

Design and Implementation of the Front End Board for the Readout of the ATLAS Liquid Argon Calorimeters

N.J. Buchanan^{a,1}, L. Chen^{a,2}, D.M. Gingrich^{a,3}, S. Liu^a, H. Chen^b, J. Farrell^b,
J. Kierstead^b, F. Lanni^b, D. Lissauer^b, H. Ma^b, D. Makowiecki^b, V. Radeka^b,
S. Rescia^b, H. Takai^b, H. Ghazlane^c, A. Hoummada^c, J. Ban^{d,4}, S. Boettcher^{d,5},
G. Brooijmans^d, C.-Y. Chi^d, S. Caughron^d, M. Cooke^d, D. Dannheim^{d,6}, A. Gara^{d,7},
A. Haas^d, I. Katsanos^d, J.A. Parsons^{d,*}, S. Simion^{d,8}, W. Sippach^d, L. Zhang^d,
N. Zhou^d, E. Ladygin^e, E. Auge^f, R. Bernier^f, M. Bouchel^f, A. Bozzone^f, D. Breton^f,
C. de la Taille^f, I. Falleau^f, P. Imbert^f, G. Martin-Chassard^f, A. Perus^f,
J.P. Richer^{f,9}, V. Tocut^f, J.-J. Veillet^f, D. Zerwas^f, F. Hubaut^{g,10}, B. Laforge^g, O. Le
Dortz^g, D. Martin^g, Ph. Schwemling^g, J. Collot^h, D. Dzahini^h, M.-L. Gallin-Martel^h,
P. Martin^h, W.D. Cwienkⁱ, J. Fentⁱ, L. Kurchaninov^{i,11}, G. Battistoni^j, L. Carminati^j,
M. Citterio^j, W. Cleland^k, B. Liu^k, J. Rabel^k, G. Zuk^k, K. Benslama^l, E. Delagnes^m,
B. Mansoulié^m, J. Teiger^m, B. Dinkespiller^{n,10}, T. Liuⁿ, R. Stroynowskiⁿ, C.-A. Yangⁿ,
J. Yeⁿ, M.L. Chu^o, S.-C. Lee^o, P.K. Teng^o

^a *Centre for Particle Physics, Department of Physics, University of Alberta, Edmonton, Alberta, Canada*

^b *Brookhaven National Laboratory (BNL), Upton, NY, USA*

^c *Faculté des Sciences, University of Casablanca, Casablanca, Morocco*

^d *Columbia University, Nevis Laboratories, Irvington, NY, USA*

^e *Joint Institute for Nuclear Research, Dubna, Russia*

^f *Laboratoire de l'Accélérateur Linéaire (LAL), Orsay, France*

^g *LPNHE, Université Pierre et Marie Curie Paris 6, Université Denis Diderot Paris 7, CNRS/IN2P3, Paris, France*

^h *LPSC, Université Joseph Fourier Grenoble 1, CNRS/IN2P3, Institut National Polytechnique de Grenoble, Grenoble, France*

ⁱ *Max-Planck-Institut für Physik, Werner-Heisenberg-Institut, Munich, Germany*

^j *Dipartimento di Fisica dell'Università di Milano and INFN, Milano, Italy*

^k *University of Pittsburgh, Pittsburgh, Pennsylvania, USA*

^l *University of Regina, Regina, Saskatchewan, Canada*

^m *CEA, DSM/DAPNIA, Centre d'Etudes de Saclay, Gif-sur-Yvette, France*

ⁿ *Southern Methodist University, Dallas, Texas, USA*

^o *Institute of Physics, Academia Sinica, Nankang, Taipei, Taiwan, ROC*

¹ *Now at Department of Physics, Florida State University, Tallahassee, Florida, USA*

² *Now at Department of Electrical and Computer Engineering, University of Saskatchewan, Saskatoon, Saskatchewan, Canada*

³ *Also at TRIUMF, Vancouver, British Columbia, Canada*

⁴ *Also at Institute of Experimental Physics, Kosice, Slovakia*

⁵ *Now at Christian Albrechts University, Institute for Experimental and Applied Physics, Kiel, Germany*

⁶ *Now at CERN, Geneva, Switzerland*

⁷ *Now at IBM T.J. Watson Research Center, USA*

⁸ *Now at LAL*

⁹ *Now at LPSC*

¹⁰ *Now at CPPM*

¹¹ *Now at TRIUMF*

ABSTRACT: The ATLAS detector has been designed for operation at CERN's Large Hadron Collider. ATLAS includes a complex system of liquid argon calorimeters. The electronics for amplifying, shaping, sampling, pipelining, and digitizing the calorimeter signals is implemented on the Front End Boards (FEBs). This paper describes the design, implementation and production of the FEBs and presents measurement results from testing performed at several stages during the production process.

*Corresponding author. Email: parsons@nevis.columbia.edu

Contents

1. Introduction	2
2. FEB Specifications	2
3. Overview of the FEB Architecture	5
4. Analog Signal Processing	7
4.1 Preamplifiers	7
4.2 Hadronic Endcap Preshapers	9
4.3 Shaper	10
4.4 SCA Analog Pipeline	13
4.5 Digitization and Gain Selection	18
5. SCA Controller and FEB Digital Control	19
6. SCA Control and Address Bus	22
7. Output Optical Link	23
8. FEB Output Data Format	24
9. Level 1 Trigger Summing	27
10. FEB Control Interfaces	31
10.1 Trigger and Timing Control	31
10.2 Clock Distribution	32
10.3 Configuration and Slow Control	33
11. FEB Power Distribution and Monitoring	34
12. FEB Layout and Topology	35
13. FEB Production	38
14. FEB Testing	40
14.1 Initial Test and HASS Test	40
14.2 Digital Testing	41
14.3 Analog Testing	44
15. Performance of the FEB	48

1. Introduction

ATLAS [1] is a large general-purpose detector designed for operation at the Large Hadron Collider (LHC) at CERN. The LHC is a proton-proton collider which will operate with a center-of-mass energy of 14 TeV. A system of liquid argon (LAr) calorimeters forms one of the major ATLAS detector systems. The LAr calorimeters include the electromagnetic barrel (EMB) calorimeter, which is housed in the central cryostat and provides coverage for pseudorapidities $|\eta| < 1.5$. The ranges of larger pseudorapidity are covered by endcap calorimeter (EC) systems, housed in separate endcap cryostats. Each EC cryostat includes an electromagnetic endcap (EMEC) calorimeter, a hadronic endcap (HEC) calorimeter, and a forward calorimeter (FCAL) providing coverage up to $|\eta| = 4.9$. More details about the design, construction, and performance of the calorimeters themselves can be found in Reference [2].

The electronic readout of the ATLAS LAr calorimeters is divided into a Front End (FE) system of boards mounted in custom crates directly on the cryostat feedthroughs, and a Back End (BE) system of VME-based boards located in an off-detector underground counting house. The FE system includes Front End Boards (FEB), which perform the readout and digitization of the calorimeter signals, Calibration boards which inject precision calibration signals, boards for producing analog sums for the Level 1 (L1) trigger, and other control and monitoring boards. The BE electronics includes Readout Driver (ROD) boards which perform DSP-based digital filtering of the signals.

A more detailed overview of the entire ATLAS LAr readout chain, and details of the BE system, can be found in Reference [3]. The overall FE system architecture is documented in Reference [4]. The purpose of this note is to describe the design and implementation of the ATLAS LAr FEB, including some information on the standalone performance of the FEB. The overall system performance of the ATLAS LAr readout electronics will be documented in a subsequent publication.

2. FEB Specifications

The electronic readout of the ATLAS LAr calorimeters faces demanding specifications, a summary of which is provided in Table 1. Given the fine segmentation of the ATLAS LAr calorimeters, a total of 182468 active channels must be read out. These are distributed among the EMB (60%), EMEC (35%), HEC (3%) and FCAL (2%). For noise optimization, the FEBs are placed in crates mounted directly on the calorimeter cryostat feedthroughs. A channel density of 128 channels per FEB is required to satisfy the space constraints of this on-detector location. With this density, and given the cabling of the calorimeters, a total of 1524 FEBs is required to read out the LAr calorimeters. Any heat dissipated by the electronics must be removed in order to not impact other detector subsystems. Therefore, maintaining a low power consumption, well below 1 Watt per channel, is important. Furthermore, since convective air cooling using fans cannot be used in this

region of the detector, the FEBs are conductively cooled using an under-pressure water cooling system.

The FEBs must sample the LAr calorimeter signals at the LHC bunch crossing frequency of 40 MHz. The online selection of events is performed by the three-layered ATLAS trigger system [5]. In the first step, the L1 trigger uses partial information from the calorimeters and the muon system to reduce the trigger rate to a maximum L1 Accept rate of 75 kHz. For each L1 trigger, the FEB readout of typically five samples per channel must be accomplished with little deadtime, and the full event data made available for subsequent evaluation by the Level 2 and Level 3 trigger systems. The L1 trigger decision is performed by a deadtimeless, synchronous pipelined system with fixed latency of up to $2.5 \mu\text{s}$ (100 bunch crossings), during which the FEB must store the time-sampled calorimeter signals. The FEB must be able, at lower trigger rates, to read out up to 32 samples per channel, in order to measure the entire waveform when desired during calibration or other special runs.

Table 1. A summary of some of the main specifications of the ATLAS LAr Front End Board.

Channel count	Total Active Channels	182,468
	Number of FEBs	1524
	Channel density	128 channels/FEB
	Power consumption	$\approx 0.7 \text{ W/channel}$
Signal sampling	Sampling frequency	40 MHz
	L1 trigger latency	$< 2.5 \mu\text{s}$ (100 b.c.)
	Max. L1 trigger rate	75 kHz
	Samples to read/channel	Typically 5, Maximum 32
	Deadtime	$< \text{few } \%$
Energy measurement	Dynamic range	$\approx 17 \text{ bits}$
	Calibration uncertainty	$< 0.25\%$
	Noise/channel	$\approx 10 - 50 \text{ MeV}$
	Coherent noise/channel	$< 5\%$ of total noise/channel
Time measurement	Resolution	$< 100 \text{ ps}$
Trigger sums	Gain uniformity	$< 5\%$

The energy resolution of the ATLAS electromagnetic (EM) calorimeters for EM showers is given by $\sigma(E)/E \approx 10\%/\sqrt{E} \oplus 0.7\%/E$, with the energy E is expressed in GeV and \oplus representing addition in quadrature. The resolution at high energies is dominated by the second, so-called constant, term, which arises due to a variety of contributions, including mechanical tolerances and variations in the calorimeter construction, time variations of the detector response, and uncertainties in the calibration. Achieving the goal of 0.7% for the constant term requires that the calibration of the readout is understood at a level of better than 0.25%.

For low energies, the noise contribution is also important. The noise for an electron shower in the EMB at low pseudorapidity is expected to be 190 MeV. Depending on the section of the calorimeter, the noise of the preamp loaded by the detector capacitance for a single calorimeter channel is typically $\approx 10 - 50 \text{ MeV}$. The low end of the FEB dynamic range is set by the requirement that the FEB noise be dominated by the preamp noise. The high end of the scale has been set to \approx

3 TeV, the largest energy that could be deposited in a single channel of the EM calorimeter due to an electron originating from the decay of a 5 TeV Z' boson. Being able to cover this entire energy range requires the FEB have a dynamic range of ≈ 17 bits.

Given the fine granularity of the calorimeter, the energy of an EM shower is obtained by summing the signals of typically 50-100 calorimeter cells; even more channels must be summed for a typical jet or for more global quantities such as missing transverse energy (MET). As a result, any coherent component of the noise is particularly critical. By studying the impact on the reconstruction, using MET, of a heavy SUSY Higgs boson decaying to tau pairs [6], a requirement has been set that the coherent noise per channel be less than 5% of the total noise per channel.

In addition to measuring the deposited energy per channel, the FEB readout must provide a measurement of the time of the deposition and also a χ^2 -type quantity related to how well the pulse shape follows the expected shape. Digital processing can be performed of the typically five time samples read out per channel to produce optimized measures of energy, time, and χ^2 . The χ^2 measurement provides sensitivity to pulses which are mismeasured due to waveform distortions produced by large energy depositions in neighboring bunch crossings. The time measurement provides, first of all, identification of the bunch crossing to which the energy deposit belongs. A timing resolution of ≈ 5 ns would be sufficient for this purpose given the time separation of 25 ns between subsequent bunch crossings. However, the calorimeter has a much better intrinsic timing resolution, and there are physics arguments to aim for excellent timing resolution. For example, some variants of Gauge Mediated Supersymmetry Breaking (GMSB) models imply existence of a heavy gaugino which would decay to an undetected Graviton and a photon which, due to the finite lifetime of the gaugino, would hit the calorimeter with a slight delay and from an angle not directed to the primary vertex. A study [6] of such a scenario demonstrated that the pointing resolution of the calorimeter would be a powerful tool to identify such “non-pointing” photons, and a timing resolution much better than the mean delay of ≈ 2 ns would be helpful in confirming this signature. For large pulses, a timing resolution better than 100 ps should be achievable. Such a performance would also allow the possibility to use timing to help select the proper vertex in Higgs decays to a pair of photons, and to achieve some rejection of “pile-up” due to other proton-proton scatters within the same bunch crossing.

Due to their on-detector location, the FEBs must be able to tolerate significant levels of radiation. Simulations of the expected radiation field [7] predict that, during 10 years of LHC operation at design luminosity, the FEBs will experience a total ionizing dose (TID) of 5 kRad and non-ionizing energy loss (NIEL) effects due to an equivalent fluence of 1.6×10^{12} 1-MeV-neutrons/cm². In addition, the possibility of single event effects (SEE), including single event upsets (SEU) of logic states, must be considered, given the exposure to 7.7×10^{11} hadrons/cm² with energies above 20 MeV. As will be described in a subsequent publication, all FEB components were subjected to an extensive radiation tolerance qualification process, with the required tolerance specified by multiplying these expected radiation levels by three safety factors (SF) which aim to account for uncertainties in the simulated radiation levels, for possible low dose rate effects, and for variations in radiation tolerance from lot to lot in production of an ASIC. The SF values were lower for the DMILL process [8], which was guaranteed by the vendor to be radiation tolerant, than for “commercial-off-the-shelf” (COTS) components or ASICs produced in a commercial process that does not guarantee radiation tolerance. Table 2 summarizes the expected radiation levels as well

as the “Radiation Tolerance Criteria” (RTC) values to which DMILL and other components had to be tested and qualified. The RTC values for commercial processes take into account the fact that the final production parts were purchased from known, homogeneous lots; otherwise, each of the commercial RTC values would have been an additional factor of two higher.

Table 2. Summary of the radiation levels to which components for the FEB had to be tested and qualified. For more details, see the text.

Radiation		Estimated Level	DMILL RTC	Commercial Process RTC
Type	Units			
TID	Gy	50	525	1700
NIEL	1 MeV equiv. n/cm ²	1.6×10^{12}	1.6×10^{13}	1.6×10^{13}
SEE	Hadrons (> 20 MeV)/cm ²	7.7×10^{11}	7.7×10^{12}	7.7×10^{12}

The radiation tolerance issues led to the development of a number of custom ASICs in specialized, radiation-tolerant semiconductor processes, and to a very limited use of COTS components. Table 3 summarizes the process technologies utilized for the main active components on the FEBs. A number of ASICs were developed in the DMILL process, and several more using a commercial 0.25 μm “deep submicron” (DSM) process, but using a special library which was radiation hardened through the use of a custom-developed enclosed transistor geometry [9].

The on-detector location of the FEBs implies that there is no access to them during operation of the LHC, and gaining access requires a shutdown of considerable length. In particular, the FEBs of the EMB detector are mounted in crates at either end of the barrel cryostat. Accessing these crates to repair or replace the FEBs requires rearranging or extracting large detector elements in the endcap regions, including the endcap calorimeter cryostats. While it is foreseen to access these crates once per year, during the long annual shutdown planned for the LHC, other opportunities for access will be very limited. Given these access difficulties, reliability is a key concern.

The limited use of COTS components, and reliance on specialized radiation-tolerant semiconductor processes, implies that very little is known about the expected lifetime of the individual components. All components were purchased with $\approx (5\text{-}8)\%$ spares, to allow for repairs. As is often the case for high energy physics instrumentation, essentially all FEB components are already obsolete and unavailable, so procurement or production of additional spares is not feasible. This fact, plus the current planning to upgrade the luminosity of the LHC in the future and extend even further its expected operations phase, makes the FEB lifetime issue even more critical. An accelerated lifetime test is being performed by operating a number of pre-production FEBs at elevated temperature in an environmental chamber. Results of this test, once available, will be reported in a subsequent publication.

3. Overview of the FEB Architecture

A block diagram indicating the main features of one four-channel element of the FEB architecture is shown in Figure 1. The raw signals from the calorimeter are mapped onto the FEB inputs as they emerge from the cryostat feedthroughs. On the FEB, the signals are first subject to several stages of analog processing. Preamplifier hybrids amplify the raw signals. To reduce the dynamic

Table 3. The acronyms of the main active components of the FEB, the number of each component per 128-channel FEB, and the component’s functionality. The components are grouped according to the semi-conductor technology used in their production.

Production Process	Component Name	Number per FEB	Description of Functionality
Hybrid	Preamp	32	First stage amplification
	Preshaper	32	Amplification, preshaping for HEC
AMS BiCMOS	Shaper	32	Amplification and shaping
DMILL	SCA	32	Sampling and analog pipeline
	SMUX	1	32:16 multiplexor
	SPAC slave	1	Serial control interface
	CONFIG	1	Configuration controller
	TTCrx	1	Trigger and timing control receiver
DSM	GSEL	8	Gain selection, data formatting
	CLKFO	7	Clock fanout
	SCAC	2	SCA controller
	DCU2	2	Temperature and voltage monitor
	QPLL	1	Quartz-crystal phase-locked loop
STm RHBip1	VREG	19	Radiation-tolerant voltage regulator
COTS	OpAmp	32	Match SCA output to ADC input
	ADC	16	12-bit digitization
	GLINK	1	1.6 Gbps Serializer
	OTx	1	VCSEL-based optical transmitter

range requirements of the sampling and digitization stages, the preamp outputs are split and further amplified by shaper chips to produce three overlapping linear gain scales which are each subject to fast analog shaping. The shaped signals are then sampled at the LHC bunch crossing frequency of 40 MHz by switched-capacitor array (SCA) analog pipeline chips, which store the signals in analog form during the L1 trigger latency.




Figure 1. Block diagram of the FEB architecture, depicting the data flow for four of the 128 channels.

For events accepted by the L1 trigger, typically five samples per channel are read out from the SCA using the optimal gain scale, and digitized using a 12-bit ADC. The digitized data are formatted, multiplexed, and then transmitted optically out of the detector to the ROD via a single 1.6 Gbps optical output link per FEB.

The FEB implements the first two stages in the summing tree producing analog sums for the

L1 trigger. The shaper chips sum their four input channels. The shaper sum outputs are then summed further on plug-in Layer Sum Boards (LSB) mounted on the FEB. The LSB outputs are sent off of the FEB to trigger boards mounted in the same crate, which perform further summing and processing.

For configuration and operation, the FEBs require a number of external control signals. The 40 MHz LHC clock, as well as the L1 trigger signal and a few other signals synchronous with the clock, are delivered via the ATLAS Trigger and Timing Control (TTC) system [10]. The FEB is configured and monitored via a dedicated “Serial Protocol for ATLAS Calorimeters (SPAC)” [11] serial control link, operating at 10 MHz.

The following sections describe each of the blocks on the FEB architecture in more detail.

4. Analog Signal Processing

As depicted in Figure 1, and described in more detail in the following, the analog processing performed by the FEB includes amplification, shaping, storage in an analog pipeline, and then digitization for events accepted by the L1 trigger.

4.1 Preamplifiers

For the ATLAS EM calorimeters (both barrel and endcaps), the raw calorimeter signals are cabled out of the detector and received on the FEBs. Preamps mounted on the FEB provide the first step of analog processing. A similar situation applies to the FCAL signals, except that there is one intermediate step in which transformers are used to sum the raw signals into the final towers, in order to reduce the channel count.

The raw calorimeter signal is an approximately triangular current pulse having a risetime of a few ns and a falling edge extending for the total drift time (≈ 400 ns) of electrons in the LAr gaps. The response of the EM calorimeters to an electron results in a current of typically $\approx 2 - 3 \mu\text{A}/\text{GeV}$. The dynamic range of interest corresponds, therefore, to currents in the range of nA up to several mA.

The LAr preamps [12] are coupled to the detector by a transmission line. As the signal duration is long compared to the shaping time, current preamps are used which provide a voltage output directly proportional to the input current. The principle of coupling a preamp to a high capacitance detector is described in Reference [13].

The preamp circuit schematic is shown in Figure 2. The main characteristic of the ATLAS preamp is the use of a local feedback in the input stage to attribute the functions of low noise and high dynamic range to two different transistors. This circuit configuration allows excellent linearity and noise performance, with relatively low power (≈ 50 mW per channel). The gain (i.e. the transresistance) and the input impedance can be chosen independently without changing the power supply voltages and power dissipation.

In order to optimize the match of the preamp characteristics to the calorimeter cell to which it is connected, three different variations, or “flavors”, of preamp are used. As summarized in Table 4, the flavors differ in terms of their input impedance, gain (ie. transimpedance), and dynamic range. These variations are achieved by using the appropriate values of components R2, R12, L1, and L2 in the schematic of Figure 2. Flavor 1 is used in the presampler and in the front sections of the EM

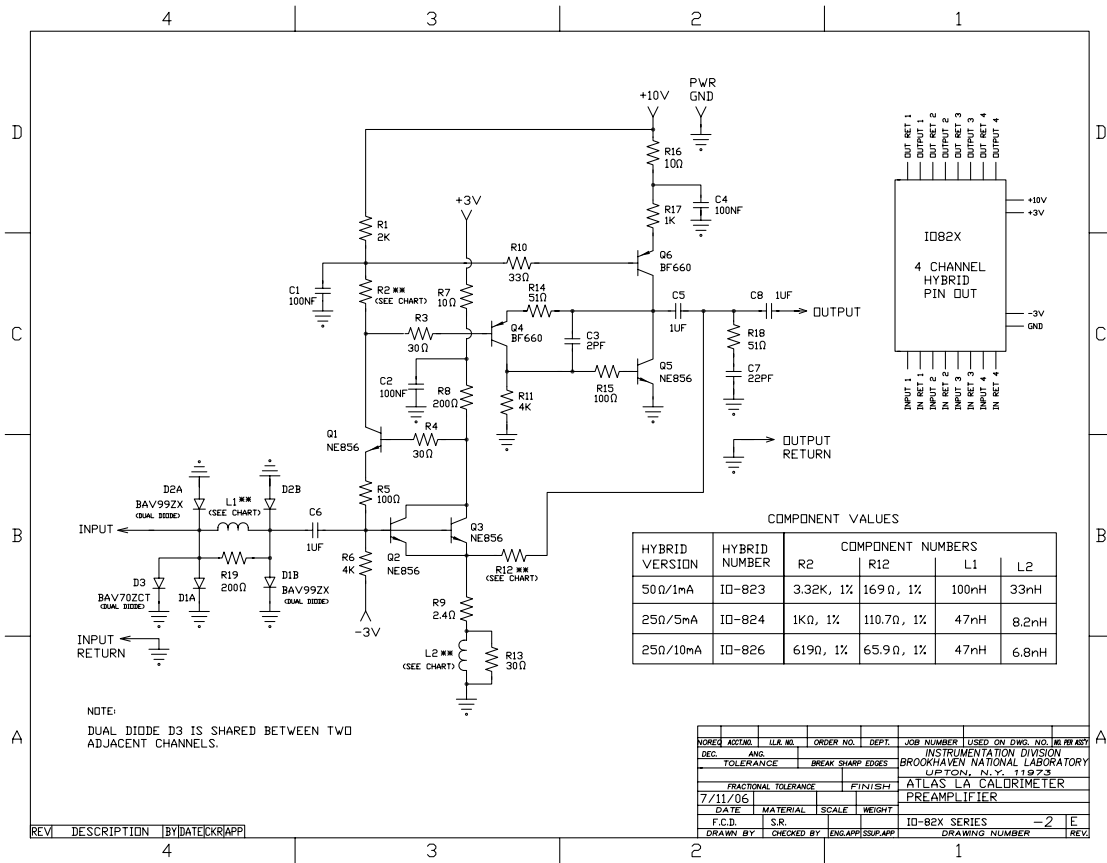


Figure 2. Schematic of one channel of the four-channel preamplifier hybrid.

calorimeter, Flavor 2 is used for the FCAL as well as the EM middle/back sections for lower $|\eta|$, and Flavor 3 is used for the EM middle/back sections for higher $|\eta|$.

Table 4. Characteristics of the three different flavors of preamps.

Preamp Flavor	Input Impedance (Ω)	Transimpedance (Ω)	Maximum Input Current (mA)	Typical Detector Capacitance (pF)
1	50	3k	1	160 - 500
2	25	1k	5	300 - 2000
3	25	500	10	400 - 2000

The preamp is realized as a four-channel thick-film hybrid, measuring 55 mm \times 23 mm. The hybrid is equipped with gold-plated pins that plug into matching sockets mounted on the FEB PCB, allowing customization of the FEBs for the different calorimeter sections by insertion of the appropriate preamp hybrids.

Production of the preamps was split approximately equally between Brookhaven National Lab (BNL) and INFN Milano. Each site procured all components for their local production, and some differences resulted. Measurements performed during the pre-testing of the hybrids revealed

slightly different performance parameters [14]. The most significant difference was in the gains. For example, Figure 3 shows the measured gains for the $25\ \Omega$ 5 mA version of the preamp, for which two different distributions are clearly visible. The lower gaussian corresponds to the INFN preamps, with a mean gain which is 1.7% below that of the BNL preamps. Each distribution separately has a Gaussian width of $\sigma \approx 0.5\%$. Within each lab, an acceptance window of $\pm 2\%$ was applied during preamp testing. For the precision readout, each channel is separately calibrated and therefore matching of the absolute gain is not relevant. However, to ensure sufficiently uniformity for the L1 trigger sums, the gains are required to be uniform within $< 5\%$. As shown in the figure, less than 2% of the hybrids would fall outside a $\pm 2\%$ acceptance window applied to the combined distribution, and all are well within the L1 specification. The slight differences, therefore, are not considered an issue.

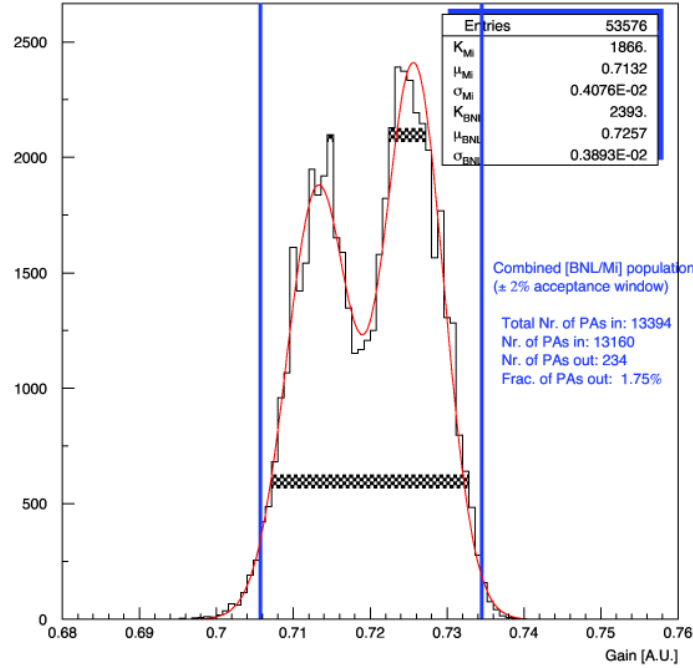


Figure 3. Distribution of gains measured for the $25\ \Omega$ 5 mA version of the preamplifier hybrids. For more details, see the text.

4.2 Hadronic Endcap Preshapers

The Hadronic Endcap (HEC) calorimeters represent the one LAr subsystem where preamps are mounted directly on the detector in the liquid argon. The HEC cryogenic electronics are described in Reference [15]. The preamp outputs are driven on cables out of the cryostat to the FEBs.

Instead of the plug-in preamp hybrids used for the rest of the LAr subsystems, the HEC FEBs are equipped with plug-in “HEC preshapers” [16] that are designed to be pin-to-pin compatible with the preamp. A schematic of the HEC preshaper is shown in Figure 4. The roles of the HEC preshapers include providing a pole-zero cancellation to adapt to the widely varying HEC detector capacitance, and inverting, amplifying and pre-shaping the signal so that the input to the shaper is

the same polarity and approximately the same shape as for the rest of the LAr calorimeters. With these adaptations, the same FEBs are used for all LAr calorimeters.

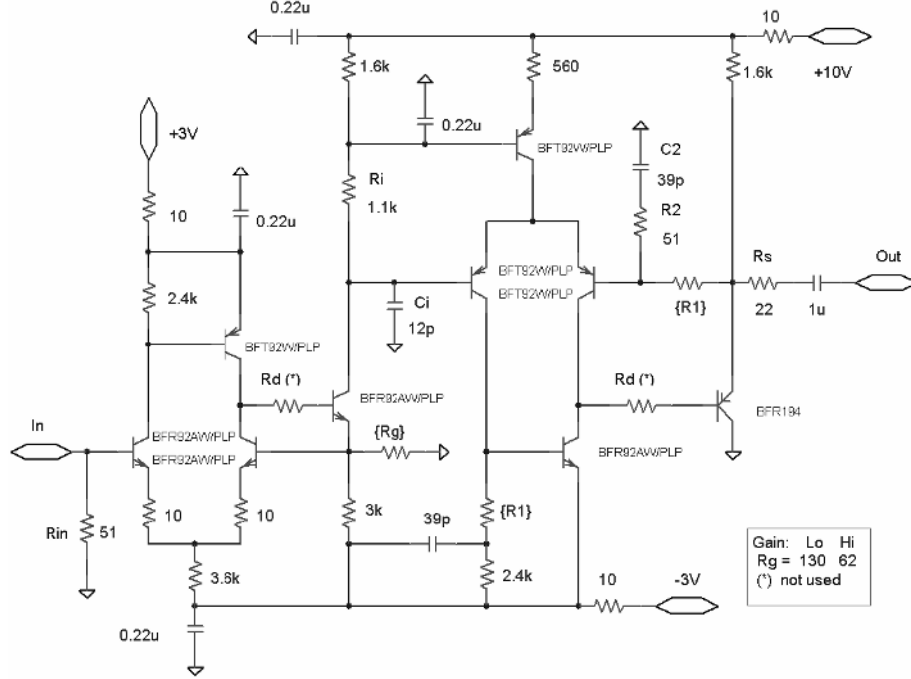


Figure 4. Schematic of one channel of the four-channel HEC preshaper hybrid.

To optimize the match of the preshaper to the HEC channel to which it is connected, 14 different preshaper pole-zero time constants are needed. In fact, since the preshapers are mounted on both sides of the FEB, a total of 28 different preshaper versions is required. Depending on the value of component R_G , the preshaper has a nominal gain of either 6 (“LO”) or 12 (“HI”). Figure 5 shows the measured gains for 168 hybrids, which are in good agreement with the design values. The dispersion among hybrids has an RMS of less than 1.5%, satisfying the uniformity requirement set by the L1 trigger. Figure 5 also shows the measured noise as a function of detector capacitance, with the results of Spice simulations superimposed on the measurements; good agreement is seen between the noise levels achieved and those expected by simulation.

4.3 Shaper

The preamp/preshaper outputs are AC-coupled into a four-channel shaper chip [17], which applies an $CR-(RC)^2$ analog filter to the signals in order to optimize the signal-to-noise ratio. The one differentiation serves to remove the long tail from the detector response, while the two integrations limit the bandwidth in order to reduce the noise. The $CR-(RC)^2$ filter architecture, with transfer function $H(s) = (\tau s)/(1 + \tau s)^3$, is a good compromise between the number of stages (and corresponding power consumption) and the performance, since it is only 10% worse than an ideal filter.

The RC time constant of the shaping function is set to 13 ns, representing a compromise between minimizing pile-up noise, which increases for slower shaping, and thermal noise, which

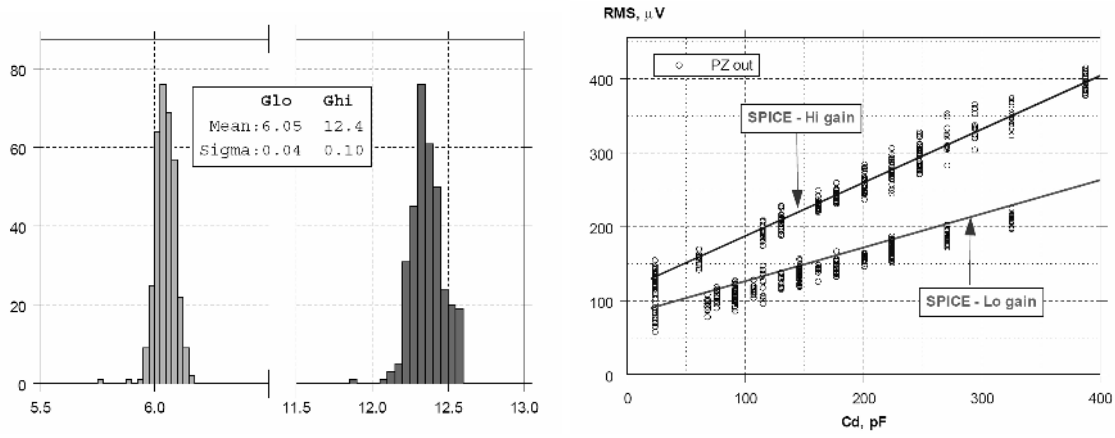


Figure 5. Results of measurements of a sample of 168 HEC preshapers. The left plot shows the measured gains, while the right plot shows the measured and simulated noise levels as a function of detector capacitance.

decreases for slower shaping. The characteristic form of the resultant output signal shape is shown in Figure 6(a). For the typical detector capacitance, the peaking time of the signal when convoluted with a 450 ns triangular input current is ≈ 35 -40 ns. Figure 6(b) demonstrates that this value of peaking time is close to optimal for the LHC design luminosity of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. For lower luminosities, further downstream digital filtering can be used to achieve an effectively longer shaping time, and subsequently better noise performance. Another consideration is that the shaped signals are sampled at the 25 ns LHC bunch-crossing rate, so even faster shaping would start to suffer from aliasing effects.




Figure 6. (a) Shape of the signal output from the shaper chip. The dots indicate the position of samples separated by 25 ns. (b) Total noise versus shaping time for luminosities of 10^{33} and $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. Also shown are the separate contributions from pile-up noise and from electronics (series and parallel) noise.

In addition to shaping, the shaper provides additional amplification, so that the preamp noise

will dominate over the contributions from later stages. To achieve the full dynamic range, the shaper splits each signal to provide three overlapping linear gain scales, with gain ratios of ≈ 10 . The absolute gains are 0.8 (LO gain), 8.4 (MED gain) and 82 (HI gain). The shaper also includes one “Dummy” output which has no input signal but the same output stage as the analog signals. The Dummy channel thereby provides a measurement of the on-chip noise. As described in Section 4.4, this signal is connected to the input Reference channels of the SCA analog pipeline in order to provide a “pseudo-differential” architecture, with the goal of reducing low frequency noise.

The shaper chip also incorporates a “Linear Mixer” (LM) stage which provides an additional output corresponding to a shaped sum of the four input channels. The channels are summed before shaping, and then the sum is subject to a CR-RC filter before being output. Depending on how the “Gain” pin on the shaper package is connected on the FEB PCB, the LM gain can be set to either one or three. As described in more detail in Section 9, the LM output signals are used in the generation of analog sums for use by the L1 trigger system. Programmable switches are used to enable or disable individual channels in the trigger sum, allowing the disabling of noisy or faulty channels. The values of the switches are set, and can be read back for verification, via the SPAC slow control link to the FEB.

The architecture of the shaper chip is depicted in Figure 7. As shown, the CS0 through CS3 inputs are connected to fuses which can be blown to connect additional capacitors in the shaping stages. Their purpose is to allow tuning of the RC time constant in the presence of unavoidable process variations in the values of the on-chip resistors and capacitors.

The shaper is realized as a four-channel ASIC implemented in the AMS 1.2 μm BiCMOS technology [18], and is packaged in a 100-pin rectangular QFP package with 0.65 mm pitch. Some of the relevant shaper parameters and performance figures are summarized in Table 5.

Table 5. Some of the parameters and performance measures for the shaper ASIC.

Parameter		Value
Number of channels		4
Die size		$4 \times 4.5 \text{ mm}^2$
Voltages	VDD	4.5 V
	VSS	-3.0 V
Power consumption		500 mW total
Output signal swing		-2 V to +3.5 V
Input impedance		$50 \Omega \pm 15\%$
Gain	HI	82
	MED	8.4
	LO	0.8
Noise (μV)	HI	850
	MED	400
	LO	250
Linearity		$\pm 0.2\%$
Crosstalk		$< 0.1\%$

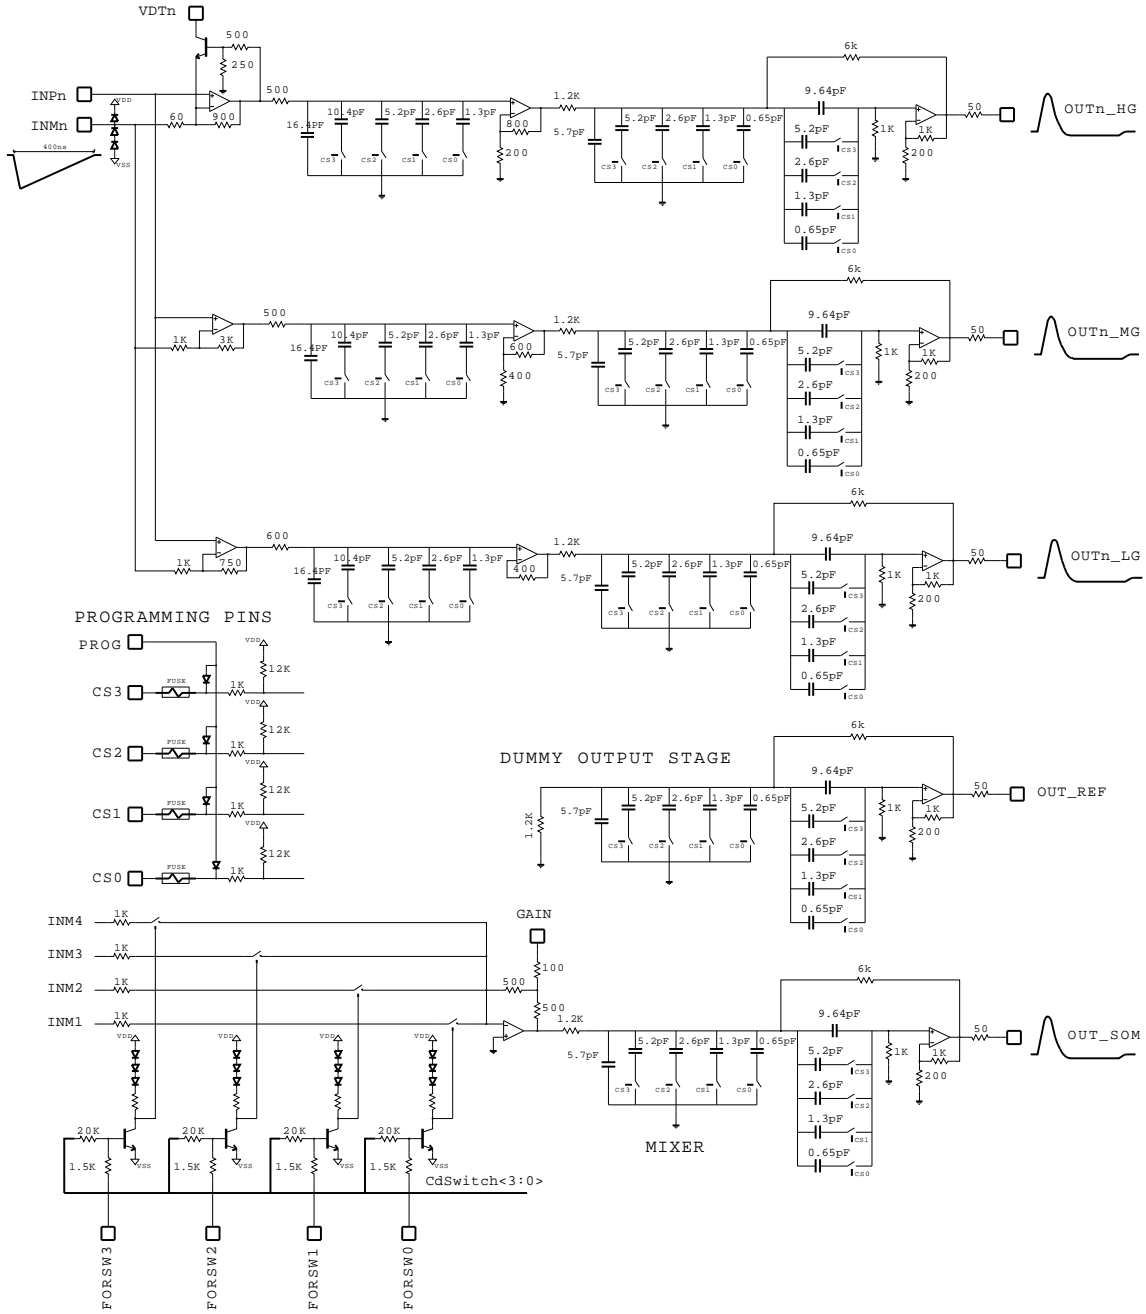


Figure 7. Block diagram of the shaper architecture. One of the four identical signal channels is shown in the upper half of the figure, while the fuse programming pins, Dummy output and Linear Mixer circuits are shown in the lower half.

4.4 SCA Analog Pipeline

The shaper output signals are sampled at 40 MHz and stored in analog form by a switched-capacitor array (SCA) analog pipeline. The SCA stores the analog signals during the L1 trigger latency in pipelines of 144 cells, and further serves as a multiplexor and de-randomizing buffer in front of the ADC for triggered events. The SCA is designed to allow simultaneous Write and Read operations.

Separate 8-bit Write addresses (WADD) and Read addresses (RADD), as well as a 40 MHz Write clock (WCLK) and 5 MHz Read clock (RCLK) must be provided from off-chip to control the operation of the SCA.

Each SCA chip processes all three gain scales for each of four calorimeter channels. In addition, an extra so-called “Reference” channel is associated with each group of three gains corresponding to a calorimeter channel, so the SCA chip contains a total of 16 analog pipeline channels. On the FEB PCB, the inputs of the four Reference channels are connected together to the “Dummy” output of the shaper. The Reference channel is treated by the SCA exactly as the other channels and its output is subtracted from the signal output during the Read operation. This subtraction is performed off-chip. This pseudo-differential operation is employed to reject the major part of the coherent noise generated before and inside the chip such as clock feedthrough and couplings through the substrate. It also improves the power supply rejection ratio (PSRR) of the SCA chip.

The output of the shaper is DC-coupled to the input of the SCA. As a result, the SCA is asymmetrically powered ($V_{SS}=-1.7$ V, $V_{DD}=+3.3$ V) to deal with input signals in the range from -0.9 V to $+2.5$ V with a baseline voltage of 0 V.

The architecture of the SCA ASIC is depicted in Figure 8. More details about the SCA design can be found in Reference [19]; here we describe the main design features and characteristics.

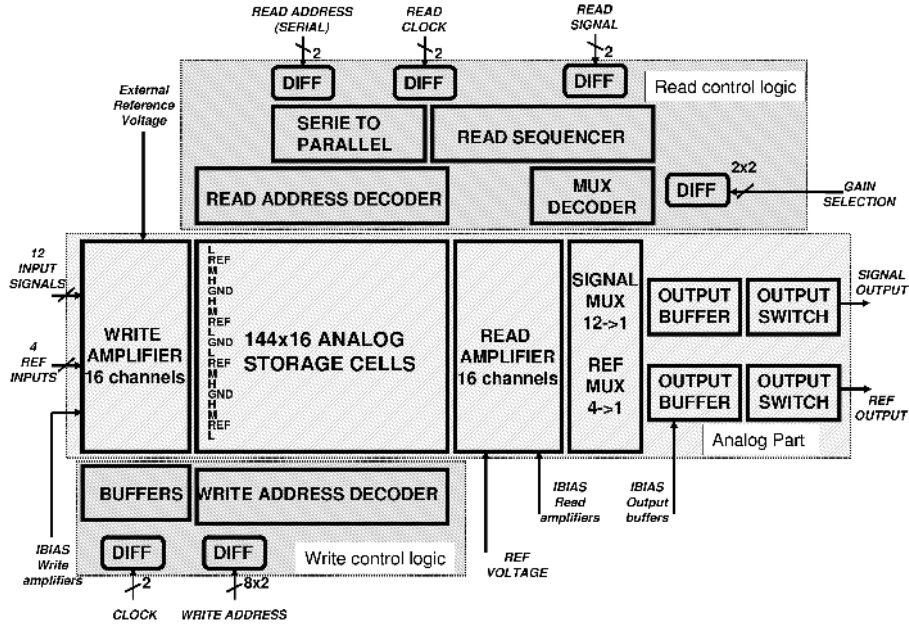


Figure 8. Block diagram depicting the architecture of the SCA.

SCA analog pipeline architecture

A more detailed depiction of the structure of one channel of the SCA analog pipeline is shown in Figure 9. Each pipeline contains 144 identical cells. Each cell comprises a ≈ 1 pF storage capacitor, and separate switches for Writing (switches S1 and S2 in Figure 9) and Reading (switches S3 and S4). Switches S1 and S3, which have to deal with the input signal dynamic range, are

CMOS switches, whereas S2 and S4 are simple NMOS switches as they are connected to a fixed intermediate level.

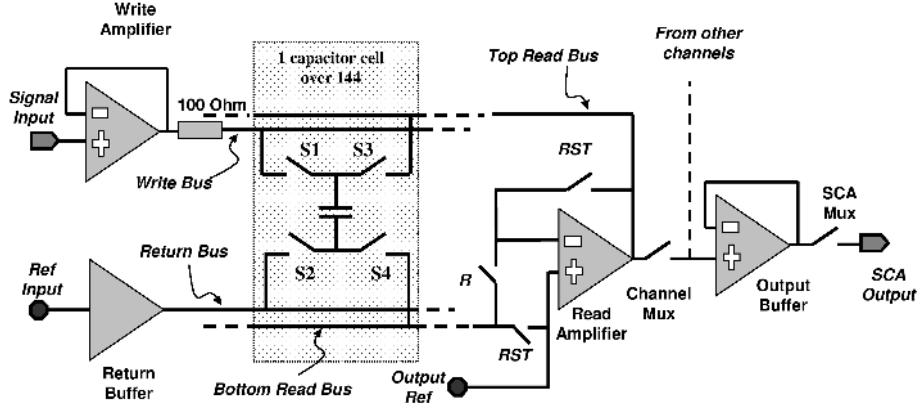


Figure 9. Block diagram of the SCA analog pipeline design.

The input signal is connected to the Write bus through an input buffer amplifier, the aims of which are to limit the voltage swing of the input signal, to present a constantly low capacitive load to the shaper output, and to decrease drastically the level of crosstalk by buffering the flow of signal currents on and off of the SCA chip. The input buffer is an operational amplifier connected as a voltage follower. To ensure stable operation, a 100 Ω p+ diffusion resistor is placed in series between the amplifier output and the Write top bus.

Similarly, another buffer, connected to the return bus, reduces cross-talk by limiting currents from the common reference pin of the chip. This buffer does not have a wide dynamic range requirement, and so a simple NMOS source-follower is used for this purpose. The common input reference voltage of these 16 followers is generated by a servo-control system which connects the Write return bus to a reference voltage, denoted by VREF.

During a Write operation, S1 and S2 are closed and the voltage across the capacitor tracks the difference of voltage between the Write bus and the return bus. The effective capacitor charging time constant is ≈ 2 ns, set by the value of the capacitance and the ON resistance of the switches. The input signal voltage is sampled and stored on the capacitor at the moment the Write switches are opened again. For a Read operation, the corresponding sampling capacitor is put in the feedback of the Read amplifier by closing S3 and S4 after the Read bottom bus has been reset. To reduce the sensitivity to noise injected in the substrate, the top plate of the storage capacitor is the one connected to the sensitive input node of the Read amplifier during the Read operation whereas the bottom plate, which presents a parasitic capacitance to back silicon, is the one connected to the output of this amplifier. As the Write and Read operations are both performed in voltage mode, the gain of the SCA is very close to unity and is, to first order, independent of the value of the storage capacitance.

The Read amplifier has the same open-loop structure as the Write amplifier but uses a PMOS transistor input pair to provide lower white and 1/f noise. Before reading a storage cell, the residual charge stored on the parasitic capacitance of the bottom Read bus needs to be cleared to avoid the mixing of two consecutive signals. During this Reset phase, the RST switches are closed, setting

the Read amplifier in voltage-follower mode and discharging the bottom Read bus, while Read switch R is opened, insulating the Read amplifier input from the bottom Read bus capacitance to avoid oscillations. This phase lasts at least 100 ns to allow discharging of the bottom Read bus with sufficient precision. At the end of the Reset phase, the RST switches are opened just before switch R is closed.

At the output of the Read amplifiers, the 12 signal channel outputs are multiplexed by CMOS switches towards the signal output buffer. In the same way, the four Reference channels are multiplexed towards the Reference output buffer. The commands of the switches are generated by a sequencer within the digital Read logic section. When no channel is selected, the inputs of the two output buffers are clamped by PMOS transistors to VREF.

The output buffers have the same structure as the Write amplifiers (NMOS input), and are designed to be able to drive up to 30 pF for the 5 MHz read-out frequency. To allow multiplexing of two SCA chips to the same external ADC, the output buffers are insulated from the output pads by CMOS switches. The ON resistance of these switches is 400 Ω and varies by about $\pm 15\%$ along the SCA voltage dynamic range. Achieving a good linearity requires that the electronics connected to the SCA outputs have a high input resistance, at least 100 k Ω . As described in Section 4.5, commercial op-amps configured as voltage followers are used for this purpose. Moreover, the outputs of the SCA are never left in a high-impedance state; one of the paired SCA chips is configured to always drive the output lines except for when the other SCA is being read out.

SCA Write and Read Control Logic

As shown in Figure 8, in addition to the analog pipeline section, the SCA chip contains separate digital sections for the Write and Read control logic. These are physically separated from each other and from the analog pipeline section. All of the building blocks were developed with full-custom designs. All the flip-flops and registers use gated-inverter-based master-slave flip-flops.

As described later in Section 6, the WADD, RADD, and other control signals are transmitted to the SCA chips via a control bus using low voltage, differential signals. The differential signals are received on the SCA chip via the “DIFF” blocks depicted in Figure 8, which use CMOS comparators to convert the inputs to differential CMOS signals. The comparators have NMOS inputs in order to be compatible with the common mode voltage of the SCA control bus drivers.

The Write control signals include the 8-bit WADD and the 40 MHz WCLK. On the falling edge of WCLK, the WADD bits are stored in a register. The 8-bit output of this register is pre-decoded into four groups of four signals, two bits at a time, by an asynchronous AND gate based decoder. Inside each group, each of these four signals flags one of the four possible values of the two corresponding bits. The 16 resulting lines are then bused to 144 synchronous four-input decoders, designed such that the timing of the Write operation is defined solely by the rising edge of WCLK and is not address dependent.

At the output of the WADD decoders, the signal is buffered to drive the gates of the NMOS transistors of the S1 and S2 Write switches. The same signals are also inverted and then buffered with a delay of about 2ns to drive the PMOS transistor in S1. Thus, the opening of switch S2 defines the time of the Write operation. Since S2 is connected to a quasi-invariant voltage (the return bus), both the time of the sampling and the charge injected into the capacitor during the operation are kept independent of the signal value, thus minimizing jitter and avoiding undesirable

non-linearity. The typical delay between the rising edge of the differential low voltage input WCLK and the sampling time is 4 ns.

The Read control signals include the 5 MHz Read clock (RCLK) and a Read signal which is used to trigger a Read operation. The 8-bit RADD is transmitted bit-serially to the SCA at the 5 MHz RCLK frequency in order to reduce the numbers of lines, and associated noise, on the SCA control bus. The data bits are clocked on the falling edge of RCLK into an 8-bit shift register (SR). Detection of a Read signal triggers the loading of the SR into the RADD register, which is subsequently decoded as described above for WADD. The SCA then internally generates a complete Read cycle which includes cycling the MUX output through each of the four calorimeter channels. For each channel, the gain scale indicated by the two Gain Selection bits (supplied by the Gain Selector chip, as described later) is output. The inverted Read signal is used internally as the Reset signal, to hold the output op-amps in reset between Read operations.

SCA Realization and Performance

The SCA chip is realized in the DMILL process, and contains ≈ 45000 transistors and 2384 capacitors on a die measuring $4.5 \text{ mm} \times 4.4 \text{ mm}$. The SCA chip is packaged in the same 100-pin QFP package used for the shaper. For the SCA, two different packaging geometries were used, one with the die mounted up and the other with the die mounted upside down. This was achieved by bending the pins on the chip frame in the appropriate direction before the plastic encapsulation stage. Using this method, half of the chips were packaged in the “standard” geometry and the other half using the “mirror” geometry. The purpose was to allow SCA chips mounted on opposite sides of the FEB PCB to directly share the vias providing the large number of connections to the addresses and control signal buses, as well as power and grounds. This greatly simplified the routing of the FEB PCB. Table 6 summarizes some of the parameters and performance figures for the SCA chip.

Table 6. Some parameters and performance measures for the SCA chip.

Parameters	Value
Power Consumption	290 mW \pm 5%
VSS, VDD	-1.7 V, +3.3 V
Input Bandwidth	50 MHz
Input Slew Rate	175 V/ μ s
Sampling time jitter	\approx 10 ps
Cell-to-cell sampling time variation	1.9 ps/cell
DC Gain	0.995
Input range with < 0.2% integral nonlin.	-0.8 V to 2.9 V
Noise	290 μ V
Dynamic Range	13.2 bits
Fixed Pattern Noise	250 μ V
Channel Offset Dispersion	12 mV RMS
Droop rate	< 0.54 V/s
Crosstalk	< 0.01%

4.5 Digitization and Gain Selection

Two commercial dual op-amp chips (AD8042 from Analog Devices [20]) couple the SCA outputs to the commercial ADC. The first pair of op-amps are connected as emitter followers to provide high impedance loads to the two pseudo-differential CMOS SCA output drivers. The second pair perform the differential subtraction in front of the single-ended ADC. In addition, they are operated with less than unity gain to map the SCA output voltage range onto the more limited (1 V) ADC input signal range. Finally, a voltage offset is added such that the pedestal value corresponds to ≈ 1000 ADC counts, allowing measurements on both the positive and negative lobes of the shaped calorimeter signals.

The commercial 12-bit ADC (AD9042 from Analog Devices [21]) is operated continuously at 5 MHz. The 12 single-ended TTL digital outputs of the ADC are interfaced to a Gain Selector (GSEL) ASIC [22]. To reduce noise, the transition edges of the digital lines are slowed down by inserting 1 k Ω resistors in series. In addition, resistive voltage dividers are used to map the 5 V ADC logic levels onto the 2.5 V range of the GSEL chip.

To reduce the event size and required readout bandwidth, typically only one of the gain scales is digitized and read out for each channel. To optimize the readout precision, one would like to read out the scale which has the highest gain and yet is not saturated. The gain selection needs to be performed individually for each calorimeter channel, and separately for each Level 1 trigger. The GSEL is designed to perform this gain selection. In addition, it formats and serializes the data before sending it to the interface to the optical link used to transmit the data off-detector from the FEB to the ROD. For the eight channels corresponding to a single ADC, the GSEL formats the data into an event fragment of 16-bit words. The format is discussed in more detail in Section 8.

The GSEL can be configured for calibration purposes to read out either one, two, or even all three gain scales. During normal physics running, the GSEL is configured to operate in “Auto Gain” mode, where the gain scale is selected dynamically for each channel and trigger by first digitizing the peak sample on a specified gain. This value is then compared by the GSEL against two 12-bit thresholds, downloadable for each channel separately, to determine the optimal gain scale to be used. The five samples of the given event are then digitized, all on the selected gain scale. The gain selection method thus requires that six digitizations be performed per channel in order to get the final five samples, but avoids systematic effects which would be encountered in combining the five samples if they were not all digitized on the same gain scale. The GSEL communicates the appropriate gain via two gain selection bits sent to the output multiplexor sections of the corresponding SCA chips.

The GSEL was designed, prototyped and tested first using the DMILL process. However, when the SCA Controller (see Section 5) was developed in the DSM process, the opportunity was taken to target the GSEL design also to DSM, using essentially the same Verilog design as for the DMILL GSEL version. The minimum DSM production run was such that the SCAC, GSEL and also CLKFO chips were all produced on the same wafers with essentially no extra cost, apart from packaging, than would have been incurred to produce only the SCA Controller chips.

The DSM GSEL is realized as a 4 mm \times 4 mm ASIC implemented in the radiation-tolerant DSM process, and was packaged in the same 100-pin QFP package as the shaper and SCA. One GSEL ASIC handles the outputs of two ADCs, corresponding to 16 channels. The GSEL design

incorporates features to protect against SEU-triggered corruption of the downloaded parameters needed to configure and operate the GSEL. For each of the eight calorimeter channels corresponding to a single ADC, a 32-bit word is assigned in the GSEL to store the relevant parameters. The mapping of these bits includes six Hamming code bits, a 12-bit upper threshold and a 12-bit lower threshold for use in the gain selection algorithm, and two bits which specify the mode in which the GSEL should operate. The six Hamming code bits per 32-bit word are sufficient to provide Error-Detection-and-Correction (EDC) functionality such that, if any single bit gets flipped, the error is detected and automatically corrected. An error flag is set in the output data in order to allow monitoring of the rate of single-bit errors. If two bits get flipped, the error can be detected, but not corrected. For double bit errors, an error flag is set to indicate the need to download the parameters again via SPAC in order to fix the problem. A total of five mode bits is sufficient to fully specify the GSEL operation, while 16 modes bits are available, distributed over the eight 32-bit channel words. The mode bits are provided even more protection against SEU corruption by distributing three redundant copies of the five mode bits over the eight words, and using majority voting logic to set the selected mode.

5. SCA Controller and FEB Digital Control

The SCA Controller (SCAC) ASIC [23, 24] provides the main digital control of the FEB. The SCAC receives the CLK and trigger information, and generates the WADD and RADD sequences for the SCA pipelines. In addition, the SCAC controls the flow of the SCA readout through its communication with the GSEL. A block diagram depicting the basic SCAC architecture is shown in Figure 10.

The main SCAC functionality, namely the SCA WADD and RADD bookkeeping, involves cycling the 144 pipeline addresses of the SCA through a system of FIFOs. The SCA addresses are gray encoded, to reduce address pattern noise on the FEB and inside the SCA. To achieve this noise reduction, the addresses must be maintained in sequential order as much as possible. As illustrated in Figure 10, the movement of SCA addresses takes the following path:

1. After a SCA cell number is used for the SCA WADD for a particular bunch crossing, it is written into the Latency FIFO, which stores the sequence of SCA cells corresponding to the bunch crossings of the L1 trigger latency. The cell numbers are stepped through the Latency FIFO every 25 ns. The depth of the FIFO is programmed, depending on the actual value of the L1 latency, such that a given cell number emerges from the FIFO at the time that the L1 trigger decision for the relevant bunch crossing arrives at the FEB.
2. Cells which emerge from the Latency FIFO and are not part of an L1 trigger are transferred to the Free FIFO, which contains addresses which are available for use as WADD in upcoming bunch crossings.
3. Those addresses emerging from the Latency FIFO which are part of a L1 triggered event are transferred to the de-randomizing buffer of samples awaiting digitization. The SCAC will typically be configured to read out five consecutive samples for each triggered event. The relevant cell addresses are stored in either the Event FIFO or the Sample FIFO. The Event

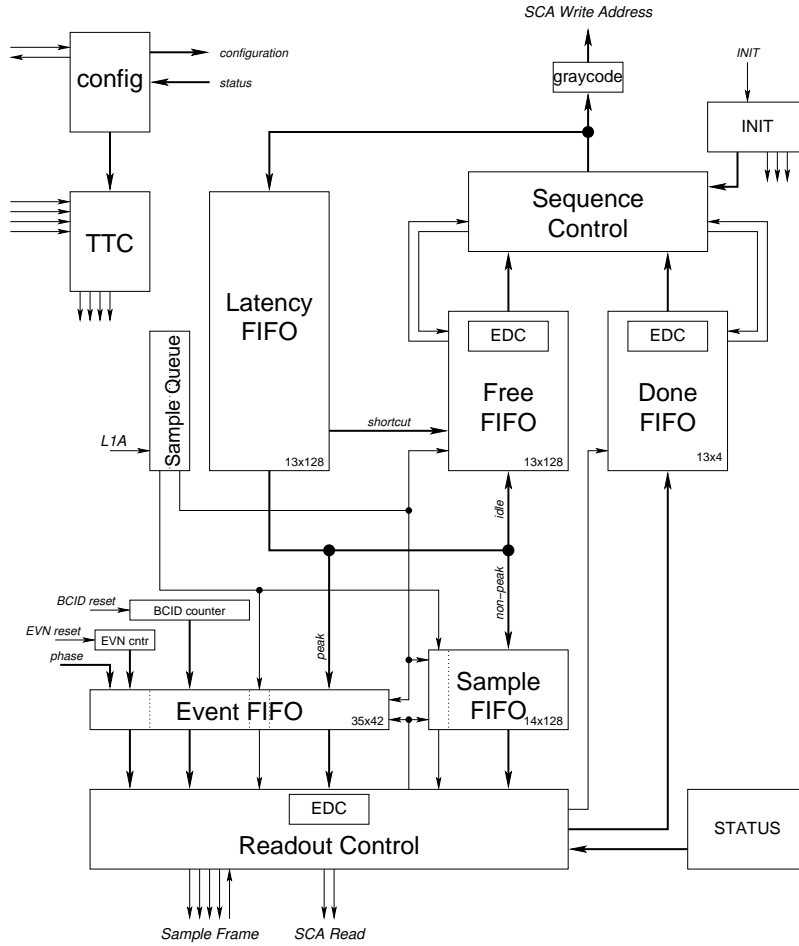


Figure 10. Block diagram of the SCA Controller architecture.

FIFO stores some event status information and the cell address of the peak sample of the event. The Sample FIFO stores the remaining non-peak sample addresses.

4. After the signal samples corresponding to a particular cell address have been fully digitized, the address is again available for re-use. However, in order to keep the gray encoded addresses as sequential as possible, these cells are not immediately put into the Free FIFO. Instead, they are transferred to the Done FIFO.
5. Every 25 ns, the Sequence Control compares the addresses at the top of the Free FIFO and Done FIFO and chooses the one which most closely preserves the ongoing sequence of WADD. In this manner, the addresses in the Done FIFO, corresponding to samples read from previous L1 triggers, will be inserted back into the correct sequence.

The readout operation must be synchronized with the GSEL chips, depending on the gain mode in use. The Readout Control block of the SCAC handles the interface between the SCAC and the GSEL. The SCAC provides the GSEL with the RADD as well as other information (such as Bunch Crossing number) required for the readout data stream.

The “config” block of the SCAC provides the interface to the SPAC system used to configure the SCAC. The L1 latency is configurable between 64 and 127 bunch crossings. In addition, the SCAC can be configured to read from 1 to 32 signal samples for each L1 trigger. These signal samples can form an arbitrary pattern (with gaps) of up to 32 bunch crossings in length. One sample is designated as the peak sample and is read first to perform the gain selection. The remaining signal samples are read in sequence.

The number of events that can be stored in the de-randomizing buffer depends on the L1 latency and on the number of samples to digitize per L1 trigger. The 144-capacitor depth of the SCA was chosen to provide at least an eight-event buffer, assuming a maximum L1 latency of $2.5\ \mu\text{s}$ (100 bunch crossings) and five signal samples per event. The deadtime caused by the FEB is then less than a few percent for the maximum allowed L1 trigger rate of 75 kHz. The FE readout system does not send a “Busy” signal to the trigger system. Instead, the L1 central trigger processor models the status of the SCAC and implements a leaky bucket algorithm to avoid issuing additional triggers when the SCA buffers are full.

The TTC block of the SCAC provides the interface of the SCAC to the TTCrx chip [10] located on each FEB. In addition to the 40 MHz clock, the SCAC receives from the TTCrx chip the L1 Accept (L1A) signal, Reset signals for its internal Bunch Crossing and Event Number counters, and a signal for initialization. The bunch counter number (BCID) is used in ATLAS to align the various detector readout elements in order to correctly combine them when building an event. More details on the trigger and timing control are provided in Section 10.1.

Status information for monitoring the performance of the SCAC and checking for possible corrupted data is obtained in two ways. First, event status and SCAC status information is communicated by the SCAC to the GSEL and included in the event data output. In addition, a 32-bit SCAC status register, containing additional information about the internal state of the SCAC, can be read out via SPAC.

The SCAC functionality was first implemented using the radiation-hard DMILL technology. However, the DMILL design resulted in a rather large die size and limited operating speed margin. It was decided, therefore, to perform a new SCAC design using the DSM technology. The increased speed and density of the DSM process allowed inclusion in the design of a number of features aimed at mitigating SEU effects. The SCAC is configured by downloading 64 bits through the serial configuration interface. Since the corruption on any of these bits would result in incorrect operation of the SCAC, all 64 configuration bits are stored using triple-redundant flip-flops, and majority voting is used to resolve any ambiguities.

Proper functioning of the SCAC depends on the entire history since the last initialization. Whenever the SCAC takes a wrong decision due to SEU, the sequence of SCA addresses is very likely to be corrupted, and recovery requires a reinitialization. Due to this nature of the SCAC, efforts at “hardening” the SCAC design against SEU-induced errors focused on reducing the rate of SEU errors that corrupt the sequence of SCA cell addresses. The largest contribution to SEU on the SCAC are the SRAMs of the FIFOs. The address pointers in the FIFO implementation are not protected against SEU errors. The SRAM blocks employ input registers at the address busses. These registers contribute about equally to the SEU cross section as the address pointers. Since the SRAM module was provided from an external source, we could not harden the input registers. Thus, protecting the address pointers would only reduce the SEU cross section of the FIFO logic

by half, which was not considered worth the considerable effort required. However, all information stored in the FIFOs includes redundancy to protect against SEU-induced errors. Every 8-bit SCA cell address is accompanied by five Hamming code bits, which allow for correction of single bit errors, and detection of double bit errors. The EDC algorithm is applied to the cell addresses before the SCAC takes a decision based on the address value, and before an address is sent off-chip. When any EDC unit encounters a single or double bit error, a bit is set in the next event's status word. The occurrence of EDC errors can also be detected by reading the SCAC status word via the serial configuration interface.

The Event FIFO also stores the BCID, event number, and read clock phase. These bits are unprotected, since they are not used for any decision within the SCAC. This is one example where a transient SEU error in the SCAC would corrupt the data for a single event.

The DSM SCAC is packaged in the same 100-pin QFP package as the shaper, SCA and GSEL. The minimum size of the DSM SCAC die was $2.7\text{ mm} \times 2.7\text{ mm}$. Due to quantization requirements on the reticle, the edges of the die were required to be a multiple of 2 mm. While the design could have fit into a $2\text{ mm} \times 4\text{ mm}$ area, it was decided to keep the geometry square to simplify packaging. The final die size was therefore $4\text{ mm} \times 4\text{ mm}$. A drawing of the layout can be seen in Figure 11. The layout is pad-limited, with the core occupying less than one-half of the available space.

In addition to the SRAM, the SCAC contains 957 D-flip-flops, including those embedded inside the SRAM blocks. With few exceptions, all I/O to and from the SCAC uses low-voltage differential signal (LVDS) levels, with differential receivers on all input pairs and current mode LVDS drivers on the output pairs.

The SCAC chips were subjected to extensive functional testing before being assembled on to FEBs. The maximum frequency for which the SCAC operated properly ranged from 54.5 MHz to 63.5 MHz, with a mean value of 59.0 MHz and an RMS of 1.1 MHz. These values provide a wide safety margin, given the required operating frequency of 40 MHz.

6. SCA Control and Address Bus

Each of the two SCAC chips on one FEB controls a total of 16 SCA chips, corresponding to 64 readout channels. For proper operation, the SCAC must send to the 16 SCA chips the associated WADD, RADD and control lines. Given the number of SCA chips and control signals required, the SCA control was implemented as a bus, with one driver per line driving all 16 SCA chips.

The SCA control bus is driven differentially using commercial MC10H116D chips [25]. The 10H116 chips are powered with voltage rails of +3.3 V and -1.7 V in order to match those of the SCA chips. The SCA control bus lines are routed on an internal layer of the FEB PCB, with grounds on neighboring layers (see Section 12). The bus is terminated at both ends with $27\ \Omega$ resistors, with the far end termination connected to VSS through a 470 pF center-tap capacitor. Simulations including the distributed SCA loads showed that the resultant signals are very clean, with fast rise and fall times and little overshoot, and an amplitude of $\approx 400\text{ mV}$. These expectations were confirmed by measurements. The delay between the signals at the first pair of SCA chips on the bus and at the last (eighth) pair is measured to be $\approx 2\text{ ns}$, in agreement with simulation. This effect implies a channel dependence of the signal sampling time that, if uncorrected, would




Figure 11. Diagram showing the layout of the DSM SCA Controller ASIC. The four double-width SRAM blocks used to implement the SCA address FIFOs are visible in the upper region of the core, while the combinatorial and sequential logic occupies the lower region of the core.

correspond to a Gaussian sigma of ≈ 580 ps. However, since the effect is perfectly systematic and predictable, it can easily be corrected.

As described in Section 4.4, pairs of SCA chips mounted on opposite sides of the PCB were packaged such that they could share vias for the control lines. This greatly reduced the routing complexity, and also reduced loading of the bus due to vias.

7. Output Optical Link

The formatted event fragments of 16-bit words from each GSEL are serialized in a sequence of steps that results in a single 1.6 Gbps bit-serial output stream, transmitted via optical link from each FEB. The basic architecture of the optical link is shown in Figure 12.

The first step in the serialization is performed by the GSEL. For each of its two event fragments of 16-bit words, the GSEL outputs two bits parallel at 40 MHz, taking eight 40 MHz clock cycles to output each partially serialized 16-bit word. The bits are transmitted from the GSEL chips over LVDS-like signal lines.

With eight GSEL chips per FEB, and with each GSEL outputting four data lines (two event fragments, each with two data lines), the SMUX chip [26] receives a total of 32 data bits at 40 MHz. The SMUX then performs a 2:1 multiplexing and a level shifting from LVDS to TTL in order to




Figure 12. Architecture of the transmission end (left side) and reception end (right) of the 1.6 Gbps FEB output data optical link. The transmitter is mounted on the FEB and the receiver on the ROD.

generate an output stream of 16 bits at 80 MHz. In addition, the SMUX generates a FLAG, which is HI (LO) for the MUX cycle when data from channels 0 - 63 (64 - 127) are sent the GLINK. The 16 SMUX data out signals and FLAG are then transmitted at 80 MHz to the commercial “GLINK” serializer chip (HDMP-1022 from Agilent Technologies [27]), which adds protocol and control bits and produces a single serial output stream of 1.6 Gbps. In the final step, this serial stream is converted to an optical signal and transmitted off of the FEB via a custom-built optical transmitter (OTx) module which includes a commercial driver chip (SY88922V from Micrel Inc. [28]) and a 850 nm VCSEL (TTR-1A43 from TrueLight Corp. [29]) for converting the signal from electrical to optical. Table 7 lists the parameters measured, and acceptance windows applied, during QC measurements of the OTx. In addition, an eye mask test, and bit error rate tests at 0 dB and 10 dB attenuated optical power were also carried out. The bit error rates in both cases were required to be less than 10^{-12} . The distributions for several OTx parameters are shown in Figure 13, while Figure 14 gives a typical eye diagram of the optical signal from the OTx. Detailed information about the QC tests of the OTx can be found in Reference [30].

Multimode 50 μm core graded-index (GRIN) fiber is used to carry the signal from the FEB to the ROD. At the ROD end of the fiber, the 1.6 Gbps optical signal is converted back to an electrical signal via a PIN-diode based custom-built optical receiver (ORx) module. It is then de-serialized by the GLINK receiver (HDMP-1024) [27] back into an image of the 16 bits at 80 MHz data stream output from the SMUX.

8. FEB Output Data Format

The GSEL formats the output data for the 8 channels digitized by a single ADC into an event

Table 7. Parameters measured, and the corresponding acceptance windows, during the tests of the OTx modules.

Parameter	Units	Min. Value	Max. Value
Average optical power	dBm	-7.5	-3.5
Extinction ratio	-	6.0	-
Rise time	ps	-	220
Fall time	ps	-	220
Deterministic jitter (peak-peak)	ps	-	125
Random jitter (RMS)	ps	-	10

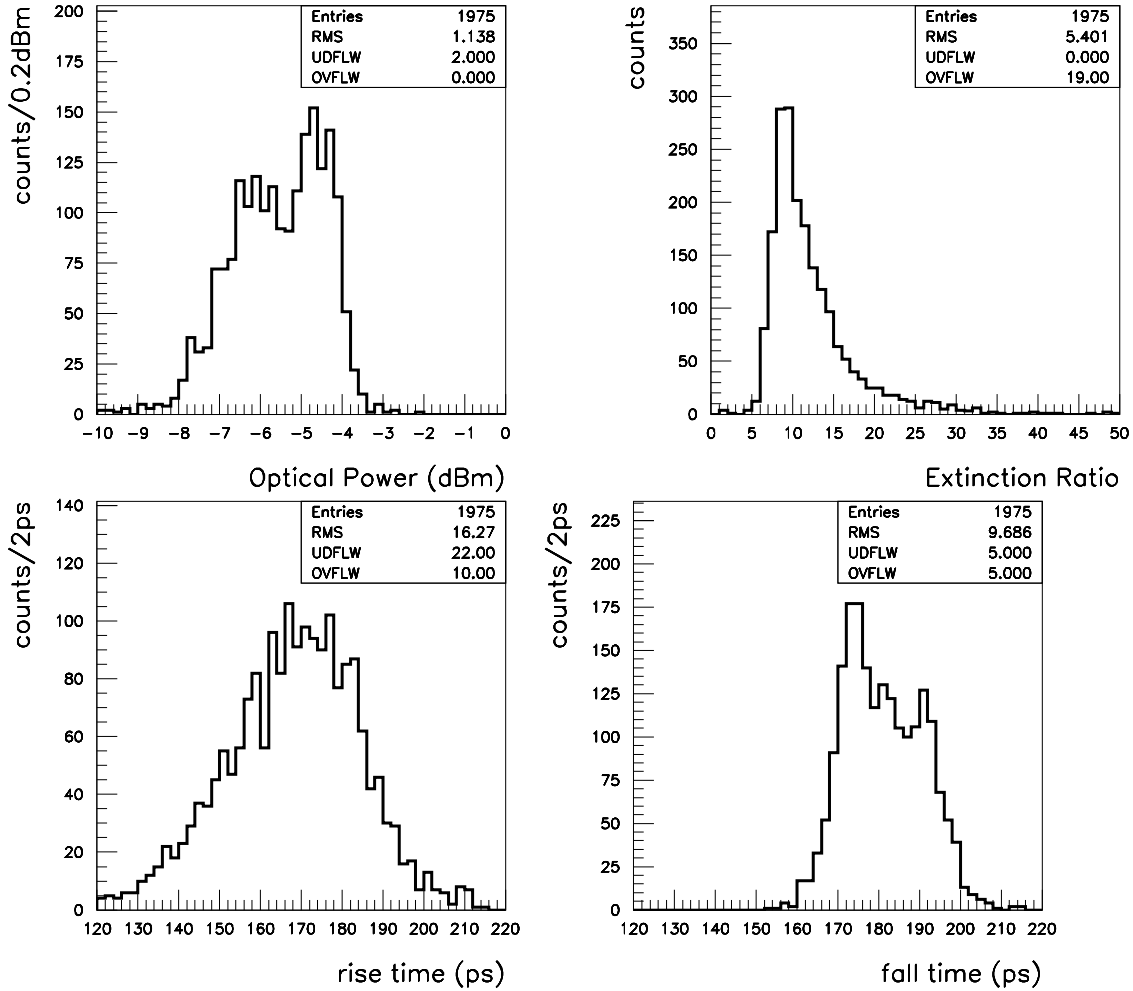


Figure 13. Distributions of optical power, extinction ratio, rise time, and fall time measured during QA tests of the OTx modules.

fragment of 16-bit words according to the format shown in Fig. 15. The event fragment contains one Frame Start word, two Event Header words, data words carrying the calorimeter signal data, one Event Trailer word, and at least one Frame End word to separate from the next event. For all words except for Frame Start and Frame End, Bit 14 serves as a parity (P) bit, and is set in the




Figure 14. An eye diagram demonstrating the eye mask test of a typical OTx module. The horizontal scale corresponds to 90 ps/division.

GSEL to the appropriate value to ensure odd parity for each 16-bit word.

Frame start tag		1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
Event header		0	P	0	0	ADCID				EVENTN						
		0	P	0	0	BCID										
Sample header		0	P	0	0	F	L	B	A	CELLN						
Sample data	per sample data	0	P	gain		ADC										
		0	P	gain		ADC										
		0	P	gain		ADC										
		0	P	gain		ADC										
		0	P	gain		ADC										
		0	P	gain		ADC										
		0	P	gain		ADC										
		0	P	gain		ADC										
		0	P	gain		ADC										
		0	P	gain		ADC										
Event trailer		0	P	0	0	1	S	E	SCAC status							1
Frame end tag		0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Figure 15. Readout event format for the ATLAS LAr FEB. See more details in the text. For each word, bit 0 is listed in the rightmost column, proceeding toward bit 15 in the leftmost column.

The first 16-bit word is the Frame Start word, defined as \$FFFF, which signals to the ROD the

start of a new event. The first of two Event Header words follows, and carries the 4-bit ID number of the corresponding ADC and an 8-bit EVENTN field that includes a 5-bit event number and the 3-bit phase of the 5 MHz RCLK with respect to the 40 MHz clock for this L1Accept. The second word contains the 12-bit BCID number for this L1Accept.

The number of signal data words per event fragment depends on the number of samples and number of gains configured to be read out. For each sample, there is first a sample header word, which contains the 8-bit SCA cell number, three “sample mode” bits, and one bit (denoted by A in figure 15) which has a value of 0 for normal data and 1 for data taken in a mode where configurable test data is transmitted instead of data from the ADC. The sample mode bits indicate whether the sample in question is the first (F) or last (L) sample of the event, or whether the so-called “Backporch” (B) bit has been set for this event. If the Backporch bit is set in auto-gain mode, the gain selection algorithm was modified to prevent any channel from using a higher gain than in the previous event. This flag can be set by the SCAC for events that occur less than some specified time (for example, the total LAr drift time of ≈ 400 ns) after the previous event, with the goal of preventing an incorrect gain selection for the second event in cases of large signals in the first event.

The sample header word is followed by $n \times 8$ words containing the ADC data for that sample, where n is the number of gains to be read out. Each ADC word includes the 12-bit ADC value plus two bits which encode the gain.

The event trailer word contains the 8-bit SCAC status word, with bits indicating SEU-induced single- and double-bit errors, as well as other SCAC error and status conditions. Bits 9 and 10 of the event trailer word are used as flags to indicate that single- or double-bit errors, respectively, have been detected in the EDC logic of the GSEL.

The final word in the event is the Event End word, defined to be \$0000. Detection of this word indicates to the ROD that the event has ended. Since words of this value are sent by the FEB between events, the ROD simply waits for the next occurrence of \$FFFF to indicate the arrival of a new event.

9. Level 1 Trigger Summing

In addition to signals from the muon detectors, the ATLAS L1 trigger system [5] considers the energies deposited in the calorimeters in regions of reduced granularity. These trigger signals are produced via a chain of analog summing, the first two stages of which are performed on the FEB.

The calorimeter signals are proportional to energy, and must be transformed to transverse energy for the L1 trigger. This transformation is accomplished through the choice of several gain factors in the L1 summing chain, including the preamp transimpedance, the LM gain, the gain of the Layer Sum Board (LSB) described below, and gains downstream of the FEB in the Tower Builder and Receiver modules (described in more detail in Reference [4]). The overall gain of the L1 summing chain must be such that a 10 mV pulseheight at the L1 preprocessor input would correspond to 1 GeV of deposited transverse energy. The sums must have a saturation level corresponding to 250 GeV transverse energy. The gain uniformity is required to lie within an allowed variation of $\pm 5\%$.

As discussed in Section 4.3, each shaper chip has a Linear Mixer (LM) section that sums its four input channels. On-chip logic, addressable via SPAC, allows the contribution of each

individual channel to be switched ON or OFF in this trigger sum. This feature allows masking of noisy or bad channels. In addition, this masking ability provides a very useful debugging feature: any single channel of the calorimeter readout can be examined off-detector with an oscilloscope by enabling only that single channel in the appropriate LM sum, and also disabling all other inputs to that particular trigger sum. This feature can be used, for example, to examine any individual channel for high frequency noise that would not be easily visible through the precision readout due to its 40 MHz sampling rate.

The LM output sums from the 32 shapers per FEB are routed to the inputs of the two Layer Sum Boards (LSBs) [31] which plug into sockets on each FEB. The LSBs use commercial op-amps to provide the next layer in the L1 trigger summing tree. The mapping of EM calorimeter channels onto FEBs has been chosen such that, except in limited regions near the barrel-endcap transition, an individual FEB processes channels from only one longitudinal section (i.e. presampler, front, middle, or back). In the EMB, all Level 1 trigger towers are of size $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$, and the analog trigger sums are made using four different types of LSBs, with 1, 2, 4, or 8 inputs (see Table 8). The LSB is designated $n \times m$, in which n is the number of inputs to each sum and m is the number of summed outputs on the LSB. In the case of the S1x16 LSB, a suffix (L, H, M) is applied to denote the channel gains, as discussed below. In the endcap calorimeters the trigger tower structure is somewhat more complex, but the same types of LSBs are used, employed in different ways.

Table 8. Trigger Tower (TT) Structure and Layer Sum Board types for the EMB calorimeter.

EMB Layer	Cell Size ($\Delta\eta \times \Delta\phi$)	Number of cells per TT	Width of LSB sum	LSB type
Presampler	0.1×0.025	16	1	S1x16H
Front	0.1×0.0003	32	8	S8x2
Middle	0.025×0.025	16	4	S4x4
Back	0.05×0.025	8	2	S2x8

The LSB functional requirements include providing sums with a gain of one or two depending on the LSB type (with exceptions for FCAL LSBs), having a noise which is small compared to the preamp noise, and providing a voltage clamp at ≈ 3 V, with rapid recovery from saturation. The LSBs are non-inverting. There are two basic designs used in the LSBs, a one stage non-inverting configuration (Figures 16a and 16b) and a two-stage summing configuration (Figure 16c). The former is used only for the S1x16 LSBs, while the latter is used for all other types. The LSB input impedance of 1.5 k Ω is chosen to be large compared to the 50 Ω output resistor of the LM to both avoid the factor of two loss in amplitude resulting from a low impedance, and reduce the effects of unavoidable variations in the LM output resistance due to ASIC production process variations.

The gains of the S8x2, S4x4, and S2x8 LSBs are all unity. The S1x16 LSBs are made in three different types: high gain H ($g=2$, used for the HEC and the presampler), low gain L ($g=1$, used for the HEC and the front section in the EMEC) and mixed gain M ($g=1$ or $g=2$ on different channels, used only in the HEC). In the FCAL, weighted sums are used, to account for relatively large variations in the value of $\sin\theta$ over the bins of width $\Delta\eta = 0.4$. In addition to their nominal gains, all boards require a small (3%) additional gain to compensate for the voltage reduction

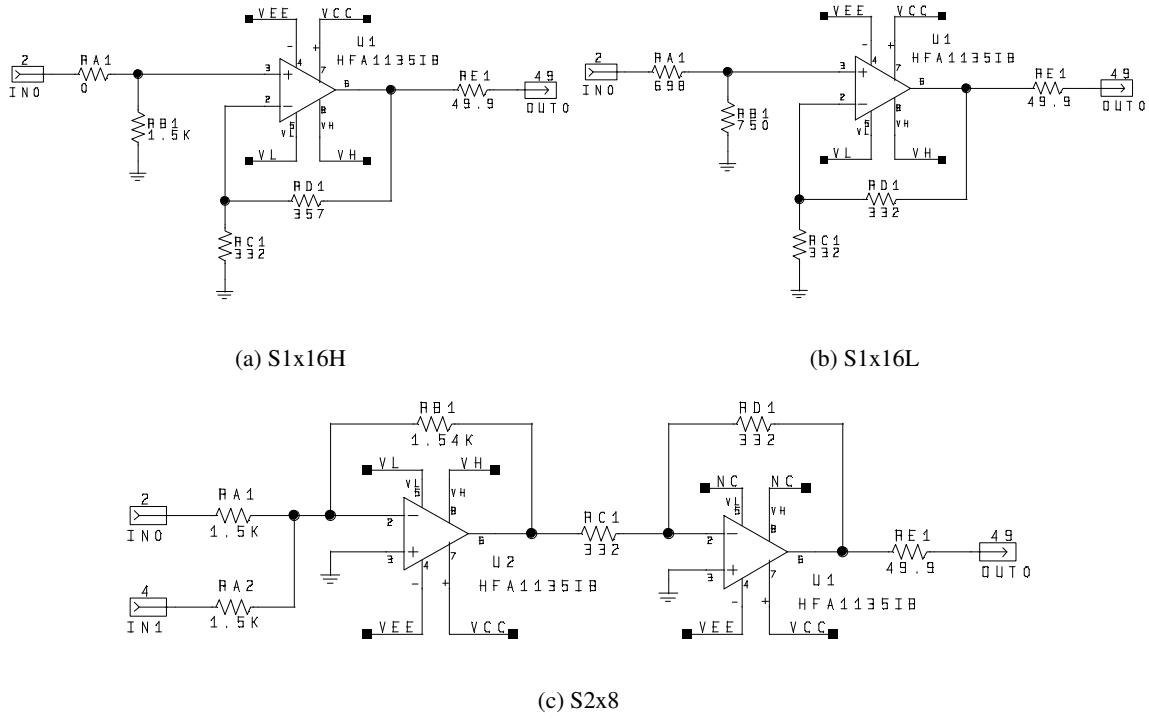


Figure 16. Schematic diagram of the circuits used in the LSBs. The upper two circuits are single stage amplifiers of (a) gain=2 and (b) gain=1. The lower drawing is a schematic of the adder circuit in the S2x8 LSB, which utilizes two stages to achieve summing without polarity inversion. The S4x4 and S8x2 LSBs use identical circuits with 4 and 8 inputs, respectively.

caused by the presence of the $50\ \Omega$ series resistor at the LM output.

The amplifier chip chosen for the LSBs is the Intersil HFA1135A [32], a fast current feedback amplifier produced using a bipolar technology. This device is relatively radiation tolerant, and contains voltage limiting circuitry which is needed for rapid and accurate recovery from saturated pulses. The maximum input common mode voltage of the amplifier is important for the low gain single stage amplifier of Figure 16(b), as both inputs to the amplifier ride to the maximum level of the pulse. Since the specification for this voltage is 1.7 V for the HFA1135A and the maximum amplitude of the LM output is 3.0 V, the input pulse is attenuated by a factor of two, and the gain of the amplifier is set to two to achieve an overall gain of unity. This precaution is unnecessary for the high gain version of Figure 16(a), as the voltage clamp is operational whenever the input is outside the operating range of the IC. In the two-stage circuit of Figure 16(c), the resistors used in feedback are 1.5 k Ω to yield unity gain. The resistors of the second stage are chosen to be the lowest values recommended by the manufacturer, in order to reduce the noise contributed by that stage.

The largest source of noise for the two-stage LSBs is the non-inverting input noise current density of the HFA1135A in the first stage. The specified value [32] of 20 pA/ $\sqrt{\text{Hz}}$, combined with the 1.5 k Ω input resistance, yields a noise voltage density of 30 nV/ $\sqrt{\text{Hz}}$ at the output. The bandwidth of the system, measured to be 12 MHz, is limited by the long cables over which the signals are sent to reach the counting room, yielding a noise level of approximately 100 μV for this

particular source. Measured values of noise at the LSB output range from $140\ \mu\text{V}$ for the S2x8 LSB to $200\ \mu\text{V}$ for the S8x2 LSB. The noise for the single stage configurations is lower, with measured values of $20\ \mu\text{V}$ and $42\ \mu\text{V}$ for the high and low gain circuits respectively. In the L1 energy scale, $200\ \mu\text{V}$ corresponds to 20 MeV of transverse energy, approximately the preamp noise of a single calorimeter cell at $\eta = 0$. A complete trigger tower, which is the sum of approximately 60 cells, has preamp noise of about $300 \sin \theta$ MeV and electronic noise from the four LSBs (one per depth layer) of 30 to 40 MeV. Other sources of electronic noise in the chain (principally the Receiver module, with a variable-gain amplifier) contribute additional noise at the level of 100 MeV.

A typical DC response curve of the S1x16H LSB (gain=2) is shown in Figure 17, where the output voltage is plotted versus the input voltage. From such a curve, several parameters are defined. The upper part of the curve can be modeled as two intersecting lines, one in the linear region, obtained by fitting all points with V_{out} below a certain limit (chosen to be 2.8 V). The second line, in the saturation region, is obtained by fitting all points with V_{out} greater than another limit (3.2 V). The value of V_{out} where the two lines intersect is defined as the upper breakpoint voltage V_{B+} , the onset of saturation in the positive voltage region. A similar procedure is used to determine the lower breakpoint voltage V_{B-} . The offset voltage V_{DC} is the intercept of the fitted line in the linear region, and the gain in each of the regions is taken to be the slope of the fitted line in that region. The integral nonlinearity is found by finding the maximum deviation between the points and the fitted line in the linear region and dividing it by the breakpoint voltage.




Figure 17. Typical response curve for the S1x16H layer sum board (left), and illustration of the parameters that are used to characterize the circuits (right).

The LSB boards were produced and assembled in industry by the University of Pittsburgh. After receipt, the boards were visually inspected for assembly or soldering faults and were then given a test for basic functionality. Any faults found in the initial tests were repaired. This process was followed by a burn-in (168 hours at 70°C while under power, with a low frequency sine wave signal applied at the input). This was followed by an acceptance test, in which the response curve described above was measured and analyzed. Boards were accepted for installation on the FEBs if the criteria listed in Table 9 were met.

The acceptance rate for LSBs was about 97% once the initial assembly faults (cold solder joints, broken resistors, chips incorrectly inserted) were corrected. No failures of components

Table 9. Acceptance criteria applied during testing of the LSBs. For more details, see the text.

Parameter	Acceptance Window
Deviation of gain from nominal value	$< 2\%$
Integral nonlinearity	$< 0.7\%$
DC offset	$< 40 \text{ mV}$
DC gain in both saturation regions	< 0.15
Breakpoint voltages	3.2 - 3.8 V

could be attributed to the burn-in process itself. All boards which failed the initial acceptance tests were repaired and retested, resulting in essentially 100% yield.

10. FEB Control Interfaces

For configuration and operation, the FEB requires a number of external control signals. These include a 40 MHz clock derived from the LHC machine clock, as well as the L1 trigger signal and a few other signals synchronous with the 40 MHz clock. The FEB is configured and monitored via the SPAC serial control link, operating at 10 MHz. These systems are described in more detail in the following subsections.

10.1 Trigger and Timing Control

As described in more detail in Reference [4], the ATLAS TTC signal is delivered to each FE crate via redundant optical fibers. Inside each FE crate, Controller boards [33] receive the optical TTC signal, convert it from electrical to optical, and then fan out the TTC signal to each of the boards to which it is connected. The electrical TTC signal fanout is performed via point-to-point “mini-B” USB 2.0 cables routed within the FE crate from the Controller to each individual board.

The FEB receives its electrical TTC signal input via a mini-B USB 2.0 connector which penetrates its front panel. As described in more detail in Section 10.2, the 40 MHz clock from which all FEB control signals are generated is derived from the input TTC signal via the on-board TTCrx chip. To avoid noise generation, the FEB uses only a few of the decoded TTC signals, namely L1A, Bunch Counter Reset (BCR), and Event Counter Reset (ECR). So, for example, instead of distributing the multi-bit single-ended TTL Bunch Counter and Event Counter outputs of the TTCrx, the SCAC has an internal 12-bit Bunch Counter and 5-bit Event Counter, which are synchronously reset using the relevant Reset signals. These signals are routed to each of the two SCAC chips on the FEB. Use is also made of two of the possible TTC Broadcast commands in order to perform control operations on all FEBs on the TTC branch in question; one is used to simultaneously initialize the two SCAC chips on each FEB, in order to have all FEBs operating synchronously, while the other resets the SPAC Slave chip.

The L1A signal can be encoded in the ATLAS trigger system logic as either one or two bits before transmission to the FEB, where it will be subsequently decoded by the SCAC. The two-bit option was implemented to reduce the rate of fake triggers due to SEU-induced transient effects of the PIN diode used to receive the optical TTC signal. The value of the BCID counter is sent with the first sample command of every event to the GSEL, and included in the event header data.

10.2 Clock Distribution

All clocks on the FEB are derived from a single input 40 MHz clock. The clock used to sample the calorimeter signals must be of excellent quality if the ultimate timing resolution achievable by the LAr calorimeter, of order a few tens of ps, is to be approached. Furthermore, the FEB serial output data link operates at 1.6 Gb/s. Since this high frequency clock must be derived by multiplying the 40 MHz clock, stable operation of the output optical link requires an input clock with very low jitter.

The input 40 MHz clock is recovered on the FEB from the input TTC signal, via the on-board TTCrx chip. The TTCrx includes an on-chip phase-locked-loop (PLL) circuit with a wide lock range, but provides a recovered clock with rather large random and TTC data-dependent jitter. During the first round of prototype FEB production, it was discovered that the jitter levels were too high to prevent stable operation of the optical link. To solve this problem, the FEB design was modified to include the QPLL chip [34]. The QPLL is a PLL based on a voltage-controlled quartz crystal oscillator (with external crystal), developed as a jitter filter for the TTCrx clock and implemented in the DSM process. The QPLL has a narrow lock range, typically less than 8 kHz, and should provide an output clock with less than 50 ps jitter peak-to-peak.

Implementing the QPLL as a jitter filter on the TTCrx clock resolved the FEB optical link stability problems. However, problems were observed in the first batch of production FEBs, and were traced to activity dips in the external crystals. The solution required screening the crystals against activity dips within the relevant frequency range. The FEB test procedure (described in Section 14) included measuring the jitter on a copy of the QPLL output clock brought for diagnostic purposes to a coaxial connector on the FEB front panel. Typical values were ≈ 10 ps, and a requirement was made that the value not exceed 30 ps. The measurement was made across the QPLL lock range in order to detect and reject any crystals with instances of activity dips.

The QPLL output clock feeds a clock fanout tree for the entire FEB which is comprised of custom CLKFO chips. The CLKFO ASIC was developed to provide several functions needed in the FEB clock distribution, and was implemented as a minimum size (2 mm \times 2 mm) DSM chip, produced on the same wafers as the SCAC and GSEL. The CLKFO receives as inputs two LVDS CLK signals, denoted CLK1 and CLK2. The CLK2 signal is clipped such that the high phase of the clock is ≈ 6 ns. Two identical copies of the clipped CLK2 differential signal are output. The differential CLK1 input is fanned out internally to provide three single-ended TTL output copies. In addition, the CLKFO has 8 LVDS outputs related to CLK1. Depending on whether one of the CLKFO pins is tied high or low, the LVDS outputs are either eight identical copies of CLK1, or four copies of CLK1 and four identical 5 MHz signals that are derived by counting down the CLK1 40 MHz signal. The down-counter for generating the 5 MHz is designed using triple redundant counters and majority logic to harden the design against SEU-induced errors.

Using these functionalities, the clock distribution and fanout system for the FEB is implemented via a tree totaling seven CLKFO chips. In the first step, the 40 MHz differential output of the QPLL is connected to the CLK1 input of the first CLKFO chip in the tree. Its three 40 MHz TTL outputs are connected together and routed to a connector on the FEB front panel that can be used to externally measure the FEB clock quality, such as its jitter. Two LVDS CLK outputs are routed to two additional CLKFO chips which are used to provide further fanout in order to pro-

vide 40 MHz clocks to operate the relevant chips (SCAC, GSEL, DCU, SPAC, CONFIG, SMUX, GLINK). Two of them also use their CLK2 input in order to generate the two clipped 40 MHz clock signals used to drive the left and right SCA Write Clock buses. The clipping is performed for SCA noise optimization. Finally, the 40 MHz signal is counted down by the first CLKFO chip in the tree to generate a 5 MHz signal for the Read logic. The four LVDS 5 MHz outputs are routed to four additional CLKFO chips which are used to provide further fanout in order to provide 5 MHz clocks for the relevant chips (ADC, GSEL, SCAC). Note that both the 40 MHz and the 5 MHz clocks are derived in a single CLKFO chip at the top of the tree. This fact, plus the attention paid in routing the various clock signals on the FEB, provide the synchronous and in-phase operation of all of the various chips on the FEB.

The skew of the TTCrx clock signal used as input to the QPLL and therefore to the entire FEB clock fanout tree can be set via SPAC through the I2C interface of the TTCrx chip. The skew can be chosen, in steps of 104 ps, to optimize the phase at which the SCAs sample the analog calorimeter signals. The choice of this skew can, therefore, be set by the desired position of the samples on the calorimeter signal waveform. Since all clocks on the FEB are derived from this one input signal, there is only one phase adjustment possible per FEB.

10.3 Configuration and Slow Control

The various types of data to be loaded to, and/or read back from, the FEB include the settings of the switches on the shaper chips which enable/disable contributions from individual channels to the L1 analog trigger sums, the parameters required to configure the SCAC (such as the L1 latency, the number of samples to digitize, and the order of digitization), the parameters needed to configure the GSEL (such as the number of gains to digitize, the order of digitization, and the thresholds to be used in the gain selection algorithm), and the skew of the ClockDes1 signal from the TTCrx. All parameters were designed to be also read back for verification. In addition, as described in Section 11, it is possible to enable/disable various voltage regulators, and also to monitor some voltages and temperatures on the FEB.

As described in more detail in Reference [4], the ATLAS LAr FE electronics is configured using a custom “Serial Protocol for ATLAS Calorimeters” (SPAC) [11] system. Each FEB has an on-board SPAC Slave ASIC to allow configuration and monitoring of the FEB via the SPAC bus. An 8-bit switch which can be seen and modified through a slot in the FEB front panel is used to control the SPAC addressing. Seven of the switches are used to set the 7-bit address of the SPAC Slave. The eighth switch is used to configure the Local Broadcast Address to which the FEB will respond.

The SPAC Slave provides two I2C interfaces, as well as an 8-bit parallel interface. Each provides both Read and Write functionality. Most of the FEB configuration is performed via the parallel interface of the SPAC Slave, which is connected to the Configuration Controller chip (CONFIG) [35]. The exceptions are the use of the SPAC Slave I2C ports connected to the TTCrx (I2C port 0) and to the two DCU chips (I2C port 1).

The CONFIG is designed to serve as the interface between the SPAC Slave parallel interface and the various custom chips on the FEB that require configuration data. Simple serial protocols are implemented in the CONFIG to connect to the SCAC and GSEL chips.

The CONFIG is realized as a DMILL ASIC with die area 31 mm^2 , packaged in the same 100-pin QFP package as the shaper, SCA, SCAC and GSEL. The interface to the shaper switches, which has logic levels of 0 V and -3 V due to the shifted shaper power rails, is implemented by taking advantage of the SOI nature of the DMILL process, which allows the use of different voltages on different sub-blocks of the same chip. While most of the CONFIG is powered with 0 V and +5 V, the CONFIG shaper interface logic is powered with 0 V and -3 V, thereby avoiding the necessity of including external level-shifting components.

11. FEB Power Distribution and Monitoring

The FEB receives its power on a 10-pin power connector along one edge of the board. Three of the pins provide Ground connections, and one additional pin is reserved for the Calibration board and is not connected on the FEB. The remaining pins are used to provide six different input voltages to the FEB. The FEB total power consumption is $\approx 80.7 \text{ W}$, corresponding to less than 700 mW per channel.

Each input voltage is protected for safety reasons with a fuse and a reverse-biased Zener diode. The current rating of the fuse is chosen at least three times larger than the nominal current, to avoid instances of fuses blowing due to aging. In the case of over-voltages or reversed voltage polarities, the diodes should be able to divert any potentially damaging currents long enough for the fuse to blow, without causing damage to the other FEB components.

After the fuses, the input voltages are fed to a number of on-board radiation-tolerant positive and negative voltage regulators (L4913 and L7913 from STm [36]). To generate all of the different voltages required, while staying safely below the 3 A current limit per regulator, a total of 19 regulators are needed per FEB. The regulators serve several functions, including reducing the sensitivity to power supply noise, particularly at low frequencies, and providing current limiting as well as a thermal shutdown function in case of over-heating. Since the regulators are equipped with an Inhibit control pin, they can be switched ON and OFF remotely. The Inhibit signals are connected to CONFIG switch outputs designed such that, upon power-on of the FEB, most of the regulators are OFF until being switched ON via SPAC. This feature allows the FEB to be powered on in a controlled and staged manner, reducing surges or other problems which might arise should the entire load of $\approx 2 \text{ kW}$ per crate be switched on simultaneously. The possibility also exists to power cycle a component without having to power off the entire crate, in case a radiation-induced latch-up condition is unexpectedly encountered.

FEB Temperature and Voltage Monitoring

The CONFIG has 8 inputs which are connected to Overcurrent Monitor (OCM) outputs from voltage regulators. Given the total of 19 voltage regulators per FEB, the voltage regulators must be grouped together for these purposes.

Two DCU [37] chips, each with eight ADC inputs, are located on each FEB, in order to monitor temperatures as well as selected voltages. The DCU chip design provides two current sources which can be connected to external thermistors, the voltage across which can then be connected to one of the inputs in order to measure additional temperatures. Given the known temperature-sensitivity of the GLINK and TTCrx chips, temperature measurements are made for

these critical components by connecting external thermistors mounted close to these components to DCU current sources. One current source of a DCU is connected to the temperature-sensitive diode input of the GLINK chip.

With the remaining DCU inputs available, it is not possible to monitor the voltages from all 19 of the voltage regulators per FEB. Since the goal of the voltage monitoring is to be able to track any possible variations in analog performance with voltage, preference was paid to the voltages in the analog sections of the FEB.

12. FEB Layout and Topology

The FEB is realized as a large, ten-layer printed circuit board (PCB). The dimensions of the PCB measure 490 mm \times 409.5 mm. Figure 18 shows a photograph of the top layer of a FEB, with superimposed labels identifying many of the main active components. To achieve the required density, there are components mounted on both sides of the FEB.




Figure 18. Photograph of the top layer of a FEB with the main active components indicated.

A notable feature of the FEB topology is that connectors are mounted on three sides of the PCB. Along the “bottom” side in Figure 18 one sees three 96-pin DIN-style signal connectors.

The outer two are used to bring in 64 channels each (with the 32 center pins of each connected to ground). The central connector is used to output the LSB trigger signals back into the baseplane, in which they are routed to the Tower Builder Board situated in another slot in the same crate. Along the “top” side of the PCB, in the front panel, are mounted the output optical link connector and the input TTC connector. Finally, along the right side are mounted the power and SPAC connectors. The power and SPAC connections must be made with “combs” which slide through buses mounted on the side of the crate and connect with the FEB. Details can be found in Reference [4].

In order to minimize couplings in this mixed analog-digital board, the general topology of the FEB layout moves progressively from the sensitive analog electronics to the digital portion of the board. Each set of components is mounted in a single row across the FEB (along the z direction in ATLAS), and the rows of different components do not overlap. As one moves outward in the radial direction in ATLAS across the FEB, one encounters first the row of preamps, then shapers, then SCAs, then ADCs, then digital logic, and finally the digital optical link. This spatial separation reduces noise coupling into the analog electronics.

A cross-section of the ten-layer FEB PCB is depicted schematically in Figure 19. While the density of the FEB does not allow the use of complete power and ground planes, the FEB design makes extensive use of extended copper pours to implement segmented power and ground layers to distribute the various voltages and grounds. As shown, separate “AGND” and “DGND” ground regions are used for the analog and digital portions of the FEB, respectively. An important function of the form of the analog ground architecture is to extend the Faraday cage for the input signals from the input connectors to the preamplifier inputs. The AGND regions extend from the signal input connector side of the FEB up until the ADC region, where the DGND regions begin. The AGND system also minimizes feedback from the digital activity from the digital sections of the FEB. The two ground systems, AGND and DGND, are tied together under the ADC chips using 32 zero Ohm resistors. These resistors connect the plane copper pours over the entire width of the FEB PCB. There is another massive ground connection at the top of the FEB, where the ground from the power connector connects also to AGND and DGND. At this point the front panel is also connected to the FEB ground system, closing the Faraday cage created by the whole crate mechanics. This connection should divert the power supply noise currents out of the sensitive region of the FEB near the signal inputs.

The quality of the ground connection between the FEB and baseplane is very important since ground motion could be seen at the preamp inputs and amplified, contributing to coherent noise in the readout. Referenced to the preamp inputs, coherent effects at the level of micro-Volts would already violate the coherent noise specification. The 96-pin DIN input connectors themselves do not provide a sufficiently high quality ground connection, due in part to the rather large inductance of the connector pins and to the assignment of one ground pin for every two signal pins. In addition, the plastic connector does not provide a shield connection which would allow a separation between signal return and ground. To address these issues, custom metal shields are used to cover both sides of the FEB connector. These shields connect to mating springs mounted on the crate baseplane. In addition, additional ground pins are used to further improve grounding between the FEB and the baseplane.

From the input connectors, the raw input signals are routed to the four-channel preamp hybrids. As depicted in Figure 19, the input signals are routed on two inner layers, each surrounded by

Layer definition for Atlas FEB											
Layer/Area	Input connector		Preamplifier and Shaper		SCA and ADC			Digital			
Layer 1	AGND		Routing		Routing			Routing			Top
Layer 2	AGND		AGND		AGND			DGND			P
Layer 3	Input signals routing	Trigger out	-3V		VSS			Routing			C
Layer 4	AGND		VDDAW		VDDD	VDDAR	DGND			P	
Layer 5	Trigger routing(in)		VSS		+5.5V		+5V_ADC	+3.3V	Routing	C	
Layer 6	PA +10V,+3V		Shaper control		VSS	Digital bus	VSS		+5V +2.5V	P	
Layer 7	AGND		VDDAW		VDDD	VDDAR	DGND			C	
Layer 8	Input signals routing	Trigger out	+4.5V		VSS			Output signal routing			P
Layer 9	AGND		AGND		AGND			DGND			C
Layer 10	AGND		Routing		Routing			Routing			P
Bottom											
total: 70.8 mils											

Figure 19. Figure depicting the cross-sectional structure of the FEB.

ground planes, to avoid pick-up. In addition, the use of two layers allows the routing without signal crossings, in order to avoid cross-talk. The preamps are enclosed in their own RF-gasketed Faraday shields to provide additional noise immunity.

Having gone through the first stage of amplification, the preamp outputs are routed on the outer layers directly to the inputs of the four-channel shaper ASICs, the outputs of which are then routed to the four-channel SCA chip. This allows a compact layout with short signal traces, and avoids potential problems in cross-talk and signal trace loads which could result if a more complex routing of signal traces over several signal layers were necessary.

Care has been taken to minimize digital coupling into the analog sections of this mixed analog-digital board. Whenever possible, all digital signals are transmitted via “LVDS-like” push-pull differential low-voltage signal pairs. Notable and unavoidable exceptions are the single-ended TTL digital outputs of the ADC chips. To minimize noise, these lines are kept short. In addition, series 1 k Ω resistors are inserted on each line to slow down the transitions to reduce dV/dt effects. Finally, additional 1 k Ω resistors are used to divide the signal swing by a factor of two in order to map the levels onto the GSEL inputs, while also providing further noise reduction.

As discussed in Section 6, the control of the SCA chips requires that a large number of digital address and control lines be delivered to the SCA chips at 40 MHz. These lines are particularly critical since they must be distributed at 40 MHz across the entire board, penetrating deep into the analog section. Given the use of the SCA as a de-randomizing buffer for L1 triggered events awaiting digitization, the gray encoding of the WADD bits reduces but does not eliminate the

occurrence of multi-bit transitions which might cause coherent noise in the readout. The SCA control bus lines, labeled “Digital bus” on Figure 19, are routed on a middle PCB layer and isolated by a Faraday cage created by a combination of VSS and VDD copper pours.

The use of an optical link for the FEB output data connection to the off-detector electronics avoids pick-up and ground loops which could arise from using copper cables. A general description of the grounding scheme of the overall FE electronics is provided in Reference [4].

13. FEB Production

A total of 1524 FEBs are required to equip the ATLAS LAr calorimeters, A total of 1627 FEBs were produced. This number included an allocation of up to 16 FEBs which, for whatever reason, might not be fully functional or usable in the final system. Therefore, the goal was to deliver at least 1611 FEBs to the experiment, corresponding to $\approx 6\%$ spares. After testing, debugging and repair, all but two FEBs were delivered. One FEB was too warped to fit properly in the crate, and attempts to sufficiently flatten the FEB through an additional heating cycle were unsuccessful. The second faulty FEB was damaged in the reflow oven during assembly and was not salvageable. Unfortunately, five FEBs were lost during shipment from LAL to CERN, leaving a total of 1620 FEBs delivered to CERN.

The FEB fabrication and testing process proceeded through a number of steps, as described in more detail below.

PCB Fabrication

The 10-layer FEB PCBs were fabricated in lots of typically a few hundred. The contract called for a delivery of ≈ 100 PCBs per week. The PCB producer labeled all of the FEB PCBs of a particular fabrication lot with a lot number. Each bare PCB was subjected by the PCB manufacturer to a complete automated “bed of nails” electrical test to verify the continuity and isolation of all nets.

Upon receipt at Nevis Labs, the lot number and quantity of PCBs in each production lot was recorded, along with the date received. To minimize handling of the bare PCBs, a subset of only a few PCBs per lot were visually inspected to check the quality of the PCB fabrication. The inspection checked for flatness as well as overall quality and cleanliness. A few vias were inspected under a microscope to check the alignment of the various layers. A limited number of Ohmmeter tests were also performed to check for shorts and for continuity. Accepted PCBs were delivered from Nevis to the assembler in lots, matching the delivery of the active components.

Component Pre-Testing

Each FEB includes ≈ 200 active components. As discussed previously, these include a large number of custom ASICs fabricated in different radiation-tolerant technologies. These ASICs were delivered from their fabrication and packaging process without having been tested. It was necessary, therefore, for the custom ASICs to be tested before FEB assembly.

Most of the custom ASICs were tested manually. In the case of the two analog and highest volume chips, namely the shapers and SCAs, for which more than 50,000 good chips each were required, a robotic test system was developed. As part of the shaper testing, the robotic test system

would determine which on-chip fuses to blow in order to obtain the nominal shaping time, and subsequently blow them. As discussed later, this method turned out to not work very well.

Table 10 summarizes the yields achieved in the testing of the various ASICs developed specifically for the FEB. During packaging of the DSM chips, the packaging firm did not properly use the wafer map provided by the ASIC manufacturer. As a result, a number of chips from the periphery of the wafer were packaged despite the fact that they were not fully processed during the IC fabrication process and were missing power vias. These chips were easily rejected during the testing process, since the chips in question did not draw any current. The yields in Table 10 neglect these chips, which amount to $\approx 6\%$ of the total number delivered.

Table 10. Yields from the pre-testing of the custom ASICs developed specifically for use on the FEB.

ASIC Type	Process	ASIC	Die Size (mm ²)	Yield (%)
Analog	AMS BiCMOS	Shaper	18	88
	DMILL	SCA	19.8	65
Digital	DMILL	SMUX	16	92
		SPAC slave	27	89.7
		CONFIG	31	84.3
	DSM	GSEL	16	95.8
		CLKFO	4	98.5
		SCAC	16	96.6

After testing, components were delivered in trays suitable for use in industrial automatic Pick-and-Place machinery, and for baking at 125° C in order to remove moisture from the components before they were used for PCB assembly.

FEB Assembly

Each FEB has $\approx 20,000$ solder joints. On average, about 60 FEBs per week have been processed during the production, with a peak throughput of about 120 FEBs per week.

The majority of the 200 active components per FEB were in surface-mount packages, and were assembled onto the FEB through an automated assembly line including a Pick-and-Place machine and reflow oven. The line included test equipment which performed full and automated visual and X-ray inspections of each FEB. The visual inspection verified the presence and correct orientation of each component, as well as visually inspecting all solder joints. The X-ray inspection checked each solder joint for opens, shorts, bridges, etc., and also verified the presence of sufficient solder and checked the solder profile to try identify cold or weak solder joints. Since the FEB has components mounted on both sides of the PCB, two passes through the automated assembly line were required.

The through-hole sockets used for installation of the preamps and LSBs had to be soldered onto the FEB after the surface-mount steps were complete. Given the size of the PCB, it was decided to avoid use of wave soldering. Instead, the sockets were soldered using a “selective soldering” machine in which a soldering tip automatically soldered each individual socket. The final FEB assembly step involved hand soldering of a number of components, including the OTx and the Faraday shields installed over the preamps and over the input signal connectors.

14. FEB Testing

As described in more detail in the following sections, the FEB testing plan [38] proceeded in a number of stages, first at the industrial assembler and then in various labs. Debugging and repair of failures has been mostly performed in time with the production.

14.1 Initial Test and HASS Test

Given the large number of custom ASICs, plus the mixed analog-digital nature of the FEB, the use of commercial test techniques such as JTAG, etc. was not practical. Instead, we two custom test stands were developed and installed at the industrial assembly house.

The first test stand could be used to test one FEB at a time, and included computer-controlled power supplies with current limiting in order to minimize the possibility of damage in case of a short or other problem. Once the FEB was successfully fully powered and the currents recorded and verified to be within pre-defined acceptance windows, a series of automated tests was performed. These tests included configuring and reading back all of the various chips on the FEB, and triggering and reading out the FEB in order to verify, in addition to the correct functionality of the FEB, the values (means and RMS values) of the pedestals. The entire test took ≈ 30 seconds per FEB and was fully automated. The various results of the test were recorded according to the unique serial number of the FEB under test. The readout of the FEBs and analysis of the data was supervised remotely from Nevis Labs.

FEBs which passed the first test were then installed in the second test stand, which was located within a large environmental test chamber (see Figures 20). A total of 16 FEBs could be tested simultaneously. Once the FEBs were installed, a pre-test was performed, where all 16 FEBs were tested in a manner similar to the single board test, in order to verify that each FEB was correctly installed and connected. Once this test was successful, the chamber was closed and an automated Highly Accelerated Stress Screening (HASS) test performed [39]. For a period of ≈ 6 hours, the FEBs were operated and triggered while the test chamber underwent a series of thermal cycles between 0°C and 55°C . The goal of the HASS test was to try trigger and identify any failures due to component infant mortality, cold solder joints, etc. The duration of the HASS test was limited in order to allow testing to proceed with sufficient throughput to match the required FEB delivery schedule. The temperature range was limited in order to avoid damage or excessive aging of the custom components, in particular of the OTx.

After completion of the thermal cycling, the FEBs were subject once again to the same test procedure as used in the HASS pre-test, in order to identify any failures which had occurred. The acceptance windows used for the various parameters were fairly wide, since the goal at this stage was to identify failures rather than to ensure the analog quality of the results.

Some failures of the OTx optical transmitter were observed after HASS screening of the first sets of FEBs. These failures were determined to be due to problems in the OTx burn-in process which was used for the first batch of OTx before delivery to the FEB assembly process. As a result, modifications were made to the OTx burn-in process, and it was decided to replace all transmitters from the first OTx batch. In total, 138 OTx had to be removed manually from FEBs and replaced. Apart from this first batch of OTx, no OTx failures occurred during the HASS screening. Other failures specifically attributed to the HASS screening were rare, with a few failures of other



Figure 20. Photograph showing the front end crate with FEBs installed in the environmental chamber used for the HASS test. The single-FEB test stand used for the initial test can be seen behind the chamber in the right part of the photograph.

components. In addition, in some instances the thermal stress revealed bad solder joints that had not been obvious in the data taken before the thermal cycling.

14.2 Digital Testing

After completion of the tests and HASS screening at the assembler, all FEBs were sent to Nevis Labs, where they were subject to more extensive testing. In addition, any FEBs which failed any test at the assembler were delivered for debugging at Nevis Labs, in order to ensure quick feedback in the case of potential manufacturing problems.

A visual inspection was performed upon receipt of each FEB. The inspection followed a checklist that included checks for overall quality, flatness, correct labeling, and proper component placement. After the inspection was completed, the FEBs were subject to an automated “Digital Test” that included repeating all of the tests performed at the assembler, now with tighter acceptance windows. Additional tests were also performed to precisely measure the frequency range over which the QPLL circuit properly locked, measure the clock jitter as a function of frequency, and to characterize both the jitter and performance of the output optical link. Measurements were also made to calibrate the current sources of the DCU chips. The Digital Test was performed on a single FEB at a time, and took about 6 minutes per FEB. Pre-defined acceptance windows were applied to all parameters, a sample of which are reported in Table 11.

Note that the “Digital Test” included tests of the pedestal values and noise. However, since this test was performed before the FEB was equipped with the plug-in preamps, the analog performance in response to the injection of signal pulses could not be checked at this stage. Figure 21 shows the

Table 11. Acceptance windows for some of the parameters measured during the Digital Test of the FEBs.

Parameter Type	Measurement	Units	Min. Value	Max. Value
Pedestals	Mean	ADC counts	880	1080
	RMS (w/o preamps)	ADC counts	0.5	1.5
Clock Quality	Jitter	ps	5	30
	Lower Edge of Lock Range	kHz	$f_{\text{LHC}} - 6$	$f_{\text{LHC}} - 2$
	Upper Edge of Lock Range	kHz	$f_{\text{LHC}} + 2$	$f_{\text{LHC}} + 6$
	Width of Lock Range	kHz	5	11
Optical Output Quality	Data-dependent Jitter	UI	0.01	0.18
	Average Optical Power	dBm	-10	-0.01
	Extinction Ratio Opt. Output	dB	10	30
DCU Quality	Current Source	μA	10	35
	Temperatures (uncalibrated)	degC	15	35
DCU Voltage Monitoring	3.3 V SCA VDDD	V	3.2	3.4
	3.3 V SCA VDDA	V	3.2	3.4
	2.5 V dig.	V	2.4	2.6
	-1.7 V SCA VSS	V	-1.875	-1.6
	3.0 V preamps	V	2.9	3.1
	-3.0 V shaper VSS	V	-3.3	-2.7
	4.5 V shaper VDD	V	4.40	4.65
Power Consumption	Current +6 V dig.	A	2.06	2.28
	Current +11 V (w/o preamps)	mA	2	7
	Current +6 V analog	A	4.23	4.53
	Current +4 V	A	2.6	2.9
	Current -4 V	A	4.84	5.36
	Total Power Dissipation	W	67.4	74.0

mean pedestal and RMS per channel as measured in the Digital Test for all FEBs. The distributions are uniform and narrow. Any significant deviation would reveal component or assembly problems.

To measure the lock range of the QPLL, the TTC system controlling the FEB was driven from an Analog Devices AD9852 DDS synthesizer, which can be programmed with 1 Hz precision. The stability and accuracy of the clock frequency setting were guaranteed by utilizing a GPS-referenced 10 MHz source (Symmetricom XLi). As a result, an absolute measurement of the QPLL lock range was made with a precision of few Hz, the residual uncertainty arising mainly due to FEB temperature variations. As shown in Figure 22(a), the lock range of the final QPLL circuits on the FEBs safely covers the expected LHC operation frequency range of $40.078966 \text{ MHz} \pm 12 \text{ ppm}$.

The jitter of the QPLL clock was measured by analyzing the copy of the clock brought for diagnostic purposes through a coaxial connector on the FEB front panel. The system clock was very clean, with a jitter of less than 5 ps RMS, and was used as the reference trigger source to a 20 GSa/s, 6 GHz analog bandwidth oscilloscope (Agilent model 54855A) used to measure the QPLL clock. The peak-to-peak and RMS jitter of the FEB clock were measured and recorded throughout the lock range of the QPLL, in steps of 200 Hz. For about 1% of the QPLL circuits




Figure 21. Mean value (left) and RMS (right) of the pedestal per channel in high gain, as measured during the Digital Test (ie. before insertion of the preamps).




Figure 22. Distributions of (a) lower and upper lock range of the QPLL circuits, as compared to the clock frequency range expected during LHC operation, and (b) the measured jitter of the FEB clock.

it was found that for small frequency bands inside the lock range the jitter increased to very high values, due to activity dips in the quartz crystal. In such cases the QPLL circuit was replaced and the measurement repeated.

The clock jitter was within acceptable limits for most FEBs (see Figure 22(b)). About 3.5% of all FEBs show jitter above 20 ps, which however can be reduced by reprogramming the PLL current of the TTCrx timing receiver chip. Fifteen FEBs with jitter above 20 ps for all of the possible settings of the PLL current were sent back to the assembler for replacement of the TTCrx chip. In all cases the jitter was found after the repair to be below 10 ps for default PLL settings. The 40 FEBs with high initial jitter above 20 ps for default PLL settings, but jitter below 20 ps for non-default PLL current settings, have been selected for a dedicated test. This test has been performed to determine the optimal PLL current settings and ensure that the jitter is stable over a wide temperature range. The jitter has been determined inside an environmental chamber for temperatures between 5° C and 45° C and as function of the PLL current setting. Optimized and stable PLL settings were found and stored in a database for all 40 FEBs.

Another set of tests investigated the performance of the GLink serializer and high speed optical transmitter. These measurements used the FEB test clock signal as a reference to measure the FEB optical data output, which was connected to the fast oscilloscope through a NewFocus 12 GHz DC-coupled receiver (Model 1554). While the FEB under test was transmitting 5-sample fixed-gain data at a 50 kHz L1A rate, the 1.6 Gb/s serial data output was acquired, for 20 sweeps, using the full memory depth of the 54855A oscilloscope, spanning approximately 13 ms (at 20 GSa/s). A dedicated analysis of the serial data was carried out, taking into account the specific features of the GLink CIMT protocol. The phase error of the GLink master transitions was calculated, allowing detection of possible problems with the PLL on the GLink serializer. The data-dependent jitter, sensitive to problems with the optical transmitter or the GLink output, was also calculated. In terms of the Unit Interval (UI) of 625 ps, it was typically 0.1 UI, including the contribution from the rise-time of the filter between the optical receiver and the oscilloscope input. Finally, the average output optical power (see Figure 23) and modulation amplitude were also measured.

The DCU readout of all monitored voltages and temperatures was tested for each FEB. In addition, the 20 μ A constant-current outputs of the two DCU chips on each FEB were measured at room temperature (see Figure 23). This current is used to drive two thermistors on the top layer of each FEB. A large spread of the measured current with respect to the nominal current of 20 μ A was observed. All current values are stored in a database used to calibrate the measured temperatures of the FEBs after installation in the experiment.




Figure 23. Distributions of measured optical power of the OTx (left) and output current of the current source of the DCU chips (right).

The overall failure rate after digital testing was 15%. Besides component failures (73% of all failures), assembly problems (23%) and PCB manufacturing failures (2%) have been observed. The failure rates observed for the active components are summarized in Table 12. The observed component failure rates ranged from approximately 0.01% to 1%.

14.3 Analog Testing

Upon successful completion of the Digital Test at Nevis Labs, the FEBs were shipped to either Brookhaven National Laboratory (BNL) in New York (USA) or to Laboratoire de l'Accélérateur Linéaire (LAL) in Orsay (France). In the two labs the 23 different flavours of the FEBs were configured by mounting the appropriate preamps and LSBs. Once configured, the FEB were subjected

Table 12. Component failures observed in tests up to and including the Digital Test. The OTx numbers do not include the failures observed in the first batch before they were replaced.

Component	Number Per FEB	Number of Failures	Failure Rate
Shaper	32	5	0.01%
SCA	32	32	0.06%
OpAmp	32	4	0.01%
ADC	16	5	0.02%
GSEL	8	11	0.08%
SMUX	1	0	0%
GLINK	1	7	0.43%
OTx	1	4	0.27%
TTCrx	1	8	0.49%
QPLL	1	19	1.17%
CLKFO	7	4	0.04%
SCAC	2	1	0.03%
VREG	19	6	0.02%
CONFIG	1	11	0.68%
SPAC slave	1	2	0.12%
DCU2	2	12	0.37%




Figure 24. Signal form (in ADC counts) in 1 ns steps for one Medium gain channel.

to a rigorous test of their analog performance, including noise and pulse measurements. Figure 24 shows a sample pulseshape measured in MED gain for one channel. Analog measurements were also performed at Nevis Labs for a control sample of $\approx 3\%$ of the FEBs in order to have quick feedback to the FEB production and assembly process.

The BNL test setup was based upon a standard LAr crate and electronics, as described in Reference [4]. One half of the crate could be populated with FEBs, allowing a test of up to 14 FEBs simultaneously. The LAL test bench was based upon a custom, locally developed setup, in which one FEB was tested at a time. The analog testing procedure for the FEB consisted of a set of different types of runs, including pedestal runs, signal shape runs, and ramp runs for the standard readout as well as for the trigger output of the FEB. Each FEB flavour had a dedicated configuration file with acceptance cuts. An example of the acceptance criteria for one of the FEB flavours is shown in Table 13. Similar tests were performed at BNL and LAL. For illustration, results from the tests at LAL are discussed below.

Table 13. Acceptance windows for some of the parameters measured during analog tests at LAL of FEBs equipped with 50 Ω preamps and S8X2 LSBs. The number of events used in each measurement is also listed.

Parameter	Gain	Min. Value	Max. Value	Events
Mean Pedestal (ADC counts)	All	900	1100	see Noise
Noise (ADC counts)	HI	6.0	8.5	100000
	MED	0.9	1.5	50000
	LO	0.6	1.2	25000
Fixed Sequence Noise (ADC counts)	HI	0.1	0.95	see Noise
	MED	0.1	0.6	see Noise
	LO	0.1	0.6	see Noise
Coherent Noise (ADC counts)	All	-	0.075	see Noise
Time Offset of Peak (ns)	HI	68	77	5000
	MED	70	77	5000
	LOW	68	75	5000
Deviation of Gain	All	-7%	+7%	6700
Trigger Pedestal (ADC counts)		400	550	10000
Trigger Noise (ADC counts)		1.5	6.	see Trigger Ped
Trigger Timing (ns)		66	71	5000
Deviation of Trigger Gain		-3%	+3%	10200

The pedestals, noise and fixed sequences noise (FSN) were measured separately for HI, MED, and LO gains. Dead channels were easily detected via too low noise. The noise measurement is also sensitive to the time constant of the shaper. The FSN measures the pedestal dispersion of the 144 capacitor storage cells of the SCA. An analysis of coherent noise was also performed, summing over all 128 channels. The overall noise distribution of all channels and FEBs equipped with 50 Ω 1 mA preamps is shown in Figure 25 for HI gain.

A calibration signal of fixed amplitude, chosen to yield a signal in the saturation region of each of the gains, was then injected and the minimum and maximum signal of the peak sample and a non-peak sample were determined for each channel as a function of the SCA cell used to store the signal. The difference between the minimum and maximum for each channel (and gain) was required to be small. This test was designed to detect premature saturation of individual pipeline cells of the SCA (peak sample) as well as timing deviations (off peak sample).

Timing runs with ten samples were performed to measure the signal shape in steps of 1 ns. The signal shape was compared to a reference signal to detect deviations. The signal shape analysis also verifies the read and write operations of the combination of SCAC chips, SCA control bus and SCA chips. The position of the signal peak was required to fall within an acceptance window of about 7 ns. A significant number of shapers were found to have a shaping time shorter than nominal, since the fuses had not been properly blown during the robotic pre-testing of the shaper chips. These shapers were either replaced or the pins of the fuses which should have been blown were cut in order to correct the problem. After installation of the FEBs in ATLAS, many more instances of this problem were found, and in fact it was seen that the fraction of shapers with such problems was growing with time, suggesting a systematic problem with the fuses. By the time this fraction reached $\approx 1\%$ and was the leading cause of bad channels, it was decided to cut the pins of all ≈ 52000 shaper chips in order to eliminate this problem.

The linearities of the three gain scales were calibrated separately, with the highest DAC value chosen appropriately for each gain scale to give an almost full height signal. Typically 67 DAC steps were used, equally spaced between the minimum and maximum signal. The signal timing was adjusted so that the maximum of the signal coincided with the third sample in time. The third sample signals were averaged and the gain was determined with a linear fit. The gain was required to be within an acceptance window of the expected value. The distribution of the gains determined for the LO gain of all channels of all FEBs equipped with $50\ \Omega$ preamp is shown in Figure 25. A ramp run was also performed with the GSEL chips configured in autogain mode, in order to verify the proper operation of the gain selection algorithm.




Figure 25. Some results measured for all channels during analog testing at LAL of 479 FEBs equipped with $50\ \Omega$ preamps. The left plot shows the noise distribution for HI gain, while the right plot shows the fitted slope of the linearity curve in LO gain.

In addition to evaluating the analog performance of the readout channels, the trigger sum outputs were also checked. The first test consisted of injecting a large input signal while turning off all shaper sums and verifying that none of the 32 trigger channels showed an output signal. The shaper sum outputs were then turned on one at a time. This test verifies the correct functioning of the shaper output and also crosschecks the correct configuration of the LSB as the channel with the expected output signal is dependent on the LSB type. The gains as well as the shapes of the trigger sum outputs were measured, to ensure a homogeneous trigger response in ATLAS.

The three most frequently detected failures were preamp problems, mostly due to breaking a pin or the preamp hybrid itself during insertion on the FEB, shapers with a short peaking time, and SCAs with a cell with large leakage currents or a limited dynamic range. These problems occurred at rates of $\approx 0.6\%$ each for the preamps and shapers, and $\approx 1.2\%$ each for the SCAs. All FEBs were repaired as necessary to ensure they passed the Analog Test requirements during subsequent retesting.

Once the Analog Test was passed, the cooling plates were mounted on both sides of the FEB and an air pressure test at 3 bar was performed to verify that the system was leakless. A partial re-test of the FEB was then performed to ensure no damage was caused during this step, after which the FEB was shipped to CERN. Upon arrival at CERN, a leak test was performed, and then the FEBs were prepared for installation on the calorimeter in the ATLAS pit.

15. Performance of the FEB

To evaluate the performance of the LAr FE readout as a *system*, a dedicated set of tests of a half-crate of electronics, including 14 FEBs, was performed. Detailed results of this FE system test and tests during installation and commissioning of the electronics on the ATLAS detector at CERN can be found in Reference [4]. Here are summarized only a few results of the measurements of the performance of the FEB.

The FEB meets or exceeds all of the specifications that were summarized in Table 1. The noise of a single channel of the FEB is ≈ 0.8 ADC counts, without the preamp. This includes contributions from the shaper, SCA, op-amps, and the ADC itself. Including the (unloaded) preamp, the noise increases for HI gain to ≈ 3 -6 ADC counts, depending on the preamp type. Therefore, as required, the HI gain noise is dominated by the preamp noise. Once the FEB is connected to the detector, this noise increases to ≈ 4 -10 ADC counts, depending on preamp type and detector capacitance. The coherent noise per channel across the 128 channels of one FEB is typically 1-3% of the total noise per channel.

During typical operation at the LHC, for each L1 trigger five samples will be digitized and read out from the FEBs per channel. For large pulses, the amplitude can be reconstructed with a resolution better than 0.1% and the time can be measured with a resolution of ≈ 20 ps. The nearest-neighbour crosstalk is $< 1\%$, dominated by the input connectors, and small compared to crosstalk and energy sharing effects in the detector itself.

16. Summary

We have discussed the design, implementation and production of the Front End Boards designed to meet the challenging specifications of the readout of the ATLAS LAr calorimeter system. Each FEB handles 128 signal channels. Implementing the specified FEB functionality and performance while meeting the radiation tolerance requirements required the utilization of 11 different custom ASICs, including eight developed specifically for the FEB. The FEB is implemented as a large 10-layer PCB with components mounted on both sides, and with $\approx 20,000$ solder joints per board. Production and testing the full set of FEBs took approximately one year, with the testing work shared among three labs.

Each FEB was subjected to a rigorous Quality Assurance test program, and debugged and repaired as necessary until it passed all tests. The main difficulties encountered during production and testing of the FEBs included the following:

- the sizeable jitter of the 40 MHz clock from the TTCrx chip required a re-design of the FEB after the first prototype production to introduce the QPLL-based jitter filter;
- the first deliveries of crystals for the QPLL demonstrated activity dips and had to be subject to a special screening process to reject any crystals with such problems;
- the first batches of OTx optical transmitters were subject before delivery to a faulty burn-in process, and subsequently experienced failures during the FEB HASS test; well over 100 had to be de-soldered from FEBs and replaced with OTx from later deliveries after the problem was solved;
- some SCA chips had to be replaced since the FEB testing revealed that they had pipeline cells that had large leakage currents or which had a response that saturated before reaching the required dynamic range;
- as discussed in Section 14.3, problems were identified with the fuses implemented on the shaper chips to allow tuning of the shaping time. Eventually it was decided to cut the appropriate pins of all ≈ 52000 shaper chips in order to eliminate the problem.

The FEBs meet or exceed the exacting specifications of the ATLAS LAr readout. For example, for large pulses the FEBs measure the deposited energy with a relative precision better than 0.1%, and with a timing resolution of ≈ 20 ps. In total, 1524 FEBs are required to read the LAr calorimeter system. A total of 1620 FEBs were delivered to CERN.

Acknowledgements

We would like to acknowledge the excellent work and essential contributions of the technical staff working in our institutions, including L. Holm, J. Schaapman and L. Wampler (Alberta), K. Sexton, P. Bijoneau, R. Burns and A. Hoffman (BNL), P. Betts, N. Bishop, J. Capone, R. Gardner, M. Hwang, and A. Teho (Columbia), O. Bohner, B. Debennerot, P. Favre, M. Lechowski, M. Quentin, W. Roudil, and S. Trochet (LAL), M. Dhellot, C. Goffin, A. Guimard, and J.M. Parraud (LPNHE), L. Gallin Martel, J.P. Richer, O. Rossetto, and C. Vescovi (LPSC), and F. Sabatini and B. Monticelli (Milano).

This work has been supported by the Natural Science and Engineering Research Council of Canada (Canada), by the US National Science Foundation under contract number NSF PHY03-01292 (Columbia), by the Commissariat à l’Energie Atomique (France), by the Bundesministerium fuer Bildung, Wissenschaft, Forschung und Technologie, Germany, under contract numbers 05HA8EX16 and 05HA1EX12 (MPI), by INFN (Milano), by the US Department of Energy grant DE-FG02-04ER41299 (SMU), and by the National Science Council, Taiwan, ROC (Taiwan). We thank all funding agencies for their financial support.

References

- [1] ATLAS Collaboration, “ATLAS Technical Proposal”, CERN/LHCC/94-43 (1994), available at <http://atlasinfo.cern.ch/ATLAS/TP/tp.html>.
- [2] ATLAS LAr Collaboration, “ATLAS Liquid Argon Calorimeter Technical Design Report”, CERN/LHCC/96-41 (1996), available at <http://atlas.web.cern.ch/Atlas/GROUPS/LIQARGEXT/TDR/>;
B. Aubert et al., “Performance of the ATLAS Electromagnetic Calorimeter Endcap Module 0”, Nucl. Instr. and Meth. A **500** (2003) 178-201;
B. Aubert et al., “Performance of the ATLAS Electromagnetic Calorimeter Barrel Module 0”, Nucl. Instr. and Meth. A **500** (2003) 202-231;
J. Colas et al., “Position Resolution and Particle Identification with the ATLAS EM Calorimeter”, Nucl. Instr. and Meth. A **550** (2005) 96-115;
B. Aubert et al., “Construction, Assembly and Tests of the ATLAS Electromagnetic Barrel Calorimeter”, Nucl. Instr. and Meth. A **558** (2006) 388-418;
M. Aharrouche et al., “Energy Linearity and Resolution of the ATLAS Electromagnetic Barrel Calorimeter in an Electron Testbeam,” Nucl. Instrum. Meth. A **568** (2006) 601-623;
M. Aharrouche et al., “Response Uniformity of the ATLAS Liquid Argon Electromagnetic Calorimeter”, Nucl. Instr. and Meth. A **582** (2007) 429-455.
- [3] A. Bazan, et al., “ATLAS Liquid Argon Calorimeter Back End Electronics”, JINST 2:P06002, 2007.
- [4] “ATLAS Liquid Argon Calorimeter Front End Electronics”, in preparation for submission to JINST.
- [5] ATLAS Trigger Collaboration, “ATLAS Level 1 Trigger Technical Design Report”, CERN/LHCC/98-14 (1998), available at <http://atlas.web.cern.ch/Atlas/GROUPS/DAQTRIG/TDR/tdr.html>.
- [6] ATLAS Collaboration, “ATLAS Detector and Physics Performance Technical Design Report”, CERN/LHCC/99-14/15 (1999), available at <http://atlasinfo.cern.ch/Atlas/GROUPS/PHYSICS/TDR/access.html>.
- [7] Details of the ATLAS policy on radiation tolerant electronics can be found at <http://atlas.web.cern.ch/Atlas/GROUPS/FRONTEND/radhard.htm>.
- [8] M. Dentan, et al., “DMILL, A Mixed Analog-Digital Radiation-Hard BiCMOS Technology for High Energy Physics Electronics”, IEEE Trans. Nucl. Sci. **43** (1996) 1763-1767.
- [9] G. Anelli, et al., “Radiation Tolerant VLSI Circuits in Standard Deep Submicron CMOS Technologies for the LHC Experiments: Practical Design Aspects”, IEEE Trans. Nucl. Sci. **46** (1999) 1690-1696.
- [10] Information about the LHC TTC system is available at <http://ttc.web.cern.ch/TTC/intro.html>.
- [11] F. Hubaut, et al., “Specifications of the Serial Protocol for the ATLAS LAr Calorimeters: The SPAC3 Protocol and Slave ASIC”, available at <https://edms.cern.ch/file/110351/3/>.
- [12] M. Citterio, J. Kierstead and S. Rescia, “The ATLAS Preamplifier: Performance, Radiation Damage, ESD Resistance, Reliability, and Manufacturing Issues”, Proceedings of the Fifth Workshop on Electronics for the LHC Experiments (1999) Snowmass.
- [13] R.L. Chase, et al., “Transmission Line Connections Between Detector and Front End Electronics in Liquid Argon Calorimetry”, Nucl. Instrum. Meth. A **330** (1993) 228-242;
R. L. Chase and S. Rescia, “A Linear Low Power Remote Preamplifier for the ATLAS Liquid Argon EM Calorimeter”, IEEE Trans. Nucl. Sci. **44** (1997) 1028-1032.

- [14] H. Chen, et al., “Atlas-LARG Preamplifier: A Summary Report on comparative studies between BNL and Milan productions”, available at <https://edms.cern.ch/document/867072/1>.
- [15] J. Ban, et al., “Cold electronics for the liquid argon hadronic end-cap calorimeter of ATLAS”, Nucl. Instrum. Meth. **A556** (2006) 158-168.
- [16] E. Ladygin, et al., “Preshaper for the Hadronic Endcap Calorimeter”, available at <https://edms.cern.ch/document/875444/1>.
- [17] J. Collot, et al., “The LAr Trigain Shaper”, available at <https://edms.cern.ch/file/871897/1>.
- [18] Information about Austria Micro Systems (AMS) processes is available at <http://www.austriamicrosystems.com/>.
- [19] D. Breton, et al., “HAMAC, a Rad-Hard High Dynamic Range Analog Memory for ATLAS Calorimetry,” Proceeding of the 6th Workshop on Electronic for LHC Experiments, Cracow, Poland (2000) 203-207.
- [20] Data sheet for the Analog Devices AD8042 dual op-amp, available at http://www.analog.com/UploadedFiles/Data_Sheets/AD8042.pdf.
- [21] Data sheet for the Analog Devices AD9042 ADC, available at http://www.analog.com/UploadedFiles/Data_Sheets/AD9042.pdf.
- [22] S. Boettcher, J. Parsons and B. Sippach, “The Gain Selector ASIC for the ATLAS LAr Calorimeter”, available at <https://edms.cern.ch/file/871897/1>.
- [23] S. Boettcher, J. Parsons, B. Sippach, and D. Gingrich, “The SCA Controller for the ATLAS LAr Calorimeter”, available at <https://edms.cern.ch/file/871897/1>.
- [24] D.M. Gingrich, et al., “Radiation Tolerant ASIC for Controlling Switched-Capacitor Arrays”, IEEE Trans. Nucl. Sci. **51** (2004) 1324-1332.
- [25] Data sheet for On Semiconductor 10H116 triple line receiver, available at <http://www.onsemi.com/pub/Collateral/MC10H116-D.PDF>.
- [26] D. Dzahini, “Development of a DMILL Radhard Multiplexer for the ATLAS GLink Optical Link and Radiation Test with a Custom Bit Error Tester”, Seventh Workshop on Electronics for LHC Experiments (2001) Stockholm, available at <http://cdsweb.cern.ch/record/588659>.
- [27] Detailed information about the GLINK chipset and protocol is posted at <http://literature.agilent.com/litweb/pdf/5966-1183E.pdf>.
- [28] Data sheet for the Micrel SY88922 laser diode driver, available at http://www.micrel.com/_PDF/HBW/sy88922v.pdf.
- [29] Information about the TrueLight TTR-1A43 VCSEL TOSA is available at <http://www.truelight.com.tw/detail/41052/41052.html>.
- [30] T. Liu, et al., “Quality Control Test of Optical Transmitters in Liquid Argon Calorimeter Readout System”, ATL-LARG-PUB-2005-004, available at <https://edms.cern.ch/document/882033/1>.
- [31] W. E. Cleland, “Layer Sum Boards for the ATLAS Liquid Argon Calorimeters”, ATL-AL-EN-0019, available at <https://edms.cern.ch/document/110377/1>.
- [32] Data Sheet for the HFA1135, available at <http://www.intersil.com/data/fn/fn3653.pdf>.

- [33] M. Dhellot, et al., “The Controller Board and SPAC Buses of the ATLAS LAr Calorimeter Frontend Crates”, available at <https://edms.cern.ch/document/867048/1>.
- [34] Data sheet for the Quartz Crystal Phase-Locked Loop (QPLL), available at <http://proj-qpll.web.cern.ch/proj-qpll/>.
- [35] P. Cros, et al., “The Configuration Controller”, available at <https://edms.cern.ch/file/871897/1>.
- [36] Information, including data sheets, for the L4913 and L7913 radiation-tolerant voltage regulators is available from <http://lhc-voltage-regulator.web.cern.ch/>.
- [37] G. Magazzu, A. Marchioro, and P. Moreira, “The Detector Control Unit: An ASIC for the Monitoring of the CMS Silicon Tracker”, IEEE Trans. Nucl. Sci. **51-4** (2004) 1333-1336.
- [38] J. Ban, et al., “ATLAS LAr FEB Qualification and Testing Procedures”, available at <https://edms.cern.ch/file/872263/1>.
- [39] K. Benslama, et al., “Highly Accelerated Stress Screening of the ATLAS Liquid Argon Front End Boards”, available at <https://edms.cern.ch/file/872263/1>.