

The ATLAS TRT electronics

This article has been downloaded from IOPscience. Please scroll down to see the full text article.

2008 J. Inst. 3 P06007

(<http://iopscience.iop.org/1748-0221/3/06/P06007>)

The Table of Contents and more related content is available

Download details:

IP Address: 152.3.153.219

The article was downloaded on 24/07/2008 at 20:30

Please note that terms and conditions apply.

The ATLAS TRT electronics

The ATLAS TRT collaboration

E. Abat,^{a,*} T.N. Addy,^j T.P.A. Åkesson,^m J. Alison,^u F. Anghinolfi,^c E. Arik,^{a,*}
M. Arik,^a G. Atoian,^z B. Auerbach,^z O.K. Baker,^z E. Banas,^f S. Baron,^c C. Bault,^c
N. Becerici,^a A. Beddall,^{a,1} A.J. Beddall,^{a,1} J. Bendotti,^c D.P. Benjamin,^g
H. Bertelsen,^d A. Bingul,^{a,1}, H. Blampey,^c A. Bocci,^g M. Bochenek,^e
V.G. Bondarenko,^p V. Bychkov,^l J. Callahan,^k M. Capeáns Garrido,^c L. Cardiel Sas,^c
A. Catinaccio,^c S.A. Cetin,^{a,2} T. Chandler,^z R. Chritin,^h P. Cwetanski,^k M. Dam,^d
H. Danielsson,^c E. Danilevich,^v E. David,^c J. Degenhardt,^u B. Di Girolamo,^c
F. Dittus,^c N. Dixon,^c O.B. Dogan,^{a,*} B.A. Dolgoshein,^p N. Dressnandt,^u C. Driouchi,^d
W.L. Ebenstein,^g P. Eerola,^m U. Egede,^m K. Egorov,^k H. Evans,^k P. Farthouat,^c
O.L. Fedin,^v A.J. Fowler,^g S. Fratina,^u D. Froidevaux,^c A. Fry,^j P. Gagnon,^k
I.L. Gavrilenko,^o C. Gay,^y N. Ghodbane,^r J. Godlewski,^c M. Goulette,^c I. Gousakov,^l
N. Grigalashvili,^l Y. Grishkevich,^q J. Grognuz,^c Z. Hajduk,^f M. Hance,^u F. Hansen,^d
J.B. Hansen,^d P.H. Hansen,^d G.A. Hare,^u A. Harvey Jr.,^j C. Hauviller,^c A. High,^u
W. Hulsbergen,^c W. Huta,^c V. Issakov,^z S. Istin,^a V. Jain,^k G. Jarlskog,^m L. Jeanty,^y
V.A. Kantserov,^p B. Kaplan,^z A.S. Kapliy,^u S. Katounin,^v F. Kayumov,^o P.T. Keener,^u
G.D. Kekelidze,^l E. Khabarova,^l A. Khristachev,^v B. Kisielewski,^f T.H. Kittelmann,^w
C. Kline,^k E.B. Klinkby,^d N.V. Klopov,^v B.R. Ko,^g T. Koffas,^c N.V. Kondratieva,^p
S.P. Konovalov,^o S. Koperny,^e H. Korsmo,^m S. Kovalenko,^v T.Z. Kowalski,^e
K. Krüger,^c V. Kramarenko,^q L.G. Kudin,^v A-C. Le Bihan,^c B.C. LeGeyt,^u K. Levterov,^l
P. Lichard,^c A. Lindahl,^d V. Lisan,^l S. Lobastov,^l A. Loginov,^z C.W. Loh,^y S. Lokwitz,^z
M.C. Long,^j S. Lucas,^c A. Lucotte,ⁱ F. Luehring,^k B. Lundberg,^m R. Mackeprang,^d
V.P. Maleev,^v A. Manara,^k M. Mandl,^c A.J. Martin,^z F.F. Martin,^u R. Mashinistov,^p
G.M. Mayers,^u K.W. McFarlane,^j V. Mialkovski,^l B.M. Mills,^y B. Mindur,^e V.A. Mitsou,^x
J.U. Mjörnmark,^m S.V. Morozov,^p E. Morris,^k S.V. Mouraviev,^o A.M. Muir,^y A. Munar,^u
A.V. Nadtochi,^v S.Y. Nesterov,^v F.M. Newcomer,^u N. Nikitin,^q O. Novgorodova,^o
E.G. Novodvorski,^v H. Ogren,^k S.H. Oh,^g S.B. Oleshko,^v D. Olivito,^u J. Olszowska,^f
W. Ostrowicz,^f M.S. Passmore,^c S. Patrichev,^v J. Penwell,^k F. Perez-Gomez,^c
V.D. Peshekhonov,^l T.C. Petersen,^c R. Petti,^b A. Placci,^c A. Poblaguev,^z X. Pons,^c
M.J. Price,^c O. Røhne,^t R.D. Reece,^u M.B. Reilly,^u C. Rembser,^c A. Romaniouk,^p
D. Rousseau,^s D. Rust,^k Y.F. Ryabov,^v V. Ryjov,^c M. Söderberg,^m A. Savenkov,^l
J. Saxon,^u M. Scandurra,^k V.A. Schegelsky,^v M.I. Scherzer,^u M.P. Schmidt,^{z,*}
C. Schmitt,^c E. Sedykh,^v D.M. Seliverstov,^v T. Shin,^j A. Shmeleva,^o S. Sivoklov,^q

**S.Yu. Smirnov,^p L. Smirnova,^q O. Smirnova,^m P. Smith,^k V.V. Sosnovtsev,^p
 G. Sprachmann,^c S. Subramania,^k S.I. Suchkov,^p V.V. Sulin,^o R.R. Szczygiel,^f
 G. Tartarelli,ⁿ E. Thomson,^u V.O. Tikhomirov,^o P. Tipton,^z J.A. Valls Ferrer,^x
 R. Van Berg,^{u*} V.I. Vassilakopoulos,^j L. Vassilieva,^o P. Wagner,^u R. Wall,^z C. Wang,^g
 D. Whittington,^k H.H. Williams,^u A. Zhelezko^p and K. Zhukov^o**

^a*Faculty of Sciences, Department of Physics, Bogazici University,
 TR - 80815 Bebek-Istanbul, Turkey*

^b*Brookhaven National Laboratory, Physics Department,
 Bldg. 510A, Upton, NY 11973, United States of America*

^c*CERN, CH - 1211 Geneva 23, Switzerland, Switzerland*

^d*Niels Bohr Institute, University of Copenhagen,
 Blegdamsvej 17, DK - 2100 Kobenhavn 0, Denmark*

^e*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*

^f*The Henryk Niewodniczanski Institute of Nuclear Physics, Polish Academy of Sciences, ul.
 Radzikowskiego 152, PL - 31342 Krakow, Poland*

^g*Duke University, Department of Physics, Durham, NC 27708, United States of America*

^h*Universite de Geneve, Section de Physique,
 24 rue Ernest Ansermet, CH - 1211 Geneve 4, Switzerland*

ⁱ*Laboratoire de Physique Subatomique et de Cosmologie, CNRS-IN2P3, Universite Joseph
 Fourier, INPG, 53 avenue des Martyrs, FR - 38026 Grenoble Cedex, France*

^j*Hampton University, Department of Physics, Hampton, VA 23668, United States of America*

^k*Indiana University, Department of Physics, Swain Hall West, Room 117,
 727 East Third St., Bloomington, IN 47405-7105, United States of America*

^l*Joint Institute for Nuclear Research, JINR Dubna, RU - 141 980 Moscow Region, Russia*

^m*Lunds Universitet, Fysiska Institutionen, Box 118, SE - 221 00 Lund, Sweden*

ⁿ*INFN Milano and Università di Milano, Dipartimento di Fisica,
 via Celoria 16, IT - 20133 Milano, Italy*

^o*P.N. Lebedev Institute of Physics, Academy of Sciences,
 Leninsky pr. 53, RU - 117 924 Moscow, Russia*

^p*Moscow Engineering & Physics Institute (MEPhI),
 Kashirskoe Shosse 31, RU - 115409 Moscow, Russia*

^q*Lomonosov Moscow State University, Skobeltsyn Institute of Nuclear Physics, RU - 119 992
 Moscow Lenskie gory I, Russia*

^r*Max Planck Institut fuer Physik, Postfach 401212,
 Foehringer Ring 6, DE - 80805 Muenchen, Germany*

^s*LAL, Univ. Paris-Sud, IN2P3/CNRS, Orsay, France*

^t*Department of Physics, University of Oslo, Blindern, NO - 0316 Oslo 3, Norway*

^u*University of Pennsylvania, Department of Physics & Astronomy,
 209 S. 33rd Street, Philadelphia, PA 19104, United States of America*

^v*Petersburg Nuclear Physics Institute, RU - 188 300 Gatchina, Russia*

^w*University of Pittsburgh, Department of Physics and Astronomy, 3941 O'Hara Street, Pittsburgh,
 PA 15260, United States of America*

^x*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*

^y*University of British Columbia, Dept of Physics,
6224 Agriculture Road, CA - Vancouver, B.C. V6T 1Z1, Canada*

^z*Yale University, Department of Physics,
PO Box 208121, New Haven CT, 06520-8121, United States of America*

¹*Currently at Gaziantep University, Turkey*

²*Currently at Dogus University, Istanbul*

**Deceased*

E-mail: rick@hep.upenn.edu

ABSTRACT: The ATLAS inner detector consists of three sub-systems: the pixel detector spanning the radius range 4cm-20cm, the semiconductor tracker at radii from 30 to 52 cm, and the transition radiation tracker (TRT), tracking from 56 to 107 cm. The TRT provides a combination of continuous tracking with many projective measurements based on individual drift tubes (or straws) and of electron identification based on transition radiation from fibres or foils interleaved between the straws themselves. This paper describes the on and off detector electronics for the TRT as well as the TRT portion of the data acquisition (DAQ) system.

KEYWORDS: VLSI circuits; Analogue electronic circuits; Detector control systems (detector and experiment monitoring and slow-control systems, architecture, hardware, algorithms, databases); Digital electronic circuits.

*Corresponding Author.

Contents

I	Front End — On detector	7
1.	ASDBLR	7
1.1	TRT signal processing	7
1.2	Design goals	7
1.3	Description of the ASDBLR	8
1.3.1	Dual preamp	9
1.3.2	Shaping stages	10
1.3.3	Baseline restorer	10
1.3.4	High and low level discriminator	11
1.3.5	Ternary driver	11
1.4	Wafer fabrication	12
1.5	ASDBLR packaging	12
1.6	ASDBLR testing	12
1.7	Radiation hardness	14
2.	DTMROC	15
2.1	Analog blocks	18
2.1.1	Ternary receiver	18
2.1.2	LVDS interface	18
2.1.3	Digital to analog converters	19
2.1.4	Testpulse generation	19
2.1.5	Temperature and voltage sense	19
2.2	Timing	20
2.2.1	DLL	20
2.2.2	Latch	20
2.3	Digital blocks	20
2.3.1	Pipeline	20
2.3.2	Derandomizer	21
2.3.3	Serializer	21
2.3.4	Command decoder	21
2.3.5	Error handling	22
2.4	Measured performance	22
3.	Board designs	24
3.1	End cap boards	24
3.1.1	ASDBLR boards	24
3.1.2	DTMROC triplet boards	25

3.1.3	Shielding	25
3.1.4	Power distribution	25
3.2	Barrel boards	26
3.2.1	Design constraints	27
3.2.2	Design rules	28
3.2.3	Barrel printed circuit boards	29
3.2.4	Power distribution	30
3.3	Testing program	31
4.	Cabling and cooling	31
4.1	Cabling	31
4.1.1	Detector to Patch Panels	31
4.1.2	Patch Panels to control room	33
4.2	Cooling	34
4.2.1	Power dissipation — Front End	34
4.2.2	Power dissipation — Patch Panels	34
4.2.3	Design and deployment	34
II	Patch Panels	37
5.	TRT-TTC Patch Panels	38
5.1	I/O compensation and filtering	38
5.2	Clock fine delay	39
5.3	Temperature read-out	41
6.	TRT-ROD Patch Panel	41
6.1	Inputs and outputs	41
6.2	Concatenation and serialization	42
7.	Interlocks, controls and DCS	43
7.1	Overview of architecture of the DCS	43
7.1.1	Components	43
7.2	Tools and methods	44
7.2.1	PVSS	44
7.2.2	Finite State Machine (FSM) methodology	46
7.3	Detector safety	47
7.3.1	DCS protections	47
7.3.2	TRT hardware interlock	48
III	Back End electronics, DAQ and power	51

8. Timing, Trigger, Control (TTC)	52
8.1 TRT-TTC segmentation	53
8.2 Modes of operation	53
8.3 Interfaces	54
9. Readout drivers (ROD)	56
9.1 ROD-TTC communication	56
9.2 ROD-Patch Panel communication	57
9.3 Data flow	59
9.4 Compression	60
9.5 Self-test features	61
10. Data acquisition	61
10.1 ATLAS trigger	61
10.2 Data acquisition software framework	61
10.2.1 ATLAS DAQ framework	61
10.3 TRT software libraries and databases	62
10.4 Readout scheme	62
10.4.1 Physics data	62
10.4.2 Calibration data	63
10.4.3 Monitoring	63
10.4.4 ATLAS tools and the TRT viewer	63
10.4.5 Voltages and temperatures	64
11. High voltage power	64
11.1 HVSys power supply	65
11.2 Implementation	65
11.3 Control	67
12. Low voltage power	68
12.1 Requirements	68
12.2 WIENER MARATON supplies	69
12.3 Overall efficiency	70
12.4 Control	70
12.4.1 Bulk supplies	70
12.4.2 Patch Panels	71
12.4.3 Control and monitoring	72
IV Installation, integration and performance	74
13. Installation, integration and testing	74
13.1 Front End board initial tests and preparation	74
13.2 Front End board installation and tests	74

13.3 TRT stand-alone noise and system tests	75
13.4 Combined noise tests with SCT	76
13.5 Cosmic ray runs with SCT and ATLAS	78

Introduction

The Transition Radiation Tracker (TRT) is composed of over 250,000 straws, ranging in length from 50cm (endcap) to 150cm (barrel). The analog readout is performed by a custom-designed Application Specific Integrated Circuit (ASIC) which performs Amplification, Shaping, Discrimination, and Base-Line Restoration (ASDBLR — section 1). These analog results are sampled by a second ASIC, the Drift Time Measuring Read Out Chip (DTMROC — section 2), which makes the time measurement of the signals and provides a digitized result to off-detector electronics for up to 16 straw channels.

Control of the DTMROCs is supplied by the TRT specific Timing, Trigger, and Control module (TRT-TTC — section 8) each of which manages up to 480 DTMROCs. The readout of the DTMROCs is performed by the TRT specific Read Out Drivers (TRT ROD — section 9). Both the TRT TTCs and the RODs are single-width 9U VME64x modules, the entire collection of which fill 10 VME crates for the whole detector (96 RODs and 48 TTCs). Each VME crate is controlled by a single board computer running a customized linux operating system.

The ATLAS TRT electronics is designed to satisfy a number of challenging requirements:

- Input signal sensitivity from about 2 to 100 femto Coulombs
- High and low threshold for transition radiation detection and tracking respectively
- Input signal rate per channel up to 20 MHz
- Time measurement accuracy to 1ns rms
- More than 350,000 input signals
- Level 1 buffering to $> 4 \mu\text{s}$
- Level 1 trigger rate up to 100 kHz
- Radiation levels $> 10 \text{ Mrad}$ and 10^{14} n/cm^2
- Tight on detector spatial constraints
- Stringent detector radiation length goals
- At source cooling - no net heat load to the ATLAS cavern
- Very high reliability (expectation of no access to detector for extended periods)
- Cable length from detector to counting room $> 80 \text{ m}$
- At maximum Level 1 Accept rates, the data rate out of the detector is over 100 GB/s
- Full compliance with the ATLAS Trigger DAQ Back End standards




Figure 1. Overview of the TRT electronics from straws to Back End hardware in USA15.

The solutions adopted by the collaboration and detailed in this document involve separating the electronics system into three geographically distinct parts - on-detector or *Front End*, intermediate *Patch Panel* and counting room *Back End* regions. The Front End electronics use full TRT custom radiation hard integrated circuits (the ASDBLR and DTMROC below). The Patch Panel boards located in the midst of the ATLAS muon system use radiation tolerant commercial parts plus several CERN designed custom chips, and the Back End electronics in the USA15 counting room are designed entirely with commercial components except for a few specialized LHC custom timing circuits. Power for the Front End and Patch Panel areas relies on LHC custom analog regulators supplied from commercial bulk voltage supplies especially designed for LHC conditions. HV power for the straw detectors is provided by semi-custom crates of precision low current HV supplies with monitors and adjustable trips. The basic TRT electronics blocks are shown in figure 1.

The basic operational mode of the TRT is the detection of avalanche currents on the anode wire initiated by clusters of primary ionization electrons from a through-going track [1]. The arrival time of the cluster depends upon the distance from the wire of the primary ionization. This cluster arrival time relative to the time of a track from a collision coupled with knowledge of the drift velocity in the gas is what allows the TRT to make a precise measurement of track position. Tracks passing near the anode produce avalanche current with a leading edge at or near the track time and a trailing edge at a time corresponding to a 2 mm drift. For avalanche currents above threshold, the ASDBLR sends a pulse to the DTMROC — the leading edge of the pulse corresponds to the distance of closest approach and the trailing edge is, roughly, fixed to the 2 mm maximum drift time. This fixed maximum drift time is useful as a flag that a given hit is, in fact, associated with

the beam crossing of interest. The DTMROC divides each 25 ns LHC bunch period into 8 equal time bins and records TRT hits to that precision — 3.12 ns (25/8). The DTMROC stores this data for up to 6 μ s until a Level 1 Accept is received. On the receipt of a Level 1 Accept from the central ATLAS trigger, the DTMROC data (24 bits of tracking and 3 bits of high threshold hit information) for 3 LHC crossing times (75 ns) is shipped to the Back End RODs and assembled into events for Level 2 and later processing. The choice of sending three crossing times' worth of data is governed by the maximum drift time in the magnetic field of about 60 ns.

Signal transmission from the Front End to the Patch Panels is, unlike most other ATLAS sub-detectors, implemented entirely with LVDS¹ signals running on 36 AWG² shielded twisted pair. This solution was chosen because of the distributed nature of the data sources (each DTMROC chip produces a 40 Mbit/sec data stream at the full L1Accept rate and the almost 22,000 DTMROCs are spread over many square meters of detector surface). In addition, the cost and power penalties associated with many low speed radiation tolerant optical links are avoided. The TRT output data is converted to optical fiber only at the Patch Panels where 30 DTMROC outputs can be merged to form a single 1.2 Gb/s fiber output using commercial optical devices located in a region that can be accessed relatively easily. This multiplexing of signals onto fiber at the Patch Panels also saves the very considerable monetary cost of running over 20,000 shielded twisted pairs nearly 100 meters while avoiding the costs of Front End mounted optical elements.

The Back End design, a TRT specific ROD and TTC plus a custom P3 backplane board set is similar to other ATLAS detectors but implements many TRT specific functions for control and data acquisition.

The TRT Detector Control System (DCS) that controls and monitors the TRT and its parts and the TRT Data Acquisition (DAQ) system that sets up and then controls the movement of data from the Front End into the larger ATLAS DAQ system are also discussed.

¹Low Voltage Differential Signaling

²American Wire Gauge

Part I

Front End — On detector

1. ASDBLR

1.1 TRT signal processing


The analog signal processing for the ATLAS TRT requires detection of avalanche signals from a few primary drift electrons in the TRT straw sensors with nanosecond timing precision at counting rates as high as 18 MHz in the presence of a high radiation background. To minimize noise and pickup it was decided to mount the readout electronics on the detector as close to the anode wires as possible. Direct connection to the wire anodes was made possible by choosing to operate the straw cathode at a high negative potential. The closely packed 4mm diameter straws represent a high density challenge for on-detector electronics. This challenge was addressed by the design of a custom analog Application Specific Integrated Circuit (ASIC), the ASDBLR, in the BiCMOS DMILL radiation hard process. The design primarily exploits NPN bipolar transistors for their intrinsically low noise, high current gain and excellent matching. This ASIC provides the complete analog signal processing chain for eight straws. It amplifies and shapes the straw signals eliminating the long ion tail and provides baseline restoration prior to presenting the signal to a dual comparator section. The output of the ASDBLR is a three level differential (ternary) current signal delivered to a custom receiver on its companion chip, the DTMROC. In the DTMROC the ASDBLR comparator output pulse width is recorded in 3.12 ns bins and stored in a pipeline. The 16 channel DTMROC provides timing, storage and control for two ASDBLR ASICS. More complete descriptions of the ASDBLR can be found in [2].

1.2 Design goals

The average electron drift velocity in the TRT straw gas is $\sim 50 \mu\text{m}/\text{ns}$. Ignoring the curvature of the primary path this velocity can be used to set a limit on the timing required to satisfy the position resolution specification. The electronics goal of 1ns timing resolution is then small compared with the position resolution goal of less than $150\mu\text{m}$. The wide dynamic range of straw track signals, up to 100 times threshold, high occupancy and the 25 ns beam crossing rate make this goal a challenging one. Electrons and gas atoms ionized in the avalanche process near the wire induce a signal current as they drift toward the anode and the cathode respectively. The time development of this current, the ion tail, is primarily determined by the type of gas and the anode wire diameter. The ASDBLR employs a traditional fixed time shaping technique to remove this unwanted, predictable, signal by building a mirror image impulse response into the signal processing electronics so that the ion tail and mirror signal cancel after the initial avalanche signal. After this cancellation process only a small fraction of the total avalanche signal is available. In our case about 5% or 0.15 fC per primary electron at a gas gain of 2.5×10^4 . Since our objective is to detect the earliest clusters of electrons arriving at the wire, the electronics must add as little noise as possible to the incoming signal. After careful study and several design iterations, an electronics peaking time of 7.5 ns with a semi-gaussian shape after ion tail cancellation was chosen. This shaping function allows integration of the avalanche signal from the first few clusters of drift electrons arriving at the wire

Table 1. Design goals for the TRT Front End electronics.

Peaking Time for Track Detection	7.5 ns
Peaking time for TR photon detection	10 ns
Double Pulse Resolution	< 25 ns
Intrinsic Electronics Noise	< 0.3 fC rms
Operational Track Disc. Threshold	2 fC
Maximum Threshold for TR photons	120 fC
10 year Neutron exposure	$3 \times 10^{14}/\text{cm}^2$
Total Dose (10 year) Ionizing Radiation	5 Mrad

**Figure 2.** Block Diagram of one channel of the ASDBLR

to give a trigger point based on a signal significantly higher than the intrinsic and systematic noise. Transition Radiation (TR) photons created by energetic electrons transiting layers of polypropylene radiator placed between the straws are stopped primarily by xenon atoms in the gas mixture. This high Z inert gas allows efficient absorption of TR photons whose energy is in the range of 5-15 keV, well above the typical 2 keV deposited by a minimum ionizing track. Extending the peaking time of the TR photon detection circuit to 10 ns allows integration of the direct and reflected signal from the far end of the unterminated 0.5–0.7 m straw. The extended peaking time reduces the variation in amplitude versus position of the avalanche along the straw [3]. The basic design goals for the ATLAS TRT Front End electronics are summarized in table 1.

1.3 Description of the ASDBLR

The ASDBLR consumes about 40 mW per channel and provides the complete analog signal processing chain for eight straw tubes from straw anode input to a three level digital output that signals arrival of the earliest ions at the wire with one level and the presence of a signal compatible with a transition radiation photon for the other level.

Figure 2 shows the signal processing blocks: Dual Preamp, Shaper, Baseline Restorer, High and Low level discriminator and driver. This largely differential design depends on good device

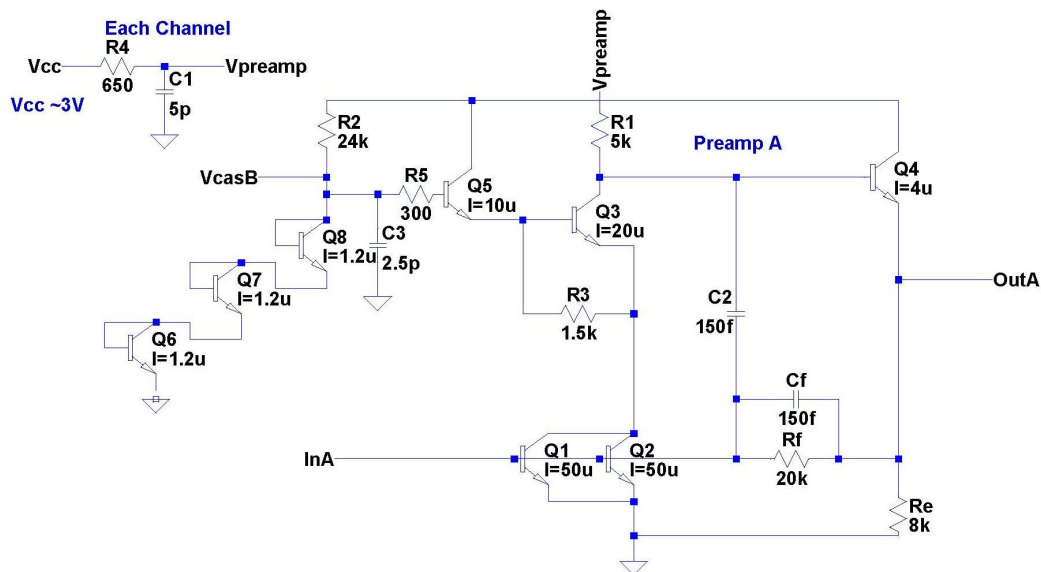


Figure 3. Schematic of one of the two matched preamps used per channel.

matching rather than the absolute values of circuit elements in this DMILL process. The differential nature of the design also has the benefit of providing a rate independent power consumption.

1.3.1 Dual preamp

The dual cascoded common emitter preamp shown in figure 3 is optimized for low-noise and radiation resistance. The duplicate circuits are interleaved in the layout to equalize topological effects and provide the best DC-balanced output to the differential shaper stage. The common emitter inputs are self-biasing at approximately +750mV. Since the preamp input is directly coupled to the wire, the common emitter input sets the anode potential. Anode current from the avalanche process flows directly into the preamp input. At high rate operation the average straw anode current is as high as $10 \mu\text{A}$ but does not significantly affect the operation of the circuit. The gain of the preamp is 1.5 mV/fC with a peaking time of 1.5 ns for an impulse input. Both preamp inputs for each channel are bonded out to package leads although only one input is attached to the straw wire anode. The second preamp input offers an important benefit in helping control channel-to-channel pickup external to the chip and reduces sensitivity to common-mode signals.

Significant attention was given to the protection of the preamp input since the input is directly attached to the straw anode wire. A breakdown or spark discharge in the straw sensor shorts the anode with the cathode which is held at a large negative voltage. Such a breakdown discharges significant current into the preamp input. To protect the input a fast response input protection diode was designed into the ASDBLR using the normally reverse biased collector-base junction of an array of large, single stripe, NPN transistors. The protection device on the preamp input that is attached to the wire is significantly larger than that on the other, typically unused, input. The “unused” input can be (and is in the board designs discussed in section 3) attached to a board level trace to help balance the common mode charge injection in the interconnect between the straw and the preamp. A total of $380 \mu\text{m}$ of emitter length provides 0.5 mJ breakdown protection when a

series 24 Ohm resistor is used between this input and the wire. This external resistor contributes 8% of the total intrinsic noise at the track comparator input. Nevertheless, this on-chip breakdown protection is not entirely sufficient so an external diode (BAW56) is mounted on the Front End boards in parallel with the ASDBLR protection diode to fully guarantee the inputs.

In low-noise amplifier designs, the input transistor is the most significant source of additional noise from thermal noise due to unavoidable resistance in the base and from the statistics of transistor current. Usually the technology dependent base resistance noise is ameliorated by increasing the size of the input transistor, but this results in reduced current density for a given power dissipation. Radiation studies (see section 1.7) of the DMILL transistors have shown that in the high radiation environment of the ATLAS detector the minimum acceptable current density is $5 \mu\text{A}/\mu\text{m}$ of emitter length. At smaller values of current density, the gain of the transistor after ten years of ATLAS operation falls about an order of magnitude from its original value to ~ 20 - 30 .

The preamp is the part of the circuit most affected by the loss of transistor gain but will still maintain acceptable performance down to a beta of 30. Noise optimization showed that a collector current of $700 \mu\text{A}$ is a reasonable tradeoff between low power consumption and low intrinsic noise in the input transistor. This optimization sets the size of the input transistor to approximately $100 \mu\text{m}$. Each input transistor is realized as two physical transistors in parallel which allows the layout of the dual preamp to use a cross quad configuration. This cross quad helps to match the thermal and topological environment of the two preamps.

The dominant pole of the preamp is created by an 300 fF feedback capacitance in parallel with the $20 \text{ k}\Omega$ feedback resistor. These components help minimize noise and create a 275 Ohm input impedance that is well matched to the straw characteristic impedance of 295 Ohms ³ and that is reasonably independent of frequency. The dynamic range of the preamp is greater than 600 fC making it the last stage to saturate from large signal depositions in the straw. The total preamp current per channel is about 2.4 mA .

1.3.2 Shaping stages

The differential three-stage shaper and preamp together provide four equivalent 1.25 ns poles of shaping to produce a nearly symmetric response with 5 ns peaking time for an impulse ionization input. The first shaper stage converts the dual preamp output to a differential signal with a gain of two. The second stage provides ion tail cancellation for either xenon or more conventional argon-based gases, as selected externally. A full scale range of 600 fC allows the tail cancellation to be effective over the widest feasible range of charge depositions. The final shaping stage contains a pole-zero network that cancels the short tail added by preamp feedback components and limits the maximum output response of the shaper to 140 fC equivalent input, the largest expected threshold setting for the TR discriminator.

1.3.3 Baseline restorer

The differential signal from the shaper is AC-coupled through 8 pF capacitors into the baseline restorer (BLR) where a bridge diode network with dynamic current control (figure 4) provides a

³Note that the anode wire resistance of about $80 \Omega/\text{m}$ must also be taken into account in understanding the signal propagation.




Figure 4. Baseline restorer functional schematic.

variable impedance shunt across the AC coupled differential signal. Current in the bridge determines the shunt impedance and is dependent on the polarity of the differential output. The shunt impedance increases as signals of the desired polarity are passed to the next stage (discriminator), and decreases when the shaper output returns to baseline. Overshoot (due to discharge of the coupling capacitors) results in an increase in current in the bridge which lowers the shunt impedance across the outputs and quickly returns the signal to baseline.

1.3.4 High and low level discriminator

The BLR is followed by two independent discriminators, one with a low threshold level for tracking, the other with a higher threshold for transition radiation (TR) detection. The low level discriminator contains additional integrations to increase the 5 ns peaking time of the shaper output to 7.5 ns and can reliably be set to trigger on signals between 1 and 10 fC. This discriminator is designed to mark the time of arrival of the avalanche signal from the earliest primary electrons liberated by an ionizing track. Since the primaries move at a predictable velocity, the time information can be used to designate the closest point of approach of the track to the wire. The high level transition radiation discriminator utilizes the same basic configuration, but has a 10 : 1 attenuation at its input and adds 5 ns to the shaping time for a shaping time of 10 ns. This longer time constant allows integration of the prompt and reflected straw signal for accurate energy discrimination of transition radiation photons.

1.3.5 Ternary driver

The low level and TR discriminators switch separate 200 μA currents between shared differential outputs to form a current sum of the combined discriminator outputs. This simple encoding scheme

Table 2. ASDBLR Single Channel Ternary Encoded output.

Signal Detected	Ternary Plus	Ternary Minus
Below Low and High Threshold	-200 μA	0 μA
Above Low Threshold	-100 μA	-100 μA
Above Low and High Threshold	0 μA	-200 μA

shown in table 2 is based on the assumption that the track discriminator output is always present when the TR discriminator is triggered due to its lower threshold [2].

1.4 Wafer fabrication

The ASDBLR was fabricated in a 0.8 μm minimum feature size, BiCMOS Silicon On Insulator (SOI) process. Two prototyping cycles allowed the understanding of issues of yield and substrate pickup that were not apparent from manufacturer provided information. Previous experience with several generations of similar analog wire chamber ASICs in fully commercial processes led to an expected a yield of 90% or greater. It was somewhat surprising to find that the yield was closer to 50 to 60% after folding in reasonable parametric limits on threshold. Measurement of test resistors on the fabricated wafers showed an unusually large number of devices whose absolute resistance varied by 20% or larger from the average value on the wafer. This variance was attributed to a high defect density in the process. Although there was no obvious circuit based way to eliminate this quality control problem it was possible to adjust the number of wafers purchased to accommodate this lower yield.

A second problem had to do with pickup between circuit blocks. Analog device models provided by the manufacturer did not account for capacitance across the insulator layer to the back substrate leading to an observable difference between calculated and measured impulse response when the input capacitance was not balanced between the two inputs. Figure 4 shows a measurement on the first ASDBLR prototype at the analog monitor with and without balanced capacitance on the inputs. Once these were observed it was a relatively simple matter to revise the models, symmetrize the pickup and eliminate the harmonics. The final production of 141, six inch wafers with 1017 useable die per wafer took place in the second half of 2003.

1.5 ASDBLR packaging

Geometric constraints of the TRT Barrel allow a space for each channel of only 30 mm² on the surface normal to the direction of the wires. In addition there is less than 2 cm in distance along the direction of the wires. This tight constraint encouraged a search for a small footprint package for both the ASDBLR and the DTMROC. A custom Fine Pitch (0.8 mm) Ball Grid Array (FPBGA), shown in figure 6 was the most cost effective choice. Due to the relatively low lead count of the ASDBLR it was possible to depopulate part of the 0.8 mm² grid near the inputs and outputs to improve the signal path and simplify the printed circuit board layout.

1.6 ASDBLR testing

Packaged ASDBLRs were delivered in 208 position JEDEC trays with laser engraved, 2-D bar coded, serial numbers. All chips were tested at the University of Pennsylvania Chip Testing Facility




Figure 5. Response at the Shaper monitor output of the first DMILL prototype ASDBLR indicated a potential for harmonic ring when the capacitance at the inputs was not balanced (upper trace). In the lower trace the capacitance on one input was 22 pF and 7 pF on the other. The model parameters were modified to reflect these results and the second prototype and production devices show no harmonic ring.




Figure 6. Two ASDBLR and one DTMROC ASIC, in their custom FPGA packages are shown in comparison with a US dime for size comparison. The solder balls are 300 micrometers in diameter on an 800 micrometer grid.

- an 80 pin, 400 MHz, IMS MSTs chip tester fed by a custom adapted Exatron robotic chip handler. Each chip was DC tested for power supply currents, input channel resistance, output switching and output switching current. In addition, each channel was parametrically characterized for 50% efficiency thresholds using a test pulse injection at values of 0; 3; 5; 30; and 50 fC and the results were written into a MySQL database. Chip testing required slightly more than 30 seconds per chip including robotic handling from the tray to the test socket and back. More than two thirds of the available chips were tested before the selection criteria were finalized. Channel to channel matching at 3 and 30 fC test pulse proved the largest hit on final yield numbers of slightly better than 50%. After good chips were identified from the database, the same Exatron robot sorted good chips into trays that were sent on to board assemblers.




Figure 7. Measured single isolated DMILL SOI transistor current gain (Beta) prior to annealing but after 3.5×10^{14} neutrons per cm^2 as a function of current density in A/m. The arrows indicate approximate collector current densities in various parts of the production ASDBLR circuit (Note that annealing at $> 100^\circ\text{C}$ for 24 hrs was found to improve the measured gain by a factor of 2.).

1.7 Radiation hardness

The ASDBLR relies on the characteristics of the DMILL NPN transistors, resistors and capacitors. The most sensitive parameter observed to change with radiation is the NPN transistor gain which is sensitive to displacement damage from exposure to neutrons and protons. Figure 7 shows one of the preliminary measurements of the gain (Beta) of the SOI process NPN transistors as a function of current density after exposure to 3.5×10^{14} n/cm² at the fast fission neutron source in Prospero, France. Although time efficient, exposure to a ten year accumulated dose in a few hours at a high rate facility is not a completely realistic test. It is well established that some gain recovery will take place over time due to self-annealing. An improvement of a factor of two in gain was measured after annealing the same transistors at $> 100^\circ\text{C}$ for 24 hrs. The results after annealing are generally believed to more realistically represent the expected result of low rate, long term exposure that the Front End electronics will actually experience. As can be seen in the plot the gain of the NPN transistors is a strong function of current density after exposure to neutrons. We chose to operate the preamp input transistor at a current density of $7 \mu\text{A/m}$ to minimize its gain loss. See Ref [4] for more details.

Radiation tests on production ASDBLRs have shown that the ASDBLR is capable of withstanding the predicted 3.5×10^{14} (1 MeV NIEL - Non Ionizing Energy Loss) neutrons/cm² and 7 Mrad ionizing radiation without significant loss of performance, although it has also been shown that unexpectedly high levels of thermal neutrons, in addition to the projected higher energy neutron dose, would compromise the useful lifetime of the ASDBLR [4]. The ASDBLR threshold input gives access to an on-chip transistor and, thus, a way to obtain a sensitive measurement of the change in NPN transistor gain after radiation damage. The threshold input is wired directly to the base of a $4 \mu\text{m}$ NPN transistor in each of the 8 channels. Measurement of the base current in those




Figure 8. Measured change in ASDBLR threshold current versus DTMROC DAC setting (one DAC count = 5mV) after exposure to 3.5×10^{14} n/cm² 1 MeV NIEL neutrons at the French Prospero facility and after annealing. Using the measured threshold current the worst case transistor gain was estimated to be 55 after neutron exposure and after annealing. The 20% variation from post annealing single transistor measurements is within the expected fabrication run to run variation.

threshold input transistors in conjunction with knowledge of the corresponding calculated collector current in those transistors allows a relatively accurate estimation of the ASDBLR current gain or beta (beta = Collector current/Base current). Figure 8 shows the measured threshold current vs setting for production ASDBLR ASICS prior to and after exposure to a neutron source. It should be noted that the post exposure threshold measurements were taken after annealing for 24hrs at 100C. These results imply a transistor gain of 55 after exposure to to 3.5×10^{14} n/cm². Simulations show that the ASDBLR performance is threshold shifted but remains acceptable for values of transistor gain as low as 30.

The DMILL NPN transistors in the ASDBLR show very little sensitivity to ionizing radiation at the doses expected after 10 years of operation at LHC. Figure 9 shows the channel by channel threshold offsets before and after exposure to 7 Mrad of Co⁶⁰.

2. DTMROC

The complementary digital readout chip for the TRT is the Drift Time Measurement/Read Out Chip (DTMROC). The initial prototype DTMROC [5] was produced in the same DMILL process as the ASDBLR but relatively low yields led to a redesign. The production chip [6] was implemented in a commercial 0.25 micron CMOS process and is designed to operate using the LHC 40 Mhz clock.




Figure 9. The plot above shows the measured ASDBLR threshold offsets before and after exposure to 7 Mrad of Gamma radiation. A small amount of broadening is evident in the distribution plotted with the solid line.

A block diagram for the DTMROC is shown in figure 10. We outline the functionality below and then describe each block in detail. Finally we provide performance information.

The DTMROC accepts 16 ternary inputs from two ASDBLRs, the analog Front End processing integrated circuit, described above. The ternary signal encodes time over threshold for the two separate discriminators on the ASDBLR. The low threshold signal is used for tracking and the DTMROC records it in 3.12 ns bins. The high level discriminator is sensitive to the presence of a transition radiation signal and the output of that high level discriminator is latched as a single bit during each 25 ns clock cycle that it is active.

The DTMROC has five differential LVDS connections to the Back End electronics:

- BX — the 40 MHz LHC clock
- Reset — a hard reset
- Command In — trigger and control data input
- Command Out — control data output
- Data Out — drift time data

BX, Reset, Command In, and Command Out are bussed to a group of DTMROC chips. Each DTMROC has a private Data Out line to the Back End electronics.

The DTMROC contains a programmable depth Level 1 pipeline to hold the front end data until the first level trigger arrives. When a trigger arrives as a serial string on the Command In line, the data (eight bits from the digitization of the low threshold input and a single bit from the state of high threshold input) for three successive bunch crossings are transferred to a derandomizing buffer. From the derandomizing buffer those 27 bits of data per channel are serialized and sent off chip over a low voltage differential signal (LVDS) twisted pair copper connection — Data Out. The data are preceded by a header giving the contents of a 3 bit trigger counter and a 4 bit clock counter. There is also a bit to indicate if the chip is in “SENDID mode,” in which case it sends its hardware address rather than the input data and a bit to indicate if any one of a number of error conditions is true — a total of 10 header bits.




Figure 10. Overall block diagram of the DTMROC chip.

In addition to the main data path described above, the DTMROC Command Out port is used for reading back the contents of downloaded registers and of status registers. In addition, the Command Out port implements, an optional asynchronous wire-or output of the 16 low level outputs of ternary receivers for use in self triggered running.

The DTMROC also contains four 8 bit Digital to Analog Converters (DACs) to set the low and high thresholds on each of the two ASDBLRs to which the DTMROC is connected. There are also a pair of one bit digital outputs that are used to control the shaper of the ASDBLR.

A testpulse generator with programmable amplitude and timing is provided for testing and calibrating the ASDBLR. Four additional DACs and four comparators are available to measure the temperature and V_{dd} (+2.5V) on the DTMROC die as well as V_{cc} (+3V) and V_{ee} (-3V) for the ASDBLR via two external connections. This array of DACs and comparators is referred to as VT-sense (Voltage/Temperature-Sense).

A Command Decoder provides register read and write functionality and well as general chip control and fast signal decoding. The Command Decoder receives commands over the Command In line and transmits readback data over the Command Out line.

Table 3. Encoding function of the current in units “U” of current that are sourced from the DTMROC inputs into the ASDBLR outputs.

Amplified Signal in the ASDBLR	True Signal	Complementary Signal
No signal above thresholds	2U	0
Signal between low and high thresholds	1U	1U
Signal above both thresholds	0	2U

2.1 Analog blocks

2.1.1 Ternary receiver

The ASDBLR output is a constant total-current, differential ternary signal (0, 1 and 2 units as shown in table 2.1.1) where a signal unit is $200 + 100 / -40 \mu\text{A}$. This current is pulled from the DTMROC Ternary Receiver inputs which sit at about +1.2 V and have an input impedance of about 200 Ohms. In this scheme, there is a constant return current flow through the negative ASDBLR supply to the analog and then digital grounds and then to the positive DTMROC supply. This non-zero return current produces a small DC offset between analog and digital grounds but obviates the need for a much more complex fully balanced output drive stage in the ASDBLR.

The ternary receivers convert the encoded currents back into two separate digital waveforms. Under quiescent conditions (no ASDBLR pulse) there is a current of $400 \mu\text{A}$ leaving a TRUE input of the DTMROC. A pulse will be viewed as a rising voltage on the TRUE input.

The Ternary Receiver circuit is capable of accurately decoding tri-level differential current pulses as short as 4ns (at base). The circuit is based on differently ratioed differential current mirrors after the common gate current receiver. The use of ternary signals permits high-density communication between the ASDBLR and DTMROC chips without driving up pin-counts and without causing self-oscillations via large signal couplings back to amplifier inputs.

2.1.2 LVDS interface

The interfaces to the Back End electronics are all implemented as true LVDS differential signals. The BX, Reset and Command In input lines are received in standard high impedance differential receivers taken directly from the CERN-RAL library with a guaranteed sensitivity for signals of 25 mV differential or greater and minimum signal widths of about 5 ns.

The Data Out driver is a custom design based on standard LVDS topology with dual current sources driving an H bridge output. For the DTMROC Command Out driver to function on a multi-source Command Out bus it was necessary to include a tri-state function to turn off the current sources except when a particular chip was addressed. Given the need for the tri-state functionality, it was straightforward to include a secondary mode of operation for the Command Out driver where it can be used as a simple current drive into a distant summing node. By connecting the input of the Command Out driver to an OR of the 16 straw tracking signals, it is possible to treat the set of Command Outs on a given Front End board as a single current sum so that a distant comparator can act as a multiplicity trigger on any minimum number of simultaneous ternary inputs. This allows self-triggering for cosmic or source running and was used during some of the commissioning tests.

2.1.3 Digital to analog converters

Four dual, eight-bit Digital to Analog Converter (DAC) blocks are used in the DTMROC to provide a total of four ASDBLR thresholds, two temperature — voltage monitor reference voltages and two Test Pulse Output References. Each dual DAC creates two eight bit reference voltages with a source impedance of 5 kOhms. Two eight bit switch arrays steer ratioed currents from current mirror slave devices to provide an eight bit current output into an internal 5 kOhm resistor. Current in the mirror master array is adjusted by an Opamp driver to provide an output voltage across an internal resistor (10 kOhm) that matches the internal band gap 1.25 V reference. The output of each DAC is a voltage in the range of 0 to 1.25 V with a least significant bit step of about 5 mV.

Two of the dual DACs are used to program thresholds for the ASDBLR chips associated with the DTMROC — two tracking thresholds and two transition radiation thresholds. The remaining two dual DACs are used to control test pulse amplitudes and measure on-chip temperature and voltage and two external voltages as explained in the next sections.

2.1.4 Testpulse generation

Each ASDBLR has two test inputs to couple an external test signal to an on-chip bus capacitively connected to each of four even or four odd ASDBLR preamp inputs. The DTMROC has a test pulse generator that provides a shaped signal to these ASDBLR test inputs. The DTMROC test signal is shaped as an integral of the expected point ionization signal in a TRT straw so that the differentiated signal at the ASDBLR preamp input closely follows the shape of a pulse from a Xe filled 4 mm straw.

Each test pulse generator is adjustable in amplitude (roughly 0 to 50 fC) and time (roughly 0 to 35 ns relative to the BX clock). A dual eight bit DAC (see above) controls the amplitude of the Odd and Even test pulse outputs although only the top six bits of the DACs are connected to the control register. An external connection (TPbias) can be used to adjust the output range — the nominal 0-50 fC range requires a 9 kOhm resistor between TPbias and 0V. The time delay of both Odd and Even outputs is controlled relative to BX through a single 32 position delay line, addressable via a five bit register field. There are separate enable bits for the Odd and Even test pulse outputs to simplify program control of the various test pulse sequences.

2.1.5 Temperature and voltage sense

A simple measuring scheme is used to monitor on-chip temperature, Vdd voltage, and two off chip sense inputs. A monitored voltage is compared to a programmed DAC value by a dual low offset voltage comparator. The comparator output value is stored in the Common Status Register and can be enabled as an input to the Error Bit. The DAC voltage range is limited to 1.25 V so the power supply voltages are scaled down suitably.

Two on-chip comparators are fed from one DAC output. These two comparators are attached to an external input via an on-chip or an external divider and to an internal current source feeding a temperature sensing diode. Two additional comparators are fed from a second DAC and compare that value to a second external voltage source and an internally divided copy of Vdd.

By sweeping the DAC setpoints it is possible to find the value of Vdd with a resolution of about 10 mV or the temperature with a resolution of about 2 degrees C given the 5 mV LSB of

the DACs. By attaching the external voltage inputs of some DTMROCs on a board to the local ASDBLR Vcc or Vee power lines it is possible to monitor power supply and temperature values at each Front End board. Because the comparator outputs can be set as an input to the Error Bit, by properly enabling the error flags it is possible to learn of significant power or temperature changes in a few tens of microseconds.

2.2 Timing

2.2.1 DLL

The Delay Locked Loop (DLL) used in the DTMROC is a standard design with eight delay stages made with current starved inverters feeding back via the last inverter to a common point where there is a comparison with the incoming 40 MHz LHC clock. The control voltage for the starved inverters is then adjusted up or down depending upon the sign and magnitude of the phase error at the comparison point. This gives eight buffered clocks, BX0 to BX7 which are used to latch incoming straw data in 3.12 ns bins.

The DLL block also generates a nearly perfect 50% duty cycle clock waveform and this may be used as the clock signal throughout the DTMROC upon register selection. This internal clock, however, does not obviate the need for a stable 40 MHz external clock, it simply relaxes the duty cycle requirement somewhat on that clock.

As a diagnostic and monitoring tool, the DTMROC includes an eight bit latch for the BX0..BX7 signals that is updated each LHC clock cycle. A latch pattern showing a $50 \pm 12.5\%$ duty cycle results in a “DLL locked” bit set in the Common Status Register.

The DLL can be reset via the command decoder or via the external Reset line or upon power up.

2.2.2 Latch

The asynchronous data from the ASDBLR arrives as pulses with ranges in duration from ~ 5 ns to many 25 ns clock cycles. For the low threshold (tracking) signal, the asynchronous data can simply be latched continuously by the BX0..BX7 clocks producing an ongoing bit stream record with 3.12 ns resolution or about 1 ns RMS. Each 25 ns clock cycle (eight of the 3.12 bins) is then recorded as an eight bit record in the Level 1 Pipeline memory. However, the high level threshold, because it only needs to be recorded as one bit per clock cycle, requires a synchronizing latch sensitive to transitions as short as 5 ns. Any high state during a given clock cycle is then stored in the Level 1 Pipeline as a high for the next clock cycle. In addition, a special “accumulate” mode has been implemented for this latch so that it does not reset each BX cycle as it would in normal data taking, but resets only upon external command. The accumulate mode is useful for chamber testing with radioactive sources or for looking at average noise rates.

2.3 Digital blocks

2.3.1 Pipeline

The Level 1 Pipeline is a logical FIFO block that stores the nine bits (one high threshold bit plus eight time digitization bits of the low threshold) of latched data from each input channel, the current value of a four bit clock counter plus a one bit error flag from the DLL. This FIFO is designed to hold the Front End data for the length of the Level 1 trigger latency.

The pipeline is implemented as two banks of 128 deep by 153 bit wide Random Access Memory (RAM) built out of 34 blocks of 128x9 bit memory cells - a CERN library part. The pipeline operates as a continuously running circular buffer, updating the read and write pointers into the memory. Writes (and therefore reads) alternate between the two banks of RAM to minimize power consumption fluctuations that could occur on parallel access to all memory blocks. The memory controller also implements a Built In Self Test (BIST) function to allow thorough memory tests without massive I/O traffic.

2.3.2 Derandomizer

The Derandomizer is an additional buffer also acting as a FIFO built from the same synchronous dual-port static RAM memory as the Pipeline, but with half the number of banks giving a storage capacity of 128 words of 153 bits. Upon receipt of a Level One Accept (L1A) the current Pipeline output word and the following two consecutive words are stored in the Derandomizer for readout. In addition to these data, the SENDID status bit, the L1ID and the Common Error status bit are stored. This gives 441 bits to be stored per event. The Derandomizer can store 42 events. In the case of memory overflow control logic provides a “full” flag and skips complete events avoiding synchronisation problems until the memory clears an event and is able to store a subsequent event. The Derandomizer and Pipeline SRAM blocks are equipped with Build-In-Self-Test (BIST) controlled via the Configuration register. The BIST result can be read out from the General Status register.

2.3.3 Serializer

As soon as the Derandomizer has any content, data is placed serially on the DataOut line at a 40MHz bit rate. Given the known constant event size it is possible to use a simple protocol with a three bit preamble “101” sent at the beginning of each event. When idle, the data output signal is “0”.

2.3.4 Command decoder

The command decoder block receives LVDS BX and a CommandIn signals from the TTC. The command decoder examines the CommandIn stream bit by bit and issues all the necessary timing signals (L1A,...), internal registers read/write strobes and data. As implemented the decoding algorithm is very simple and is built from a shift register, a look-up table with valid command codes and comparator logic. This architecture is, in this case, lower cost in power and resources than a finite state machine. In order to insure that the circuit is robust against single event upset, additional logic was incorporated. The decoder for the critical ‘fast’ commands (L1A, SoftReset and BunchCrossingReset) and the command length counter were implemented as three copies for automatic error correction. A “surveillance” counter was implemented to guarantee the Register Transfer Level (RTL) state coverage and to release any access lasting longer then 171 clock cycles. Upon a read request the command decoder serially transmits the contents of the selected register on the differential CommandOut line, three clock cycles after the last bit of the command,. This CommandOut line is common to multiple chips on a board and therefore has a ‘tri-state’ capability. A three bit preamble (‘101’) is used for each data transmission and the idle state of this line is “HiZ”. The chip address is defined by setting external pins.




Figure 11. DTMROC’s time-measuring performance with the nominal 2.5 V and 2.0 V power supply. A 4.0 ns wide tracking pulse was injected at 100 ps intervals across three full clock periods, 75 ns in total. The picture shows the leading (red) and falling (blue) edges deviations from fit and differential non-linearity from a single channel.

2.3.5 Error handling

The implemented protocol is based on the ABCD chip specification [7], which lacks any advanced protection against transmission errors. The bit patterns for the commands are, however, chosen such that a single bit error should not cause the acceptance of a wrong command. An erroneous bit pattern in the middle of a command causes a flush of that unrecognized field and an attempt to decode the following command.

2.4 Measured performance

The preproduction fabrication of the design was fully successful and so the full production run of 48 wafers was done using the original masks which included lithographic adjustments to simulate two “faster” and two “slower” (from 85% to 125% of nominal propagation time) variations about nominal process parameters. The DTMROC preproduction tests looked at logical performance, time linearity, DAC linearity, and performance over voltage and process variations. As shown in figure 11, the time linearity is at the expected 1 ns level, differential and integral, not only at the nominal 2.5 V Vdd but also down to at least 2.0 V. The basic logic of the chip such as the command decoder, was shown to work properly at over 100 MHz at 2.0 V, but the RAM blocks did show some degradation at elevated clock speeds, lowered process speed and lowered supply voltage as shown in figure 12. As noted above, the RAM blocks can be tested with a built in self test routine as access through the DTMROC command structure is indirect and full coverage tests would be very time consuming.




Figure 12. RAM access rate vs. Vdd and over various process variations. Note that for the DTMROC production run, the slowest (labeled 125% on this plot) process variation was not packaged.

Radiation tolerance was tested at the CEA Saclay Pagure facility up to 7 Mrad of Co⁶⁰ with the only degradation being about a 10% decrease in maximum DAC output voltage. SEU sensitivity was tested at the CERN PS irradiation facility using a 24 GeV proton beam with an integrated fluence of 1.8×10^{14} p/cm² and shown to be of order 10^{-14} cm² [8]. This rate of SEUs produces a per TRT error rate of much less than 1 Hz. The event counter reset signal is available at about 1 Hz and so the TRT can take advantage of ATLAS scheduled deadtimes to reset DTMROC chips that have suffered from SEU.

Production tests of the chips, conducted at the University of Pennsylvania Integrated Circuit Test Facility, concentrated on checking I_{DD} , exercising the built in memory self test, and then loading preset pulse patterns into the sixteen input channels and verifying proper time encoding within one LSB. The ternary receivers were checked for proper performance over their full range of amplitude and offset as were the LVDS I/O blocks. The DAC outputs and VT sense DACs were also exercised over their range and the test pulse output was checked at several values of amplitude and time delay. About 40,000 test vectors per chip were completed in about 30 seconds including robotic handling. All production chips were laser engraved with 2-D barcode serial numbers and test results were stored in a database. Chips passing all test criteria (about 85% yield) were pulled from the JEDEC trays and placed into final trays for the assembler by the same robot.

3. Board designs

It is necessary to design the Front End boards, the detector grounding and shielding, and the power and signal connections to the outside world very carefully in order to achieve optimal performance. For instance, the 30 to 70 cm long straw tubes in the TRT are close to being ideal 1/4 wave antennas tuned to the ASDBLR peak bandwidth of about 33 MHz. In addition, the ASDBLR inputs must be carefully isolated from the 40 MHz digital logic of the DTMROC,

The geometry and other constraints of the TRT End Cap [9] and Barrel [10] are very different, therefore the designs of the End Cap and Barrel printed circuit boards need different optimizations. The basic schematic is, however, the same in each case — input protection (24 Ω series resistor with a diode clamp to ground) at the detector and then the ASDBLR, DTMROC, and power filtering on the Front End boards. All signals in the system are differential except that the straw anode input is matched only pseudo-differentially by a trace carrying the ASDBLR dummy input as close to the detector as possible. The End Cap and Barrel printed circuit designs and justifications for the optimizations are discussed in detail below.

3.1 End cap boards

The End Cap straws are positioned radially, therefore while the 4 mm straws nearly touch with a center to center distance of 5.2 mm on the inner radius of the TRT, the straw to straw center distance along the ϕ direction is 8.2 mm at the outer radius where the electronics are located. This spacing allows almost 3/4 of a cm² per channel for the electronics located on the tread of the End Cap wheels. This location for the electronics, at the outer radius of the TRT up against the ATLAS solenoid cryostat, means that the radiation length of the electronics package is not as critical a concern as it is in the Barrel case where the electronics must be located at the end of the Barrel just before the End Cap. There is also, by design, sufficient room in r to allow separate, albeit closely packed, layers of printed circuit to handle the analog (ASDBLR) and digital (DTMROC) parts of the system.

3.1.1 ASDBLR boards

There are two schematically identical but physically different designs for the ASDBLR boards shown in figure 13 for the Type A and Type B End Cap Wheels. The Type B wheels are constructed with half the straw density in the Z direction of the Type A wheels. The Type A boards are about 60 mm square and serve 64 channels of detector (1/96 of a wheel in ϕ) with eight ASDBLR chips. The Type B boards are about 120 mm in the Z direction to match the wheels. Both types are four layer, 1.5 mm thick, through-via designs. Connection to the straw anodes and cathodes is made via miniature connectors to the End Cap Wheel WEBs, complex rigid flex printed circuits that serve to attach the electronics boards to the detector both physically and electrically [9]. The high voltage decoupling capacitors for the cathode connections and 24 Ohm input protection resistors are mounted on the WEB as are the HV protection fuses which also serve as blocking or isolating resistors to the HV supply.

ASDBLR ternary outputs are sent to the DTMROC chips mounted on the Triplet Boards via four high compliance miniature connectors. ASDBLR control signals from the DTMROC and power for the ASDBLR chips also travel over these four connectors.




Figure 13. Two A Wheel (l) and one B Wheel (r) ASDBLR boards top view showing the ASDBLR chips and the four multi-pin white connectors going to the Triplet Board described below. The semi-circular cutouts in the board edges are designed to clear plumbing and HV connections to the wheels.

3.1.2 DTMROC triplet boards

The Triplet boards are identical for Type A and Type B wheels. One Triplet board serves three ASDBLR boards — 32 Triplet boards per End Cap Wheel. Each Triplet board, shown in figure 14, has 12 DTMROC chips and is constructed as three separate four-layer through-via printed circuits joined by flexible soldered on jumpers. The jumpers allow flexibility so that the Triplet can follow the curve of the End Cap Wheels. An initial design effort using rigid flex technology demonstrated good performance, but was abandoned because of cost considerations and because the stiffness of the flex circuitry was too great to allow easy assembly on the wheels.

3.1.3 Shielding

For the TRT End Cap, the Faraday shield around the straws is made complete through the WEB structure on one side of the wheel up through the ASDBLR board and down the opposite WEB as shown in figure 15. The analog ground plane of the ASDBLR board acts not only as the cathode connection and reference for the ASDBLR input but also as a part of the Faraday shield. As a secondary shield, it was learned that completing the connection of the digital ground plane of the Triplet boards with a low impedance phosphor bronze clip between the edge plated segments of the Triplet boards provided additional cross talk protection from the digital logic on the Triplet. A full external Faraday shield constructed at the level of the TRT cable trays surrounds the detector and electronics.

3.1.4 Power distribution

Power is brought to the Triplet via the same 60 pin connector that also carries the four control lines and the 12 data output lines. The +2.5V power for the DTMROCs is filtered at the 60 pin




Figure 14. End Cap Triplet board showing the four DTMROC chips per segment with the power and data connector on the right. The four white jumpers between segments carry control and data signals and power and power returns. Note that the breakaway bars at top and bottom of the Triplet are removed after testing.




Figure 15. Cross-section drawing of the construction of an End Cap wheel. The DTMROC Triplet board is at the very top, then the ASDBLR board, the cooling structure and the WEBS.

connector, is returned via the digital ground plane on the triplet board, and is distributed only on the Triplet board. The $\pm 3V$ power for the ASDBLRs is filtered to the analog return at the 60 pin connector and then distributed to the three ASDBLR boards. Analog and Digital power returns are commoned at the connector. Cooling is described in section 4.2.

3.2 Barrel boards

The Barrel readout system must function in a much more constrained space and, in addition, there




Figure 16. End-on view of the TRT Barrel during above ground testing at CERN before installation in the ATLAS detector hall. Each of the three types of quadrilateral modules (types 1,2, and 3 from inner to outer radius) is bisected by a member from the Barrel Support Structure. These six different triangular shapes define the geometry of the Front End boards.

is a stringent radiation length limit as the boards must be situated in the middle of a tracking region. The available board area per straw is also about a factor of two less than in the End Cap case. The mounting of the Barrel modules in the Barrel Support Structure (BSS) defines a set of triangular spaces shown in figure 16 into which the electronics must fit in order to connect to the anode and cathodes brought out on the module tension plates.

After exploring a number of possible design concepts all of which failed to meet the design constraints discussed below, the collaboration settled on a single board mounting both ASDBLR and DTMROC chips [11] as opposed to the two board solution used in the End Cap.

3.2.1 Design constraints

The connectors designed into the Barrel Tension Plates (individual sockets for 0.5 mm pins arranged in groups of 16 anodes and 6 grounds) shown in figure 17 [10] are not suitable for simultaneous mass insertion as would be required by a triangular printed circuit board serving from 10 to 27 of these 16 anode patterns. In addition, the areal electronics density required, especially in the smallest (Type 1) modules, is somewhat greater than practical. Therefore, an additional small




Figure 17. Cross-section view of a barrel module tension and HV plate region. The Protection Board (labeled ‘Electronics Card’) is shown plugged into sockets in the tension plate. The active electronics plugs into a high density, high compliance connector on the protection board and is parallel to the Protection Board.

board, the Protection Board, was added to the scheme to allow a transition from the 0.5 mm pin connectors to the high compliance miniature connectors used on the End Cap — the protection board improves the mechanical compliance to allow, in the worst case, making 600 simultaneous connections and provides a small additional surface area to house the clamp diodes and input protection resistors for the ASDBLR inputs.

The existence of a Protection Board coupled with the constraint that the overall depth in Z available for Barrel electronics including cooling and connectors is about 1.5 cm, a single substrate printed circuit solution becomes almost inevitable. Because of the obvious mechanical difficulties in using a liquid cooled plate under the ASDBLRs, a set of two piece machined Al plates were designed to be attached to the outer or DTMROC side of the triangle boards. In that way the cooling connections could inhabit space reserved for other plumbing connections.

3.2.2 Design rules

Although variations in geometry by end and by module layer require a total of 12 different board and cooling plate footprints, the design approach is the same for each. The basic rule is to separate the board into analog and digital domains splitting a total of 14 layers equally between them. Four levels of blind vias allow isolation of analog and digital signals within their separate domains. All fast signals are low level differential. Barrel boards are arranged in 16 channel ASIC triplets consisting of two eight channel ASDBLRs and one DTMROC. Analog connectors and ASDBLRs face inwards towards the barrel module. Two signal layers separated by an analog ground plane utilize short vias to isolate the inputs from the rest of the board. An empty layer below the input signal layers and a power layer with cutouts corresponding to the location of input signals minimize the capacitance to the other board layers. The analog side is enclosed by a second analog ground plane with a power layer finally facing outwards towards the digital side. The layer stackup is shown in figure 18

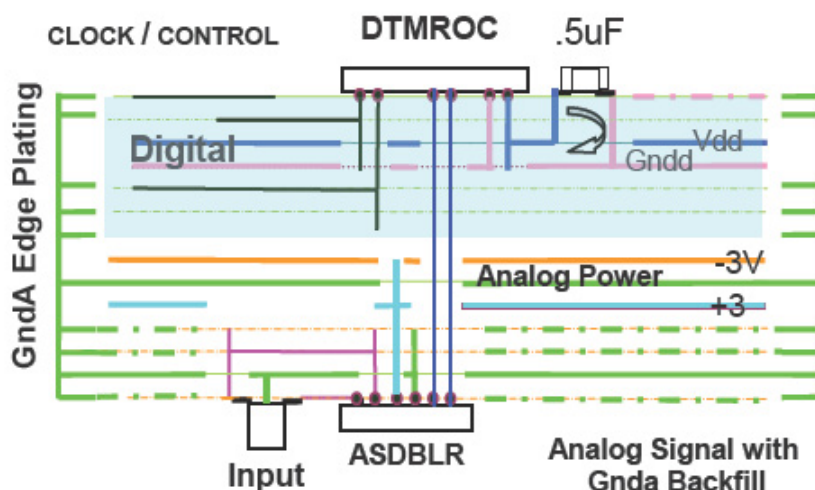


Figure 18. Cross-section of the printed circuit planes showing the digital and analog domains within the 14 layer stackup. Four different classes of blind vias, two for the digital domain and two for the analog are used for almost all connections except direct connections for ASDBLR outputs to the DTMROC and control signals from the DTMROC. Power is provided from a through-hole connector. The edge plating noted on the left of the figure wipes against the RF connector fingers mounted on the Barrel Support Structure (BSS).

The DTMROC ASICs are placed on the outermost layer away from the barrel. Analog and digital grounds are completely separate but provision is made on the digital side to connect them using surface mount resistors placed around the perimeter of the triangular footprint. At these locations direct access to analog and digital ground is provided with large pads to allow low inductance connection from digital to analog ground and from analog ground to the conductive support structure of the barrel module. The support frame of the barrel module, the RF contacts (“fingers”) around the perimeter of each board, the analog layers of the board, and the inner and outer cylinders of the Barrel comprise a Faraday shield which surrounds the “analog zone” of the detector — the straws and the signal paths to the inputs of the ASDBLR. Special care is taken at the two sides of each module where HV connections to the cathodes are made via wide Kapton foils — the foils are notched every few centimeters and custom grounding clips bridge the gap between the RF fingers on the BSS and the edge plated analog board ground.

3.2.3 Barrel printed circuit boards

The 12 final designs of the boards shown in figure 19 are designated ARxyz (where AR is an historic reference to an Active Roof), x is the module number (1-3), y is Front or Back corresponding to C or A sides of ATLAS, and z is Large or Small triangle of the pair (the smaller triangle has its base at smaller radius). Front and Back boards are not mirror images of each other as the Tension Plate designs reflect a translation, not a mirror operation from end to end in the modules — this geometry forces twelve distinct printed circuit designs. General characteristics of these designs are listed in table 4.

Table 4. AR Board Designs.

Board	No. DTMROCs	No. Clk Groups	No. Power Connects
AR1F(B)S	10	1	1
AR1F(B)L	11	1	1
AR2F(B)S	15	1	1
AR2F(B)L	18	2	1
AR3F(B)S	23	2	2
AR3F(B)L	27	2	2




Figure 19. The photograph above shows a module 3 barrel board (two power and clock groups) with one through hole and one pig-tailed power connector and two small white data connectors near the center of the board. The DTMROC Fine pitch Ball Grid Array (FBGA) chips can be seen on the top of this board.

3.2.4 Power distribution

Whereas the End Cap uses a single connector for power and data, that single connector is much too large for use in the Barrel case. All AR boards have separate power and data connectors. The data connectors are 60 pin surface mount versions of the high compliance connectors used to connect to the Protection Boards. The power connectors are 6 pin industrial style through hole connectors. As seen in figure 19 some of the power connections are implemented as pig-tails to ease routing of the power connections through the plumbing. Unlike the End Cap case in which electrical connections were made to the boards before plumbing connections and shielding were completed, the Barrel cabling was installed after the full plumbing manifold system was in place. Cooling is described in section 4.2

3.3 Testing program

All TRT Front End boards were assembled by industrial printed circuit assemblers and then shipped to The Niels Bohr Institute for visual inspection, burn in and testing. Each board was run on a current and voltage monitored stand at an elevated temperature for one week or more prior to detailed testing. Boards failing the initial tests were repaired either at NBI or sent to the University of Pennsylvania for diagnosis and repair. Tested boards were then sent to CERN, retested in the surface assembly building, mounted on the detector, and then tested again. Results of the on-detector tests are covered in section 13.2.

4. Cabling and cooling

4.1 Cabling

The organization of the cabling plant in the inner detector of a collider experiment is always a logistical and engineering struggle. The compromise between low mass and minimal volume versus high quality signal and services transmission can differ in each example. The following paragraphs describe briefly the solutions arrived at for the TRT.

The cabling of the TRT has been split into two separate parts in the two different regions of the detector — Front End to Patch Panel (Type II cables in ATLAS nomenclature) and then Patch Panel to USA15 (Type III cables) [12]. The regions differ in their environmental design constraints. The detector contains the following types of cable:

- Signal cables for data transmission, read-out, timing and temperature sensing.
- High voltage cables to bias the straw cathodes
- Low voltage power cables to supply the FE electronics and auxiliary electronics on Patch Panels
- Control cables for several functions implemented on the Patch Panels

4.1.1 Detector to Patch Panels

For the innermost part of the TRT (Front End to Patch Panel), the main design constraints are mass and volume. The connection between the Front-End electronics on the detector and the Patch Panels routes all the services and data through the very limited space between the TRT and the cryostat wall — a region shared in ϕ with the SemiConductor Tracker (SCT) services. These ‘Type II’ cables terminate at the TRT Patch Panels (see section II). The ‘Type II’ cables were produced as prefabricated bundles customized for each part of the detector. For each function, the cable was chosen for the minimum possible diameter — usually set by the insulator thickness. For power cables high quality Kapton (polyimide) insulation was selected. There are three types of cables used in this section of cabling:

- Read-out, control, monitoring (temperature) cables — individually shielded twisted pairs, tin plated copper conductor, 0.13 mm conductor diameter (36AWG), four drain wires of silver plated copper, a wrapped aluminized polyester foil as shield, 0.95 mm overall diameter.

- Miniature HV cables — 2 kV mini-coaxial cable, 0.16 mm conductor diameter, tin plated silver-copper alloy, aluminized polyester foil shield, two silver plated copper drain wires.
- low voltage shielded cables procured in triplets or pairs according to the voltage sets carried on them.
 - Analog — wheel type A — three conductors 1.0 mm², extra flexible tinned copper insulated in MULRAD, 0.12 mm tinned copper braid with a 85 % minimum coverage, Kapton insulated, OD 4.6 ±0.3 mm
 - Analog — wheel type B and barrel — three conductors 0.6 mm² (37 strands of 0.15 mm), extra flexible tinned copper insulated in MULRAD, 0.12 mm tinned copper braid with a 85 % minimum coverage, Kapton insulated, OD 4.00 ±0.2 mm
 - Digital — two conductors 1.0 mm², extra flexible tinned copper insulated in MULRAD, 0.12 mm tinned copper braid with a 85 % minimum coverage, Kapton insulated, OD 4.3 ± 0.2 mm

In both the end cap and barrel cases, the HV cables were fabricated as separate bundles for each ϕ sector of the detector. For the end cap, low voltage power cables were bundled at the detector end with the data and control twisted pairs since a common connector is used for both types of connections (as shown in figure 20). However, in the barrel case, the low voltage and data-control lines were bundled separately. The data and control lines are such a fine gauge (36AWG) that connection had to be made via hand soldering in all cases — about 360,000 manual soldering operations. At the detector side, in both the End Cap and Barrel cases, the data and control twisted pairs go to a common connector, but at the Patch Panel end the data cables go to Data Patch Panels while the control lines go to the TTC Patch Panels. Thus the bundles are of a fairly complex one connector to two connector form. For the Type II cables, there are 736 complex harnesses (some 2,800 separate bundles of different types of wire, multiple connectors, and multiple lengths). The LV cables for the barrel used crimped and soldered connections at each end, the End Cap LV cables were soldered to an adapter board on the detector end but used the same crimp connections as the barrel power cables on the Patch Panel end. HV cables used commercial crimped multi-way connectors at the Patch Panel ends and were soldered and custom strain relieved into three way commercial connectors at the end cap or into custom fuse boxes at the barrel end.

All cables were thoroughly tested for continuity and against shorts on custom built cable test jigs at the bundle assembly site. In addition, the HV cables were subjected to HV leakage tests. Each cable bundle was labeled for its intended position in the detector as cable lengths depended in detail upon the route from the detector face, across the cryostat flange, and through the tile calorimeter fingers and first layer of muon chambers to the Patch Panel boxes.

The bundled cabling (of all types for a given ϕ section) was placed into cable trays at CERN, retested, and, as an assembly, mounted either on the cryostat wall (barrel part) or on the mechanical structure of the detector (endcap part). The barrel cables were run in one piece from the detector up to the Patch Panel area. The endcap cables have been split at the cryostat flange where the Patch Panel Front 1 (PPF1) has been located, allowing easier and separate installation of the two cable lengths. On average, the Type II cables are 12 m long.



Figure 20. Partially assembled End Cap showing two full cable trays at the top and one empty cable tray at the bottom. Each tray carries cables for 1/32nd of the detector and acts as part of the outer Faraday cage of the detector. The blue cables are the 36 AWG data and control shielded twisted pairs, the heavier orange-brown cables are low-voltage supply cables. The light colored HV cables are just barely visible in the upper tray on the left. In this view the Triplet Boards serving the A type Wheel nearest the interaction point are on the left and the PPF1 connectors are just out of the picture on the right.

4.1.2 Patch Panels to control room

From the PP2 location on to the control room there is enough space to use standard cables for all functions. The data from thirty DTMROCs are merged at PP2 into a single 1.6 Giga Baud opto-fiber carrying 1.2 Gb/s of TRT data. The temperatures of the NTC board level thermistors and RTD detector and manifold mounted sensors are measured on the TTC Patch Panel boards and those cables also stop at PP2. Only timing/control cables and bulk HV and LV cables are installed to connect PP2 with the control room and the bulk LV power supplies on the HS structure. The average length of these Type III and Type IV cables is around 75 m with longest being 102 m. The PP2 crates (see section II) contain boards at which many DCS control and monitoring operations are performed. The communication between those cards and the control room is performed via a standard industrial bus — CAN bus. This communication need adds CAN bus cables to the inventory (NG18P, 9-pair shielded cable, 120 Ohms). The PP2 crates also contain the low voltage regulators with their control (see section 12) and status monitoring. The “bulk” power lines which deliver LV to PP2 are simple multi-strand copper cables of 35 and 50 mm². The signals controlling the data Patch Panels (where the data signals from the front-end boards are converted from electrical to optical) are carried from the control room to PP2 via standard Ethernet-style CAT5 shielded cables. The timing (TRT-TTC) system uses the following cable:

- 25 pair twisted, round shielded and jacketed discrete wire cable, made of 28 AWG tinned stranded copper core with an overall aluminized polyester foil and tinned copper braid shield. The cables are halogen free, flame retardant, 110 ohm balanced impedance, color coded.

At the PP2 position the HV cables also change their type via a simple passive disconnect. From the PP2 up to the power supplies in the control room following cable is used:

- 3 kV DC, round overall shielded/jacketed 56 discrete coaxial cables, made of 0.12 mm² [AWG26] tinned stranded copper core, overall tinned copper braid (shield), halogen free, flame retardant

4.2 Cooling

The heat produced in the Front End and Patch Panel electronics and that produced by the voltage drop in the LV power cables must be removed from the cryostat and muon system areas in order to keep the TRT and the rest of the Inner Detector at reasonable operating temperatures. It is also necessary to avoid local hot spots within the muon tracking system which could distort the geometry of that system. The overall ATLAS requirement is that each subsystem within the cavern must be thermally neutral. This need for thermal neutrality mandates the need for an efficient cooling system. The space constraints imposed by the detector geometry dictate a liquid cooling system. A monophasic system relying only on the heat capacity of the fluid has been shown to be sufficient for the needs of the TRT.

4.2.1 Power dissipation — Front End

Power dissipation in the front-end electronics is ~ 60 mW per channel (40 mW — ASDBLR and 20 mW — DTMROC) or about 320 mW/chip for ASDBLR, and 320 mW/chip for the DTMROC. This gives a total TRT Front End power dissipation in the range of 25 kW. The power loss in the cables running from the PP2 regulators to the Front End is an additional 4 kW.

4.2.2 Power dissipation — Patch Panels

The PP2 crates contain three LV distributor boards which dissipate about 60 W each per crate. For safety reasons the PP2 crates are closed with front/back panels and there is no air convection which might remove the heat. It is necessary to remove the heat via liquid cooling. The PP2 LV cabling running from the power supplies on the HS platforms outside the ATLAS detector has been routed in between high precision muon chambers which need to be kept at a constant temperature. Therefore this stretch of the LV cables must also be cooled.

4.2.3 Design and deployment

For all the above loads a simple monophasic cooling system was chosen. A liquid is circulated under sufficient pressure and with adequate mass flow to remove heat from the cooled elements. The liquid needed to have a number of special properties - to be radiation hard (i.e. not decompose under ionizing radiation or create aggressive compounds which might endanger the integrity of piping or the liquid itself). It should also be a good insulator, a dielectric. Due to the inaccessibility of the detector electronics, a limited leak developing during the lifetime of the experiment could not be allowed to cause collateral damage through leakage currents or short circuits — a requirement that excludes water, despite its excellent heat capacity. After a series of radiation tests a fluorinert, C_6F_{14} , was selected as the working fluid for the TRT cooling system.

A cooling station of the design shown in figure 21 was constructed with a cooling capacity of 70 kW. This station supplies four distribution racks where the main flow is split in a manifold to barrel and endcap detectors via 34 rack output lines and 50 rack return lines. A similar cooling station provides cooling for the PP2 boxes and the Type III and Type IV cables. The cooling system operates at six bar with an adjustable supply liquid temperature 15-17 °C. The pressure is high enough to cope with the dynamic drop in the piping as well as the hydrostatic conditions in an




Figure 21. Block diagram of the monophase cooling system.

extended and high volume system. Fluid operating temperature in the experimental hall is required to be maintained well above the dew point.

The cooling elements in the barrel detector are “cooling plates” machined from aluminum. These plates are attached to the electronics boards by means of a thermally conductive silicone adhesive and are supplied with the cooling liquid through a manifold of pipework running over the BSS face. Each 1/32 in ϕ of the barrel is a separate cooling loop.

The endcap electronics is cooled using a logically similar solution with a very different mechanical design. In this instance a thermally conductive rubber mat is used as the thermal connector between ASDBLR boards and aluminum plates attached to SS pipes that carry the fluid. There are two cooling loops per end cap wheel each covering 180 degrees in ϕ , but with entrance and exit points rotated in ϕ for each wheel to distribute the cooling inlet and outlet pipes around the end cap support structure [9]. Both the barrel silicone adhesive and the endcap conductive rubber mats were tested for radiation tolerance. Neither material lost flexibility or thermal conductivity at neutron doses above 10^{14} n/cm².

At PP2, each Patch Panel board is bolted to a 2 mm thick aluminum cooling plate with heat transfer conductive rubber between the cooling plate and the high dissipation components. This cooling plate is then thermally clamped to the PP2 box top and bottom plates with expansion

clamps. The top and bottom plates of the PP2 boxes have embedded copper tubing carrying the cooling liquid for that box. At each PP2 location, the PP2 boxes are connected to a local input and output manifold.

The cooling of the cables is realized by routing LV cable bundles along side the return pipes of the electronics cooling system. Cabling between PP2 and the periphery of ATLAS is cooled by a separate network of piping common to several inner detectors. This network is driven by second cooling station of identical design to the TRT Front End electronics cooling station.

The system has proven, during tests, capable of keeping all barrel Front End board components under 40 degrees and, in the end cap case, keeping even the worst case triplet board components under 45 degrees even though they are cooled only via conduction through the electrical connectors to the ASDBLR boards.

Part II

Patch Panels

The TRT (and Pixel) Patch Panels (formally, Patch Panel 2, abbreviated PP2) are located just outside the first layer of Muon Barrel chambers in 64 different crates (one crate per 1/32 per end) on five different platforms on both the A and C sides of ATLAS as shown in figure 22. These locations were chosen as being accessible during a short shutdown but as close to the Inner Detector as possible in order to reduce dispersion and signal loss in data cables and voltage drop in the low voltage power cables. All the TRT electrical services are either repeated or are at least disconnectable, in the case of the High Voltage, at these Patch Panels. Treatment of the TTC signals, and Data Signals is described in the sections below — the HV “Patch Panel” is simply a set of multiconductor connector pairs located at the Patch Panel location and is not further described.

The TRT Low Voltage DC power is subject to a final level of regulation, control, and monitoring at the Patch Panels using custom radiation tolerant linear regulators for both the Front End electronics ± 3 V and +2.5 V requirements and the PP2 located logic for TTC and ROD. The Low Voltage regulator PP2 boards are described in detail in the Low Voltage power section (12), but it is worth noting that these boards dissipate most of the power at the PP2 location and dominate the cooling requirements of the PP2 crates.

The data and control signals transmitted between the Front End Boards and Backend Modules are all received and repeated at the Patch Panel boards. Each PP2 location houses between five and eight Patch Panel boxes, each of which contains all the Patch Panel boards which are associated




Figure 22. ATLAS Patch Panel 2 locations.

with 1/32nd of one side of the detector. Each box houses eight Patch Panel boards in total; three for 1/32nd of one side of the barrel (one power PP, one ROD (or “Data”) PP, one TTC PP), and five for 1/32nd of one endcap (two power, one TTC, two ROD). All eight boards in one box share common cooling, power and CAN bus.

Because of the differences in readout granularity and cabling between the Barrel and the End Caps, a single Patch Panel design could not be used for both parts of the detector. As a result, the Patch Panels have been designed using a modular scheme consisting of an active part which contains all the functionality and a passive part which handles the routing of signals to and from the Front End boards. This avoids the need for two separate active designs which would have complicated not only the design and production, but also the testing and installation processes. The specifics of each board design will be described in the following sections.

5. TRT-TTC Patch Panels

The primary function of the TTC Patch Panel board, shown in block form in figure 23, is to receive and repeat signals between the TTC and the Front End boards. 20 TTC lines, each consisting of one group of Clock, Command In, Command Out, and Reset signals, are handled by this board. The Command In and Command Out signals are received and repeated directly, with some shaping in the inputs and outputs (discussed below). The Clock and Reset signals, on the other hand, arrive once for every 10 lines, and are fanned out on the Patch Panel to each of the Front End boards. In addition to being fanned out, the Clock signal can also be independently delayed for each line (after the fan-out) in increments of 0.5 ns. Finally, the TTC Patch Panel is responsible for driving temperature read-out of the detector and front-end boards, which is accomplished by means of an Embedded Local Monitoring Board (ELMB) [13] mounted on the Patch Panel board.

5.1 I/O compensation and filtering

All of the signals entering and exiting the TTC Patch Panel board are LVDS traveling over small copper twisted pair wire. As mentioned in the cabling section, the signals going to and coming from the TRT-TTC travel over 60-100 m of 28 AWG copper twisted pair, while the signals going to and coming from the Front End boards travel over 10-14 m of 36 AWG shielded twisted pair. The shaping that is applied to each input or output depends on the cable type, as well as the particulars of the signal.

There are two different filters that are used for signals coming onto the TTC Patch Panel board. The Command-In and Reset signals coming from the TTC, as well as the command-out signals coming in from the Front End boards are received with a passive eleven component differential RLC filter shown in figure 24. This filter has been shown to recover good signal performance for signals traveling up to 100 meters over 28 AWG twisted pair wire as shown in figure 25. The Beam Crossing Clock (BX) signals coming from the TTC are received with an active equalization circuit that can be tuned for cable lengths between 60 and 90 meters in steps of 10 meters by moving a simple jumper. This active equalization uses a commercially available IC and provides a lower jitter and better signal quality than the passive filters used for the other signals albeit at higher cost. This step was taken for the clock in particular because a well-shaped, reliable clock edge is imperative to the proper performance of the drift time measurement in this detector.




Figure 23. Block diagram of the TTC Patch Panel showing connections to and from the Back End (the TRT-TTC) on the left and the Front End electronics on the right. The ELMB mounted on the Patch Panel board is for detector and Front End electronics temperature monitoring only, it does not control any Patch Panel functionality.

For the outputs as well, two different approaches were taken. The signals going out over the 36 AWG shielded twisted pair wire towards the Front End boards (BX Clock, Reset, Command in) are shaped with a three component passive RL pre-compensation filter to ensure good signal shape when they are received on the Front End boards. The signals going out of the TTC Patch Panel towards the TRT-TTC (Command Out only) have no pre-compensation, but instead the LVDS drivers have been doubled (with the second driver capacitively coupled) to provide the extra drive needed to traverse the 100m cables. The compensation in this case is handled on the TRT-TTC board itself with an active equalization similar to the one described above for the BX Clock on the Patch Panel board.

5.2 Clock fine delay

Aside from receiving and repeating signals, the TTC Patch Panel also houses a few important functions for detector control and monitoring. One of these functions is the fine delay of the clock which is supplied to the Front End boards. This delay is used to compensate for differences in




Figure 24. RLC Cable compensation network for 100 m of 28 AWG twisted pair.

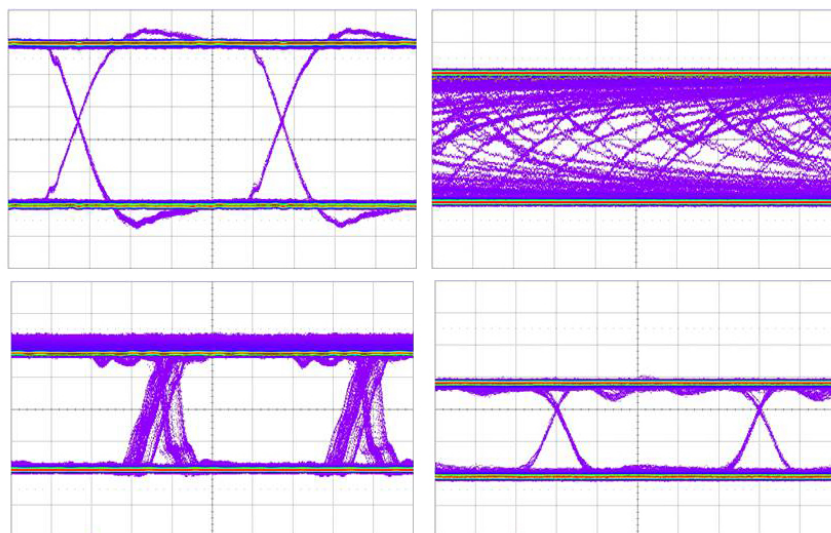


Figure 25. Oscilloscope traces for the generator signal (upper left), received signal at the end of 100 m of 28 AWG shielded twisted pair (upper right), recovered signal after passive cable compensation (lower left) and after active cable compensation (lower right). All traces are shown at 1 V, 5 ns per division.

cable length, as well as differences in particle flight time, so that the rising edge of the clock is synchronous with the earliest arrival of particles at every part of the detector. Because the clock is fanned out to the Front End boards at the TTC Patch Panel, the fine delay of the clock has to be implemented at the Patch Panel level. This function is realized using the CERN-developed Delay25 chip, which can provide from 0 to 25 ns of delay for any signal in 0.5 ns steps. The Delay25 chip [14] has 5 delay channels and is controlled by an I2C interface. The I2C interface in this case is driven from the TTC board using spare lines in the same 50-pair cable that carries the Front-end control signals.

5.3 Temperature read-out

Another function of the TRT-TTC Patch Panel is temperature read-out for the detector and Front End boards. NTC and PT1000 sensors mounted on the Front End boards and the detector are connected to the Patch Panel board via shielded twisted pair wires that run along with the Front End control lines in the Type 2 cables. These are read out by means of the ELMB mounted on the Patch Panel board. The ELMB has a multiplexed radiation tolerant precision ADC and this device is used extensively in ATLAS for DCS and monitoring purposes. The ELMB communicates with its host computer via CAN bus (automotive industry standard) carrying the data over CANOpen protocol. The temperature readout implemented on the TTC Patch Panel covers 28 NTC sensors and 36 PT1000 sensors per ELMB. A reference voltage generated on the ELMB is applied to all of the sensors and the ELMB monitors the mid-point of a voltage divider formed from the remote temperature sensor and an on-board reference resistor. Power for the digital part of the ELMB along with its CAN control is supplied from the CAN bus and is isolated from the rest of the Patch Panel by optocouplers internal to the ELMB.⁴ The analog part of the ELMB is supplied from the same power source that supplies the Patch Panel so that the analog part of the ELMB has the same ground potential as the TTC Patch Panel board. There is also an external connector on the TTC Patch Panel board which allows external connection to the NTC temperature sensor voltages. This external connection is used by the temperature interlock system which is discussed in section 7.3.2 of this paper.

6. TRT-ROD Patch Panel

The Patch Panel board associated with the TRT ROD receives data over copper lines from a set of Front-end boards and passes that data to a ROD over optical fibers as shown in figure 26. Up to 120 40 MHz LVDS signals come in from the Front End boards over 36 AWG shielded twisted pair copper wire, and that data is sent out to the TRT ROD via four 1.6 Gb/s optical links using an 850nm VCSEL from Honeywell which has been radiation qualified by the ATLAS muon group. The optical power coupled into the fibre is about -3 dBm and the receiver works up to -15dBm. The data are serialized by the GOL (Gigabit Optical Link) chip which was developed by the CERN microelectronics group [15] specifically for high radiation LHC applications. Serialized data are then driven over each optical link to the ROD. we are

6.1 Inputs and outputs

The inputs to the ROD Patch Panel from the Front End boards consist of up to 120 LVDS signals (one per DTMROC) coming in over 10-14 m of 36 AWG shielded twisted pair wire. These signals are shaped with a passive eleven component differential RLC filter (see section 5.1) that provides near optimal cable compensation, and are then latched by a comparator before being sent to the GOL. Data from up to 31 of the digital inputs are concatenated into and serialized as a 1.2 Gb/s stream onto a single optical link,⁵ then sent from the Patch Panel area to the ROD in the USA15

⁴Note that the ELMB with all its components, including the optocouplers, has been extensively tested for radiation tolerance and is fully qualified for use in the ATLAS cavern except in the actual collision region [13].

⁵Note that the GOL link itself runs at a significantly faster bit rate — 1.6 Gb/s — partly because the GOL is designed for 32 inputs and partially because of the link overhead.




Figure 26. Block diagram of the ROD Patch Panel showing connections to and from the Back End (the TRT-ROD) on the left and the Front End electronics on the right.

counting area. Each of these links is driven by a laser diode, which is itself driven by the CERN designed GOL serializer chip. In addition to the data path, there is a set of two ethernet cables running between the ROD and the ROD Patch Panel which deliver a clock for the GOL chips, as well as inhibit and I2C control signals for all the chips on the board. The clock and control signals are received by an active equalization circuit, with the expected cable length tunable in steps of 10 m by a jumper on the board.

6.2 Concatenation and serialization

The concatenation and serialization stage on the TRT ROD Patch Panel is achieved by using the GOL (Gigabit Optical Link) chip. The design idea for the chip was to allow serialized transmission of data at 800MHz or 1.6 GHz over copper wire or optical fiber. Depending on the operating mode, the chip accepts up to 16 or 32 simultaneous 40 MHz inputs with the output being chosen dependent on the type of transmission line that will be used. In addition to the GOL chip itself, there is a set of custom chips provided by the CERN Microelectronics group [16] that go along with the GOL to ensure its proper functionality. These chips are the CRT4T for power switching (required for proper GOL startup due to a bug which can cause the chip to lock up if power is not applied in a precise order), a custom rad-hard QPLL [17], tuned to LHC frequencies for stabilizing the clock, and the Delay25 chip for setting a proper phase of the GOL control clock with respect to the incoming data. All of these chips are controlled through an I2C interface which is driven by the ROD board.

7. Interlocks, controls and DCS

7.1 Overview of architecture of the DCS

The TRT Detector Control System (DCS) controls, monitors and supervises operation of the TRT and related apparatus. The main DCS design aims are:

- monitoring, loading, logging and setting of parameters
- receiving commands from the central DCS system
- issuing commands for certain actions — mostly operational
- correlating parameters from different parts of the detector
- collaborating with the DAQ system via the ‘Run_control’ layer
- supervising the safety of the detector
- triggering alarms, emergency procedures etc

It is necessary that the status of all the detector components is always available to users and that user intervention is straightforward where and when permitted. Because the DCS of the TRT is a part of the overall ATLAS DCS, it will also accept general ATLAS commands and return the summary status of the TRT subsystem to the central ATLAS DCS. Additionally, the DCS was designed to supply a user interface in order to allow separate operation during the construction, testing, commissioning and calibration of the TRT detector. To facilitate the overall integration and operation of the central DCS this interface needed to be identical to the one used by the central operator, with variations only as required for the specifics of the TRT detector.

7.1.1 Components

The TRT detector control system monitors and controls the following subsystems.

1. Detector temperature measurement system
2. Gas systems
 - (a) Active gas delivery and control
 - (b) Gas Gain Stabilisation System (GGSS)
3. Power supply systems
 - (a) LV system for the TRT detector
 - (b) HV system for the TRT detector
4. Cooling systems
 - (a) Gas cooling of endcap detectors and barrel ventilation
 - (b) Monophase liquid cooling (Front End electronics and cables)
5. Infrastructure
 - (a) VME racks




Figure 27. Global architecture of the ATLAS DCS including control computers, connections to the infrastructure, and the ATLAS Detector Control System.

- (b) Can bus power supply system
- 6. Interlock system
- 7. Connection to centrally monitored systems
 - (a) Racks
 - (b) Environment in the experimental areas
 - (c) LHC

Figure 27 shows the DCS global architecture, assignment of the control computers, and connections to the infrastructure and ATLAS DC system.

The system is designed to use commercial and industrial standards in both the hardware and software layers as much as possible. Such an approach facilitates the design and construction phase and later allows for reasonably easy maintenance, servicing, and upgrading during the long lifetime of the experiment. Another useful tool is the so called FrameWork released by CERN IT department. FrameWork contains templates and skeletons for popular and widely used equipment all of which greatly ease the integration process.

7.2 Tools and methods

7.2.1 PVSS

PVSS II [18] is a SCADA system. SCADA stands for Supervisory Control And Data Acquisition. PVSS is used to connect to hardware (or software) devices, acquire the data they produce and use that data for task supervision, i.e. to monitor device behavior and to initialize, configure and operate those devices. In order to do this PVSS provides the following main components and tools:

A run time database in which data coming from the devices is stored, and which can be accessed for processing, visualization, and other purposes.

Archiving of data in the run-time database for long term storage, and later retrieval by user interfaces or other processes.

Alarm generation and handling such that alarms can be generated if new data arriving in PVSS exceeds predefined conditions. The alarms are stored in an alarm database and can be selectively displayed by an Alarm display. Alarms can also be filtered summarized, etc.

A Graphical Editor (GEDI/NG) which allows users to design and implement their own user interfaces (panels).

A Scripting Language which allows users to interact with the data stored in the database, either from a user interface or from a “background” process. PVSS scripts are called **CTRL** (read control) scripts and follow the C syntax and include many SCADA-specific functions.

A Graphical Parameterization tool (PARA) which allows users to:

- define the structure of the database.
- define which data should be archived.
- define which data, if any, coming from a device should generate alarms
- etc.

Drivers which provide the connection between PVSS and the hardware or software devices that are to be supervised. Common drivers that are provided with PVSS are OPC, ProfiBus, Can-Bus, Modbus, TCP/IP and Applicom. A DIM driver is provided as part of the Framework.

PVSS has a highly distributed architecture. A PVSS application is composed of several processes called “Managers” in PVSS nomenclature. These Managers communicate via a PVSS-specific protocol over TCP/IP. Managers subscribe to data and data is sent over the link (only on a change value) by the Event Manager. The Event Manager is the heart of the system shown in figure 28.

The Event Manager (EVM) — is responsible for all communications. It receives data from Drivers (D) and sends it to the Database Manager to be stored in the data base. However, it also maintains the “process image” in memory, i.e. the current value of all the data. In addition it ensures the distribution of data to all Managers which have subscribed to this data. The DataBase Manager (DBM) — provides the interface to the (run-time) data base. User Interface Managers (UIM) — can get device data from the database, or send data to the database to be sent to the devices, they can also request to keep an “open” connection to the database and be informed (for example to update the screen) when new data arrives from a device. There is a UI for Linux and a Native Vision (NV) for Windows. The UIM can also be run in a development mode; PARA for parameterization of DPTs/DPs, GEDI for the Graphical Editor (GEDI for Linux and NG for Windows). Ctrl Managers (Ctrl) — provide for any data processing as “background” processes, by running a scripting language. This language is like “C” with extensions.

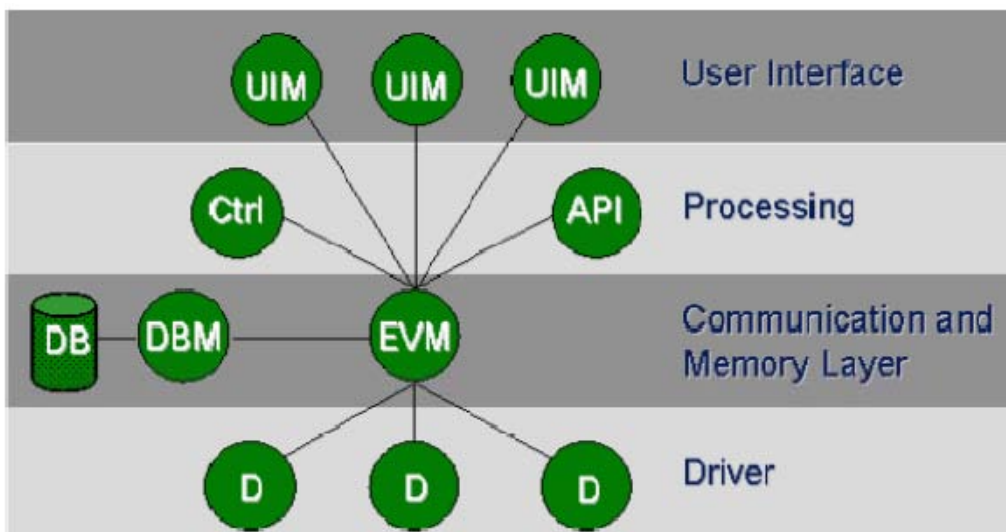


Figure 28. Event manager architecture showing the various layers.

API Managers (API) — Allow users to write their own programs in C++ using a PVSS API (Application Programming Interface) to access the data in the database. Drivers (D) — Provide the interface to the devices to be controlled. These can be PVSS provided drivers like Profibus, OPC, etc. these will be described later in more detail, or user-made drivers. Archive Managers — Allows users to archive data for later retrieval and viewing. A project, which is the name for a PVSS application, may have one or more Archive Managers and one can configure which data is stored in which manager.

7.2.2 Finite State Machine (FSM) methodology

We have built a system where the system behaviour is described in terms of “objects” and the “states” of these objects. Such an approach is useful since it abstracts the problem from the details of the system and the underlying hardware. This separation allows a clear definition of the logical states of the various parts of the system. An additional parameter describing system health, a quality qualifier, is a “Status” object associated with every “State”. Each granule of the system can be seen from outside as an object — an entity with certain behaviour and characteristics. An object can be in one of its allowed states which can be set either by command or by a change of an external input. The meaning of states with the same names can be different depending on the object definition. The “Status” object describes different quality factors of a given state. Some combinations are obviously excluded (e.g. “READY” state with “FATAL” status). The set of commands used in the system is hierarchical and is adapted to the level of access. Some commands are restricted to a certain categories of user and are password protected. Each command calls for certain actions, which should at lowest level reach the smallest addressed objects of the system. Not all commands change the object state e.g. those changing parameters or reading status.

As shown in figure 29 the HV system states which easily can be mapped on any of the objects of which TRT detector is composed. It is clear that some objects will not attain certain states due to their internal nature. Some of the states are “mandatory” i.e. every object has to be in one of

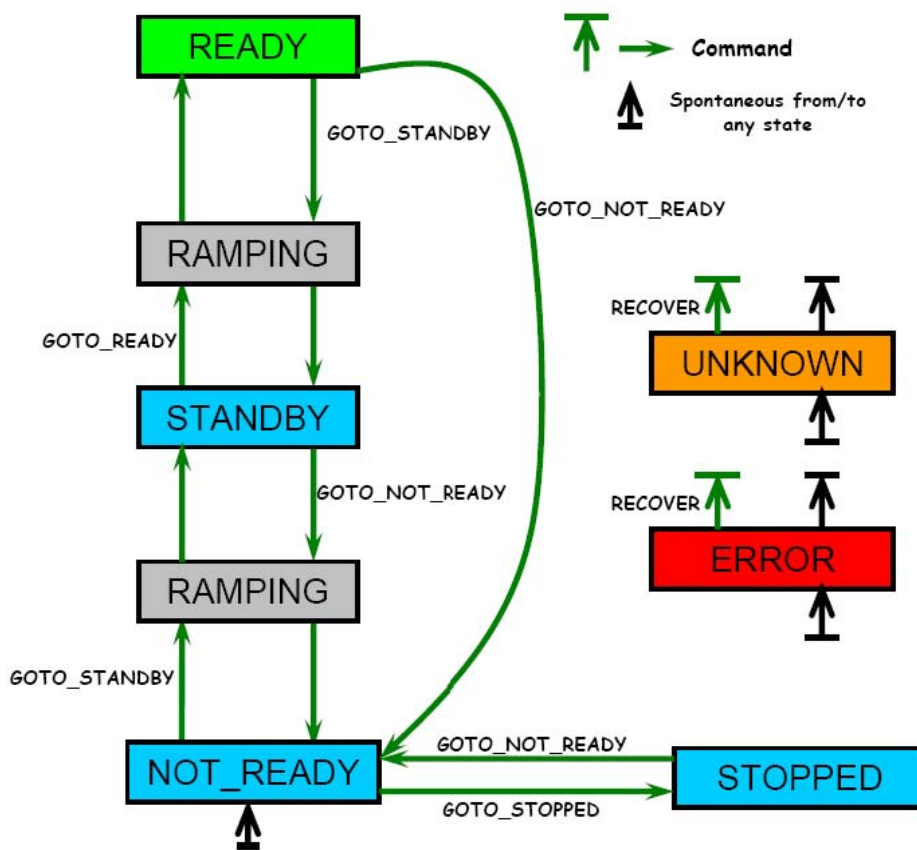


Figure 29. The TRT High Voltage system state machine.

this states; others are “optional” i.e. depending on the characteristics and nature the object can be in one of this states.

The control hierarchy of the objects is composed of the control units, logic units and device units. The first two have very similar properties except for the lack of a stand-alone operation mode for logic units. The device units represent direct connections to the hardware and need to reflect all possible behavior of the apparatus. Whenever a “parent” in a hierarchy receives a transition command it sends that command to all the “children” in its branch. Any “child” may propagate the command and so on throughout the system. A “parent” (at any level of the hierarchy) confirms the transition only when all (or, in some cases, a majority of) the “children” have confirmed the new state. Therefore any errors occurring during the performance of a state change are effectively propagated to the root node

7.3 Detector safety

7.3.1 DCS protections

The DCS system monitors all slowly changing parameters which have some influence on detector performance. However, the dense concentration of electronics in a confined volume puts strin-

gent requirements on the cooling system. Any unexpected or uncontrolled rise in the temperature of either electronics or detector mechanics could lead to severe detector damage. Therefore, the detector control system implements limits on the all of measured temperature values. There are two separate limit values. A “warning” (set at 2-4 degrees above the normal working temperature) serves as a signal to the operator that something is wrong. An “alarm” (set 7-10 degrees above the working temperature) turns off the low voltage to the affected region of the detector electronics. Any such incident is also handled by the FSM system and by the central DCS alarm screen. A similar philosophy is applied to the power supply voltages and currents where any channel exceeding the alarm value is switched off.

7.3.2 TRT hardware interlock

Backup protection against overheating of the TRT front-end electronics, beyond that provided by the DCS, is guaranteed by the TRT Hardware Interlock (TRT-HWI). This system of custom electronics monitors temperatures on the front-end boards read out using Negative Temperature Coefficient Thermistors (NTC), which also provide inputs to the DCS system. If a specified number of NTCs register temperatures above pre-determined thresholds for a user-adjustable period of time, the system will switch off the Low Voltage (and also High Voltage, if desired) to the entire TRT. In addition, auxiliary inputs to the TRT-HWI are also provided to allow other conditions (for example, the proposed DCS watchdog status) to cut power to the TRT. The TRT-HWI is designed to be as fail-safe as possible: using no programmable logic or remotely settable parameters, while being monitored but not controlled, by the DCS.

A block diagram of the major elements of the TRT-HWI is shown in figure 30. The system itself is comprised of two elements:

1. *Comparator Boards*, sitting on the PP2 boxes in UX15, which compare NTC outputs in 1/32 of each end of the detector to pre-determined thresholds and send out current approximately proportional to the number of NTCs over threshold in the barrel and the endcap
2. *the Logic Box*, located in USA15, which is itself composed of two elements: *Receiver Boards* that receive all of the outputs of the Comparator Boards for the A and C sides of the detector, and “count” Comparator Boards with more than a specified number of NTCs over threshold; and an *Alarm Board* that uses the “counts” from each Receiver Board, as well as auxiliary inputs to generate a KILL signal based on a simple algorithm.

The TRT-HWI KILL signal is sent to the LV and HV control racks in USA15, disabling their outputs and thus switching off power to the entire TRT. The system is designed so that a failure in any component — NTC, cable, board — will contribute to (in the case of a single NTC failure) or produce a KILL signal.

Comparator Boards. Each of the 64 Comparator Board receives input from a total of 28 NTCs mounted on TRT front-end electronics boards: eight from the barrel and 20 from the endcaps. These NTC signals are routed to the Comparator Boards from the TTC PP2 boards where they are passively split off from the main TTC DCS path. On the Comparator Boards, the output signal from the NTCs is compared to thresholds using rad-tolerant comparators (LM339). The input to




Figure 30. Block diagram of the TRT-HWI system.

the comparators from the NTCs is pulled down to ground, so that both shorted and disconnected NTCs appear, to the comparator, to be in a high-temperature state, thus potentially causing a KILL signal to be generated.

Thresholds for all barrel and endcap NTC comparators can be set independently using two socketed resistors on the Comparator Board. This scheme allows the thresholds to be adjusted if conditions change, but prevents them from being set remotely to incorrect values. The state of each comparator (above or below threshold) is indicated by an LED.

The current outputs of all eight barrel comparators and all 20 endcap comparators are summed separately to form the two primary board outputs: BR_SUM and EC_SUM. These sums contain a small current offset to allow disconnected cables to be detected at the Logic Box.

If a given NTC is known to be bad, its comparator output can be individually removed from the sum (disabled) using a set of switches mounted on the board.

The BR_SUM and EC_SUM outputs can be used in two ways, although both methods will not be used simultaneously.

1. They can be sent to the *Logic Box* for use in constructing the global KILL signal.
2. They can be sent, via another set of comparators, to disable inputs on the PP2 MARATON power supplies (see section 12 switching off LV power to the front-end boards controlled by the PP2 in question (LOCAL_KILL)).

Outputs to the Logic Box are sent differentially using two wire pairs each (for BR_SUM and EC_SUM) on ~100 m ethernet-style cables. The alternate LOCAL_KILL signal, which is used during commissioning, is generated by two other comparators. The threshold for each of these can be set using socketed resistors. Two LOCAL_KILL signals (for barrel and end-cap) are sent out

through LEMO connectors. The end-cap signal is split with a simple “T” to accommodate the two end-cap LV boards in a PP2 crate.

The Comparator Boards, which are fabricated as 75×125 mm, four layer PCBs, are collected at each PP2 location into mechanical housings, which can hold five or eight boards. This housing is screwed onto the side of the PP2 structure. The Comparator Boards receive power, via a daisy-chain connector, from the same LV supply used to power the PP2 TTCs.

The Logic Box. After traveling from UX15 to USA15, Comparator Board output signals are received by instrumentation amplifiers on one of two *Receiver Boards* that are part of the *Logic Box*: one of which deals with signals from side A, the other for side C. After reception each Comparator Board output sum serves as input to two comparators. One comparator detects whether the cable from that Comparator Board is disconnected, and the other compares the signal to a multiplicity threshold, set separately, by socketed resistors, for barrel and endcap sums. The results of the two comparators for each Comparator Board input are summed to give two output signals, BR_MULT_SUM and EC_MULT_SUM corresponding to the number of Comparator Boards in the barrel and end-cap that have more NTCs over threshold than that specified by the multiplicity threshold.

Individual Comparator Board inputs can be excluded from this sum (disabled) using switches on the Receiver Board. Because this action would leave a significant portion of the electronics unprotected, several levels of monitoring have been added to the Receiver Boards including LED visual indicators and independent monitoring by DCS to ensure that disabled boards do not go unnoticed.

The BR_MULT_SUM and EC_MULT_SUM signals from each Receiver Board are daisy chained (effectively summing them from each board) and sent by a short ribbon cable to the Alarm Board, where they are compared against barrel and end-cap board-count thresholds, again set by socketed resistors. If any of the barrel, end-cap, or auxiliary inputs are over their respective thresholds, an ALARM condition is generated. To reduce sensitivity to transient noise in the system, this ALARM condition is only translated into the KILL signal if it persists for a specified period of time (set by an RC circuit, and currently ~ 45 s). The KILL signal activates a relay, which produces the signal to turn off power to the TRT, via a cable to the LV (and possibly HV) racks in USA15.

Critical voltages and signals (the various SUM, ALARM, KILL, DISABLED, etc.) are monitored by a set of five ELMBs. One ELMB monitors the Alarm Board states and settings and for each side (A and C) of the detector. Two ELMBs monitor the _SUM voltages and disabled states of the 32 sectors. A PVSS panel integrated into the DCS framework monitors and logs all abnormal conditions.

In order to be able to power up the TRT after a power failure or a trip, it is necessary to temporarily disable the TRT-HWI. This reset can be accomplished in two ways. First (and normally) a push-button disable is provided on the front of the Alarm Board. When pressed, the button disables the KILL output signal for approximately five minutes, after which time the signal is automatically re-enabled and can again cause TRT power to be cut. For those occasions when a longer period of inactivity is desired, a “hidden” switch is provided that disables KILL until the switch is returned to its enabled position.

Part III

Back End electronics, DAQ and power

The TRT Back End electronics is housed in ten 9U and two 6U crates located in the USA15 counting hall. As shown in figure 31, the 9U crates which hold the TRT specific TRT-TTC and TRT-ROD modules are arranged in four racks corresponding to the four logical (and physical) partitions of the detector - End Cap side A, Barrel side A, Barrel side C, and End Cap side C. The central rack houses the general ATLAS trigger hardware with two partitions per 6U crate.

Each 9U crate contains multiple three card groups of readout electronics. Each “group”, as shown in figure 32 consists of one TRT-TTC and two TRT-ROD cards. All cards are configured and monitored via a single board computer (SBC) over the standard VME bus but communicate trigger and timing information within their local group via a custom P3 backplane. Each group controls up to 40 logical Front End boards and receives DTMROC data from those same 40 boards. The crates are organized as shown in tables 5 and 6 below.



Figure 31. Photo of the two leftmost of five TRT Back End racks in the USA15 counting hall. These two racks house the three 9U crates for End Cap C and the two 9U crates for Barrel side C. The next rack in line house the two 6U ATLAS TTC crates and then two more racks for the side A Back End crates. The electrical and optical cable plant obscures most of the electronics.




Figure 32. TRT 9U crate-based backend electronics “group” block diagram. .

Table 5. End cap crate configuration showing the number of TRT-ROD and TRT-TTC cards per crate and the number of DTMROC readout links received and the number of Front End boards controlled per crate. The configurations are identical for sides A and C - six crates total for the end cap readout.

Crate	RODs	TTCs	DTMROCs	FE Boards
1	10	5	2400	200
2	10	5	2400	200
3	12	6	2880	240

Table 6. Barrel crate configuration - identical for sides A and C. Note that “FE Board” here means logical board. The three largest triangle boards have two TTC Groups per physical board.

Crate	RODs	TTCs	DTMROCs	FE Boards
1	8	4	1664	144
2	8	4	1664	144

8. Timing, Trigger, Control (TTC)

The ATLAS timing, trigger and control signals⁶ are distributed by the RD12 [19] optical TTC system. In the inner tracker, this system is used up to the level of the read-out drivers (ROD). At that point there is a protocol change in order to minimize the functionality required in the Front-

⁶The LHC bunch crossing clock [BX], level-1 accept trigger signal [L1A], synchronisation signals [bunch and event counter resets, BCR and ECR], tests and calibration pulses, resets, etc.




Figure 33. The TRT-TTC VME Module.

End. This protocol conversion is accomplished in the TRT-TTC module [20] housed in the Back End VME crates. The TRT-TTC module shown in figure 33 interfaces the backbone TTC system to:

- The TRT-RODs, providing them with the BC, L1A, BCR, ECR, event identifier, bunch crossing identifier, and the trigger-type word;
- The Front-End electronics read-out chips (DTMROCs), providing them with a hard-reset signal and dedicated versions of the ATLAS timing trigger and control signals.

The TRT-TTC module also receives commands via VME bus. It accepts triggers via the VME bus, the TTC optical link and the module front panel.

8.1 TRT-TTC segmentation

The TRT-TTC signals are sent to each of the local geographic zones of the detector. Zones consist of 1/32 of an 8-plane wheel (a single “Triplet” board) in the end-cap region and either one half or all of a triangular Front End board in the barrel. Each “zone” receives three signals (Clock [BX], Command_In, and Reset) from the TRT-TTC and returns one signal (Command_Out) to the TRT-TTC. This collection of four signals is sometimes referred to as a “TTC Group”.

Barrel. For the Barrel, one TRT-TTC board communicates with the Front-End boards of four modules of each type (M1, M2 and M3). As each (M1+M2+M3) is 1/32nd of one end of the Barrel (9 TTC Groups), one TRT-TTC board controls 4/32nd of one end of the barrel (36 TTC Groups total).

End-cap. For the End Caps, one TRT-TTC communicates with the Front-End boards of 2/32nd of each wheel A and each wheel B of one side of the detector (40 TTC Groups total).

8.2 Modes of operation

The TRT-TTC module provides a number of different modes for writing and reading back parameters and status bits to and from the Front-End. These modes are initiated by VME commands and, in some cases, use LHC beam structure signaling provided from the ATLAS TTC system.

BASIC. The BASIC mode uses direct access from VME. This mode which is used for system tests.

INIT. The INIT mode writes all parameters to the Front-End on all TTC links simultaneously. These parameters are stored in on-board memory.

POLL. The POLL mode reads all parameters from the Front-End to on-board memory. These parameters are optionally compared to values written with the INIT mode and any discrepancies are flagged for further action by the DAQ. The POLL mode can be used to read the parameters during the $3 \mu\text{s}$ LHC beam gap. This mode works for all parameters except for the 144 bit test/mask register. All other registers take $1.5 \mu\text{s}$ or less to read out.

REFRESH. The REFRESH mode is similar to the INIT mode in that all of the Front-End parameters are written from memory. However, the writing is done serially, one parameter at a time, but in parallel across all forty links and the “writes” take place during the beam gap periods.

The POLL and REFRESH modes are used to detect and correct any Front-End register content corruption which might be expected from single event upsets during high luminosity running.

DTMROC chips can be switched to the Fast OR mode for cosmic ray trigger purposes. In this mode, the TRT-TTC modules receive Fast OR signals on the parameter read-back lines and form a trigger using simple on-board logic. For more complicated trigger building the TRT-TTC propagates the Fast OR lines to the VME P2 connector where a dedicated extension module can be installed.

8.3 Interfaces

The TRT-TTC has four main interfaces: a VME interface for software configuration; an optical link with the ATLAS-TTC system; an electrical connection with the DTMROC Front End; and finally a VME (P3) connection with the TRT RODs. In addition there are front panel auxiliary interfaces useful for testing and setup.

VME interface. The TRT-TTC module is a 9U VME64x board, using the J1, J2 and J3 connectors. The user pins of the J2 connectors are dedicated to the 'FAST-OR' signals described in the previous section. The J3 connector is fully dedicated to the communication between TRT-TTC and its two adjacent RODs. The TRT-TTC board uses the VME 32 bit address and 32 bit data mode to communicate with the VME crate processor. The VME interface allows for communication with the Front-End boards via a change of protocol inside the TRT-TTC, for reading and writing of the status and control registers of the TRT-TTC and for generation of test signals such as trigger, test pulse or resets for calibration or system test purposes.

ATLAS-TTC interface. The TRT-TTC board is equipped with a photodiode and a TTCrx chip [21], configurable via the VME bus. The TTCrx chip decodes the incoming ATLAS TTC signal and delivers the Bunch Clock, Level-1 Accept Trigger, Bunch Counter Reset, Event Counter Reset, data (eg, trigger type) and commands (eg TestPulse) to be transmitted to the Front-End boards and to the ROD modules.

Front-End interface The Front-End electronics (and in particular the DTMROC chips) need to receive the following information from the TRT-TTC module: Clock, Level-1 Accept Trigger, Event Counter Reset, Bunch Counter Reset, Test Pulse (for calibration purposes), Hard Reset and read and write commands for the parameters to configure the chip's registers.

These Front End signals and values are transmitted via a serial protocol which is carried on four lines (Command_In, Command_Out, Clock and hard Reset - a TTC-group). Forty separate Front-End boards (each having between 9 and 15 DTMROC chips) can communicate independently with one TRT-TTC. The clock and hard-reset lines are common for ten Front-End boards and are fanned out at the TTC Patch Panel (see section 5). However, each Front-End board has its own dedicated command and read-back lines. The TRT-TTC also controls the I2C interfaces to the TTC Patch Panel in order to remotely adjust the Clock phase for each Front-End board via TTC Patch Panel mounted Delay25 chips.

The communication to the Front-End boards is carried on four 25 pair cables from 60 to 100 m long. The large difference in length of the installed cables requires a set of timing and transmission adjustments. In addition to the 0.5 ns step clock phase adjustment on the TTC Patch Panel, the duty cycle of clock signal can be adjusted on the TRT-TTC as can the clock polarity. Each command line can also be phase-shifted, in 0.5 ns steps, and the read-back lines can be latched at any of four different phases, each 6 ns apart.

The duty cycle and polarity of the clock are chosen to optimize Front-End stability. The duty cycle and polarity can be set independent of the physics constraints. The remaining parameters (remote clock phase, command data phase, and read-back latching phase) are tightly coupled, and must be chosen in a specific order.

First, the choice of the remote clock phase shift is driven by the need to see the full pulse of a minimum ionizing particle as digitized by the DTMROC. Shifting the clock at the DTMROC shifts the DTMROC readout window, allowing one to align the 75ns window such that both the leading and trailing edges of the analog pulse are seen.

Next, the choice of the command data phase and read-back latching phase is driven by the choice of the remote clock phase. An appropriate choice of the command data phase allows the DTMROC to latch incoming commands with respect to the clock supplied by the Patch Panel. Similarly, the read-back latching phase is chosen so that data sent by the DTMROC to the TRT-TTC on the command-out line are latched correctly by the TRT-TTC. Both the command data phase and the read-back latching phase are set locally on the TRT-TTC.

A more thorough description of the tuning of clock and command phase shifts, and the physics constraints that guide that tuning, can be found in [22]. An algorithm for the prediction of specific phase shifts based on those constraints is described in [23].

For data alignment purposes, the trigger and reset signals can be delayed by an integer multiple of clock periods. These adjustments are done independently for each Front-End board.

ROD interface. The TRT-TTC board communicates with its two adjacent ROD boards, via the J3 connector of the VME backplane. The TRT-TTC transmits the trigger, the synchronization related signals, (Bunch Counter Reset and Event Counter Reset) and for every Level-1 Accept the corresponding trigger ID, Trigger Type, and Bunch Clock ID. The TRT-TTC receives from each of the two RODs a BUSY signal to be retransmitted to the Central Trigger Processor via the partition

BUSY module. A single coaxial cable runs from each TRT-TTC to the BUSY module for that partition.

Auxiliary interfaces. For test beam and system test purposes the TRT-TTC can receive NIM level signals from a front panel connector. For dead time monitoring purposes TRT-TTC measures the duration of the BUSY signals generated by the SLINKs which send data from the RODs to the ATLAS Level 2 buffers. The TRT-TTC on-board buffer control logic generates a buffer signal at near full conditions. The combined BUSY signal is sent to the Central Trigger Processor through the ATLAS VME BUSY module.

9. Readout drivers (ROD)

The TRT Readout Drivers (RODs) are the data processors of the TRT Data Acquisition system. A TRT-ROD connects to the Front-End electronics via data Patch Panels, to TRT-TTCs for triggers and clocks, to the Level 2 Readout Systems and to VME-based readout single board computers that control each readout crate, as shown in figure 32. The main function of the TRT-ROD is to receive data from the Front-End boards via the data Patch Panel (see section 6), check the validity of the data, compress the data, format and build the data into mini event fragments and output the event to the Readout System (ROS). In addition, the RODs allow sampling of the events via VME-based single board computers (SBCs) for in-crate monitoring. The RODs also send clock and control information to the data Patch Panels, receive trigger information from local TRT-TTCs (see section 8) and assert BUSY when the readout system cannot accept any more data.

The TRT DAQ contains 96 RODs in total, 32 for the TRT Barrel (16 per end), each handling 2/32 of one end of the detector, and 64 for the Endcaps, with each ROD serving 1/32 of an endcap unit.

Each ROD is connected to two identical Patch Panels, and much of the ROD is divided into two logical sections, one for each data source. The ROD interfaces to each Patch Panel via four 1.6 Gb/s optical transceivers and two Cat5 ethernet cables containing clock and control information, for a total of eight optical links and four ethernet cables per physical TRT-ROD. The optical transceivers are used for local loopback tests, but during normal data taking only the receiver portion is used.

The TRT-ROD is a 16 layer, 9Ux400mm deep single-width VME-based board as shown in figure 34. All external interfaces are located on the front panel. Most of the ROD functionality (see figure 35) is implemented within 5 Xilinx FPGAs. An XC3S400 Spartan 3 FPGA handles the VME-bus protocol and translates it to a simple local bus for communication with the other FPGAs on board. A second XC3S400 implements the interface with the TRT-TTC, and handles other miscellaneous tasks. A Virtex II Pro XC2VP7 is responsible for decoding the eight optical input data streams, providing test data for checking the data processing and interfacing to the output SLINK daughtercards. Finally, two Virtex 4 XC4VLX40 chips each handle one half of the data processing (one for each data Patch Panel connected), control the data buffering and compression, check for errors and build the data into formatted event fragments.

9.1 ROD-TTC communication

Each TRT-TTC communicates to two RODs via a custom J3 backplane. The timing, trigger and busy signals are sent and received from each ROD using dedicated point-to-point connections (see




Figure 34. Photograph of the TRT ROD board.

section 8). The ROD receives the LHC clock from the TRT-TTC, which is then passed through the on-board QPLL to reduce jitter. Upon receipt of an L1A, the TRT-TTC sends the ROD the L1A signal, L1 identification counter (L1ID), Bunch crossing identifier (BCID) and Trigger Type via a 16-bit data bus. These signals are used to ensure that each subdetector's data is synchronized with the rest of the detector, and to check each Front-End chip is reading out the correct bunch crossing's data. If the ATLAS readout system asserts BUSY via the SLINK, the ROD asserts BUSY to the TRT-TTC which alerts the trigger system to suppress further triggers. In addition, the ROD can assert a separate BUSY if it needs to halt triggers for internal reasons.

9.2 ROD-Patch Panel communication

Each ROD is connected to two independent data Patch Panels that service up to 120 Front-End DTRMOC chips. The data from up to 31 DTMROCs⁷ are concentrated onto one optical fiber, resulting in four fibers per Patch Panel, each carrying 1.2 Gbits/s of data. This data is latched and 8b/10b encoded by four GOL [15] chips operating at 40 MHz, so the actual data rate of each fiber is 1.6 Gb/s. The GOLs require low-jitter clocks, and each group of up to 31 data lines has slightly different timing. To maximize the system stability, the ROD transmits the phase-locked 40MHz

⁷The widely varying number of DTMROCs per board in the Barrel case resulted in one GOL with 31 inputs with the remaining 73 inputs spread over the remaining GOLs in order to concentrate all DTMROCs from a single TTC group in a single GOL. In the End Cap case each GOL has exactly 30 inputs.




Figure 35. TRT ROD functional block diagram.

LHC clock via one of the Cat5 ethernet cables to the Patch Panel. Each ethernet cable has 4 twisted pairs of wire for carrying signals. The clock from the ROD is passed through a 4-channel Delay25 chip on the Patch Panel, allowing each optical link to have an individual phase adjustment of its clock for maximum timing headroom. The configuration of the GOL and Delay25 chips on the Patch Panel is implemented via an I²C interface carried on the ethernet cables. In addition, voltage inhibits for power sequencing, a patch-panel reset signal and an optical IDLE pattern enable are also transmitted over the cables.

The Patch Panel logic is implemented within the Spartan 3 FPGA that also handles the TRT-TTC interface.

9.3 Data flow

Each data Patch Panel reads out data from either 120 (Endcaps) or 104 (Barrel) DTMROCs, so each ROD handles either 240 or 208 inputs. In either case the RODs are identical — 32 channels are simply unused when operating in Barrel mode. Each DTMROC transmits 12 bits of header information and the data from 16 straws, at 25 ns/bit, meaning each ROD handles the data from up to 3840 straws.

The data from the eight optical inputs are received using Finisar FTRJ8519 transceivers and passed to a single Xilinx Virtex II Pro which has eight RocketIO⁸ blocks. Each RocketIO block can independently align, decode and perform a serial to parallel conversion on its data stream. After deserialization, data from each Patch Panel is split and sent to one of two Virtex 4 FPGAs for further processing.

The Virtex 4 FPGAs must first extract the header information from each of the 120 data streams from the Front-End DTMROCs, and further deserialize the following 16 straw's data into 16 27-bit wide words. At this point the data is stored into DRAM for buffering, as the time taken to compress this data varies from event to event depending on its complexity.

The DRAM consists of four 16-bit wide 256 Mbit DDR (dual data rate) memory modules with a maximum frequency of 200 MHz organized into a 64-bit wide bus. Raw data is written two 27-bit words at a time. The controller is implemented in the Virtex 4 FPGA, with read and write scheduling optimized to minimize latency penalties due to mode switching, while ensuring priority for data writes when needed. As the data is read/written on both the rising and falling clock edges, the bandwidth in/out of the memory buffer is 800 M 27-bit words/s, less the time spent in refresh cycles. Since each ROD half must process approximately 200 M straw words/s at the maximum Level 1 trigger rate there is ample memory bandwidth.

After retrieval of the event data from DRAM, the header information is checked against the L1ID and BCID received from the TRT-TTC to ensure data synchronization across the detector. Errors, e.g. channels without data, are flagged and built into an error block which is appended onto the event data during the event building phase. The data words are first sent to the compression engine described below before being passed to the event builder.

Once the full event fragment, with header, data block and error block has been assembled, it is sent back to the Virtex II Pro chip which contains the code to interface to the SLINK readout into the ROS. There are two SLINKs daughtercards on board (one for each half-ROD's data) that can each transmit at 160 MBytes/s via a duplex optical fiber. The ROS can inhibit data output by asserting Link Full, in which case the ROD stops sending data and asserts BUSY to the TRT-TTC module. Data flow is immediately resumed once the ROS de-asserts Link Full. In addition, a fraction of the events may be stored into DRAM as "spy" data, which can be read out on demand over VME to the single board computer in the crate for monitoring.

⁸Xilinx Inc., proprietary transceiver logic blocks embedded in Virtex II and 4 FPGA parts, capable of 622 Mbps to 6.25 Gbps rates.

9.4 Compression

The data from each of the 350k straws is read out from the DTMROCs on every Level 1 trigger, at a rate of up to 100 kHz. Each straw's data consists of 27 bits for a total of 1 Tbit/s of data. A Tbit/s is beyond the bandwidth of the subsequent parts of the DAQ system, as well as the storage system, and so a data compression/zero suppression scheme must be implemented. By studying the entropy of the data, i.e. the actual information content of the digitized data, it was found that a completely lossless compression scheme is possible, and fits well within the available bandwidth of the DAQ chain. Since the entropy of the data stream varies logarithmically with the complexity of the data, there is only modest growth in the data volume from low to high luminosity conditions, and one compression works scheme for both cases. This compression also allows the monitoring of the detector performance via the VME-based readout at a higher rate than otherwise would be possible.

The entropy-based compression scheme is a Huffman encoding of the data [24]. While each straw's data is a 27-bit pattern, the number of bit patterns that actually occur is a very small fraction of the total number of 2^{27} possible patterns. Ideally, one expects the tracking data from each straw to be quite simple. For example, the data is all zeros if there was no particle crossing the straw, and no noise pickup. Since the per-straw occupancy varies from a few percent at low LHC luminosity to approximately 50% at high-luminosity in the inner regions, this is the most common pattern. If a particle did intersect the straw, the data bits are zero for the time-slices before the closest ionization reaches the sense wire. Once the charge does reach the sense wire, the signal goes above threshold and the data bits are 1 until the farthest ionization (at the straw wall) reaches the sense wire, after which the data bits are again 0 (see section 2). In practice many more complicated patterns are present (approximately 100,000 in total), though the least likely may only occur at a rate of less than 1 in 10^8 . These bit patterns of data are ordered by their likelihood of occurring, and assigned a variable-length encoding based on this ranking. Very likely patterns get short encodings, and rare patterns get longer ones. The Huffman encoding algorithm generates the table of these encodings in a way that results in an output data stream that is very close to the theoretical minimum length to losslessly transmit the data.

The implementation of this scheme at the high rates needed (200 M lookups/s) is based upon an IDT75K62100, a specialized Network Search Engine (NSE) ASIC that was developed for internet routers. Two of these NSE chips, operating with custom firmware in the ROD's Virtex 4 FPGAs, allow for deadtimeless, lossless compression of the TRT data that approaches the theoretical limit.

The NSE is programmed via VME with the 128k most common 27-bit data patterns. A parallel 128k deep SRAM is loaded with the encoding for each pattern as well as the code length. In operation, the 27-bit data pattern from a straw channel is presented to the NSE, which searches for the table address holding the data pattern. This address is used as a lookup into the SRAM which then drives the encoding and code length for this data pattern. The variable-length code is read by the Virtex 4 chip, then packed into 32-bit words and passed to the event-builder subblock. Patterns not in the encoding table (typically very anomalous patterns that are not possible under normal behaviour of the Front-End electronics) are given a fixed prefix code followed by the full 27-bit data, or optionally mapped to preset codes marking data patterns that are unusable for physics.

9.5 Self-test features

The ROD is designed so that a good deal of its internal functionality can be tested with only a VME crate and computer. The ROD contains a QPLL and oscillator that can function as a stand alone LHC clock source for the board. The TRT-TTC interface chip can mimic Level 1 triggers in various modes. The Virtex II Pro, through which the data flows, contains firmware that can inject known test data into the system in various ways. The Virtex II can generate standard data patterns, or trigger an on-board FIFO to transmit custom data, such as data that reflects actual TRT detector data, which is loaded via VME. These data can bypass the RocketIO deserializers and test the functionality of the board's data processing. In addition, the RocketIO blocks have the ability to serialize and transmit data in addition to receiving it. The data can then be sent out the optical transmitters on board, looped back to the receivers, deserialized and checked for bit errors.

10. Data acquisition

The primary function of the TRT data acquisition system (DAQ) is the collection of data from Front-End electronics at a maximum trigger rate of 100 kHz. To this end, it is composed of hardware and software components which form an interface between TRT-specific electronics and the ATLAS trigger and data-acquisition system. In addition to the collection of physics data, the DAQ is also part of a calibration system designed to exercise the full functionality of the front- and back-end electronics. Finally, the DAQ provides the ability to monitor the detector on a variety of levels, from basic temperature and voltage sensing to real-time calibrations of noise rates and efficiencies.

10.1 ATLAS trigger

The ATLAS trigger is a three-tiered system, designed to reduce the incoming event rate of 40 MHz to a sustainable output rate of 200 Hz [25]. The first layer is hardware based, followed by two software based layers that perform region-of-interest based analyses. In the first layer, only calorimetry and parts of the muon system are used to make the Level 1 decision. The inner tracker information, including that of the TRT, is first used in the High Level Trigger (Levels 2 and 3), along with the information from the calorimeters and the full muon spectrometer. To transmit essential signals, such as the beam clock, triggers, and reset signals, a series of trigger modules have been developed by the collaboration. At the top level is the Central Trigger Processor (CTP) [26], which makes the Level 1 decision using information from the calorimeters, the muon trigger system, and the busy signals supplied by the subsystems. The CTP passes information to the Local Trigger Processors (LTP) [27], which are responsible for a single "partition". There are four partitions for the TRT, and 36 for the whole of ATLAS. The LTP information is then passed to a serializer (the TTCvi [28]), which transmits the signals to the subsystem-specific hardware via the TTCex optical encoder module [28].

10.2 Data acquisition software framework

10.2.1 ATLAS DAQ framework

The ATLAS experiment has created a data acquisition framework to facilitate the integration of the different subsystems. The system is a C++ based Finite State Machine, and uses a combination

of gigabit ethernet and fiberoptic connections to join the different parts of the experiment. Each subsystem is required to maintain its own dedicated hardware, and to supply a collection of libraries which guide the hardware through the state transitions before, during, and after a physics or calibration run. As the interface between the ATLAS DAQ and the TRT, the TRT DAQ has the responsibility of receiving the Level 1 trigger decisions, processing those triggers, and returning the corresponding data buffers to the High Level Triggers with minimal latency.

10.3 TRT software libraries and databases

The collection of libraries provided by the TRT to steer the TRT hardware make use of a second tier of low-level hardware libraries. There is one hardware library designed for each piece of configurable hardware in the TRT (DTMROC, TTC, ROD). The libraries are implemented in C/C++, and provide access to all of the functionality of the corresponding component. These libraries are used on the single board computers, and interface with a VME driver (and the associated API) to communicate with the hardware. The high-level libraries define the order of low-level function calls at each state transition, and are responsible for retrieving the configuration parameters needed to configure the hardware components. The configuration parameters for all of the electronics are stored in a database. The “configuration database” for the TRT is two-tiered. In the first tier, the TRT is presented to the ATLAS DAQ framework by way of an Object Kernel Source (OKS) database [29]. This database is an ATLAS-specific implementation of a completely relational database, which contains lists of identifiers for all of the configurable TRT hardware in a given run. These identifiers are used to generate software objects, which then access the second tier of the configuration database. This second tier is composed of an Oracle back-end and a CORAL [30] Front-End, and contains channel specific parameters that need to be archived for future analysis. These parameters include tracking thresholds, timing parameters, and status elements inserted into the data stream for calibration purposes.

10.4 Readout scheme

10.4.1 Physics data

When a Level 1 Accept is sent out by the ATLAS Central Trigger Processor, one of the four TRT Local Trigger Processors (LTP) receives a copy of this signal. This LTP then passes the trigger on to the next LTP, as well as to the ROD crates for which it is responsible. At the ROD crate, each of the TRT-TTCs will receive a copy of the trigger, and pass this trigger on to the Front-End electronics. Any delays necessary to compensate for trigger latency or Front-End pipeline depth are made in the TRT-TTC. The DTMROC at the Front-End, upon receipt of the L1A, will proceed to transmit its straw data from a pre-determined place in its pipeline. The data will then travel over small-twisted-pair wires to a series of Patch Panels, which receive the electrical signals and transmit them, over optical fibers, to the RODs in the ROD-crates. The ROD then organizes the straw data for up to 120 DTMROCs into a single buffer, and packages those data with: the event, run, and ROD identifiers; any error conditions detected in the DTMROC data; and markers to indicate the beginning and end of the event. The data are then sent to local, rack-mounted computers via the S-Link protocol [31], where they will be available for the High Level Trigger farms, and for collection into complete ATLAS events. In high trigger rate conditions, the ROD will also perform on-board compression

of the straw data before packaging. The compression scheme is a lossless, entropy based Huffman encoding, which has been designed to reduce the event size by a factor of four at a luminosity of 10^{34} (which becomes a factor of 16 at low luminosities). Without compression, the size of a single event for the entire TRT is 1.5 MB.

10.4.2 Calibration data

A principal design challenge for the TRT was the creation of high-speed, low-noise Front-End electronics. While these goals were met in the design phase, a thorough calibration of operating thresholds and straw efficiencies needs to be made on a routine basis in order to maintain optimal noise performance. Several tools have been developed to perform these calibrations. Some are suited to detailed testing, as described in section 13. To test and calibrate the full detector, a collection of tools has been created within the data acquisition framework. The calibration data are acquired in the same way as physics data. A run consists of a scan over some number of points in configuration space. To control the event rate, and to ensure adequate statistics for a given test, the trigger electronics are controlled directly to provide a specific number of triggers per scan point. The DTMROC event data, once collected by the RODs, is packaged with the necessary information to reconstruct the configuration at the time of that event. The ROD data is then collected by a dedicated computer, which either stores the data for later processing, or processes the data in real time.

10.4.3 Monitoring

Monitoring the activity and health of the detector has been integrated with the data acquisition system. In addition to monitoring changes in detector performance, the DAQ also watches environmental conditions such as chip temperature and low voltage levels on the Front End electronics. This DAQ based monitoring complements the monitoring done by the DCS. Part of this monitoring data is transmitted to the DCS (for comparison and crosscheck) via a data channel (Information Services) which connects the two (DAQ-DCS) subsystems.

10.4.4 ATLAS tools and the TRT viewer

At the ATLAS level, there are several tools available to the subsystems to facilitate detector monitoring. One such tool is an event display, which incorporates all of the subsystems into a single interface. This tool, called Atlantis [32], is designed to show event-level data in an easy to navigate format, with a variety of different projections available for all systems. This interface works well on an event level, but in order to see detailed information about detector performance over an entire run, another collection of tools is required. This suite, called AthenaPT [33], is similar in nature to the ATLAS level-3 trigger, and is capable of performing full reconstruction for all subsystems. It can be run in real time, sampling complete events from the DAQ event builder and storing the results into ROOT-based histograms for future analysis. At the TRT level, a special tool has been developed that combines the event display capabilities of Atlantis with the detector performance monitoring of AthenaPT. The TRTViewer extracts the TRT information from the raw data and provides both a real-time event display and an analysis package that does track reconstruction, efficiency calculations, noise occupancy maps, and error monitoring. This tool can also provide an interface to the configurations database, and has the ability to re-calibrate detector thresholds and timing based on information extracted from the monitoring process.

Table 7. Number of HV channels (feed lines) for the End Cap and Barrel.

End-cap			
Per wheel A	64	Per wheel B	32
Total wheels A	768	Total wheels B	512
Total End-cap		1280	
Barrel			
Per barrel Module-1	6	Per barrel Module-2	6
Per barrel Module-3	10		
Total Barrel		704	
Total		1984	

10.4.5 Voltages and temperatures

The on-detector ASICs were designed with the ability to report their temperature and operating voltages, which are then monitored for changes over time. The temperature or voltage can be found by searching for the threshold over which the DTMROC reports an error for the quantity being measured. Similarly, crossing the threshold from the other direction causes the error condition to disappear. As temperatures and voltages tend to behave consistently among chips that share a common low voltage connection, pairs of adjacent chips are used to bracket a common threshold over which the error condition appears. One chip will be set above threshold, and its neighbor will be set below the same threshold. If either chip changes state, the change will be registered for both chips, and a new threshold will be set. The error conditions are recorded during the beam gaps using the polling mechanism in the TRT-TTCs, which the read-back of a DTMROC register at some future time when there will be no triggers (i.e. the next orbit gap). This voltage and temperature information is collected by the DAQ, and is stored directly in an online conditions database. Alternatively, it can be passed to the DCS via a DAQ-DCS Communication [34] protocol. The voltages and temperatures can then be monitored by the DCS, as well as be archived in a conditions database.

11. High voltage power

The High Voltage power supply system which delivers 1.4 to 2 kV to each straw is based on a commercial multichannel HV power supply from HVSys [35]. Each HV channel delivers power to several straws (up to 192). In order to avoid the loss of a large fraction of the detector in case of a problem on one straw (e.g. in case of a short between a wire and the straw), each group of eight straws is individually protected with a fuse at the level of the detector (on the WEB for the end-cap and through fuse boxes in the barrel). The total number of HV channels for the end-cap and the barrel is given in table 7.

Maximum DC current per HV channel. The maximum current per line is defined by the number of powered straws and the average current consumption per straw at maximum luminosity.

Table 8. Estimated maximum current per HV channel at full luminosity.

Detector element	Number of straws	Number of lines	Number of straws per line	Current per straw at full luminosity (mA)	Total current per line (mA)	Maximum DC current to be delivered by the HV source (mA)
Wheel-A	12288	64	192	0.0045	0.86	3
Wheel-B	6144	32	192	0.0045	0.86	3
Barrel-1	329	6	55	0.020	1.1	3
Barrel-2	520	6	87	0.010	0.9	3
Barrel-3	793	10	80	0.006	0.5	3

Table 8 summarises the requirements. The maximum current to be delivered by a HV channel includes a safety factor for fluctuations in luminosity.

Maximum AC current per HV channel. The HV distribution circuit includes a filter consisting of a 100 kOhm resistor and a 1 nF capacitor. In case of a straw discharge, the HV source must recharge the 1 nF capacitor through the 100 kOhm equivalent resistor fast enough so that a discharge does not appear as a DC fluctuation. A maximum current of 16 mA during 100 ms must be delivered by the power supply.

Trip requirements. When an anode wire is in contact with the cathode, the contact resistance is in the range 0-3 kOhm and some damage can appear in the straw (e.g. holes) under some conditions. Continuous discharges (when the wire is not in direct contact with the cathode) can also lead to damaging the straw conductive layer and in some cases a hole in the straw can appear. It has been shown that such damage appears only after more than 5000 discharges have taken place. Limiting the number of discharges is, therefore, a very important requirement. The system has a programmable trip level and time duration for how long the trip level can be exceeded before actually tripping the HV output.

11.1 HVSys power supply

The system selected for the TRT is a multi-channel power supply from HVSys. Packaged in a 6U euromechanics crate, the system consists of up to twenty 21-channel HV modules and a crate controller. A backplane distributes control signals and low voltage power.

The main characteristics of the power supply are given in table 9.

11.2 Implementation

The overall TRT HV power system is divided into four geographical regions: each side of the Barrel and each End Cap. Each Barrel side and each End Cap requires, respectively, 352 and 640

Table 9. HVSys power supply nominal characteristics.

Parameter	Value	Condition
No. of independent channels	420	
Output voltage per channel	Programmable in the range 500 - 2000 V	In steps of 2 V
Current per channel	3 mA	Maximum continuous rating.
Ramp-up and Ramp-down	Programmable in the range 1 -500 V/s in steps of 10 V/s	
Noise and ripple	≤ 50 mV pp	0-100 MHz at output with a 5k - 1 nF filter
Voltage Monitor versus Output Voltage Accuracy	± 0.5 V $\pm 0.1\%$ of reading	
Voltage Set versus Voltage monitor Accuracy	± 1 V $\pm 0.1\%$ of setting	
Current Monitor versus Output Current Accuracy	$\pm 5\mu\text{A}$	
Maximum Current Set versus Output Current Monitor Accuracy	$\pm 2.5\mu\text{A}$	
Over-current trip condition	-	Output voltage set to 0V if the maximum output current value is exceeded for more than a programmable time
Efficiency	80%	

channels. The power supply crates are located in the counting room (USA15). Multiwire HV cables are used from USA15 to the PP2 region (36 or 52 wires per cable, length in the range 60 - 100 m) and then miniature cables are used from the PP2 location to the detector (coaxial cable of 1-mm outer diameter, length in the range 6 - 15 m). Several passive patch-panels are needed as seen on figure 36:

- One Patch Panel in the control room adapting the power supply 3-channel miniature LEMO connectors to 52- or 36-way connectors attached to the multiwire HV cable;
- One in the PP2 region converting between multiwire cables and miniature coaxial cables;
- On at the end of the cryostat flange for the End Cap (PPF1) or close to the barrel face for the Barrel (PPB1) to facilitate detector installation.




Figure 36. HV distribution scheme

11.3 Control

The HVSystem crate is controlled by a set of microprocessors (one main and four branch controllers) in the crate controller module and microcontrollers at each module board. Built-in control firmware accepts external commands and communicates with the individual channels (cells). An external host computer can access the crate via a serial RS-232 communication line or via optocoupled CAN-bus or USB. Two lines of the internal system bus are used for communication with the high voltage cells. The RS-232 communication is performed at a bit rate of about 10 kHz. High voltage cells (HVCs) operate as slaves on the branch. The main controller also controls the central power supplies (low voltages and branch voltages), monitors them and provides emergency procedures. Two 8-channel analog to digital converters are provided for the monitoring of the central power supplies and monitoring temperatures of critical elements. Commands for the controller are divided into several main groups according to their functional assignments.

- Commands for obtaining system unit status.
- Commands controlling the central power supply.
- Commands controlling cell(s).
- Cell detection commands.
- Commands addressed simultaneously to all cells.
- Service commands.
- Test commands.

The HVSystem USB and CAN communication modes are limited and do not exploit all the virtues of either USB or CAN but simply translate the 'character' based RS-232 protocol to either CAN (message) or USB (package) oriented media. Careful studies of the system performance lead to a decision to abandon any control method other than RS-232. To allow for easier integration of the system into the overall TRT DC system an OPC server for system controls was constructed as well as a Framework type for the hardware devices. This now allows the use of all of the Framework tools for mass parameterization and access to configuration and conditions databases. The system




Figure 37. View of the low voltage distribution system

performance is limited by the speed of the internal bus and the required overhead in the OPC server necessary for robust error handling. The average response time for a broadcast command sent to the 420 channels contained in a crate is ~ 30 secs.

12. Low voltage power

The low voltage distribution to the TRT Front-End electronics is based on a two level scheme shown in figure 37 and described in greater detail in [36].

- Bulk power supplies deliver power to the PP2 power boards;
- PP2 power boards housing radiation tolerant voltage regulators deliver the power to the Front-End.

In order to reduce the length and therefore the size of power cables, Wiener MARATON [37] radiation and magnetic field tolerant multi-output bulk supplies have been installed in the experimental cavern (UX15).

12.1 Requirements

A set of MARATON bulk supply channels provide power to a PP2 Patch Panel Box which serves 1/32 of one side of the detector. There are six PP2 locations, housing three, five or eight PP2 Boxes each. Each PP2 Box contains three power boards, one for 1/32 of the barrel, one for 1/32 of the wheels A and one for 1/32 of the wheels B. In addition the PP2 box contains some data and TTC electronics which need to be powered. Three different voltages are needed for the Front-End supply (+3V, -3V, +2.5V) and an additional one (5V) for the PP2 electronics. The total current needed per PP2 Box and voltage source is given in table 10.

Each bulk supply channel has a separate return line in order to give partial isolation of the loads. However, given that the analog and digital grounds must be connected at the Front End boards and since those grounds are connected to the power returns for their respective analog and digital voltages and because a single MARATON channel provides power for both barrel and end cap boards, there are, in fact, multiple current return paths to a MARATON channel. In normal operation this produces no ill effects, but if one section of the detector is powered off while another section is operating, the change in current flow can produce a change in effective threshold on the barrel Front End boards where the analog and digital planes are separated by multiple low value

Table 10. Current needed per supply.

	+3V	-3V	+2.5V	PP2(5V)
Barrel	13A	11A	13.5A	
Wheels A	18A	16A	19A	
Wheels B	12A	10.5A	12.5A	
PP2				< 5 A
Total	43A	38A	45A	< 5 A

resistors but not on the end cap boards where the analog and digital planes are hard connected together.⁹ This voltage offset must be accounted for in setting operating thresholds and automated software scripts are in place to adjust thresholds for changes in operating conditions. The VT voltage measurement data coupled with high threshold noise scans allow measuring and validating the actual threshold conditions in a simple way.

Each PP2 box requires four separate power supply channels. Because of the voltage drops introduced in all the cables (bulk to PP2 and PP2 to Front-End) one needs a much higher voltage at the bulk supply level than what is needed on the Front-End; an 8V/300W MARATON channel type has been used for all of loads. As the needed power for the PP2 electronics is low, one MARATON channel is shared between two PP2 boxes. The 64 powered PP2 boxes requires 224 MARATON channels (three channels per box for powering the Front-End and one channel every second PP2 box for the PP2 electronics). Given twelve channels per MARATON bulk supply, the TRT uses twenty rack mounted supplies located in five different racks in the UX cavern.

12.2 WIENER MARATON supplies

There are three main components in the WIENER system (as seen in figure 37):

1. 3U power crate with radiation and magnetic tolerant DC-DC each of them housing 12-channel 8 V/300 W elements;
2. 4U crate housing up to six 3.6 kW AC-DC converters (380 V DC output). One AC-DC is needed to power a 12-channel 8 V/300 W element;
3. 6U crates housing up to 21 control boards. One control board controls a 12-channel 8 V/300 W element.

Position 1 is in the experimental cavern (UX15), while positions 2 and 3 are in the control room (USA15).

In UX15 there are racks for the TRT power supplies at five locations on the structure around the detector. Table 11 gives the summary of the different PP2 boxes powered from each location, as well as a summary of the MARATON equipment installed at each location.

⁹Note that the two different schemes — direct ground connection in the end cap and low value resistive isolation in the barrel case are the result of noise optimization studies carried out on the bench and in test beams during the electronics development and reflect the differing geometries and connectivities of the barrel and end cap detectors.

Table 11. MARATON devices at each location in UX15.

Location	PP2 boxes location	# PP2 boxes	MARATON equipment	
			# channels	# 3U power crates
1	1&2 side-C	13	46	4
2	1&2 side-A	13	46	4
3	3&4 side-A	13	46	4
4	3&4 side-C	13	46	4
5	5&6 side- A & C	12	38	4
Total		64	256	24

Table 12. Voltage drops.

Element	Voltage Drop
Type III cable	0.5 - 1 V
Regulators at PP2	0.8 - 1.5 V
Type II cables	0.5 - 1.3 V
Total maximum drop	3.8 V

In order to maintain a clear path for the current returns of the power lines feeding the Front-End, each DC-DC is floating and two wires (power and return) are used to feed the PP2 boxes, except for the + and - 3V which are using a common return line.

12.3 Overall efficiency

There are voltage drops and inefficiencies at several levels in the system. Table 12 gives a summary of the voltage drops along the cables from the MARATON DC-DC converter and at the linear regulators in PP2. The minimum and maximum values correspond to different cable lengths and cross section. The table takes into account the voltage drop on the return lines.

In order to have 2.5 V on the Front-End boards, one has to deliver up to 6.3 V at the output of the DC-DC converters. This gives an efficiency of $\sim 40\%$. Taking into account the efficiency of the AC-DC (main to 380 Vdc) and of the DC-DC (and neglecting the very small drop on the type IV cable) leads to an overall efficiency of about 30%.

12.4 Control

12.4.1 Bulk supplies

The MARATON power supplies are delivered with an OPC server package which provides the interface for remote control of the unit. The server operates via Ethernet. The OPC standard is an interfacing method to avoid the need of device specific drivers. The CERN control group of the CERN IT department provides users with a skeleton software solution called FRAMEWORK. The package contains many useful tools for construction, integration and configuration of the complex multi-channel systems build of the standard, widely available equipment. The MARATON system has been made part of FRAMEWORK which makes the integration of the supplies straightforward. The control commands of the MARATON system are limited. Basically the only important action

which can be undertaken by user is switching the given output on or off. For monitoring purposes the system gives access to the values of the output voltages and currents. Control panels have been designed to allow the user to monitor and control the system.

12.4.2 Patch Panels

Groups of custom, rad hard, linear regulators supplying geographically clustered parts of the detector Front-End electronics have been placed on a printed circuit board together with the control and monitoring electronics. One card of 440x200 mm size houses up to 36 regulators (positive and negative) delivering voltage sets: ± 3 V for analogue electronics and +2.5 V for the digital part. Three boards are used per PP2 crate to power a 1/32 phi slice of the detector: one board for the barrel and two for the end-cap. The high power dissipation on the board (60-70 W) necessitates an effective liquid cooling system, described in section 4.2. The circuitry on the board performs the functions described in the following paragraphs.

Voltage regulation and setting. The regulators used are adjustable. Changing the voltage "adjust" input allows the output to be set at the proper value for operation of the Front End electronics. This value may change during the life cycle of the experiment as semiconductor devices change characteristics due to radiation. The variable set-point voltage is delivered to the regulator by the radiation hard DAC embedded in the DTMROC chip (see section 2. The settings of the DAC's are controlled via a simple 3-wire serial protocol. The voltage swing of the DAC output (from zero to +1.25V) allows for the regulator outputs to be varied by 0.5 V, a range which fully covers the required worst case change [38]. Because of the wiring harness layout, the Wheel A supplies need slightly more than the rated current for individual regulators. For these channels parallel operation of the regulators has been implemented. A carefully designed biasing network ensures correct sharing of current between the two regulators.

Voltage and current measurements in each individual output line. The PP2 power supply board contains an embedded controller - the ELMB (Embedded Local Monitoring Board [13]), standard control unit designed for ATLAS. All voltage outputs from the PP2 power supply board are connected to the ELMB's ADC inputs allowing for output voltage measurements. The same ADC also measures the output currents, by monitoring the voltage drop on 22 m Ω serial resistors inserted in the output lines.

Over current protection. The regulators have internal over-current protection which, during an overcurrent condition, operates the regulator in a constant current mode. The over-current state of the regulator is signaled by a level that is latched by PP2 board logic and can be monitored by the ELMB. The on-board logic will switch the affected channel off on an over-current condition if so enabled.

Enabling and disabling individual output lines (on/off). The DTMROC also has digital outputs which, via the proper command, can be set 'true' or 'false'. One bit of the DTMROC digital output is connected to the 'inhibit' input of each regulator and thus provides a means to remotely switch each regulator output on/off. This on/off control signal is "ORed" with the latched over-current signal described above. This control bit allows switching off and on the individual parts of the

Front End electronics supplied by a given voltage regulator. It should be noted that the DTMROCs are powered by the Patch Panel bulk supply, not the outputs of the controlled regulators.

12.4.3 Control and monitoring

The TRT power supplies control system has been built upon the commercial SCADA (System Control And Data Acquisition) software package chosen by CERN as standard for control systems of LHC experiments - PVSSII. Detailed descriptions of the available tools and methods are available elsewhere [18].

CAN bus and ELMB. ELMB is a basic hardware unit of the ATLAS detector control system and can operate in the magnetic field and radiation environment of the experimental cavern. The ELMB contains 64 channels of multiplexed 16 bit ADC input and several I/O digital ports. The basic ELMB functionality and detailed description can be found elsewhere [13]. ELMB commands and monitoring readback values are sent over CAN bus [39] - one of the CERN recommended commercial fieldbuses. The CAN bus, designed for and widely used in the automobile industry, is a robust communication layer, and both hardware and software components are available from a wide range of manufacturers. Different autonomous actions can be taken when ELMB error conditions are detected: node reboot, microcontroller firmware reload and disconnect of the affected node from CAN bus. Merging the NodeID with MessageID codes to COBID (CAN Object ID) field allows an efficient message oriented priority mechanism. A set of broadcast messages reduces the bus load in complex systems.

LVPP board control. The main development effort spent on the TRT low voltage system has been on the design, programming and debugging of the Low Voltage Patch Panel boards (LVPP) containing the voltage regulators and main controlling/monitoring circuitry. As mentioned previously, the heart of the LVPP is an ELMB communicating with the host controller over CAN bus. The read-out of the analogue values (current and voltage) is standard for ELMBs in ATLAS and is based on the CANOpen [39] software layer where the OPC standard for data access has been implemented. The ‘analog’ functionality of the ELMB has already been included in the CERN FRAMEWORK package so integration and control is reasonably easy and straightforward for those functions. A custom design was done for the connection to the DTMROCs that control the voltage regulators. An extension has been designed and built to the standard PVSS CTRL scripting language [18]. This extension allows user defined functions to be interpreted by PVSS in the same way as PVSS functions and consists of a set of class methods compiled into new shared library DLL, providing the input / output parameters and return value. The functions access directly the CANbus driver. The DLL contains the following functionalities.

- Initialization of the CANbus, ELMB, DTMROC
 - Operational:
 - Set DAC’s
 - Read back DAC’s
 - Set inhibits in DTMROC’s

- Read back inhibit state
- Enable/disable and read out the over-current monitor state
- Diagnostics:
 - Reset (soft and hard) of DTMROC's,
 - Send given number of clocks to DTMROC's
 - Get state of a given DTMROC,
 - Set ELMB in the requested state
 - Read back ELMB state
 - Close connection

Part IV

Installation, integration and performance

13. Installation, integration and testing

The Front End electronics for the TRT have undergone an extensive testing regime throughout the course of the installation process, testing functionality, individual noise performance, and performance inside the larger system. There were three distinct phases of this testing regime: functionality tests after board assembly and burn-in, connectivity and noise tests before and after mounting the boards on the detector, and system level tests after the electronics installation was complete. These three phases are described below.

13.1 Front End board initial tests and preparation

The initial testing of the TRT Endcap Front End boards was done at the Niels Bohr Institute (NBI) in Copenhagen, and the testing of the Front End boards for the TRT Barrel was done at CERN, with effort from the University of Pennsylvania (Penn), Lund University, and Hampton University. A unified testing procedure was used, with testers at both locations using the same test suite, writing results into a single MySQL database. The test suite [40] was designed to exercise the complete functionality of each board and its corresponding chips. All of the registers of each DTMROC were exercised, a fine timing scan was run to check the health of the command lines, and a noise scan was run with the internal test pulse firing to characterize the analog performance and identify dead channels. For the Endcap boards, a test was also run with an external pulse injector as another way to identify dead or problematic channels. The data from these tests were stored in the previously mentioned MySQL database, and a web interface to this database was developed for later viewing, as well as for tracking board shipments and repair histories.

The criteria for passing these tests were: no register/memory failures or fine timing anomalies were tolerated in the chips on any boards. For the endcap boards, no dead or otherwise non-functional channels were allowed. For the Barrel boards, which had a much less forgiving granularity (up to 432 channels on a single PCB, vs. 64 for the ASDBLR endcap boards), a requirement of no more than roughly 0.5 percent dead or otherwise non-functional channels for a single board was used. All boards which passed the initial testing phase were burned in for one week at ~ 80 degrees C at a special facility which was constructed for this purpose at NBI. The burn-in facility continuously measured the current consumption of each board while maintaining a constant temperature through the burn-in process. After burn-in, all boards were tested a final time. Those boards which failed the tests (either before, or after burn-in) were sent to Penn for repairs, and then returned to CERN or NBI to be inserted back into the testing program.

13.2 Front End board installation and tests

After passing the first round of tests, The FE Boards for both the TRT Barrel and Endcap were transferred to the SR-1 Integration Facility at CERN to be installed on the detector. A series of tests was run on the boards at this point with two main goals: characterizing the noise profile of each Front End board, and ensuring that the boards were properly connected to the detector.

In order to characterize the noise performance, three tests were run: one which recorded the low (tracking) threshold noise response for all low threshold values; a second, similar test which looked at the response to a calibrated internal test pulse rather than noise; and a third test which recorded the high (transition radiation) threshold response to the internal test pulse for all high threshold values. This series of tests was run both before and immediately after each board was mounted on the detector. See figure 38. Anomalies in the noise response of a given board seemed to be associated with problems with the mechanical insertion and contact of the board with the Faraday shield, so particular attention was paid to the results of these tests.

In addition to the threshold scans, other tests were done after the board was mounted to determine if the channels were properly connected to their respective straws. In the barrel, a ‘connectivity test’ was run where an external pulse was applied to the high voltage system and data from the corresponding pulse induced on the anode wires was read out to show connectivity. For the End Caps, where the design of the high voltage system made this approach impossible, the determination of connectivity was made through a twofold process. First, analyzing the difference between on and off detector noise rates which are sensitive to anode wire capacitance and second, running a scan which used the ‘accumulate mode’ of the DTMROC to measure a background rate of high threshold hits when high voltage and active gas are applied to the detector [41]. The combination of this “accumulate mode” scan with the difference between on and off detector noise rates provided an equally powerful tool for the identification of disconnected channels.

When the electronics integration was complete, the numbers of dead channels/straws were as follows: For the Barrel, there were 1755 dead straws out of 105,088 total (1.7%), of which only 119 were from dead or otherwise non-functional electronics channels. The vast majority of dead straws were due to wires that had to be removed due to high voltage instability or physical non-conformity of the straw. For the Endcaps, Endcap C had 932 dead straws out of 122880 total (0.8%), while endcap A had 493 (0.4%). All of these failures arose from mechanical problems, chief among them failed high voltage lines. The feat of having zero non-functional electronics channels was achieved mainly as a result of the forgiving granularity of the endcap board design. The noise levels on both the barrel and endcap sections of the detector were well within specifications after all electronics were installed and the Faraday cages completed. The noise levels stayed constant after the subdetectors were integrated into ATLAS.

13.3 TRT stand-alone noise and system tests

After the Front-End board installation phase was complete, a phase of system level tests began, looking for noise or readout issues in the larger detector. The first of these tests was simply to look for increasing levels of noise as larger and larger fractions of the detector electronics were powered and read out in unison. The Front-End electronics (and especially that of the barrel) is known to have a sensitivity to clock pick-up as a result of having the 40MHz control lines in close proximity to the sensitive analog inputs, and there was concern that as more and more clock signals were running on the detector this pick-up would increase. When this test was first run, some regions of increased noise were found. However, all of them pointed to gaps in the outer Faraday shield, and they all disappeared when the gaps in the shield were filled as shown in figure 39. Once the outer Faraday shield was completed, there was no increase in noise as a result of powering and clocking large sections of the detector.




Figure 38. Per channel 300kHz threshold distributions for barrel electronics on-detector (far right curve, green), off-detector (second from right, red), difference of on-off-detector (third from right, blue) and then difference from chip average (far left curve, purple). The increase in the 300 kHz threshold when the electronics are mounted in place is due to the detector capacitance which raises the equivalent noise charge figure for the ASDBLR. The smaller capacitance of the first nine layers of ‘short’ straws is clearly evident in the difference (third from right, blue) distribution.

Another concern was that the Front End electronics would be affected by pickup arising from its own data read out. To test for this susceptibility, a special trigger was set up to send pairs of triggers, the first one being a normal random noise event, and the second one taking data during the time when the first event was being read out. A larger occupancy in the data from the second trigger would point to a sensitivity to data read-out pick-up. The result of this study also depended heavily on the state of the outer Faraday shield. When this test was first done, and the Faraday shield was not yet complete, some chips showed a slight sensitivity to data read-out pick-up. However, the magnitude of this effect was very small, and in the barrel, it decreased to zero once the Faraday cage was complete.

13.4 Combined noise tests with SCT

After internal noise issues, the main worry for the TRT electronics performance was interference or pick-up between the TRT and the SCT. Because the two subdetectors are in such close proximity, and because their services often run one on top of the other in the cable trays leading to the detectors, great care was taken to isolate the two systems in the design. This approach seems to have succeeded, since no crosstalk or interference between the two systems has been observed to

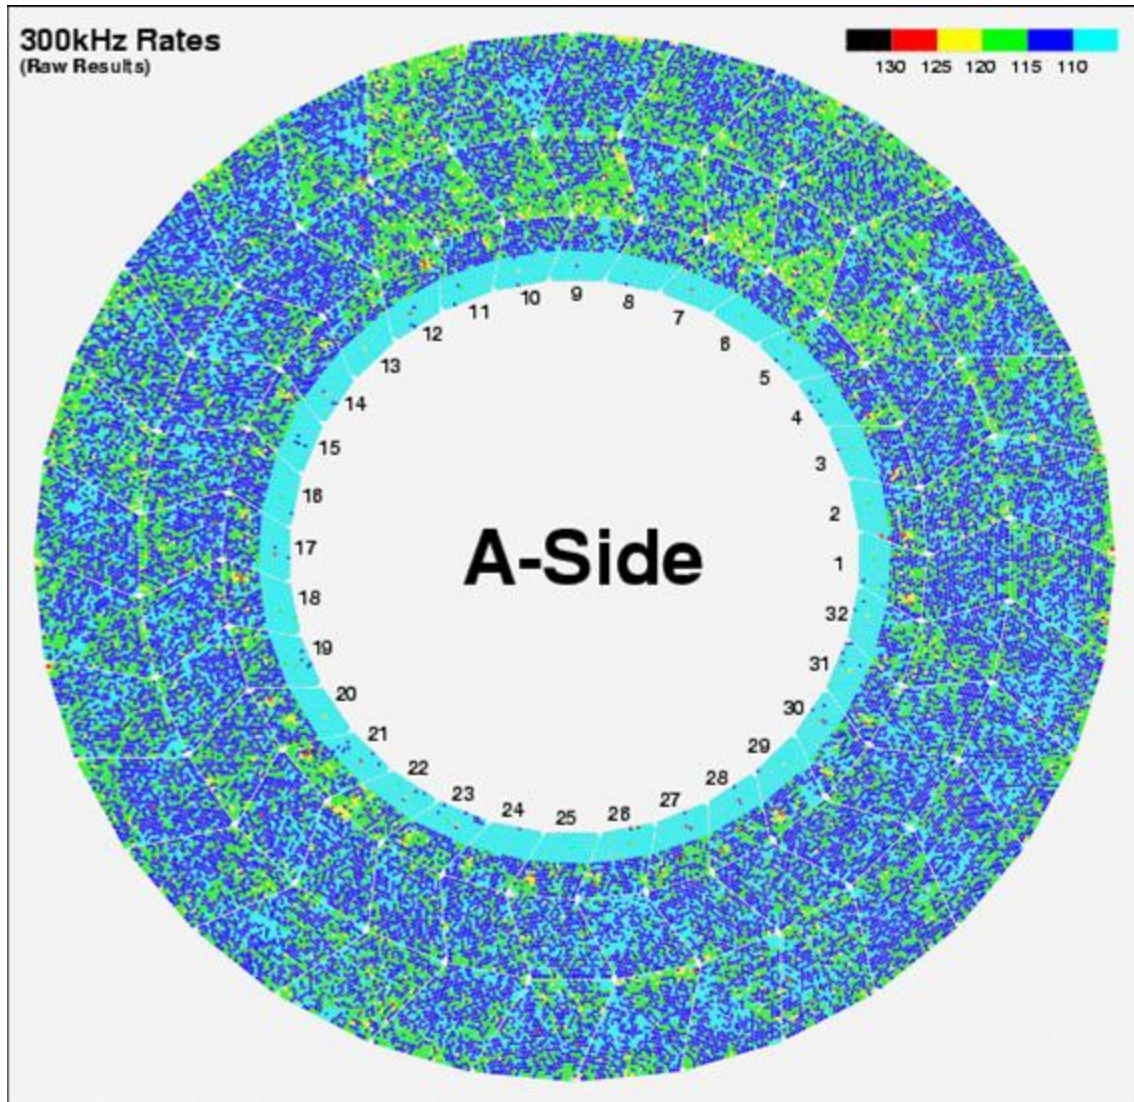


Figure 39. A side of the TRT Barrel showing the threshold setting required to obtain a 300 kHz average noise rate in each channel.

date [42]. The first test in this domain was to simply run the TRT with a noise trigger with the SCT powered off, and then make a second noise run with the SCT powered on. The comparison of these two states shows no difference in noise occupancy in the TRT as a result of the SCT being on or off. A second test was to look for a sensitivity in the TRT electronics to the SCT read-out. This test was similar to the internal TRT test for read-out pick up, except that in this case, the TRT trigger was timed to coincide with the SCT read-out cycle, rather than that of the TRT. No increase in noise was found when the SCT was reading out as compared to when it was not. Finally, the effect of a possible incorrect ground connection was investigated. In this test, the TRT analog ground plane was attached to the SCT digital power return ground, and a noise run was taken. Once again, there was no noticeable increase in noise rates, indicating that there is little sensitivity in the TRT electronics to mistaken or poorly made grounding connections between the two subsystems.




Figure 40. Cosmic ray tracks in the TRT Barrel from the “M3” run in ATLAS using the Tile Cal trigger.

13.5 Cosmic ray runs with SCT and ATLAS

In 2006, a series of cosmic ray runs was taken with the TRT and SCT combined, in the surface building (SR-1) at CERN where the two detectors were assembled and integrated [42]. For the run including the TRT Barrel, 1/8 of the barrel was read out (1/16th on top and 1/16th on bottom), comprising just over 13,000 channels. For the endcap run, 1/16 of the endcap was read out, comprising roughly 7500 channels. In both cases, a simple scintillator trigger was set up to highlight the active regions, and the standard ATLAS DAQ and monitoring infrastructure was used to record and monitor the data. These runs provided useful experience in commissioning, running, and monitoring the detector. In addition, the data sets that were generated were also useful for tracking and alignment studies because these were the first data showing tracks in the two jointly assembled sub-detectors.

Once the detector was fully installed and integrated into ATLAS, the TRT participated in several global commissioning runs (the series of which is still ongoing as of this writing), again taking cosmic ray data (see figure 40), with varying portions of the detector being read out. These commissioning runs have provided running experience with the final hardware setup and the global ATLAS DAQ - something that was not possible before the detectors were installed. In particular, methods for setting and tuning the Front-End voltages, setting and tuning the tracking thresholds, and for tuning the fine timing required for proper drift-time measurement were developed. In particular, the voltage and temperature sense functionality in the DTMROC has proved essential in correctly setting up and monitoring the voltages on the Front-End electronics since this is the only direct measurement available of the voltage at the Front-End electronics, deep within the detector. Also, commissioning runs provide critical experience for operation of the DAQ system, detector monitoring and data analysis.

Acknowledgments

The research described in this publication was partly supported by following funding agencies: the European Union (DGXII), the International Science Foundation, the Swedish Research Council, the Knut and Alice Wallenberg Foundation, the Ministry of Science and Higher Education, Poland, the International Science and Technology Centre, the Civil Research and Development Foundation

and grants from the U.S. Department of Energy and the U.S. National Science Foundation, the Natural Science and Engineering Research Council of Canada, the Natural Science Research Council of Denmark, the Ministry of Education and Science of the Russian Federation, the Academy of Science of the Russian Federation, the International Association for the Promotion of Cooperation with Scientists from the new Independent States of the former Soviet Union, The Turkish Atomic Energy Authority.

References

- [1] E. Abat et al., *The ATLAS Transition Radiation Tracker (TRT) proportional drift tube: design and performance*, 2008 JINST 3 P02013.
- [2] N. Dressnandt, N. Lam, F.M. Newcomer, R. Van Berg, H.H. Williams, *Implementation of the ASDBLR straw tube readout ASIC in DMILL technology*, *IEEE Trans. Nucl. Sci.* **48** (2001) 1239.
- [3] ATLAS-TRT collaboration, T. Akesson et al., *Straw tube drift-time properties and electronics parameters for the ATLAS TRT detector*, *Nucl. Instrum. Meth.* **A 449** (2000) 446.
- [4] N. Dressnandt, F.M. Newcomer, O. Rohne, S. Passmore, *Radiation hardness: design approach and measurements of the ASDBLR ASIC for the ATLAS TRT*, *IEEE Nucl. Sci. Symp. Conf. Rec.* **1** (2002) 151.
- [5] C. Alexander et al., *Progress in the development of the DTMROC time measurement chip for the ATLAS Transition radiation tracker (TRT)*, *IEEE Trans. Nucl. Sci.* **48** (2001) 514.
- [6] F. Anghinolfi et al., *DTMROC-S: Deep submicron version of the readout chip for the TRT detector in ATLAS*, Prepared for 8th Workshop on Electronics for LHC Experiments, Colmar, France (2002), <http://cdsweb.cern.ch/record/1056713>.
- [7] W. Dabrowski et al., *Design and performance of the ABCD chip for the binary readout of silicon strip detectors in the ATLAS semiconductor tracker*, *IEEE Trans. Nucl. Sci.* **47** (2000) 1843.
- [8] 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. Symp. Conf. Rec.* **1** (2002) 549.
- [9] E. Abat et al., *ATLAS TRT End-Cap Detectors*, to be submitted to JINST.
- [10] E. Abat et al., *ATLAS TRT Barrel Detector*, 2008 JINST 3 P02014.
- [11] C. Baxter et al., *Progress on the development of a detector mounted analog and digital readout system for the ATLAS TRT*, *IEEE Nucl. Sci. Symp. Conf. Rec.* **1** (2003) 202.
- [12] H. Blampey et al., *Services for the read-out of the ATLAS TRT cables types I and II, from detector to PP2*, prepared for 11th Workshop on Electronics for LHC and Future Experiments, Heidelberg, Germany (2005), <http://cdsweb.cern.ch/record/922717>.
- [13] <http://elmb.web.cern.ch/ELMB/>.
- [14] H. Furtado, *Delay25, an ASIC for timing adjustment in LHC*, Prepared for 11th Workshop on Electronics for LHC and future Experiments, Heidelberg, Germany (2005), <http://cdsweb.cern.ch/record/920425>; see also http://proj-delay25.web.cern.ch/proj-delay25/delay25_-_home.htm.
- [15] P. Moreira et al., *A radiation tolerant gigabit serializer for LHC data transmission*, 7th Workshop on Electronics for LHC Experiments, Stockholm, Sweden (2001), <http://cdsweb.cern.ch/record/588665>.

- [16] CERN EP Microelectronic Group, <http://web-micfe.web.cern.ch/web-micfe/index.html>.
- [17] P. Moreira *QPLL: a Quartz Crystal Based PLL for Jitter Filtering Applications in LHC*, 9th Workshop on Electronics for LHC Experiments, Amsterdam, Holland (2003), <http://cdsweb.cern.ch/record/722067>; see also <http://www.lecc2003.nikhef.nl>.
- [18] S. Schmeling, *Common tools for large experiment controls-a common approach for deployment, maintenance, and support*, *IEEE Trans. Nucl. Sci.* **53** (2006) 970.
- [19] RD12 PROJECT collaboration, B.G. Taylor, *TTC distribution for LHC detectors*, *IEEE Trans. Nucl. Sci.* **45** (1998) 821.
- [20] P. Lichard et al., *Evolution of the TRT backend and the new TRT-TTC board*, Prepared for 11th Workshop on Electronics for LHC and Future Experiments, Heidelberg, Germany (2005), <http://cdsweb.cern.ch/record/920977>.
- [21] J. Christiansen, A. Marchioro, P. Moreira and A. Sancho, *Receiver ASIC for timing, trigger and control distribution in LHC experiments*, *IEEE Trans. Nucl. Sci.* **43** (1996) 1773.
- [22] M. Hance, *Readout of the ATLAS Transition Radiation Tracker: Timing Parameters and Constraints*, ATL-COM-INDET-2008-002.
- [23] S. Fratina, B. LeGeyt and J. Saxon, *A Method for Determining First-Pass Fine Delays for the ATLAS TRT Detector*, ATL-INDET-PUB-2008-007.
- [24] D. Huffman, *A Method for the Construction of Minimum Redundancy Codes*, *Proc. of the IRE* **40** (1952) 1098.
- [25] P. Jenni et al., *ATLAS high-level trigger, data-acquisition and controls: Technical Design Report*, ATLAS-TDR-016, CERN-LHCC-2003-022.
- [26] P. Borrego Amaral et al., *The ATLAS Level-1 Central Trigger System*, 10th Workshop on Electronics for LHC and Future Experiments, Boston, MA, USA (2004).
- [27] P. Borrego Amaral et al., *The ATLAS Local Trigger Processor (LTP) 018*, 10th Workshop on Electronics for LHC and Future Experiments, Boston, MA, USA (2004), ATL-DAQ-2004-018.
- [28] ATLAS TTC System, available online at <http://tc.web.cern.ch/ttc/>.
- [29] I. Soloviev, et al., *The Configurations Database Challenge in the ATLAS DAQ System*, *Computing in High Energy Physics and Nuclear Physics*, Interlaken, Switzerland (2004), ATL-COM-DAQ-2004-021.
- [30] I. Papadopoulos, et al., *CORAL, A Software System for Vendor-Neutral Access to Relational Databases*, 15th International Conference on Computing In High Energy and Nuclear Physics, Mumbai, India (2006).
- [31] H.C. van der Bij, R.A. McLaren, O. Boyle and G. Rubin, *S-LINK, a data link interface specification for the LHC era*, *IEEE Trans. Nucl. Sci.* **44** (1997) 398.
- [32] N.P. Konstantinidis et al., *The Atlantis event visualisation program for the ATLAS experiment*, *Computing in High Energy Physics and Nuclear Physics*, Interlaken, Switzerland (2004), <http://cdsweb.cern.ch/record/865603>.
- [33] AthenaPT wiki page, available online at <https://twiki.cern.ch/twiki/bin/view/Atlas/AthenaPT>.
- [34] R. Hart and V.Khomoutnikov, *ATLAS DAQ-DCS Communication Software: User Requirements Document*, Nov. 2000, http://atlasinfo.cern.ch/ATLAS/GROUPS/DAQTRIG/DCS/DDC/ddc_urd.pdf.

- [35] HVSystems, Dubna, Russia, <http://www.hvsys.dubna.ru/>.
- [36] Z. Hajduk et al., *Distributed low voltage power supply system for Front End electronics of the TRT detector in ATLAS experiment*, Prepared for 12th Workshop on Electronics for LHC and Future Experiments, Valencia, Spain (2006), <http://cdsweb.cern.ch/record/1028143>.
- [37] W-IE-NE-R, Plein & Baus GmbH, Mullersbaum 20, D-51399 Burscheid Germany, www.wiener-d.com.
- [38] F. Anghinolfi, B. Kisielowski and V. Ryjov, *Control of the voltage regulators with a DTMROC*, ATL-IT-EN-0052, EDMS ID 468670 (2006).
- [39] B. Hallgren and H. J. Burckhart, *Frontend I/O via CANbus of the ATLAS detector control system*, 4th Workshop on Electronics for LHC Experiments, Rome, Italy (1998).
- [40] O.M. Rohne, University of Oslo, private communication.
- [41] N. Ghodbane, X. Pons, O.M. Rohne, *A Method to check the Connectivity for the ATLAS TRT Detector*, physics/0510130, ATL-COM-INDET-2005-009 (2005).
- [42] E. Abat et al., *Combined performance tests before installation of the ATLAS Semiconductor and Transition Radiation Tracking Detectors*, submitted to JINST.