment, *IEEE Tr. Appl. Superconductivity* 10 (2000) 357.
- [24] D. E. Baynham et al., ATLAS End Cap Toroid Cold Mass and Cryostat Integration, *IEEE Tr. Appl. Superconductivity* 16 (2006) 537.
- [25] D. E. Baynham et al., ATLAS End Cap Toroid Integration and Test, *IEEE Tr. Appl. Superconductivity* 17 (2007) 1197.
- 11055 [26] P. Miele et al., ATLAS magnet common cryogenic, vacuum, electrical and control systems, *IEEE Tr. Appl. Superconductivity* 14 (2004) 504-508.
- [27] ATLAS Collaboration, Muon Spectrometer Technical Design Report (Section 10.1), CERN/LHCC/97-22 (1997).
- 11060 [28] ATLAS Collaboration, Muon Spectrometer Technical Design Report (Section 10.2.3), CERN/LHCC/97-22 (1997).
- [29] ATLAS Collaboration, Muon Spectrometer Technical Design Report (Section 10.2.1 and, in particular, Figs. 10-15 to 10-18), CERN/LHCC/97-22 (1997).
- [30] T. Nikitina and F. Bergsma, A Program to Calculate the ATLAS Magnetic Field, ATLAS Note ATL-MAGNET-2001-002 (2001).
- 11065 [31] TOSCA, program developed by Vector Fields Limited, England: www.vectorfields.com.
- [32] ATLAS Collaboration, Muon Spectrometer Technical Design Report (Section 10.2.2), CERN/LHCC/97-22 (1997).
- [33] F. Bergsma, Calibration of Hall sensors in three dimensions, Proc. 13th International Magnetic Measurement Workshop, Stanford (2003).
- 11070 [34] S. Snow et al., Field mapping of the ATLAS central solenoid, *JINST* **888** (2008) to be submitted to *JINST*, ATLAS Note ATL-MAGNET-PUB-2007-001 (2007).
- [35] M. Bosman et al., Estimation of Radiation Background, Impact on Detectors, Activation and Shielding Optimization in ATLAS, ATLAS Note ATL-GEN-2005-001 (2005).
- 11075 [36] A. Romaniouk, TRT straw counting rate from photon background at LHC, ATLAS Note ATL-INDET-96-149 (1996).
- [37] S. Ahlen et al., Measurement of ATLAS MDT neutron sensitivity and development of a simulation model, ATLAS Note ATL-MUON-98-236 (1998).
- [38] R. Engel, J. Ranft and S. Roesler, Hard diffraction in hadron-hadron interactions and in photoproduction, *Phys. Rev.* **D52** (1995) 1459-1468.
- 11080 [39] R. Engel, Photoproduction within the two-component dual parton model: amplitudes and cross-sections, *Z. Phys.* **C66** (1995) 203-214.
- [40] R. Engel and J. Ranft, Hadronic photon-photon interactions at high energy, *Phys. Rev.* **D54** (1996) 4244-4262.
- [41] T. Sjostrand et al., Pythia 6.1, *Comp. Phys. Comm.* **135** (2001) 238.
- 11085 [42] C. Zeitnitz and T.A. Gabriel, The GEANT-CALOR Interface and Benchmark Calculations for Zeus Calorimeters, *Nucl. Inst. Meth.* **A349** (1994) 106.

- [43] T.A. Gabriel et al., CALOR89: A Monte Carlo Program Package for the Design and Analysis of Calorimeter Systems, Oak Ridge National Laboratory Note ORNL/TM-11185.
- 11090 [44] R. Brun and F. Carminanti, CALOR89: A Monte Carlo Program Package for the Design and Analysis of Calorimeter Systems, CERN Program Library Long Writeup W5013, September 1993.
- [45] A. Ferrari, P.R. Sala, A. Fasso and J. Ranft, FLUKA: A multi-particle transport code (Program version 2005), CERN report CERN-2005-010 (2005).
- [46] N.V. Mokhov, The MARS Code System Users's Guide, Fermilab-FN-628 (1995).
- 11095 [47] C. Cheplakov et al., Irradiation test of readout chain components of the ATLAS liquid argon calorimeters, ATLAS Note ATL-LARG-99-018 (1999).
- [48] C. Leroy, Irradiation tests of ATLAS liquid argon forward calorimeter (FCAL) electronics components, ATLAS Note ATL-LARG-02-003 (2002).
- [49] A. Amorim et al., Study of the effect of the radiation on the TILECAL Barrel Hadron calorimeter to be used in ATLAS, ATLAS Note ATL-TILECAL-93-021 (1993).
- 11100 [50] A. Karyukhin et al., Study of molded scintillator radiation hardness, ATLAS Note ATL-TILECAL-94-025 (1994).
- [51] M. David et al., Low dose rate effect in scintillating and WLS fibers by ionizing radiation, ATLAS Note ATL-TILECAL-96-078 (1996).
- [52] M. David et al., Radiation damage in WLS fibers, ATLAS Note ATL-TILECAL-98-151 (1998).
- 11105 [53] G. Lindstrom, Radiation damage in silicon detectors, Nucl. Inst. Meth. **A512** (2003) 30–43.
- [54] R. Ichimya, Radiation qualification of electronics components used for the ATLAS level-1 muon endcap trigger system, ATLAS Note ATL-DAQ-CONF-2005-003 (2005).
- [55] V. Cindro, et al., The ATLAS Beam Conditions Monitor, JINST **88** (2007) preprint JINST 004P 1207, submitted to JINST.
- 11110 [56] The CERN-RD42 Collaboration, CVD Diamond Radiation Detector Development, <http://rd42.web.cern.ch/RD42/> and CERN Note CERN-DRDC-94-21 (1994).
- [57] Element Six Ltd., King's Park Ride, Ascot, Berkshire SL5 8BP, UK.
- [58] H. Frais-Kolbl, E. Griesmayer, H. Kagan and H. Pernegger, A Fast Low-Noise Charged-Particle CVD Diamond Detector, IEEE Tr. Nucl. Sci **51** (2004) 3833.
- 11115 [59] M. Campbell et al., Measurement of Spectral Characteristics and Composition of Radiation in ATLAS by MEDIPIX2-USB Devices, CERN EDMS document 815615, <https://edms.cern.ch/document/815615>.
- 11120 [60] T. Holy et al., Proposal to measure spectral characteristics and composition of radiation in ATLAS by MEDIPIX2-USB Devices, CERN EDMS document 815615, <https://edms.cern.ch/document/815615>.
- [61] ATLAS Collaboration, Inner Detector: Technical Design Report, CERN/LHCC/97-016/017 (1997).
- [62] The ATLAS Collaboration, ATLAS Pixel Detector: Technical Design Report, CERN/LHCC/98-013 (1998).
- 11125 [63] P. A. Coe, D. F. Howell and R. B. Nickerson, Frequency scanning interferometry in ATLAS: remote, multiple, simultaneous and precise distance measurements in a hostile environment, Meas. Sci. Technol. (2004) 2175-2187.

- [64] M. Gilchriese (editor), Construction and Performance of the ATLAS Pixel Detector, JINST **888** (2008) to be submitted to JINST.
- 11130 [65] A. Ahmad et al., The Silicon Microstrip Sensors of the ATLAS SemiConductor Tracker, Nucl. Inst. Meth. **A578** (2007) 98–118.
- [66] M.S. Alam et al., The ATLAS silicon pixel sensors, Nucl. Inst. Meth. **A456** (2001) 217–232.
- [67] F. Campabadal et al., Beam tests of ATLAS SCT silicon strip detector modules, Nucl. Inst. Meth. **A538** (2005) 384–407.
- 11135 [68] A. Abdesselam et al., The Barrel Modules of the ATLAS SemiConductor Tracker, Nucl. Inst. Meth. **A568** (2006) 642–671.
- [69] A. Abdesselam et al., The ATLAS SemiConductor Tracker End-Cap Module, Nucl. Inst. Meth. **A575** (2007) 353–389.
- [70] A. Romaniouk and K. McFarlane (editors), TRT Drift Tubes, JINST **888** (2008) to be submitted to JINST.
- 11140 [71] T. Akesson et al., Operation of the ATLAS Transition Radiation Tracker under very high irradiation at the CERN LHC, Nucl. Inst. Meth. **A522** (2004) 25–32.
- [72] M. Capeans et al., Recent ageing studies for the ATLAS Transition Radiation Detectors, IEEE Tr. Nucl. Sci **51** (2004) 960–967.
- 11145 [73] G. Sprachmann et al., Ageing effects in the ATLAS Transition Radiation Tracker and gas filtration studies, IEEE Nucl. Sci. Symposium Conference Record **2** (2005) 1185–1190.
- [74] I. Gorelov et al., A Measurement of Lorentz Angle and Spatial Resolution of Radiation Hard Silicon Pixel Sensors, Nucl. Inst. Meth. **A481** (2002) 204–221.
- [75] F. Campabadal et al., Design and Performance of the ABCD3TA ASIC for readout of silicon strip detectors in the ATLAS semiconductor tracker, Nucl. Inst. Meth. **A552** (2005) 292–328.
- 11150 [76] R.J. Apsimon et al., Application of advanced thermal management technologies to the ATLAS SCT barrel module baseboards, Nucl. Inst. Meth. **A565** (2006) 561–571.
- [77] Y. Unno et al., Application of Cu-polyimide flex circuit and Al-on-glass pitch adapter for the ATLAS SCT barrel hybrid, Nucl. Inst. Meth. **A541** (2005) 286–294.
- [78] J. Grosse-Knetter, The ATLAS Pixel Detector, Nucl. Inst. Meth. **A568** (2006) 252–257.
- 11155 [79] I. Peric et al., The FEI3 Readout Chip for the ATLAS Pixel Detector, Nucl. Inst. Meth. **A565** (2006) 178–187.
- [80] R. Beccherle et al., MCC: The Module Controller of the ATLAS Pixel Detector, Nucl. Inst. Meth. **A492** (2002) 117–133.
- [81] H. Williams (editor), TRT Electronics, JINST **888** (2008) to be submitted to JINST.
- 11160 [82] T. Akesson et al., Implementation of the DTMROC-S ASIC for the ATLAS TRT Detector in a 0.25 μ m CMOS Technology, IEEE Nucl. Sci. Symposium Conference Record **1** (2002) 549–553.
- [83] A. Abdesselam et al., The Optical Links of the ATLAS SemiConductor Tracker, JINST **2** (2007) 09003.
- 11165 [84] M.L. Chu et al., The Off-Detector Opto-Electronics for the Optical Links of the ATLAS Semiconductor Tracker and Pixel Detector, Nucl. Inst. Meth. **A530** (2004) 293–310.

- [85] K.E. Arms et al., ATLAS Pixel Opto-Electronics, Nucl. Inst. Meth. **A554** (2005) 458–468.
- [86] P. Moreira et al., G-link and Gigabit Ethernet Compliant Serializer for LHC Data Transmission, IEEE Tr. Nucl. Sci Symposium conference record, Lyon (2000) 9.6-9.9.
- 11170 [87] A. Abdesselam et al., The Data Acquisition and Calibration System for the ATLAS Semiconductor Tracker, JINST **888** (2007) accepted by JINST.
- [88] H. Pernegger, Integration and Test of the ATLAS Semiconductor Tracker, Nucl. Inst. Meth. **A572** (2007) 108–112.
- [89] P. Lichard et al., Evolution of the TRT Backend and the new TRT-TTC Board, proceedings of the 2005 LECC, Heidelberg, CERN/LHCC-2005-038 (2005) 253-257.
- 11175 [90] P. Phillips, The ATLAS SCT Power Supply System, proceedings of TWEED, Prague (2007).
- [91] M. Olcese et al., Mechanics and Integration of the ATLAS Inner Detector, JINST **888** (2008) to be submitted to JINST.
- [92] A. Abdesselam et al., The Integration and Engineering of the ATLAS SemiConductor Tracker Barrel, editor: E. Perrin et al., JINST **888** (2008) to be submitted to JINST.
- 11180 [93] A. Andreazza et al., Effect of accidental beam losses on the ATLAS pixel detector, Nucl. Inst. Meth. **A565** (2006) 50–54.
- [94] R. Nickerson et al., Robotic Mounting of ATLAS Barrel SCT Modules, Nucl. Inst. Meth. **A568** (2006) 686–691.
- 11185 [95] A. Abdesselam et al., Engineering for the ATLAS SemiConductor Tracker (SCT) End-cap, JINST **888** (2008) submitted to JINST.
- [96] P. W. Phillips, Functional Testing of the ATLAS SCT Barrels, Nucl. Inst. Meth. **A570** (2007) 230–235.
- [97] A. Abdesselam et al., Combined SCT and TRT performance tests before installation, JINST **888** (2008) to be submitted to JINST.
- 11190 [98] S. M. Gibson et al., Coordinate Measurement in 2-D and 3-D Geometries Using Frequency Scanning Interferometry, Optics and Lasers in Engineering (2005) 813-815.
- [99] C. Wang and H. Ogen (editors), TRT Barrel, JINST **888** (2008) to be submitted to JINST.
- [100] C. Rembser (editor), TRT End-Cap, JINST **888** (2008) to be submitted to JINST.
- 11195 [101] M. Olcese et al., Evaporative Cooling System for the ATLAS Inner Detector, JINST **888** (2008) to be submitted to JINST.
- [102] C. Bayer et al., Development of Fluorocarbon Evaporative Cooling Recirculators and Controls for the ATLAS Inner Silicon Tracker, 2000 IEEE Nuclear Science Symposium. Conference Record, Vol. 2, Lyon, France, 15-20 Oct. 2000 (2000) 10.
- 11200 [103] V. Vacek et al., Perfluorocarbons and their use in cooling systems for semiconductor particle detectors, Fluid Phase Equilibria 174 (2000) 191-201.
- [104] A. Abdesselam et al., The Detector Control System of the ATLAS SemiConductor Tracker during Macro-Assembly and Integration, JINST **888** (2008) submitted to JINST.
- [105] ATLAS Collaboration, Liquid Argon Calorimeter Technical Design Report, CERN/LHCC/96-041 (1996).

- 11205 [106] ATLAS Collaboration, Tile Calorimeter Technical Design Report, CERN/LHCC/96-042 (1996).
- [107] D. Axen et al., Signal feedthroughs for the ATLAS barrel and endcap calorimeters, *Rev. Sci. Instrum.* **76** (2005) 063306.
- [108] B. Botchev et al., The High Voltage Feedthroughs for the ATLAS Liquid Argon Calorimeters, *JINST* **2** (2007) T10002.
- 11210 [109] B. Aubert et al., Development and construction of large size signal electrodes for the ATLAS electromagnetic calorimeter, *Nucl. Inst. Meth.* **A539** (2005) 558–594.
- [110] B. Aubert et al., Construction, assembly and tests of the ATLAS electromagnetic barrel calorimeter, *Nucl. Inst. Meth.* **A558** (2006) 388–418.
- [111] M.L. Andrieux et al., Construction and test of the first two sectors of the ATLAS barrel liquid argon presampler, *Nucl. Inst. Meth.* **A479** (2002) 316–333.
- 11215 [112] N.N. et al., Construction, Assembly and Tests of the ATLAS Electromagnetic End-Cap Calorimeter, *JINST* **888** (2008) to be submitted to JINST.
- [113] B. Aubert et al., Performance of the ATLAS electromagnetic calorimeter end-cap module 0, *Nucl. Inst. Meth.* **A500** (2003) 178–201.
- 11220 [114] B. Belhorma et al., Evaporating short-circuits in the ATLAS liquid argon barrel presampler 006, ATLAS Note CERN-ATL-LARG-PUB-2005-003.
- [115] J. Abdallah et al., Design, Construction and Installation of the ATLAS Hadronic Barrel Scintillator-Tile Calorimeter, ATL-TILECAL-PUB-2008-00xx, to be submitted to JINST.
- [116] J. Abdallah et al., The production and qualification of scintillator tiles for the ATLAS hadronic calorimeter, ATLAS Note ATL-TILECAL-PUB-2007-010 (2007), to be submitted to JINST.
- 11225 [117] J. Proudfoot et al., Master Plate Production for the Tile Calorimeter, Argonne National Laboratory Note ANL-TR-99-04 (1999).
- [118] P. Adragna et al., The ATLAS hadronic tile calorimeter: from construction toward physics, *IEEE Tr. Nucl. Sci* **53** (2006) 1275–81.
- 11230 [119] M.J. Varanda, The Tile Hadronic Calorimeter for the ATLAS Experiment, *IEEE Tr. Nucl. Sci* **48** (2001) 367.
- [120] F. Bosi et al., A device to characterize optical fibers, *Nucl. Inst. Meth.* **A485** (2002) 311–321.
- [121] J.G. Saraiva et al., The aluminization of 600-k WLS fibers for the TileCal/ATLAS/LHC, *IEEE Tr. Nucl. Sci* **51** (2004) 1235.
- 11235 [122] C. Cardeira et al., A Robot for Fibre Insertion in a Profile, ATLAS Note ATL-TILECAL-PUB-2005-007 (2005).
- [123] D.M. Gingrich et al., Construction, Assembly and Testing of the ATLAS Hadronic End-cap Calorimeter, *JINST* **2** (2007) P05005.
- [124] J. Ban et al., Cold electronics for the liquid argon hadronic end-cap calorimeter of ATLAS, *Nucl. Inst. Meth.* **A556** (2006) 158–168.
- 11240 [125] A. Artamonov et al., The ATLAS Forward Calorimeters, *JINST* **A888** (2007) JINST 015P 1107, submitted to JINST.
- [126] N.N. et al., ATLAS liquid argon calorimeter front end electronics system, *JINST* **888** (2008) to be submitted to JINST.

- 11245 [127] J. Ban et al., ATLAS liquid argon calorimeter back end electronics, *JINST* **2** (2007) P06002.
- [128] V. Radeka, Shielding and grounding in large detectors, Proceedings of the Fourth Workshop on Electronics for LHC Experiments, LEB 98, Roma (1998).
- [129] G. Blanchot, Grounding of the ATLAS Experiment, ATL-ELEC-PUB-2007-002 (2007).
- [130] J. Colas et al., Electronic calibration board for the ATLAS liquid argon calorimeter, *Nucl. Inst. Meth.* **A888** (2008) submitted to NIM.
- 11250 [131] N.J.Buchanan et al., Design and implementation of the front-end board for the readout of the ATLAS liquid argon calorimeters, *JINST* **888** (2008) submitted to JINST.
- [132] P. Adragna et al., A PMT-Block test bench, *Nucl. Inst. Meth.* **A564** (2006) 597–607.
- [133] M. Crouau et al., Characterization of 8-stage Hamamatsu R5900 Photomultiplier for the TILE Calorimeter, ATLAS Note TILECAL-NOTE 97-129 (1997).
- 11255 [134] K. Anderson et al., Design of the front-end analog electronics for the ATLAS Tile Calorimeter, *Nucl. Inst. Meth.* **A551** (2005) 469–476.
- [135] S. Berglund et al., The ATLAS Tile Calorimeter Digitizer, ATLAS Note ATL-TILECAL-PUB-2007-007 (2007), JINST 015P 1107, submitted to JINST.
- 11260 [136] K. Anderson et al., ATLAS Tile Calorimeter Interface, Proceeding of the 8th Workshop on Electronics for LHC Experiments, Colmar, France (2002) 269-273.
- [137] A.S. Cerqueira et al., Analog System for Building the First-Level Triggering Signal Provided by the Hadronic Calorimeter of ATLAS Detector, *Nucl. Inst. Meth.* **A570** (2007) 117–125.
- [138] E. Fullana et al., Digital Signal Reconstruction in the ATLAS Hadronic Tile Calorimeter, *IEEE Tr. Nucl. Sci* **53** (2006).
- 11265 [139] W.E. Cleland and E.G. Stern, Signal Processing Considerations for Liquid Ionization Calorimeters in a High Rate Environment, *Nucl. Inst. Meth.* **A338** (1994) 467.
- [140] V. González et al., Development of the Optical Multiplexer Board Prototype for Data Acquisition in the TileCal System, *IEEE Tr. Nucl. Sci* **53** (2006).
- 11270 [141] M. Aharrouche et al., Response Uniformity of the ATLAS liquid argon electromagnetic calorimeter, *Nucl. Inst. Meth.* **A582** (2007) 429–455.
- [142] J. Carvalho, Calibration and monitoring of the ATLAS Tile Calorimeter, Proceedings of the X11 International Conference on Calorimetry in High Energy Physics, Chicago (2006) 385.
- [143] E. Starchenko et al., Cesium Monitoring system for ATLAS hadronic calorimeter, *Nucl. Inst. Meth.* **A494** (2002) 508–523.
- 11275 [144] G. Blanchot et al., ATLAS Cooling System LCS v.2 Full Scale Test, ATLAS Project Document ATC-TL-ES-0001 (2001).
- [145] A.M. Henriques-Correia, A. Karioukhine, G. Schlager, Further performance tests of the cooling system for the calibration of the ATLAS Tilecal modules, ATLAS Note, ATL-TILECAL-2002-014 (2002).
- 11280 [146] J. Colas et al., Position resolution and particle identification with the ATLAS EM calorimeter, *Nucl. Inst. Meth.* **A550** (2005) 96–115.
- [147] M. Aharrouche et al., Energy linearity and resolution of the ATLAS electromagnetic barrel calorimeter in an electron test-beam, *Nucl. Inst. Meth.* **A568** (2006) 601–623.

- 11285 [148] M. Aharrouche et al., Characterization of the ATLAS electromagnetic liquid argon calorimeters with muons, *JINST* **888** (2008) to be submitted to JINST.
- [149] M. Cooke et al., In situ commissioning of the ATLAS electromagnetic calorimeter with cosmic muons, *ATL-LARG-PUB-2007-013* (2007).
- 11290 [150] M. Aharrouche, The ATLAS liquid argon calorimeter: Construction, Integration, Commissioning and Combined Test beam results, contribution to XIth Vienna Conference on Instrumentation, Vienna (2007).
- [151] N.N. et al., Electromagnetic Energy Measurement in LAr Calorimeter at the CTB 2004, *JINST* **888** (2008) to be submitted to JINST.
- 11295 [152] B. Dowler et al., Performance of the ATLAS hadronic end-cap calorimeter in beam tests, *Nucl. Inst. Meth.* **A482** (2002) 94–124.
- [153] A.E. Kiryunin et al., GEANT 4 physics evaluation with testbeam data of the ATLAS hadronic end-cap calorimeter, *Nucl. Inst. Meth.* **A560** (2006) 278–290.
- [154] C. Cojocaru et al., Hadronic calibration of the ATLAS liquid argon end-cap calorimeter in the pseudorapidity region $1.6 < |\eta| < 1.8$ in beam tests, *Nucl. Inst. Meth.* **A531** (2004) 481–514.
- 11300 [155] The ATLAS Liquid Argon EMEC/HEC Collaboration, Muon Results from the EMEC/HEC Combined Run corresponding to the ATLAS Pseudorapidity Region $1.6 < |\eta| < 1.8$, *ATLAS Note ATL-LARG-2004-006* (2004).
- [156] J. Pinfeld et al., ATLAS Liquid Argon EMEC/HEC/FCAL Collaboration, Hadronic Calibration of the ATLAS Liquid Argon End-Cap Calorimeter in the Pseudorapidity Region $2.5 < |\eta| < 4.0$ in Beam Tests, *Nucl. Inst. Meth.* **A888** (2008) to be submitted to NIM.
- 11305 [157] J.C. Armitage et al., Electron Signals in the Forward Calorimeter Prototype for ATLAS, *JINST* **2** (2007) P11001.
- [158] J.P. Archambault et al., Energy Calibration of the ATLAS Liquid Argon Forward Calorimeter, *JINST* **888** (2008) accepted by JINST.
- 11310 [159] M. Hurwitz, Module-to-module uniformity at 180 GeV in 2002-2003 tile calorimeter calibration test-beams, *ATLAS Note ATL-PUB-TILE-2006-008* (2006).
- [160] M. P. Casado and M. Cavalli-Sforza, H1-inspired analysis of the 1994 combined test of the Liquid Argon and Tilecal calorimeter prototypes, *ATLAS Note ATL-TILECAL-96-075* (1996).
- 11315 [161] A. Akhmadaliev et al., Response Uniformity of the ATLAS liquid argon electromagnetic calorimeter, *Nucl. Inst. Meth.* **A449** (2000) 461.
- [162] D. Pomarede, M. Virchaux, The Persint visualisation program for the ATLAS experiment, *Proceedings of 2003 Conference for Computing in High-Energy and Nuclear Physics (CHEP 03)*, La Jolla (2003) CHEP-2003-MOLT009.
- 11320 [163] S. Palestini, The muon spectrometer of the ATLAS experiment, *Nucl. Phys. Proc. Suppl.* **125** (2003) 337–345.
- [164] W. Riegler et al., Resolution limits of drift tubes, *Nucl. Inst. Meth.* **A443** (2000) 156–163.
- [165] M. Deile et al., Dependence of drift tube performance on the anode wire diameter, *Nucl. Inst. Meth.* **A449** (2000) 528–536.
- 11325 [166] W. Riegler et al., Front-end electronics for drift tubes in a high-rate environment, *Nucl. Inst. Meth.* **A446** (2000) 555–559.

- [167] C. Adorisio et al., Aging studies on Atlas muon spectrometer drift tubes, *IEEE Tr. Nucl. Sci* **52** (2005) 2971–2976.
- [168] C. Adorisio et al., Recent and ongoing ageing studies for the Atlas muon spectrometer drift tubes, *Nucl. Inst. Meth.* **A535** (2004) 186–190.
- 11330 [169] C. Adorisio et al., Ageing studies for the Atlas muon spectrometer drift tubes, *Nucl. Phys. Proc. Suppl.* **150** (2006) 168–171.
- [170] G. Avolio et al., Monitored Drift Tubes aging under intensive gamma irradiation, *Nucl. Inst. Meth.* **A568** (2006) 624–633.
- 11335 [171] S. Ahlen et al., ATLAS MDT neutron sensitivity measurement and modeling, *Nucl. Inst. Meth.* **A515** (2003) 446–457.
- [172] M. Aleksa et al., Rate effects in high-resolution drift chambers, *Nucl. Inst. Meth.* **A446** (2000) 435–443.
- [173] M. Deile et al., Performance of the ATLAS precision muon chambers under LHC operating conditions, *Nucl. Inst. Meth.* **A518** (2004) 65–68.
- 11340 [174] M. Deile et al., Resolution and efficiency of the ATLAS muon drift-tube chambers at high background rates, *Nucl. Inst. Meth.* **A535** (2004) 212–215.
- [175] J. Dubbert et al., Modelling of the space-to-drift-time relationship of the ATLAS monitored drift-tube chambers in the presence of magnetic fields, *Nucl. Inst. Meth.* **A572** (2007) 50–52.
- [176] D. Orestano et al., Calibration of the MDT chambers in the ATLAS experiment, *JINST* **888** (2008) to be submitted to JINST.
- 11345 [177] S. Kircher et al., Influence of aqueous vapour on the drift properties of MDT gases, ATLAS Note ATL-MUON-97-214, ATL-M-PN-214 (1997).
- [178] H. Kroha (editor), MDT construction paper, *JINST* **888** (2008) to be submitted to JINST.
- [179] J. Dubbert (editor), Barrel station integration, *JINST* **888** (2008) to be submitted to JINST.
- 11350 [180] J. Bensingler (editor), End-cap big wheel integration, *JINST* **888** (2008) to be submitted to JINST.
- [181] The MDT parameter book:
<http://atlas.web.cern.ch/Atlas/GROUPS/MUON/layout/MDT-parameter-book.html>.
- [182] F. Bauer et al., Construction and test of MDT chambers for the ATLAS muon spectrometer, *Nucl. Inst. Meth.* **A461** (2001) 17–20.
- 11355 [183] F. Bauer et al., The first precision drift tube chambers for the ATLAS muon spectrometer, *Nucl. Inst. Meth.* **A478** (2002) 153–157.
- [184] F. Bauer et al., Large-scale production of monitored drift tube chambers for the ATLAS muon spectrometer, *Nucl. Inst. Meth.* **A518** (2004) 69–72.
- [185] G. Avolio et al., Test of the first BIL tracking chamber for the ATLAS muon spectrometer, *Nucl. Inst. Meth.* **A523** (2004) 309–322.
- 11360 [186] P. Bagnaia et al., Construction of the inner layer barrel drift chambers of the ATLAS muon spectrometer at the LHC, *Nucl. Inst. Meth.* **A546** (2005) 481–497.
- [187] A. Baroncelli et al., Assembly and test of the BIL tracking chambers for the ATLAS muon spectrometer, *Nucl. Inst. Meth.* **A557** (2006) 421–435.

- 11365 [188] R. Avramidou et al., The accuracy of the ATLAS muon X-ray tomograph, Nucl. Inst. Meth. **A496** (2003) 83–101.
- [189] S. Schuh et al., A high precision X-ray tomograph for quality control of the ATLAS muon monitored drift tube chambers, prepared for 8th International Conference on Advanced Technology and Particle Physics (ICATPP 2003): Astroparticle, Particle, Space Physics, Detectors and Medical Physics Applications, Como (2003).
- 11370 [190] H. L. Groenstege, The RASNIK/CCD 3D alignment system, ATLAS Note ATL-MUON-94-063, ATL-M-PN-63 (2004).
- [191] J. Dubbert et al., Integration, Commissioning and Installation of Monitored Drift Tube Chambers for the ATLAS Barrel Muon Spectrometer, Nucl. Inst. Meth. **A572** (2007) 53–56.
- 11375 [192] Y. Arai et al., On-chamber readout system for the ATLAS MDT muon spectrometer, IEEE Tr. Nucl. Sci **51** (2004) 2196–2200.
- [193] C. Posch, S. Ahlen, E. Hazen, and J. Oliver, CMOS front-end for the MDT sub-detector in the ATLAS muon spectrometer - development and performance, Prepared for 7th Workshop on Electronics for LHC Experiments, Stockholm (2001).
- 11380 [194] Y. Arai, Development of front-end electronics and TDC LSI for the ATLAS MDT, Nucl. Inst. Meth. **A453** (2000) 365–371.
- [195] P. Bagnaia et al., Charge-dependent corrections to the time response of ATLAS muon chambers, Nucl. Inst. Meth. **A533** (2004) 344–352.
- [196] G. Brandenburg, J. Chapman et al., The readout electronics of the MDT chambers in the ATLAS experiment, JINST **888** (2008) to be submitted to JINST.
- 11385 [197] H. Boterenbrood, P. Jansweijer, G. Kieft, A. Konig, J. Vermeulen, T. Wijnen, The Read-Out Driver for the ATLAS MDT Muon Precision Chambers, IEEE Tr. Nucl. Sci **53** (2006) 741–748.
- [198] C. Guyot et al., The alignment of the barrel part of the ATLAS muon spectrometer, JINST **888** (2008) to be submitted to JINST.
- 11390 [199] J. R. Bensinger, ATLAS muon endcap alignment, Nucl. Phys. Proc. Suppl. **150** (2006) 136–139.
- [200] C. Amelung et al., Reference bars for the alignment of the ATLAS muon spectrometer, Nucl. Inst. Meth. **A555** (2005) 36–47.
- [201] H.v.d. Graaf et al., RASNIK Technical System Description for ATLAS, NIKHEF Note ETR 2000-04 (2000).
- 11395 [202] C. Adorisio et al., System Test of the ATLAS Muon Spectrometer in the H8 Beam at the CERN SPS, ATLAS Note ATL-MUON-PUB-2007-005 (2007).
- [203] C. Adorisio et al., Study of the ATLAS MDT Spectrometer using High Energy CERN Testbeam Data, ATLAS Note ATL-MUON-PUB-*****, ATL-COM-MUON-2007-013 (2007), to be submitted to NIM.
- 11400 [204] V. Polychronakos et al., Construction of the ATLAS CSC chambers, JINST **888** (2008) to be submitted to JINST.
- [205] J. Dailing et al., Performance and Radiation Tolerance of the ATLAS CSC On-Chamber Electronics, prepared for 6th Workshop on Electronic for LHC Experiments, Cracow (2000).
- 11405 [206] J. Dailing et al., Off-Detector Electronics for a High-Rate CSC Detector, prepared for 6th Workshop on Electronic for LHC Experiments, Cracow (2000).

- [207] G. Aielli et al., The RPC first level muon trigger in the barrel of the ATLAS experiment, Nucl. Phys. Proc. Suppl. **158** (2006) 11–15.
- [208] A. Aloisio et al., The trigger chambers of the ATLAS muon spectrometer: Production and tests, Nucl. Inst. Meth. **A535** (2004) 265–271.
- 11410 [209] ATLAS Collaboration, First-Level Trigger Technical Design Report, CERN/LHCC/98-14 (1998).
- [210] L. Levinson, O. Sasaki et al., The TGC trigger system, JINST **888** (2008) to be submitted to JINST.
- [211] P. Giusti (editor), RPC construction paper, JINST **888** (2008) to be submitted to JINST.
- [212] G. Aielli et al., SF-6 quenched gas mixtures for streamer mode operation of RPCs at very low voltages, Nucl. Inst. Meth. **A493** (2002) 137–145.
- 11415 [213] G. Aielli et al., Electrical conduction properties of phenolic-melaminic laminates, Nucl. Inst. Meth. **A533** (2004) 86–92.
- [214] G. Aielli et al., A high-voltage test for the ATLAS RPC qualification, Nucl. Inst. Meth. **A533** (2004) 199–202.
- [215] G. Aielli et al., Ageing test of the ATLAS RPCs at X5-GIF, Nucl. Inst. Meth. **A533** (2004) 98–101.
- 11420 [216] G. Aielli et al., Further advances in aging studies for RPCs, Nucl. Inst. Meth. **A515** (2003) 335–341.
- [217] G. Aielli et al., Test of ATLAS RPCs front-end electronics, Nucl. Inst. Meth. **A508** (2003) 189–193.
- [218] F. Giannini, E. Limiti, G. Orengo, and R. Cardarelli, An 8 channel GaAs IC front-end discriminator for RPC detectors, Nucl. Inst. Meth. **A432** (1999) 440–449.
- [219] G. Mikenberg et al., Construction of the ATLAS TGC chambers, JINST **888** (2008) to be submitted.
- 11425 [220] The TGC parameter book: <http://www.cern.ch/Atlas/project/TGC/www/design/tgc.xls>.
- [221] ATLAS Collaboration, ATLAS Forward Detectors for Measurement of Elastic Scattering and Luminosity Determination (ATLAS TDR, in preparation); ATLAS Forward Detectors for Luminosity Measurement and Monitoring, CERN-LHCC/04-10, LHCC I-014 (2004); Zero Degree Calorimeters for ATLAS, CERN-LHCC/07-01, LHCC I-016 (2007).
- 11430 [222] M.G. Albrow et al., FP420 R&D proposal, CERN-LHCC/05-025 (2005).
- [223] C. Royon, High mass diffractive physics at the LHC, Presentation at the DIS2007 conference, Munich (2007).
- [224] T. L. Cheng and P. Teixeira-Dias, Sensitivity of ATLAS to FCNC single top quark production, ATLAS Note ATL-PHYS-PUB-2006-029 (2006).
- 11435 [225] M. Dittmar, F. Pauss, D. Zurcher, Towards a precise parton luminosity determination at the CERN LHC, Phys. Rev. **D56** (1997) 7284.
- [226] A.G. Shamov and V.I. Telnov, Precision luminosity measurement at LHC using two photon production of $\mu^+ \mu^-$ pairs, Nucl. Inst. Meth. **A494** (2002) 51.
- [227] D. Acosta et al., The CDF Cherenkov luminosity monitor, Nucl. Inst. Meth. **A461** (2001) 540.
- 11440 [228] S. Agostinelli et al., GEANT4 - A Simulation Toolkit, Nucl. Inst. Meth. **A506** (2003) 250–303.
- [229] H. Stenzel, Luminosity calibration from elastic scattering, ATLAS Note ATL-LUM-PUB-2007-001 (2007).

- [230] U. Amaldi et al., The Real Part of the Forward Proton Proton Scattering Amplitude Measured at the CERN Intersecting Storage Rings, *Phys. Lett.* **B66** (1977) 390.
- 11445 [231] S. Ask et al., Luminosity measurement at ATLAS: Development, construction and test of scintillating fibre prototype detectors, *Nucl. Inst. Meth.* **A568** (2006) 588–600.
- [232] F. Anghinolfi et al., Hadron beam test of a scintillating fibre tracker system for elastic scattering and luminosity measurement in ATLAS, *JINST* **2** (2007) P07004.
- [233] T. Sjöstrand, S. Mrenna, P. Skands, *PYTHIA 6.4 Physics and Manual*, *JHEP* **0605** (2006) 026.
- 11450 [234] O. Adriani, et al., LHCf experiment: Technical Design Report, CERN-LHCC-2006-004 (2006).
- [235] Calorimeter Trigger Groups, The ATLAS Level-1 Calorimeter Trigger, *JINST* **888** (2008) to be submitted to JINST.
- [236] J. Garvey et al., Use of an FPGA to identify electromagnetic clusters and isolated hadrons in the ATLAS level-1 calorimeter trigger, *Nucl. Inst. Meth.* **A512** (2003) 506.
- 11455 [237] Barrel Muon Trigger Groups, The Barrel Muon Trigger system of the ATLAS First Level Trigger, *JINST* **888** (2008) to be submitted to JINST.
- [238] Endcap Muon Trigger Groups, The Endcap Muon Trigger system of the ATLAS First Level Trigger, *JINST* **888** (2008) to be submitted to JINST.
- [239] Central Trigger Groups, The Central Trigger system of the ATLAS First Level Trigger, *JINST* **888** (2008) to be submitted to JINST.
- 11460 [240] P. Borrego-Amaral et al., The ATLAS Local Trigger Processor, *IEEE Tr. Nucl. Sci* **52, 4** (2005) 1202.
- [241] B. G. Taylor, TTC Distribution for LHC Detectors, *IEEE Tr. Nucl. Sci* **45** (1998) 821.
- [242] H. von der Schmitt et al., A Configuration system for the ATLAS trigger, *JINST* **1** (2006) P05004.
- 11465 [243] H.C. van der Bij, R.A. McLaren, O. Boyle, G. Rubin, S-LINK, a data link interface specification for the LHC era, *IEEE Tr. Nucl. Sci* **44, 3** (1997) 398–402.
- [244] R. Cranfield et al., The ATLAS ROBIN, *JINST* **888** (2008) submitted to JINST.
- [245] R. Blair et al., The ATLAS High Level Trigger Region of Interest Builder, *JINST* **888** (2008) submitted to JINST.
- 11470 [246] H.P. Beck, Performance of the final Event Builder for the ATLAS Experiment, ATLAS Note ATL-DAQ-CONF-2007-013 (2007).
- [247] ATLAS Collaboration, High-Level Trigger, Data Acquisition and Controls Technical Design Report, CERN/LHCC/03-022 (2003).
- [248] A. Barriuso Poy et al., The Detector Control System of the ATLAS Experiment, *JINST* **888** (2008) to be submitted to JINST.
- 11475 [249] B. Hallgren, H. Boterenbrood, H.J. Burckhart, H. Kvedalen, The Embedded Local Monitor Board in the LHC Front-end I/O Control System, Proceedings 7th Workshop on Electronics for LHC Experiments, Stockholm (2001).
- [250] ETM Professional Control: <http://www.etm.at>.
- 11480 [251] W. Salter, The LHC Experiments' Joint Controls Project (JCOP), Proceedings of Computing in High Energy Physics, Beijing (2001).

- [252] K. Karam Galvão, K. Pommès and J. Molina-Perez, Management of equipment databases at CERN for the ATLAS experiment, Proceedings of the 10th ICATPP Conference, Villa Olmo, Como (2007).
- 11485 [253] C. Lasseur, ATLAS Placement Strategy and Survey, CERN EDMS document 821475, <https://edms.cern.ch/document/821475/1>.
- [254] C. Lasseur and D. Lissauer (editors), Synthesis of the survey measurements and positioning strategy for the components of the ATLAS detector, JINST **888** (2008) to be submitted to JINST.
- [255] F. Butin (editor), The assembly of the ATLAS experiment at CERN: infrastructure provided at CERN for the ATLAS experiment, JINST **888** (2008) to be submitted to JINST.
- 11490 [256] G. Blanchot, ATLAS EMC Policy, ATLAS Note ATL-ELEC-PUB-2007-003 (2007).
- [257] R. Veness et al., ATLAS beam vacuum system interfaces, CERN EDMS document LHC-VC1-ES-0002, <https://edms.cern.ch/document/442219/1.0>.
- [258] O. Beltramello and V. Hedberg, Radiological control of material from the experimental cavern, ATLAS EDMS document, ATC-TY-EY-00XX, to be submitted to EDMS and CDS.
- 11495 [259] O. Beltramello et al., The ATLAS Detector Safety System, JINST **888** (2008) to be submitted to JINST.
- [260] ATLAS Collaboration, ATLAS Detector and Physics Performance Technical Design Report, Vol. 1, CERN/LHCC/99-14 (1999).
- [261] D. Costanzo et al., ATLAS detector simulation: status and outlook, ATLAS Note, ATL-SOFT-PUB-2005-004 (2005).
- 11500 [262] T. Cornelissen et al., Concepts, Design and Implementation of the ATLAS New Tracking, ATLAS Note ATL-SOFT-PUB-2007-007 (2007).
- [263] P.F. Akesson et al., ATLAS Tracking Event Data Model, ATLAS Note ATL-SOFT-PUB-2006-004 (2006); T. G. Cornelissen et al., Updates of the ATLAS Tracking Event Data Model, ATLAS Note ATL-SOFT-PUB-2007-003 (2007).
- 11505 [264] A. Salzburger, S. Todorova and M. Wolter, The ATLAS Tracking Geometry Description, ATLAS Note ATL-SOFT-PUB-2007-004 (2007).
- [265] A. Salzburger, The ATLAS Track Extrapolation Package, ATLAS Note ATL-SOFT-PUB-2007-005 (2007).
- 11510 [266] V. Kartvelishvili, Electron bremsstrahlung recovery in ATLAS, Nucl. Phys. B (Proc. Suppl.) **172** (2007) 208–211.
- [267] R. Frühwirth, Track fitting with non-Gaussian noise, Comp. Phys. Comm. **100** (1997) 1.
- [268] R. Frühwirth and A. Strandlie, Track fitting with ambiguities and noise: a study of elastic tracking and nonlinear filters, Comp. Phys. Comm. **120** (1999) 197–214.
- 11515 [269] ATLAS Collaboration, xKalman and iPatRec, ATLAS Inner Detector Technical Design Report, CERN/LHCC/97-16 (1997) 37-40.
- [270] S. Gonzalez-Sevilla et al., Alignment of the pixel and SCT modules for the 2004 ATLAS combined test-beam, ATLAS Note ATL-INDET-PUB-2007-014 (2007).
- 11520 [271] P. Bruckman De Renstrom, A. Hicheur and S. Haywood, Global χ^2 approach to the alignment of the ATLAS silicon tracking detectors, ATLAS Note ATL-INDET-PUB-2005-002 (2005).

- [272] A. Bocci and W. Hulsbergen, TRT alignment for SR1 cosmics and beyond, ATLAS Note ATL-INDET-PUB-2007-009 (2007).
- [273] F. Heinemann, Track-based alignment of the ATLAS silicon detectors with the robust alignment algorithm, ATLAS Note ATL-INDET-PUB-2007-011 (2007).
- 11525 [274] F. Bauer et al., ATLAS 2004 Combined Test Beam results: Muon Chamber Alignment and Muon Reconstruction, ATLAS Note ATL-MUON-PUB-2007-003.
- [275] S. Hassani et al., A muon identification and combined reconstruction procedure for the ATLAS detector at the LHC using the MUONBOY, STACO and MuTag reconstruction packages, Nucl. Inst. Meth. **A572** (2007) 77.
- 11530 [276] ATLAS Collaboration, Muon Spectrometer Technical Design Report, CERN/LHCC/97-22 (1997), 296-297.
- [277] T. Bold et al., Pile-up studies for soft electron identification and b-tagging with DC1 data, ATLAS Note ATL-PHYS-PUB-2006-001.
- 11535 [278] F. Djama, Using $Z^0 \rightarrow e^+e^-$ for Electromagnetic Calorimeter Calibration, ATLAS Note ATL-LARG-2004-008 (2004).
- [279] F. Hubaut and P. Pralavorio, Commissioning of the ATLAS electromagnetic calorimeter with minimum-bias events, ATLAS Note ATL-LARG-PUB-2007-007 (2007).
- [280] G. Blazey et al., Run II Jet Physics: Proceedings of the Run II QCD and Weak Boson Physics Workshop, arXiv:hep-ex/0005012v2 (2000).
- 11540 [281] M. Cacciari and G. Salam, Dispelling the N^{*3} myth for the $k(t)$ jet-finder, Phys. Lett. **B641** (2006) 57-61.
- [282] B. Andrieu et al., Results from pion calibration runs for the H1 liquid argon calorimeter and comparisons with simulations, Nucl. Inst. Meth. **A336** (1993) 499-509.
- 11545 [283] I. Abt et al. (H1 Collaboration), The tracking, calorimeter and muon detectors of the H1 experiment at HERA, Nucl. Inst. Meth. **A386** (1997) 348.
- [284] R. K. Ellis et al., Higgs Decay to $\tau^+ \tau^-$: A Possible Signature of Intermediate Mass Higgs Bosons at the SSC, Nucl. Phys. **B297** (1988) 221.
- [285] L. DiLella, Proceedings of the Large Collider Workshop, Aachen (1990), CERN 90-10/ECFA 90-133, Vol. II, 530.
- 11550 [286] E. Richter-Was and T. Szymocha, Hadronic tau identification with track based approach: the $Z \rightarrow \tau\tau, W \rightarrow \tau\nu$ and dijet events from DC1 data samples, ATLAS Note ATL-PHYS-PUB-2005-005 (2005).
- [287] M. Heldmann and D. Cavalli, An improved tau-Identification for the ATLAS experiment, ATLAS Note ATL-PHYS-PUB-2006-008 (2006).
- 11555 [288] S. Corréard et al., b-tagging with DC1 data, ATLAS Note ATL-PHYS-2004-006.