


Draft version 0.10



ATLAS NOTE

ATL-INDET-PUB-2007-XXX

March 7, 2008



The ATLAS Pixel Detector

The Pixel Collaboration

Abstract

Status of received/implemented comments :

Ivan - implemented

Gil - implemented

John - implemented

Maurice - implemented

Emlyn - implemented

Attilio - implemented (where I could)

KK - implemented

4 Electronics Systems

4.1 Overview

The first comprehensive proposal of the pixel electronic system was described in 1997-98 in the ATLAS Inner Detector and Pixel Detector Technical Design Reports [1, 2]. The complete system underwent several revisions in the subsequent years. The required radiation tolerance is 50 Mrad, corresponding to 10 years of operation at the nominal LHC luminosity for the external layers and 3 years of operation for the innermost layer (B-layer). The total number of instrumented channels is 80 million, each containing approximately 1,000 transistors and consuming a maximum power of $100 \mu\text{W}$ (power for on-detector circuitry only).

4.1.1 System Architecture

A block diagram that illustrates the system architecture with the principal links is shown in Fig. 1.

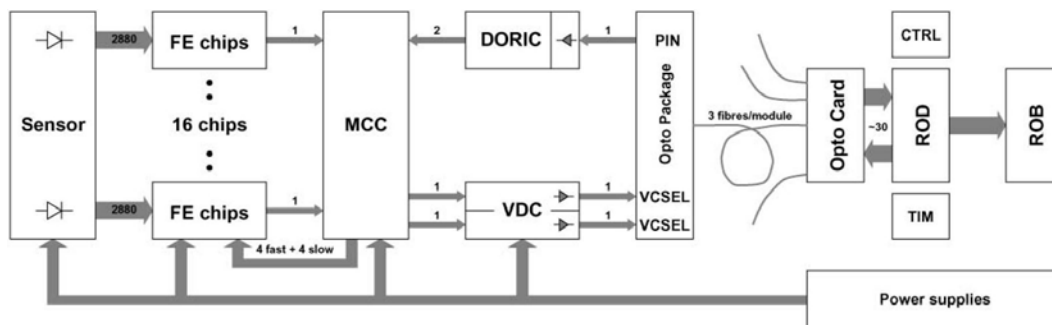


Figure 1: Block diagram of the pixel detector System Architecture. (Original figure needed)

Charge released by ionizing particles in the cells of the sensor array are collected by 16 front-end chips (FE) per module, arranged in two rows of eight chips. The 16 FEs are read out by a Module Control Chip (MCC). Data are transmitted from the FE to the MCC using Low Voltage Digital Signalling (LVDS) serial links, configured in a serial star topology. The serial protocol minimises the number lines to be routed, while the star topology maximizes bandwidth and reliability. Each module is then connected to the off-detector Read-out Drivers (ROD) through opto-links. One down link is used to transmit clock, trigger, commands and configuration data, while one or two up-links are used for event readout. The B-layer uses two up-links to increase the aggregate bandwidth needed for the higher average hit occupancy at the minimum radius. The readout (R/O) architecture is "data-push". This means that each component in the chain (FE, MCC) always transmits at the maximum rate and there is no busy mechanism to stop transmission when buffers are full. Each upstream component in the R/O chain (MCC, ROD) constantly monitors the number of events received and compares the results with the number of triggers sent. If the difference of the two is bigger than a predefined value, triggers downstream are blocked and empty events are generated.

The power supply system uses commercial components, adapted to the requirements of the pixel detector, for the low (electronics) and high (sensor bias) voltages. The use of deep sub-micron electronics and long resistive cables with significant voltage drop, required the use of low voltage regulator boards, approximately 10 meters from the pixel detector. The maximum voltage rating for the pixel electronics is 4 V. To interface between the electrical and the optical sides of the opto-links, special receiver and driver chips (DORIC and VDC) have been implemented. An opto-card (Back of Crate or BOC) is also used for each ROD.

4.2 Front-end Chip (FE)

4.2.1 Front-end Chip History

Small scale Front-end chips that demonstrated the required analog and digital architecture were first developed in the second half of 1990s (M72b [3], Lepton [4], Marebo [5, 6] and Beer & Pastis [3, 7]). The first rad-soft functional prototypes of full size chips were submitted in 1998: FE-B at LBNL, FE-A/C (Pirate) at Bonn/CPPM. The FE-B chip was designed using $0.8\ \mu\text{m}$ HP CMOS technology and had the same basic readout architecture used for the final chips. The FE-B charge amplifier uses a direct cascode¹⁾ and source follower, feedback capacitance of 4 fF. The DC feedback is based on the Marebo design. The discriminator used a dual threshold, with a low threshold for precise timing and a high threshold.

FE-A was made on $0.8\ \mu\text{m}$ BiCMOS technology from AMS, whereas FE-C was a full CMOS version. The charge amplifier used a folded cascode input stage with feedback capacitance of 3 fF and a new, improved DC feedback. The discriminator was AC coupled, with an input fully differential bipolar pair in the A version and CMOS in the C version. The column readout architecture used a shift register to transport the hit address to the bottom of the chip. Hits are associated with the level 1 trigger (L1) by counting the number of clock cycles needed for the hit to reach the bottom of the column. FE-A/B/C demonstrated all the basic ATLAS pixel performance goals in the laboratory and in beam. The subsequent chip was developed using the basic concept of the amplifier/discriminator from FE-A/C and the column readout architecture from FE-C. The European and US front-end design teams joined forces to combine all of the experience gained with radiation-soft chips into a common layout for the DMILL²⁾-”Durci Mixte sur Isolant Logico-Lineaire” technology (known as FE-D). FE-D1 was submitted in July ’99 together with the DORIC and VDC chips and a prototype MCC-D0. A new production run was submitted in Aug ’00 with two versions of FE-D2: one with dynamic and the other with static memory cells. This run included the full MCC-D2 and new DORIC and VDC chips as well. Yields of both FE and MCC were unacceptable and work with this vendor was terminated. Work on FE-H began in Dec 1999 [8]. The chip was almost ready but was never submitted also because of massive cost increases from Honeywell. The failure of both traditional rad-hard vendors left the collaboration with the Deep Sub-micron (DSM) approach, based on commercial process $0.25\ \mu\text{m}$ CMOS and a rad-tolerant layout. A major design effort was initiated in Sep 2000. Three versions (FE-I1, FE-I2 and FE-I3) were submitted. The final chip (FE-I3) became available in late 2003. Table 1 gives a summary of the front-end designs developed for the ATLAS pixel detector.

4.2.2 Design

Chip Architecture The readout chip for the ATLAS pixel detector [13, 14], shown in Fig. 2, contains 2880 pixel cells of $50 \times 400\ \mu\text{m}^2$ size arranged in a 18×160 matrix. Each pixel cell contains an analogue block where the sensor charge signal is amplified and compared to a programmable threshold using a comparator. The digital readout part transfers the hit pixel address, a hit leading edge (LE) timestamp, and a trailing edge (TE) timestamp to the buffers at the chip periphery. In these buffers a time over threshold (ToT) is calculated by subtracting the TE from LE timestamp. These hit buffers monitor the time of each stored hit by inspecting the LE time stamp. When a hit time becomes longer than the latency of the L1 trigger ($3.2\ \mu\text{s}$) and no trigger signal is recorded, the hit information is deleted. Hits marked by trigger signals are selected for readout. Triggered hit data are then transmitted serially out of the chip in the same order as the trigger arrival.

¹⁾The cascode is a two-stage amplifier composed of a transconductance amplifier followed by a current buffer.

²⁾DMILL technology was developed by CEA, France, and then produced under license by TEMIC Matra MHS

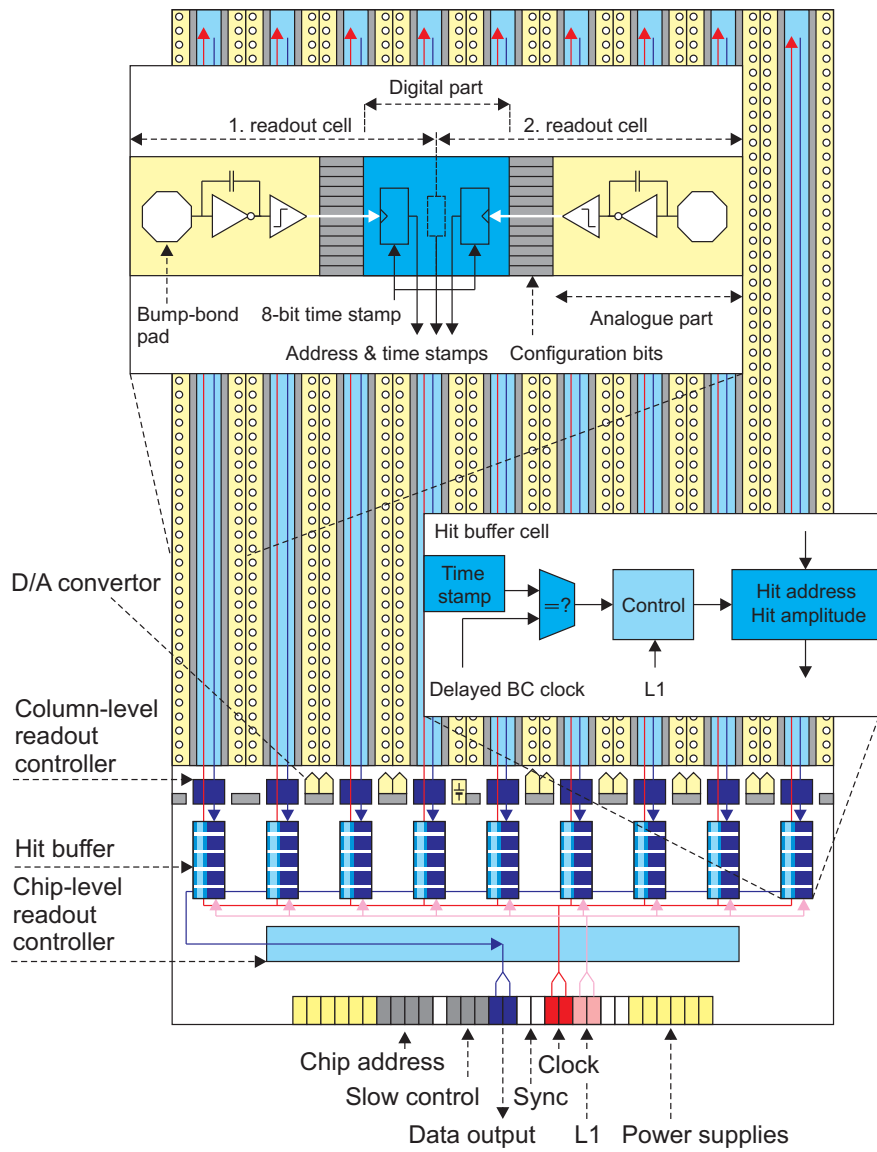


Figure 2: Floorplan of the front-end chip (FE-I3) with main functional elements.

Chip	Year	Cell size (μm^2)	Col \times Row	Transis- tors	Technology	References
Beer & Pastis	1996	50 \times 436	12 \times 63	–	AMS 0.8 μm BiCMOS, 2M	[3, 7]
M72b	1997	50 \times 536	12 \times 64	–	HP 0.8 μm CMOS, 2M	[3]
Marebo	1997	50 \times 397	12 \times 63	0.1 M	DMILL 0.8 μm BiCMOS, 2M	[5, 6]
FE-B	1998	50 \times 400	18 \times 160	0.8 M	HP 0.8 μm CMOS, 2M	[8–10]
FE-A/C	1998	50 \times 400	18 \times 160	0.8 M	AMS 0.8 μm BiCMOS, 2M	[7, 10]
FE-D1	1999	50 \times 400	18 \times 160	0.8 M	DMILL 0.8 μm BiCMOS, 2M	[10]
FE-D2	2000	50 \times 400	18 \times 160	0.8 M	DMILL 0.8 μm BiCMOS, 2M	–
FE-I1	2002	50 \times 400	18 \times 160	2.5 M	DSM 0.25 μm CMOS, 6M	[11]
FE-I2/I2.1	2003	50 \times 400	18 \times 160	3.5 M	DSM 0.25 μm CMOS, 6M	[12]
FE-I3	2003	50 \times 400	18 \times 160	3.5 M	DSM 0.25 μm CMOS, 6M	[13–16]

Table 1: Summary of the ATLAS pixel front-end chips

Charge Sensitive Preamplifier The charge sensitive amplifier uses a single-ended folded-cascode topology, which is a common choice for low-voltage and high gain amplifiers. The amplifier is optimised for a nominal capacitive load of 400 fF and designed for negative signal expected from n^+ -on- n -bulk detectors. The design of the charge amplifier was particularly influenced by requirement pertaining to sensors irradiation, which produced leakage currents up to 100 nA. The preamplifier has a 5 fF DC feedback, 15 ns risetime and operates at about 8 mA bias. Since the input is DC-coupled, a compensation circuit is implemented which drains the leakage current and prevents it from influencing the bias current of the fast feedback circuit, used to discharge the feedback capacitor. The feedback system, shown in Fig. 3 uses two PMOS devices, one (M2) providing leakage current compensation and the other (M1) continuously resetting the feedback capacitor.




Figure 3: Charge preamplifier: feedback circuit.

An important property of this feedback circuit is that the discharge current provided by the reset device saturates for high output signal amplitudes. The return to baseline is therefore nearly linear and a pulse width proportional to the input charge is obtained. The width of the discriminator output, Time-over-Threshold (ToT), can therefore be used to measure the signal amplitude. The duration of the ToT is measured by counting the cycles of the 40 MHz master chip clock. The feedback current is 4 nA for a 1 μs return to baseline with 20 ke input. The feedback circuit has an additional diode-connected

transistor M3, which acts as a level shifter so that the DC levels of input and output nodes are nearly equal. It also simplifies the DC coupling between amplifier and the comparator, described below.

Comparator Signal discrimination is made by a two stage-circuit: a fully differential low gain amplifier, where threshold control operates by modifying the input offset, and a DC-coupled differential comparator. The first stage has a bias of about $4\ \mu\text{A}$ whereas the second uses a current of about $5\ \mu\text{A}$. In order to make the threshold independent of the local digital supply voltage in each pixel and on the amplifier bias current I_f , a local threshold generator is integrated in every pixel. Seven-bits are used in each pixel to adjust the discriminator threshold.

Pixel Cell Control Logic A complete block diagram of the analogue part with several additional circuit blocks is shown in Fig. 4. Each pixel has several parameters that can be tuned through a 14-bit control

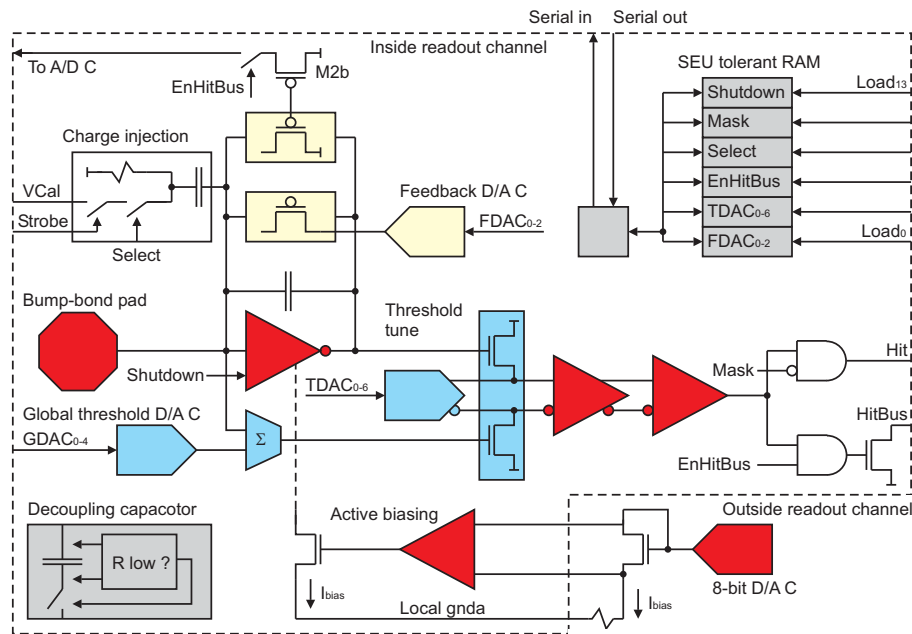


Figure 4: Pixel cell block diagram.

register. These bits are:

- *FDAC 0-2*: 3-bits to trim the feedback I_f current for tuning the ToT response.
- *TDAC 0-6*: 7-bits to trim the threshold in each pixel.
- *MASK*: the digital output of the analogue part can be switched off locally by setting this bit.
- *EnHitBus*: the digital outputs of all readout channels can be directly observed using a wired OR which is locally enabled with this bit. This bit also controls, through transistor M2b, the summing of a current proportional to the feedback plus leakage current in the preamplifier, allowing for monitoring of the feedback current and of the leakage current from the sensor.
- *Select*: enables the pixel for test charge injection. The amplitude is generated from V_{Cal} (Voltage proportional to the injected calibration charge) whereas the timing comes from an external *Strobe* signal.

- *Shutdown*: disables the charge amplifier so that no output is generated from the pixel.

Pixel Cell Readout Logic A block diagram of the column-pair readout is shown in Fig. 5. LE and

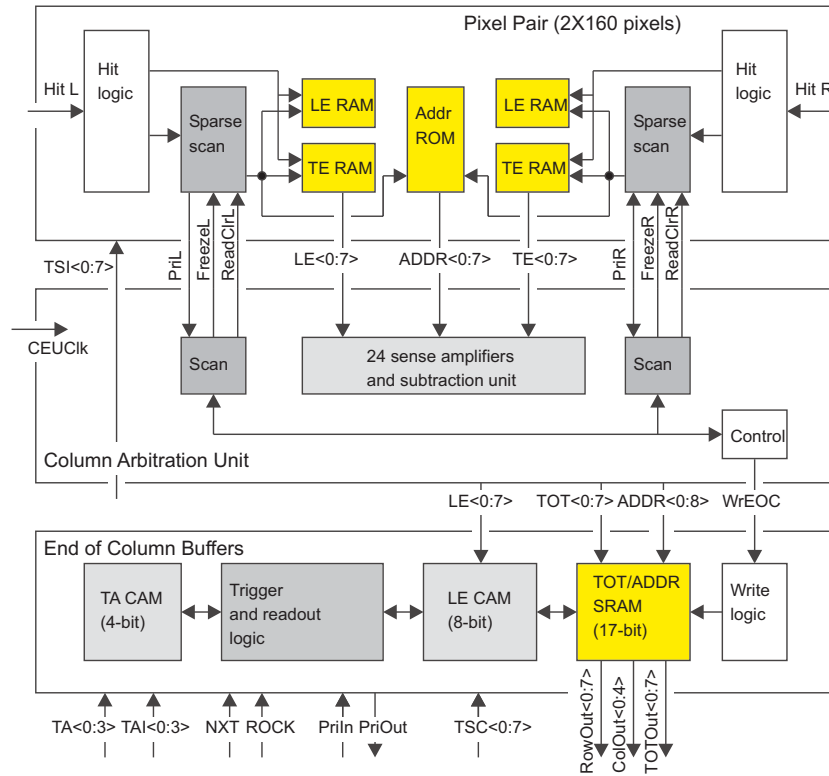


Figure 5: Block diagram of the column-pair readout.

TE timestamps are temporarily stored in local memories before being transferred to the hit buffers at the chip periphery. A digital circuitry generates two short ($1ns$) strobes at the LE and TE comparator edges, respectively. These signals are used to store the 8-bit Gray-coded time stamp into two memories. The time stamps, generated at the chip periphery, running at 40 MHz are distributed differentially, in order to decrease the digital crosstalk to analogue circuits and sensor electrodes. The complete hit information is available after the TE of the comparator signal and data transfer start. The time stamp of the LE (8-bits), of the TE (8-bits) and the row number (8-bits) are transferred to the end-of-column (EoC) buffers. Transfer happens by a priority mechanism that selects cells with data starting from the top row. The top most cell with a hit transfers its data to the bus and all the cells below it are inhibited. When the cell is read out, it releases the priority encoder bus and subsequent hits are selected and put on the readout bus. The readout speed is limited by the time the priority logic needs to ripple down. Hits can ripple through at a programmable speed that is obtained from the 40 MHz clock division. In the actual chip, the maximum speed to transfer a single hit to the EoC is 20 MHz.

Column Readout Controller The readout is column based, and two columns are readout from the same controller. The first task of the controller is the generation of the readout sequence to transfer hit information: LE and TE timestamp, plus pixel row address into an EoC buffer. This operation begins when data is complete, which is after discriminator TE is activated. The transfer of hits from a column pair is synchronized by the Controller end-of-column Unit (CEU), which operates at a speed of 5, 10,

or 20 MHz. A total of 64-hit buffers are available for each double-column. The second task involves digital processing of the hit data. Hit information is formatted by the CEU. Formatting includes the ToT calculation: subtraction of TE time stamp from LE timestamp. Optionally, a digital threshold may be applied to the ToT, and a timewalk (time slewing for small charges with respect to high charge) correction may be applied (write hit twice if below correction threshold, once with LE and once with $LE - 1$), or both. These operations are pipelined to minimize deadtime, but the EoC writes cannot occur faster than 20 MHz. Hit information is written to the EoC buffer, and waits there for a corresponding L1 trigger. If a trigger arrives at the correct time, checked using the LE time stamp for the hit, the data is flagged as belonging to a particular 4-bit trigger number. Otherwise, it is reset and the buffer is freed. Once the chip has received L1 triggers, the trigger FIFO will no longer be empty. This initiates a readout sequence in which the EoC buffers are scanned for the presence of hits belonging to a particular trigger number. If hits are found, they are sent to the output serializer block, which encodes and transmits them to the Module Controller Chip (MCC). After all hits for a given trigger number have been sent, an End-of-Event (EoE) word is appended to the data stream. All of these operations occur concurrently and without deadtime, with all column pairs operating independently and in parallel.

Event readout from the EoC buffers happens concurrently with the column readout. When the chip-level readout controller starts processing a particular L1 event, it first broadcasts the corresponding L1 readout address to all buffers. All cells with hits waiting for readout compare their stored L1 address with the request value. The readout of the selected L1 hits is controlled by a priority network, which sorts them in column and row order.

Chip Level Readout Controller The chip-level readout controller collects hit data from the EoC buffers and transmits the results of the chip serially. All hits belonging to the same L1 are grouped together into a single event, and events are transmitted out of the chip in consecutive trigger order. When a L1 trigger arrives, the current bunch crossing time and a buffer overflow bit are stored in a FIFO memory, which has a depth of 16 locations. This allows the chip to keep track of 16 pending L1 signals. The write pointer of the FIFO is used as the L1 identification, which is sent to the hit buffers. The readout sequence is started as soon as the FIFO receives an L1 trigger. If the L1 priority scan in the hit buffers flags cells with matching trigger numbers, the data of the first cell in the hierarchy is sent to a global data bus, where the information is copied to a shift register. The content of the shift register is then transmitted serially. This is repeated until the priority scan shows no more hits. An End-of-Event (EoE) data word, which includes error flags, is then added to the event.

Chip Configuration FE-I3 has 231 global configuration bits plus 14 local bits for each of the 2880 readout channels. The global bits are the settings for eleven 8-bit bias current DACs, for one 10-bit calibration voltage DAC, for the global threshold bits, for the L1 latency, for the ToT filter thresholds, column enable bits, as well as for others. The configuration is loaded into the chip using a serial protocol running at 5 MHz. This protocol uses three chip input pads: a data input, a clock and a load. Each write operation begins with a 4-bit address, which permits the 16 chips in a module to receive independent configurations. The address of each chip is encoded with wire-bonds during module assembly.

4.2.3 Requirements and Measured Performance

The design requirements for the pixel detector come from its operation at high radiation doses, from the time resolution of 25 ns to separate two contiguous bunch crossings, from noise, from the minimum operation threshold and dispersion and from the overall power budget. The calibration relies on a 7-bit adjustment of individual pixel thresholds (tuning). The un-tuned (tuned) threshold dispersion is $\sigma = 800$ (70) electrons equivalent input charge (e). The noise is $160e$ and the typical operating threshold

is $4000e$, which results in hits $> 5500e$ appearing in the correct 25 ns time bucket (described as in-time threshold) [15, 16]. Neither the dispersion nor the noise depends on the choice of threshold. The tuned thresholds have been observed to re-disperse with moderate radiation dose in prototypes, and it is expected that periodic threshold re-tuning will be needed. However, the actual dispersion rate in the real operating environment will need to be measured. A selectable option internally duplicates near-threshold hits in two adjacent time buckets in order to allow recovery of in-time threshold inefficiencies. Measurements made on a few modules irradiated to 60 Mrad (in excess of the full LHC-life dose) show a negligible tuned threshold dispersion and a 20% increase in the noise, despite the very high induced leakage current (60 nA for normal pixels). For a configured chip the typical digital current is 45 mA at 2 V and the analogue current is 75 mA at 1.6 V for a total power of 220 mW. Chip production was made in batches of 48 wafers. There are 288 chips on each 12-inch wafer. Six production batches were purchased with the 6 wafers from the engineering production run. The total number of wafers is 294. The average wafer probing yield was about 80%. The ATLAS pixel detector contains a total of 27904 front-end chips.

4.3 Module Control Chip (MCC)

4.3.1 MCC History

The prototype sequence leading up to the Module Controller Chip (MCC) is shown in Table 2. The very

Chip	Year	Std.Cells	Transistors	Chip size (mm^2)	Technology
MCC-AMS	Apr 1998	17922	363 000	10.3×6.3	AMS $0.8\mu m$ CMOS, 2M
MCC-D0	Aug 1999	–	–	6.1×3.5	DMILL $0.8\mu m$ BiCMOS 2M
MCC-D2	Aug 2000	13446	328 000	11.9×8.4	DMILL $0.8\mu m$ BiCMOS, 2M
MCC-I1	Nov 2001	33210	650 000	6.38×3.98	DSM $0.25\mu m$ CMOS, 5M
MCC-I2	Feb 2003	67919	880 000	6.84×5.14	DSM $0.25\mu m$ CMOS, 5M
MCC-I2.1	2003	67919	880 000	6.84×5.14	DSM $0.25\mu m$ CMOS, 5M

Table 2: Summary of the ATLAS pixel MCC chips.

first version of the chip, submitted in 1998, was fabricated in the rad-soft AMS using $0.8\mu m$ CMOS technology [17]. The chip had been extensively used when building rad-soft modules. The technology was chosen as it was very close to the $0.8\mu m$ BiCMOS DMILL technology which, at the time, was the chosen rad-hard technology for the ATLAS pixel detector.

A first prototype of the rad-hard chip (MCC-D0) was built in 1999. It contained only one Receiver, but all the remaining circuitry was implemented providing a good test of the DMILL technology. The final version of the chip (MCC-D2) was submitted in Aug 2000. The chip worked fine but had an unacceptably low yield, both in MCC-D2 and FE-D2. Consequently, this technology was ruled out.

At this point the MCC was ported to the DSM $0.25\mu m$ CMOS technology, and the MCC-I1 chip was submitted in Nov 2001. A new version of the chip, MCC-I2, was built in 2003 in order to provide better Single Event Upset (SEU) hardening to the chip. It turned out that the chip had a small bug that could be corrected by modifying only one metal line. Six additional wafers, containing the correction in the layout, were produced in 2003 leading to the final MCC-I2.1 chip.

4.3.2 Design

This section briefly describes the actual implementation of the production MCC chip, labeled MCC-I2.1. A simplified block diagram of the MCC internal architecture is shown in Fig. 6. The MCC has three main system tasks: (1) loading parameter and configuration data into the FEs and into the MCC itself,

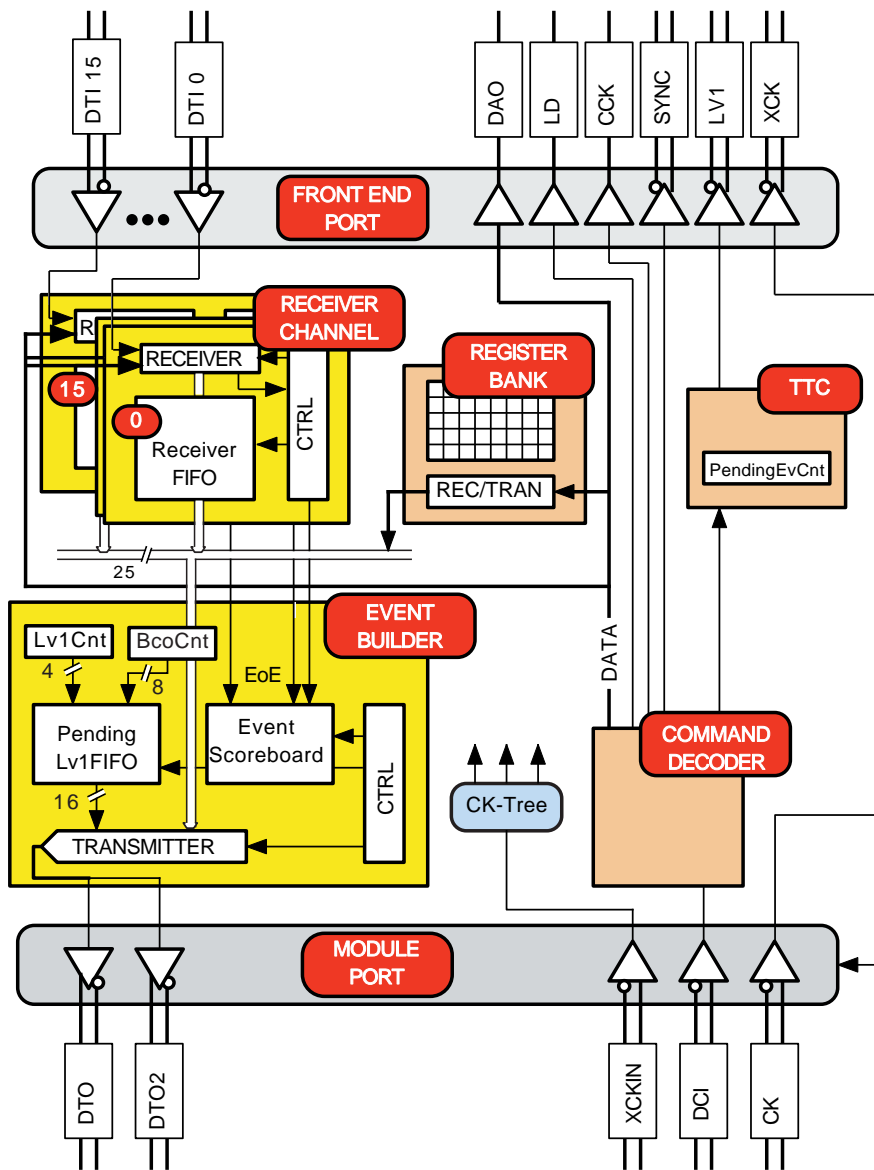


Figure 6: MCC block diagram.

(2) distributing timing signals such as bunch-crossing, L1 trigger and resets, and (3) reading out the FE chip and event building.

System Configuration The FE chips and the MCC must be configured after power-up or before starting a data-taking run. It is possible to write and read to all the MCC registers and FIFOs. This is used to configure, to read status information or to test the functionality of the chip. For this last function we provide a special set of commands that allows one to write simulated events into the FIFOs and two to run the Event Builder with the stored values in order to check the complete functionality of the chip. Once the MCC is embedded in the pixel detector, it will be important to test whether the event building works with known simulated events. Global FE chip registers and parameters in each of the pixel cells are written and read back through the MCC.

Trigger, Reset and Timing The second task of the MCC is the distribution of L1 triggers, resets and calibration/timing signals for the FE chips. In *Data Taking* mode, each time a L1 trigger command is received by the MCC the Trigger, Timing & Control (TTC) logic, it issues a Trigger to the FEs, as long as there are less than 16 events still to be processed. In case of an overflow the L1 trigger is not generated by the MCC and the corresponding event is lost. The information is sent to the ROD together with the number of missing events in order to keep up with event synchronization. In addition to the triggers, the TTC logic generates a hierarchy of reset signals that can be applied either to the MCC or to one or more FE's.

The last function of the TTC logic is the ability to issue calibration strobes to the FEs. This is used to calibrate the FE analogue cells on a pixel by pixel basis.

Event building The read-out that was chosen for the pixel detector is a data-push architecture with two levels of buffering: EoC buffers in the FE chips and 16 individual 128×27 -bit depth FIFOs (Receiver-FIFO) at the MCC inputs.

Event readout and building is by far the most complicated task and it occupies most of the chip area. Data received from the FEs, in response to a L1 trigger, are deserialized and buffered into 16 FIFOs, one FIFO for each receiving FE line. These FIFOs are used to derandomise the 16 data flows from the FEs and are used by the event builder to extract ordered hits and to prepare them for transmission out of the pixel module. Event building is performed by two concurrent processes running in the MCC. The first (Receiver) deals with the filling of the 16 input FIFOs with data received from the corresponding FE chip, while the second (Event Builder) extracts data from the FIFOs and builds up the event. Each FE sends data as soon as they are available with two constraints. Event hits must be ordered by event number and for each event an end-of-event (EoE) word is generated. The EoE is also sent for the case of an empty event to maintain the event synchronization.

The event transmitted to the ROD is organized by the Event Builder process on an event-by-event basis, instead of a hit-by-hit basis. If the FIFO becomes full while storing incoming hits, all subsequent hits are discarded and only the EoE word is written into the FIFO. In this case, a truncated event flag will be stored in the ReceiverFIFO and then recorded to the MCC output data stream. The mechanism insures that reconstructed events are not corrupted by FIFO overflows.

As soon as the Event Builder finds that an event is completely received from all of the 16 FEs, it starts building up and transmitting the event. The Event Builder learns from the Scoreboard when the events are complete. The first information written to the output data stream is the bunch crossing identifier (BCID) and the L1 identifier (L1ID). At this point the Event Builder starts fetching data from the ReceiverFIFOs, until it finds an EoE in the data. Once the Event is finished, a Trailer word is sent out to inform the ROD that the Event has ended.

I/O Protocols Several serial protocols were defined for communication to/from the ROD/MCC and the MCC/FE. All protocols that are active during data taking use only LVDS type signals, whereas signals used in the configuration time have single-ended CMOS to reduce the number of interconnection lines. Communications from the ROD to the MCC use a 40 Mb/s data line (Data Command Input - DCI) validated by the rising edge of the 40 MHz clock (CK).

The MCC to ROD link may use 40, 80 or 160 Mb/s data rates. For the case of 40 Mb/s, a new bit is transmitted at every rising edge of the CK. For the 80 Mb/s, bits are sent at both CK transitions. Finally for 160 Mb/s, both lines and clock edges are used. This can be considered as a 2-bit wide serial link. Only the event read-out uses the two higher bit rates. Read-out of configuration data is always at 40 Mb/s. Having data passing from the MCC to the ROD improved robustness by providing a bit-flip safe Header and by adding synchronization bits after known numbers of clock cycles.

Communications from the MCC to the FE chips are accomplished using a serial CMOS data bus (Data Address Output - DAO), a CMOS control line (Load - LD) and a 5 MHz validation CMOS clock (Control Clock - CCK). Both configuration and event data from the FE to the MCC are transmitted using 16 individual LVDS serial links (Data Input - DTI).

Special care is needed in the implementation of the Data-Taking protocols in order to minimize the effect of possible Single Event Upset (SEU) events. In particular, while in data taking mode, there are only two possible 5-bit commands: 'trigger' and 'exit data-taking modes'. All permutations of the trigger command, obtained by flipping one single bit, are also interpreted as a trigger command with the correct timing.

4.3.3 Performance

The design requirements needed to address the pixel detector operation at high radiation dose, the time resolution of 25 ns separating two contiguous bunch crossings, the expected bandwidths at the highest luminosity, the L1 trigger rate of 100 kHz and the number of FE chips that are controlled by a module.

The 16 FIFO's in the MCC were designed to handle the expected data rate of the FE chips operating at full luminosity with a L1 trigger rate of 100 kHz. In addition, the circuit were designed to be robust against a Single Event Upset (SEU), which can result from the harsh radiation environment that the pixel detector will have to sustain (50 Mrad in three years of operation for the innermost pixel layer). The problem was addressed using either triple redundancy majority logics or error detection and correction schemes. We irradiated several modules up to the full LHC-life dose, continuously reading out the data during irradiation. From SEU studies, we extrapolated to stable operation at the LHC without a significant loss in the configuration data coming from bit-flip in the memory elements.

For a configured chip, the typical digital current is 145 mA at 2.0 V and at a clock frequency of 40 MHz, the total power is 290 mW. All MCC-I2.1 chips were produced in a single batch of 6 wafers. The number of potentially good chips per wafer is 536. The measured yield was of 83%, providing a total of 2666 good chips. A total of 1744 chips are used in the ATLAS pixel detector.

4.4 Optical Communication

4.4.1 Optolink Architecture

The communication between the detector modules and the off-detector electronics is via optical links. The optolinks were selected to obtain electrical de-coupling and to minimize the material budget. The architecture was inherited from the ATLAS SCT [18]. Modifications were made to adapt to the data-rates, modularities and radiation hardness needs of the pixel detector.

A block diagram of the optical-link system architecture is shown in Fig. 7. The two main components in the optical link system are the opto-board, on the detector side, and the Back of Crate Card (BOC),

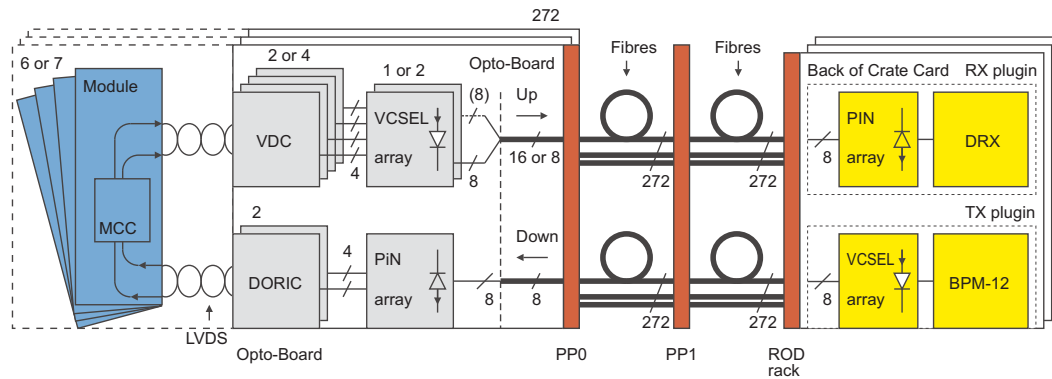


Figure 7: Optical link System Architecture

on the off-detector end. To keep the material budget low, to accommodate fiber routing requirements, to control radiation exposure, and to permit the use of optical arrays, the opto components and the related receiver/driver IC's were not implemented on the detector modules. The components have been put instead on the opto-boards at Patch Panel 0 (PP0 - see section 7), at a distance of about one meter from the modules and at relatively large radius in the the Pixel Support Tube (PST).

The transmission of the signals from the detector modules to the opto-boards uses LVDS electrical connections. These serial connections link the MCC with the VCSEL Driver Chip (VDC) and the Digital Optical Receiver IC (DORIC) sited on the opto-boards. DORIC and VDC designs were derived from the SCT project, but have been adapted to survive the higher radiation dose expected in the pixel detector. These chips have been fabricated on the same silicon wafers used to produce the MCC chips.

The communication with each detector module uses individual fibres: one for down-link and one or two for up-links. Trigger, clock, commands and configurations travel on the down-link, while event data and configuration read-back data travel on the up-link(s). On the down-link, a bi-phase mark (BPM) encoding is used to send a 40 Mbits/s control stream on the same channel as the 40 MHz Bunch Crossing (BC) clock. Decoding of the BPM channel within the DORIC recovers both the data stream and the clock signal. The use of individual links for every module permits the adjustment of the timing, used to associate the hit to the bunch crossing. The timing adjustment is accomplished by changing the delay of the transmitted signal with respect to the phase of the LHC machine reference clock received in the BOC.

The readout bandwidth required to extract the hits from the detector modules depends on the LHC instantaneous luminosity, on the L1 rate and on the distance between the module and the interaction point. Simulation of the readout architecture using generated physics events [17] shows that a bit rate of 40 Mb/s for the Layer-2 modules, 80 Mb/s for the Layer-1 or Disk modules and 160 Mb/s for the B-layer modules are needed to keep the number of lost hits due to bandwidth saturation sufficiently low. The data transmitted in up-links are encoded in non-return-to-zero (NRZ) format. Electrical to optical conversion happens in the opto-boards on the detector side and in the RX- and TX-plug-ins in the BOC.

There are two flavours of opto-boards: Disk/L1/L2-boards (D-board) with 8 down-link and 8 up-link channels and B-layer-boards (B-board) with 8 down-links and 16 up-links. The B-boards use two 80 Mb/s channels to obtain the aggregated bandwidth of 160 Mb/s. Because of the modularity of staves (13 modules) and of sectors (6 modules) D-boards use either 6 channels for the disk-sectors or 6/7 channels for the half-staves in Layer-1 and Layer-2.

In the off-detector part of the links, one BOC exists for each ROD. BOCs come with a variety of hardware options that are implemented by equipping the card with a larger or smaller number of optical

receiver (RX-plugin) or transmitter (TX-plugin) plug-ins. Each plug-in, in principle, can serve up to 8 modules, but, in practice, only 6 or 7 are used due to the modularity of the detector. Each BOC has space for 4 TX-plugin-ins and 4 RX-plugin-ins. BOCs come with 4 TX and 4 RX plug-ins, where the maximum bandwidth requirement is of 40 Mb/s, with 2 TX and 4 RX for 80 Mb/s and with 1 TX and 4 RX, for 160 Mb/s. Two custom chips have been designed by the SCT collaboration and used in the optical plug-ins: the DRX (12-channels Data Receiver ASIC in the RX) and the BPM-12 (12-channels Bi-Phase Mark encoder ASIC in the TX). In the BOC there is also the optical S-Link interface used to send the ROD output to the ATLAS Readout Buffer (ROB) units, which are the next level up in the event readout chain.

4.4.2 Opto-Board

The opto-board is the opto/electro-interface on the detector side. It consists of a beryllium-oxide (BeO) printed circuit board measuring $2 \times 6.5 \text{ cm}^2$. As discussed in Section 4.4.1 two flavour of opto-boards (D-boards and B-boards) exist and six or seven detector modules are connected to them. The D-boards are equipped with one PiN diode array and one VCSEL array (Vertical-Cavity Surface-Emitting Laser), while the B-boards have a second VCSEL array. Each opto-board loads two 4-channel DORIC chips, whereas two and four 4-channel VDC chips are loaded onto the D-board and B-boards, respectively. The optopack, which holds the PiN/VCSEL arrays and the connector to plug the optical fibres, is custom designed to fulfill requirements of low mass, non-magnetic and radiation tolerant. The total number of opto-boards in the detector is 288. This is more than the minimum 272 (44 B-boards and 228 D-boards) needed to read out the detector so that spares are available to recover from problems during integration. Ultimately, only one spare board was used. The remaining spares are mounted on the patch panels, but not connected.

PiN Diode Array Arrays of silicon PiN diodes are used to receive the data sent by the VCSELs. Epitaxial silicon PiN diodes are used because their intrinsic layer provides a thin active layer allowing for fast operation at low PiN bias voltage. The active area of each individual PiN diode is circular with a diameter of $130 \mu\text{m}$, and the depth of the intrinsic region is $35 \mu\text{m}$ [18].

DORIC The Digital Optical Receiver Integrated Circuit (DORIC) amplifies the optical signal detected in the PiN diode and extracts the clock and data from the BPM encoded signal. Data and clock are transmitted in LVDS format, to the MCC. Each DORIC chip contains four identical channels. The specification for the current from the PiN diode is in the range of $40 \mu\text{A}$ to 1 mA. The requirements for the clock are a duty cycle of $(50 \pm 4) \%$ and a time jitter better than 1 ns.

The DORIC has been designed to have a bit error rate of less than 10^{-11} after a radiation lifetime dose, for a PiN-diode bias current amplitude of $40 \mu\text{A}$. The PiN diode current amplifiers use a single ended scheme [19], avoiding the direct application of the diode bias voltage (10 V), which is much higher than the rating of the DSM technology. The DORIC must withstand up to 17 Mrad over the 10 years of ATLAS operation. It is, therefore, implemented in the same proven ($0.25 \mu\text{m}$) CMOS technology as used for FE, MCC and VDC.

VDC The VDC converts the LVDS input signal, received from the MCC, into a suitable single-ended signal to drive the VCSELs in a common cathode array configuration. The VDC chips come with four channels and drive one half of the VCSELs array. An external current used to drive the VCSEL operates up to 20 mA. The nominal current to operate the VCSEL is 10 mA. A standing (dim) current of $\sim 1 \text{ mA}$ is provided to improve switching speed in the VCSEL. The dim current is remotely controlled by an external voltage. The requirement for the rise/fall time (20 to 80 %) is in the range of 0.5 to 2.0 ns, where

1.0 ns is nominal. A further requirement, relevant to reducing the pick-up noise pertains to the constant power consumption of the VDC, that is independent of VCSEL being bright (on) or dim (off). A voltage (V_{Iset}), remotely controlled, determines the current I_{set} that sets the amplitude of the VCSEL current (bright minus dim current).

The VDC converts the LVDS input signal, received from the MCC, into a single-ended signal suitable to drive the VCSELs in a common cathode array configuration. Each VDC chip has four channels and drives one half of an VCSELs array. The requirement for the rise/fall time (20 to 80%) is in the range of 0.5 to 2.0 ns, where 1.0 ns is nominal. The VDC provides a nominal (maximum) current of 10 (20) mA to drive the VCSEL. A standing (dim) current of ~ 1 mA is provided to improve the switching speed of the VCSEL. The dim current is controlled by an external voltage; however, this feature is not implemented on the opto-board as the chip is designed to produce the nominal dim current. The VCSEL drive current provided by the chip is remotely controlled via a voltage (V_{Iset}) which determines the current I_{set} that sets the amplitude of the drive current (bright minus dim current). The chip is designed to have constant current consumption, independent from VCSEL being bright (on) or dim (off), to avoid generating ripple (noise) on the power supply which is being shared with the two DORICs on an opto-board.

VCSEL Array Vertical-Cavity Surface-Emitting Laser (VCSEL) arrays are used to transmit the data optically. The main advantages of VCSELs are that they provide large optical signals at very low currents and have fast rise and fall times. In order to maintain a low laser threshold current, VCSELs use ion-implants to selectively produce a buried current-blocking layer to funnel current through a small area of the active layer. The VCSELs [18] used in the pixel and SCT systems have an oxide implant to achieve the current confinement, which is becoming the standard VCSEL technology since it produces lower current thresholds at higher bandwidth. VCSELs are produced in arrays of 8 diodes. The typical fibre coupled power per channel is greater than 1 mW at a drive current of 10 mA. The optical power at 10 mA is sufficient to give a noise immunity of 6.2 dB. Using a slightly higher current is possible, if one adds another 1.8 dB of noise immunity [18]. The down-link, where the current is not a critical issue, can profit from this improved margin corresponding to an higher immunity to SEU and to a lower Bit Error Rate (BER). A PiN current amplitude of $100 \mu\text{A}$ ensures a BER less than 10^{-9} .

4.4.3 Back of Crate Card (BOC)

Each BOC [18] is connected to one ROD through the crate back-plane. The BOC has two functions: it interfaces between the ROD and opto-links, and it controls the distribution of the timing to the on- and off-detector electronics. Each BOC receives a system clock signal and redistributes it to the pixel detector modules and ROD. Each detector module needs a precise phase adjustment of its 40 MHz clock relative to the bunch-crossing time reference. The adjustment of this phase can be done for each module independently. This is made by an ASIC (PHOS4) that contains four delay channels, which are programmable over the range of 0 to 25 ns. The PHOS4 chip is also used to adjust the sampling clock that controls data received from the pixel modules. The opto-electrical conversion and the connection to the fibres is located in two plug-in cards: TX-plug-in and RX-plug-in, respectively, for transmission and reception of optical data. The TX-plug-in has an 8-channel VCSEL array and a BPM-12 ASIC. The RX-plug-in has an 8-channel PiN diode array and a DRX ASIC.

DRX The DRX ASIC amplifies, discriminates and converts the signal from the PiN diode into an LVDS signal. The comparator is DC coupled and the threshold can be controlled over a current range up to $255 \mu\text{A}$ by an external voltage reference generated by a DAC. The DRX chip was originally designed for the ATLAS SCT detector and contains 12 channels. Only 8 channels per chip are used in the pixel BOC.

BPM-12 The Bi-Phase Mark, BPM-12, ASIC encodes clock and data in the Bi-Phase Mark format for the fibre optic transmission. This chip was originally designed for the SCT detector and only 8 of the 12 channels are used for the pixels. A critical specification for this component is to have a short delay between the input signals and the encoded outputs to minimize the overall L1 trigger delay. The measured delay value is 60 ns.

4.4.4 Opto-fibres

The connection between the BOC and the opto-boards uses optical fibres. Two different kinds of fibres are used, Stepped Index Multi-Mode fibres (SIMM) and GRaded INdex multi-mode fibres (GRIN). SIMM fibres have been tested to be radiation tolerant but have lower bandwidth per unit length than GRIN fibres [20]. To optimise the bandwidth and radiation tolerance, splicing of 12 m long SIMM and 70 m long GRIN fibres have been used. The fibres are ribbonised into 8-way ribbons, and 8 ribbons are bundled together to form an optical cable. The 12 m length of the SIMM fibre is segmented by a MT16 connector at ~ 2.5 m from PP0 (at PP1). A total of 84 cables were installed.

4.4.5 Production and Testing of Opto-Link Components.

The opto-boards were required to pass a stringent quality assurance (QA) procedure, including burn-in and thermal cycling, at the two production sites. In addition, the boards were required to pass the reception test at CERN before installation on the service panels and subsequent tests were then performed. Each opto-board was required to produce good optical power, similar to those observed during the production QA, and over a reasonable DAC operating range in the DRX. Three major problems encountered during the test are discussed below [21].

- *Common Serial Resistance (CSR)*. During the reception test it was discovered that some of the VCSELs on the opto-boards produced very little or no optical power on all channels. Moreover the optical power on one channel was found to also depend on the current from other channels. This could be understood as the development of a common resistance in the array. The voltage drop on the CSR results in an inadequate voltage to drive the VCSELs. The only fault that could be identified in the production process of the opto-pack was that the thickness of the conductive epoxy under each VCSEL array was $\sim 5\mu\text{m}$, as opposed to $\sim 15\mu\text{m}$ as recommended by the manufacturer. A procedure was formulated to estimate the CSR by measuring the Current vs Voltage (IV) characteristics of one VCSEL channel with and without current in the other channels. Opto-boards with $> 2.25\Omega$ of CSR in the VCSELs were rejected, corresponding to $\sim 7\%$ of total production.
- *Slow Turn On (STO)*. The SCT group discovered that some of their VCSEL arrays took a few microseconds to produce full optical power after a few microseconds of inactivity. A random sample of opto-boards were tested for STO at the production sites and it was found that there was no indication of the problem until the test was performed on a prototype service panel with the production readout chain, including the fibers. It turned out that this subtle STO behaviour depended on the distance between the VCSEL surface and the fiber in polished mechanical transfer (MT) ferrule. The production fiber with the bevelled edges on the MT ferrule allowed the fiber to be pushed closer to the VCSEL, picking up transverse modes that might be time dependent. The exact cause of the slow turn-on behaviour is still under investigation. A new testing procedure was introduced, resulting in the rejection of $\sim 7\%$ of the opto-boards with severe STO VCSELs.
- *Fluctuations in the optical power (noise)*. It was discovered that the optical signal had more noise than was observed during the production testing, because of the long electrical cables. These cables

allow noise to enter into the VCSEL bias voltage via the VCSEL current control circuitry in the VDC. A bypass capacitor on the bias voltage was purposely not mounted due to the concern that the capacitor might leak after exposure to radiation, rendering the opto-board inoperable. There was no data to support the above concern but the decision was taken because the opto-boards on the production test system had low noise without the bypass capacitor. Fortunately, the capacitors could be readily retrofitted and this greatly reduced the fluctuation in the optical power.

In addition, another problem was discovered when the prototype patch panel was operated with the cooling system. The temperature of the opto-boards at a certain region on the patch panel was much lower than anticipated. We required the opto-packs on the opto-boards to produce good power at 10°C as part of the QA requirements. However, there were a significant fraction of the opto-packs that did not produce much optical power at -25°C, another temperature where we collected the data as part of the QA procedure. The optical power of these opto-packs were below the specification of 350μW on the prototype patch panel. To overcome the problem it was decided to add a heater to the opto-boards, giving the possibility to operate at 20°C. One VCSEL and one PIN channel failed during the detector integration. It is believed that the former is due to an ESD (electrical discharge damage) and the latter due to a detached solder (cold solder) on a lead of the opto-pack. In both cases the affected modules were recovering by switching to a spare opto-board in one case and by moving one module to an unused (7th channel of a board serving only 6 modules) channel, in the other.

Optical fibres, fabricated and assembled by external companies, have been tested during production by measuring light coupling and attenuation. Two 8-way ribbons in the 86 cables (8 ribbons each) showed failures. Fibres were also tested after installation using a Time Domain Reflectometer and then replaced by spares if they failed the test.

4.4.6 Optolink Performance

The selection and qualification of the components for use in the opto-link system was done by extensive laboratory tests and irradiation campaigns. From the measurements made on single components or parts of the system, one expects a stable performance for the opto-links over 10 year of operation at LHC [19, 20, 22–24]. The measurement of the BER in opto-link ring-loops running at 40 Mb/s (80 Mb/s) gives an upper limit $BER < 1.45 \times 10^{-14}$ ($BER < 3.62 \times 10^{-14}$). In fact no single errors were found during the tests [25]. The method to adjust the timing in the BOC to time the pixel detector with a bunch crossing is reported in [26]. Automatic tuning of the opto-link parameters for the entire detector (the system laser forward currents, PiN-diode photo-current thresholds, etc.) should be achievable in under 10 min.

4.5 Data Acquisition System

The pixel detector Data Acquisition System (DAQ) has been designed following the specification of the ATLAS global DAQ architecture [27].

4.5.1 Architecture overview

The off-detector readout architecture of ATLAS consists of two parts: a sub-detector specific part, where the Readout Buffers (ROD) are the main building blocks, and an ATLAS common design that is referred to as the Read Out System (ROS) [28].

The pixel ROD [29] is a 9U-VME module. The ROD handles the data transfer from the on-detector electronics on one side to the ROS system on the other side. Data from the detector arrive at the RODs through the BOCs. Data passes through the BOCs and is then received at the ROS by custom designed

interface modules. The ROS is a PC based system. These PCs temporarily store readout events into their memory and transfer only those accepted by the L2 trigger to the next level up in the readout chain.

ROD modules are plugged into ROD crates. There are 9 crates with up to 16 ROD modules per crate. In total, there are 44 modules or 3 crates for the B-layer, 38 modules for Layer-1 plus 28 modules for Layer-2 or 4 crates and 24 modules or 2 crates for the disks. A Trigger, Timing & Control Interface Module (TIM) [30–32] and a Single Board Computer (SBC) [33] complete the ROD crate. The TIM is the interface to the Trigger System. The DAQ software running in the SBCs controls the modules in the ROD crate. There is no event traffic on the ROD VME-bus during normal data-taking; data are routed directly from the ROD to the ROS PCs [28] via fast optical links (S-Links). The VME-bus is, however, heavily used during calibration of the pixel detector. RODs are controlled, via the VME bus, by the SBC which also acts as interface to the global DAQ system.

Calibration data are treated differently from collision data. The procedure to calibrate the pixel detector consists of a sequence of injections of a known charge into the pixel's front-end amplifiers. The response of each pixel is measured as a function of the injected charge and of other parameters (thresholds, preamplifier feedback currents, trigger delay) that can be varied during the calibration procedure. The typical result of a calibration scan consists of a set of occupancy histograms corresponding to different values of the scanned parameters. In order to achieve maximum precision, it is important to extract data from the FEs at the maximum speed supported by the detector links. This makes it difficult to extract calibration data using the normal data path, as the read-out chain from the ROS to the Event Builder is designed to transfer only L2-trigger accepted events, while the detector links are designed to support the full L1 trigger rate. For this reason, during calibration runs, the ROD decodes the data stream sent by the front-end electronics, fills occupancy histograms and stores them in memory. The histograms are then extracted via the VME bus by the SBC and sent to an analysis farm for further manipulation and archiving.

Single Board Computer (SBC) The Single Board Computer is a commercial VP-315³⁾ 6U VME card with a Pentium-M⁴⁾, having 1.6 GHz clock and 1 GB memory. The card uses a Universe II⁵⁾ PCI-VME bridge. It has three gigabit-ethernet interfaces, of which two are used, one to connect to the ATLAS control network and the second to the analysis farm, where histograms generated in the ROD are collected. Up to 40 MB of internal RAM memory is used to cache the configuration data needed in the pixel detector modules for a complete crate of RODs. The configuration data, cached in the SBC memory, are stored offline in an Oracle database⁶⁾ server. The memory is also used as a transfer buffer for the histograms moved from the RODs to the analysis farm.

Trigger, Timing & Control Module (TIM) The TIM is the interface between a ROD crate and the ATLAS trigger system. It receives a TTC fibre-link from a Local Trigger Processor (LTP), carrying LHC bunch crossing (BC) and orbit signals, the trigger signals like Level 1 Accept (L1A) and the trigger type, control/synchronisation signals such as the event counter reset (ECR) and the synchronisation (SYN). These signals are distributed to the RODs via a custom backplane installed in the lower part of the VME crate. On the same backplane the busy signals generated by the RODs pass to the TIM. The TIM makes a collective ROD card Busy and sends a signal to the LTP. The LTP on reception of the Busy signal stops the L1A to the detector electronics, thus allowing the front-end and ROD buffers to be emptied. Several features are implemented in the TIM to operate on the trigger signal. These include programmable delays on distributed triggers, generation of trigger bursts and strobe signals having a fixed delay from a L1A.

³⁾From Concurrent Technologies Corporation, <http://www.ctc.com>

⁴⁾From Intel Corporation, <http://www.intel.com>

⁵⁾From Tundra Semiconductor Corporation, <http://www.tundra.com>

⁶⁾From Oracle Corporation, <http://www.oracle.com/>

Moreover, the TIM can be used as a local trigger generator with a programmable rate. This has been very useful for studying ROD and DAQ behaviour for simulated event rates.

Read Out Driver (ROD) The structure of the ROD is outlined in Fig. 8. Three main sections of the design are the control path, data flow path, and the Digital Signal Processing(DSP) Farm. The control

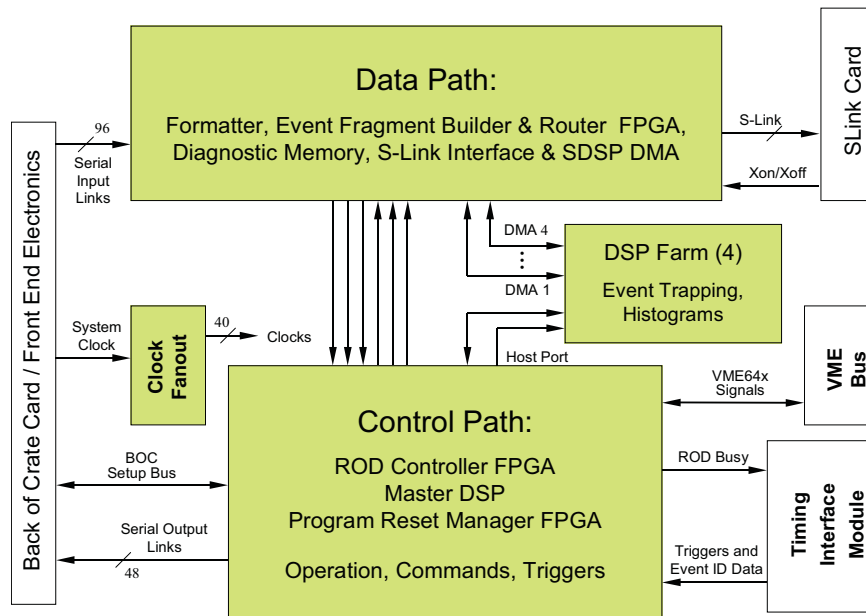


Figure 8: Block diagram of the ROD.

path section consists of two Xilinx Field Programmable Gate Arrays (FPGA) and a Texas Instruments Fixed Point Digital Processor (TI 320C6201 operating at 160 MHz with 32 MB SDRAM). The Program Reset Manager (PRM) FPGA functions as a VME slave controller, allowing read and write access to all ROD and BOC registers and a configuration controller for the data path FPGAs. To enable the users to easily upgrade the firmware on the ROD, the PRM FPGA allows VME access to an on-board Flash memory chip that stores the configuration data for all of the data flow path FPGAs. The Master DSP receives commands and transmits replies to the VME host and coordinates the configuration, calibration and data-taking modes of the ROD. The ROD Controller FPGA is used in the control path as an interface for the Master DSP to the DSP farm, the BOC, and all of the internal ROD registers in the data flow FPGAs. It also controls all of the required, data flow path specific, real-time functions on the ROD, including serial transmission of commands to the FE modules (two independent command streams can be sent to two modules or group of modules), calibration mode trigger generation, and transmission of TIM generated triggers and fast commands. In summary, these are the main actions performed by the control path block:

- full control of ROD reset and FPGA configuration;
- receives and executes command from the SBC via VME;
- receives module configurations via VME and stores them in Master DSP memory;
- transmits configuration data to the modules;
- control of calibration procedures, transmitting triggers and configuration data to FE modules;

- control of FE module data histograms;
- propagation of trigger commands from the TIM to the FE modules.

The structure of the ROD Data Flow Section is outlined in the block diagram of Fig. 9. The data flow

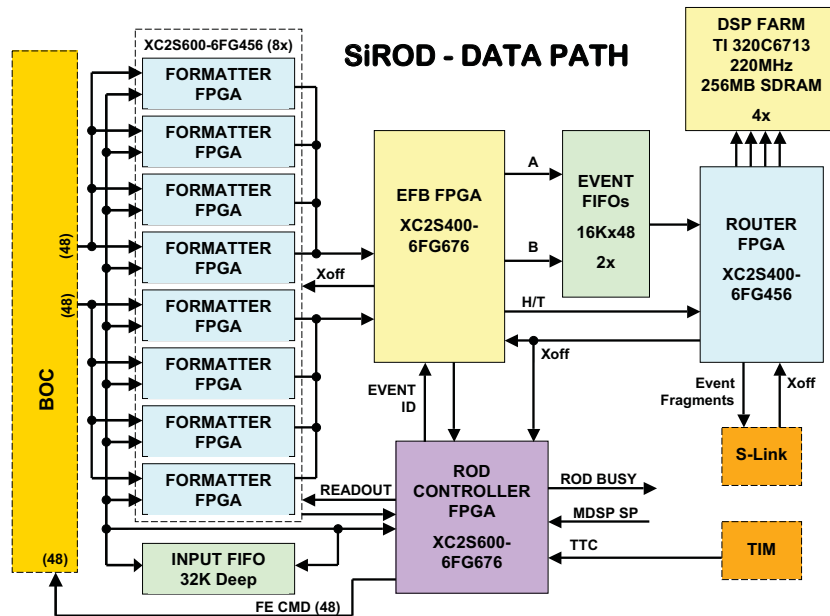


Figure 9: Block diagram of the ROD Data Flow Section.

section receives serial data from the FE modules, packs the individual module fragments into a single ROD fragment and sends it to the ROS via a custom designed optical link called the S-Link. Four blocks can be identified: the Input Link Interface, the Formatter, the Event Builder and the Router. Normal event data flows through the ROD via the Input Link Interface, which leaves the data unchanged. It can, however, trap the serial data stream in FIFOs (used in module configuration or to trap an event for diagnostics). The FIFOs can also be loaded with events for analysis by the ROD for diagnostics. After the Input Link Interface, the event data enters the Formatters. The Formatters convert the serial data streams to parallel format, and fill the de-randomizing buffers used to queue events for transmission to the Event Fragment Builder (EFB) FPGA. An event is transmitted from the Formatters to the Event Fragment Builder after the Controller FPGA sends a command notifying that a Level 1 Accept has been sent to the modules. The ROD Event Fragment is constructed in the EFB using the ATLAS Event ID data that was transmitted from the controller FPGA. In normal data taking, the primary source of the ATLAS Event ID data is the TIM with the ROD providing some additional information. After the header and mode information is sent to the EFB, the ROD Controller FPGA issues one token to the Formatters, and event data is pushed to the EFB. The EFB checks L1ID and BCID values and records errors. It also records any errors that were decoded or flagged by the Formatters. The event data is then stored in 16K de-randomizing FIFOs (two each). There are two identical engines in the EFB transferring data at up to 40 MHz (total bandwidth 80 MHz). When an event is ready (header, data body and trailer in the FIFOs), it is transmitted to the Router. The Router has two main functions. The first one, that is for the main physics data path, is to transmit 32 bit data words to the S-Link at 40 MHz. If the S-Link is receiving data faster than it can transfer to the ROS, the S-Link can assert Xoff to apply back pressure to the ROD data path. When back pressure is applied, read out of data from the EFB FIFO is stopped.

When the EFB memories are almost full, back pressure is applied to the Formatters. This will stop event data transmission from the formatter link FIFOs. The second function of the Router is to trap data for the DSPs. This is performed with no effect on the S-Link data during normal running. When the ROD is in calibration mode the DSPs can assert back pressure to pause the ROD data flow.

Finally, the ROD is equipped with four 'Slave' DSP processors (TMS320C671314) with 256 MB memory each. They are connected to the Router FPGA from which they can sample the produced ROD fragments. Different tasks can run on the DSP processors to analyze captured events: monitoring task, used during normal data taking to compute average occupancy and detect noisy pixels or data transmission errors. Calibration tasks accumulate histograms during the multi-dimensional scan procedure and perform an analysis to reduce the data volume to be transferred to the SBC. During data taking the DSPs spy on the data-flow at the maximum possible rate without introducing dead time or applying back pressure on the data flow path. During calibration, on the other hand, the slave DSPs analyze fragments, so they actually become the most important limiting factor on the data rate. For this reason the code of the calibration task must be optimised to maximum efficiency using the 128 KB internal DSP fast memory to fill the occupancy histograms.

4.5.2 ROD Crate Software and Calibration Analysis Farm

The ROD Crate software is the interface between the ATLAS Run Control and the pixel detector DAQ. Its structure is outlined in Fig. 10.

Figure 10: ROD software architecture.

For each ROD in the crate, a ROD Interface thread is created. This interface give access to the basic functionalities that an external applications can perform on the modules using the ROD. The implemented functions range from very basic commands (like ROD module reset and configuration) to complicated scan procedures. The ROD software interfaces are based on Remote Procedure Calls (RPC). They use a Common Object Request Broker Architecture (COBRA) layer called Interprocessor Communication (IPC) which is used in most ATLAS DAQ applications. These interfaces can be accessed either locally in the ROD, from another process running in the SBC, or from a process running on a remote CPU. Only one application at the time can be allowed to access a given ROD; for this reason, each SBC runs a Crate Broker. Each process accessing a ROD must first ask the Crate Broker, verify if the requested resource is free and allocate it. Only at this point is access to the ROD Interface granted. The last element of the ROD Crate software is the Run Controller. This process is a local receiver of the commands issued by the central ATLAS Run Control.

During normal data taking, the ROD Crate Run Control allocates all the ROD Interfaces and executes the transitions (INITIALIZE, CONFIGURE, START, STOP) as indicated by the global Run Control. During calibrations, the Run Control disconnects the RODs, which are then controlled by a Calibration Console, controlling the calibration procedure. The interface/broker mechanism gives the possibility to run a calibration or a debugging session on a ROD while the others are taking data. Occasionally the amount of data produced during a calibration may be too large to fit in the SBC memory. The histograms are then immediately moved (again using IPC) from the SBC to a remote analysis farm, which takes care of the final data analysis, of generating new configuration sets based on the tuning/calibration procedures, and of archiving the results. Consequently, the memory of the SBC is not saturated, and a new scanning procedure is immediately started, while the analysis farm is analyzing the previous data set.

4.6 Detector Control System (DCS), Power Supplies, and Interlock System

The operation of the pixel detector modules and of the on-detector opto-components requires a complex power supply [34, 35] and control system [36, 37]. The following supplies are required at module and opto-board level:

- V_{DDA} : analog low-voltage supply for the FE chips;
- V_{DD} : digital low-voltage supply for the FE chips and the MCC;
- V_{DET} : high-voltage supply to bias the sensor;
- V_{VDC} : low-voltage supply for the VDC and DORIC chips;
- V_{PIN} : PiN diode bias voltage;
- V_{ISET} : digital voltage to to adjust the VCSEL bias;

Electrical budget request for pixel modules and opto-boards is summarised in Table 3. The adjustment

Supply Type	Supply Voltage (V)	Supply Current (mA)	Nominal Voltage (V)	Nominal Current (mA)	Nominal Power (mW)	Worst Voltage (V)	Worst Current (mA)	Worst Power (mW)
Module								
V_{DDA}	10	1500	1.6	1100	1760	2.1	1300	2725
V_{DD}	10	1200	2.0	750	1500	2.5	1000	2500
V_{DET}	600	2	1.6	600	1	600	2	1200
					<i>Module total</i>			3860
D-board								
V_{VDC}	10	800	2.0	280	560	2.5	490	1225
V_{PIN}	20	20	5.0	–	–	20	20	400
V_{ISET}	5	20	1.0	–	–	2	20	60
					<i>D-board (Disk/Layer-1/Layer-2) total</i>		1685	
B-board								
V_{VDC}	10	800	2.0	420	784	2.5	770	1925
V_{PIN}	20	20	5.0	–	–	20	20	400
V_{ISET}	5	20	1.0	–	–	2	20	60
					<i>B-board (B-Layer) total</i>		2385	

Table 3: Specifications for module and opto-board power supplies.

of the operation conditions of the system requires a large modularity. Robust software packages are used to monitor and control the hardware. There is, in addition, an independent interlock system that focuses on safety for the equipment and human operators.

4.6.1 The Hardware of the DCS

The scheme of the powering, control and interlock system is shown in Fig. 11. The main components under the pixel DCS are:

- the power supplies to operate the sensors, the front end chips and the opto-boards;
- the Regulator Stations;

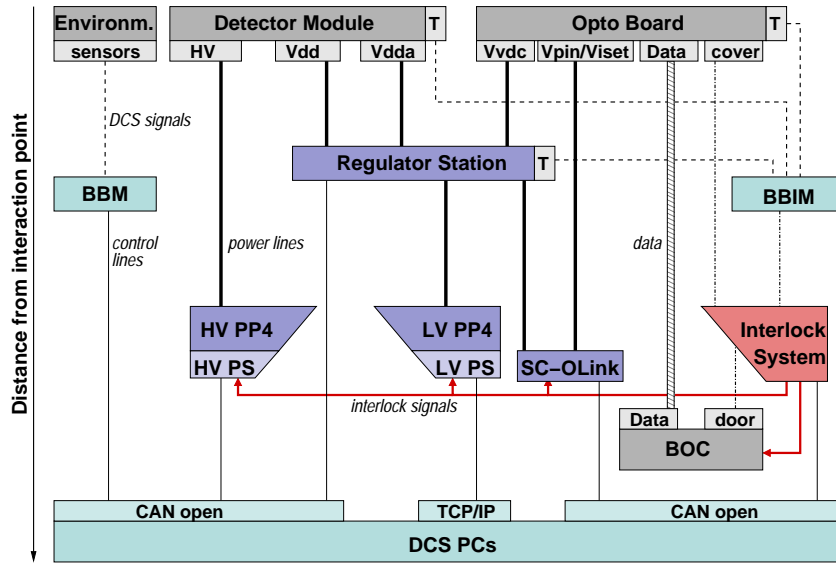


Figure 11: Overview on the hardware of the pixel detector control system.

- temperature and humidity sensors plus monitoring devices for their readout;
- multi channel current measurement units;
- the Interlock System;
- the DCS computers to control the hardware.

To comply with the ATLAS grounding scheme, all power supplies and monitoring systems must be floating. Radiation damage during lifetime operation of sensor and on-detector electronics requires that all power supplies have adjustable voltage outputs. For operational safety, over-current protection and interlock input signals are available for all the power supplies. The pixel power supply system has five main components: low voltage power supply (LV-PS), high voltage power supply (HV-PS), Regulator Station, Supply and Control for the Opto Link (SC-OLink), and the opto-board heater power supplies. Two low voltages supply the analog (V_{dda}) and digital part (V_{dd}) of the front-end read out electronics. Both are delivered by the LV-PS, which is a commercial component – the PL512M from WIENER⁷⁾. To protect the sensitive front end electronics against transients, remotely-programmable Regulator Stations are installed as close as possible to the detector. These Regulator Stations allow an individual tuning of the power lines through digital trimmers.

The pixel sensors are depleted by the high voltage V_{det} (up to 700 V) by the HV-PS. The HV-PS is assembled with EHQ-F607n_405-F modules provided by Iseg⁸⁾. The LV-PS and HV-PS are respectively connected to the low voltage patch-panel-4 (LV-PP4) and to the high voltage patch-panel-4 (HV-PP4) that are used to distribute the power and monitor the currents of the individual lines.

The SC-OLink, a complex channel consisting of three voltage sources and a control signal, delivers adequate levels for the operation of the on-detector part of the optical link. Monitoring of temperatures and of humidity is performed by the Building Block Monitoring (BBM) and the Building Block Interlock and Monitoring (BBIM) crates. While the first just provides a reading of values, the second additionally

⁷⁾WIENER, Plein & Baus GmbH, Burscheid, Germany

⁸⁾iseg Spezialelektronik GmbH, Rossendorf, Germany

creates logical signals, which are fed into the Interlock System. All components of the LV-PS, HV-PS, SC-OLink in addition to the BOC boards are connected to the hardware based Interlock System, that acts as a completely independent system. It consists of several units that guarantee safety for human operators as well as for detector parts. The Interlock System has high modularity; more than a thousand individual interlock signals are distributed. This high modularity has been chosen to maintain a low number of detector modules out-of-service, resulting from some component failure. The Regulator System and parts of the Interlock System, installed inside the ATLAS detector, had to pass requirements pertaining to radiation tolerance; this required extensive qualification for many of the components.

Besides the LV-PS and HV-PS all other components in the system are custom designs, adapted to the pixel specific needs and using the Embedded Local Monitoring Board (ELMB) [38] for the monitoring through the DCS. ELMB is the ATLAS standard front end I/O unit for the slow control signal. The Control Area Network (CAN) interface of the ELMB and its CAN-open protocol ensure that the communication is reliable and robust. Different Openness, Productivity, Collaboration (OPC) servers are used to integrate the hardware into the higher level of the software. All together 630 CAN nodes on 43 CAN busses and 48 TCP/IP nodes for the LV power supplies are used to build the pixel control network. In total, more than 44000 variables need to be monitored.

4.6.2 The Software of the DCS

The DCS software establishes the communication to the hardware, to support the operator with all required monitoring and control tools, and to provide automatic safety procedures as well as easier operation of the detector for non DCS experts. Permanent operation and reliable data taking is required. Additionally, detector operation requires good coordination between the DAQ and DCS actions. Data relevant to the offline analysis must be recorded and stored in the conditions database. The core of the DCS software is the Prozess- Visualisierungs- und SteuerungsSoftware (PVSS), ⁹⁾. These projects run as a distributed system on eight control stations. Since each part of a distributed system has its own control and data managers (processes inside PVSS), an independent development and operation of the different projects is possible. The core of the control software is the Front-end Integration Tools (FIT) which establish the communication with the hardware components. For each hardware component, like the HV-PS, the LV-PS and the different devices using the ELMBs, dedicated FIT exist. Each FIT consists of an integration and a control part. The integration part initialises each given hardware component and creates the data structures required to control it. The control part of the FIT supervises the operation of the same component. The FITs are mainly foreseen for the DCS expert who needs to check the correct behaviour of the hardware. For persons who run shifts in ATLAS, a detector oriented view of the hardware is provided by the System Integration Tool (SIT). The mapping of the read-out channels to the detector devices is done by the SIT. The SIT creates a virtual image of the detector inside the DCS. It combines all information which is relevant to the operation of the detector unit such as a half stave, a disk or even the full detector. Furthermore, the SIT is responsible for storing the data in the conditions database.

The software used to operate on the detector is the Finite State Machine (FSM). This software is developed in common for the four largest LHC experiments [39]. The FSM uses geographical organised data structures, created by the SIT, and provides the user with a set of commands, to act on small or large fractions of the detector simultaneously. The detector status is returned by the FSM. Furthermore, proper settings and special power-on sequences are performed automatically by the FSM. In addition, the FSM builds the link into the ATLAS wide control system. As part of the overall ATLAS control system, the pixel FSM will receive commands from the ATLAS FSM during normal data taking. The communication between DAQ and DCS is provided by DAQ-DCS Communication (DDC), which provides command

⁹⁾PVSS is made by ETM, Eisenstadt, Austria

transfer from the DAQ system to DCS, publishes DCS values to the DAQ and vice versa. For the tuning of the optical link, the DDC is critical.

The DCS hardware and software system has been fully exercised in various configurations during the prototyping and construction phase of the pixel detector, namely in test beams, with cosmic ray tests and during the integration of the pixel detector at CERN [40, 41].

References

- [1] ATLAS Collaboration. *ATLAS Inner Detector: Technical Design Report I*, volume TDR-4. CERN, Geneva, April 1997.
- [2] ATLAS Collaboration. *ATLAS Pixel Detector: Technical Design Report*, volume TDR-11 of *Technical Design Report ATLAS*. CERN, Geneva, May 1998.
- [3] Damon Fasching. The atlas pixel detector. *Nucl. Instr. and Meth. A*, 408(1):229–234, 1998.
- [4] CPPM. Unpublished.
- [5] L. Blanquart, A. Mekkaoui, V. Bonzom, and P. Delpierre. Pixel analog cells prototypes for atlas in dmill technology. *Nucl. Instr. and Meth. A*, 395(3):313–317, 1997.
- [6] L. Blanquart, D. Calvet, and P. Fischer. MAREBO, a full Radhard Pixel Detector prototype for ATLAS. In *4th Workshop on Electronics for LHC Experiments - LEB 1998*, volume CERN-LHCC-98-036, Rome, Italy, Sept. 21-25 1998. CERN.
- [7] L. Blanquart, V. Bonzom, G. Comes, P. Delpierre, P. Fischer, J. Hausmann, M. Keil, M. Lindner, S. Meuser, and N. Wermes. Pixel readout electronics for lhc and biomedical applications. *Nucl. Instr. and Meth. A*, 439(2-3):403–412, 2000.
- [8] K. Einsweiler, A. Joshi, S. Kleinfelder, L. Luo, R. Marchesini, O. Milgrome, and F. Pengg. Dead-time free pixel readout architecture for atlas front-end ic. *Nuclear Science, IEEE Transactions on*, 46(3):166–170, 1999.
- [9] K. Einsweiler, A. Joshi, R. Marchesini, F. Pengg, and G. Zizka. On the performance and limitations of a dual threshold discriminator pixel readout circuit for lhc. *Nuclear Science, IEEE Transactions on*, 46(4):792–799, 1999.
- [10] John Richardson. The atlas pixel front-end readout chips. *Nucl. Instr. and Meth. A*, 473(1-2):157–162, 2001.
- [11] L. Blanquart, J. Richardson, P. Denes, K. Einsweiler, E. Mandelli, G. Meddeler, P. Fischer, M. Ackers, G. Comes, and I. Peric. Analog front-end cell designed in a commercial 0.25 /spl mu/m process for the atlas pixel detector at lhc. *Nuclear Science, IEEE Transactions on*, 49(4):1778–1782, 2002.
- [12] L. Blanquart, J. Richardson, K. Einsweiler, P. Fischer, E. Mandelli, G. Meddeler, and I. Peric. Fe-12: a front-end readout chip designed in a commercial 0.25-/spl mu/m process for the atlas pixel detector at lhc. *Nuclear Science, IEEE Transactions on*, 51(4):1358–1364, 2004.
- [13] Ivan Peric. *Design and Realisation of Integrated Circuits for the Readout of Pixel Sensors in High-Energy Physics and Biomedical Imaging*. PhD thesis, Bonn University, Bonn University, 2004. BONN-IR-2004-13.

- [14] Ivan Peric, Laurent Blanquart, Giacomo Comes, Peter Denes, Kevin Einsweiler, Peter Fischer, Emanuele Mandelli, and Gerrit Meddeler. The fei3 readout chip for the atlas pixel detector. *Nucl. Instr. and Meth. A*, 565(1):178–187, 2006.
- [15] Jorn Grosse-Knetter. The atlas pixel detector. *Nucl. Instr. and Meth. A*, 568(1):252–257, 2006.
- [16] F. Hugging. The atlas pixel detector. *Nuclear Science, IEEE Transactions on*, 53(3):1732–1736, 2006.
- [17] R. Beccherle, G. Darbo, G. Gagliardi, C. Gemme, P. Morettini, P. Musico, B. Osculati, P. Oppizzi, F. Pratolongo, and E. Ruscino. Mcc: the module controller chip for the atlas pixel detector. *Nucl. Instr. and Meth. A*, 492(1-2):117–133, 2002.
- [18] M. L. Chu, S. C. Lee, D. S. Su, P. K. Teng, M. Goodrick, N. Kundu, T. Weidberg, M. French, C. P. Macwaters, and J. Matheson. The off-detector opto-electronics for the optical links of the atlas semiconductor tracker and pixel detector. *Nucl. Instr. and Meth. A*, 530(3):293–310, 2004.
- [19] K. E. Arms, K. K. Gan, P. Jackson, M. Johnson, H. Kagan, R. Kass, A. M. Rahimi, C. Rush, S. Smith, and R. Ter-Antonian. Atlas pixel opto-electronics. *Nucl. Instr. and Meth. A*, 554(1-3):458–468, 2005.
- [20] Ingrid-Maria Gregor. *Optolinks for The Atlas Pixel Detector*. PhD thesis, University of Wuppertal, March 2001. WUB DIS 2001-03.
- [21] W. Fernando. Overview and status of atlas pixel detector. In *8th International Conference on Large Scale Applications and Radiation Hardness of Semiconductor Detectors (RD07)*, Florence, Italy, June, 27-28 2007.
- [22] G. Mahout, M. Pearce, M-L. Andrieux, C-B. Arvidsson, D. G. Charlton, B. Dinkespiler, J. D. Dowell, L. Gallin-Martel, R. J. Homer, P. Jovanovic, I. R. Kenyon, G. Kuyt, J. Lundquist, I. Mandic, O. Martin, H. R. Shaylor, R. Stroynowski, J. Troska, R. L. Wastie, A. R. Weidberg, J. A. Wilson, and J. Ye. Irradiation studies of multimode optical fibres for use in atlas front-end links;. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 446(3):426–434, 2000.
- [23] K. K. Gan, W. Fernando, P. D. Jackson, M. Johnson, H. Kagan, A. Rahimi, R. Kass, S. Smith, P. Buchholz, and M. Holder. Optical link of the atlas pixel detector. *Nucl. Instr. and Meth. A*, 570(2):292–294, 2007.
- [24] S. Kirichu Nderitu. *Atlas Pixel Opto-Board Production and Optolink Simulation Studies*. PhD thesis, University of Wuppertal, February 2007. WUB-DIS 2007-03.
- [25] Tobias Flick, Karl-Heinz Becks, Peter Gerlach, Susanne Kersten, Peter Mattig, Simon Nderitu Kirichu, Kendall Reeves, Jennifer Richter, and Joachim Schultes. Optical readout in a multi-module system test for the atlas pixel detector. *Nucl. Instr. and Meth. A*, 565(1):85–89, 2006.
- [26] T Flick, P Gerlach, K Reeves, and P Mättig. Atlas pixel detector timing optimisation with the back of crate card of the optical pixel readout system. *Journal of Instrumentation*, 2(04):P04003, 2007.
- [27] ATLAS Collaboration. *ATLAS high-level trigger, data-acquisition and controls: Technical Design Report*, volume TDR-16 of *Technical Design Report ATLAS*. CERN, Geneva, 2003.

- [28] J. Vermeulen, M. Abolins, I. Alexandrov, A. Amorim, A. Dos Anjos, E. Badescu, N. Barros, H. P. Beck, R. Blair, D. Burckhart-Chromek, M. Caprini, M. Ciobotaru, A. Corso-Radu, R. Cranfield, G. Crone, J. Dawson, R. Dobinson, M. Dobson, G. Drake, Y. Ermoline, R. Ferrari, M. L. Ferrer, D. Francis, S. Gadomski, S. Gameiro, B. Gorini, B. Green, M. Gruwe, S. Haas, W. Haberichter, C. Haerberli, Y. Hasegawa, R. Hauser, C. Hinkelbein, R. Hughes-Jones, M. Joos, A. Kazarov, G. Kieft, D. Klose, S. Kolos, K. Korcyl, K. Kordas, V. Kotov, A. Kugel, A. Lankford, G. L. Miotto, M. J. LeVine, L. Mapelli, B. Martin, R. McLaren, C. Meirosu, M. Mineev, A. Misiejuk, G. Mornacchi, M. Muller, R. Murillo, Y. Nagasaka, J. Petersen, B. Pope, D. Prigent, Y. Ryabov, J. Schlereth, J. E. Sloper, I. Soloviev, R. Spiwoks, S. Stancu, J. Strong, L. Tremblet, G. Unel, W. Vandelli, P. Werner, F. Wickens, M. Wiesmann, M. Wu, and Y. Yasu. Atlas dataflow: the read-out subsystem, results from trigger and data-acquisition system testbed studies and from modeling. *Nuclear Science, IEEE Transactions on*, 53(3):912–917, 2006.
- [29] T. Vickey. A read-out driver for silicon detectors in atlas. In *12th Workshop on Electronics for LHC and Future Experiments - LECC 2006*, 2006.
- [30] J M Butterworth, D A Hayes, J B Lane, and M Postranecky. Timing, trigger and control interface module for atlas sct read out electronics. Technical Report ATL-INDET-99-018, CERN, Geneva, Oct 1999.
- [31] M. Postranecky, J. Butterworth, D. Hayes, J. Lane, and M. Warren. Tim (ttc interface module) for atlas sct & pixel read out electronics. In *7th Workshop on Electronics for LHC Experiments - LEB 2001*, volume CERN/2001/005, CERN/LHCC/2001/034, page 222, Stockholm, Sweden, Sept. 10-14 2001. CERN.
- [32] Matthew Warren, Jonathan Butterworth, John Lane, and Martin Postranecky. Ttc interface module for atlas read out electronics: final production version based on xilinx fpga devices. In *10th Workshop on Electronics for LHC and Future Experiments - LECC 2004*, page 320, Boston, MA, USA, Sept. 13-17 2004.
- [33] Markus Joos. The vmebus processor hardware and software infrastructure in atlas. In *11th Workshop on Electronics for LHC and Future Experiments - LECC 2005*, page 331, Heidelberg, Germany, Sept. 12-16 2005.
- [34] J. Schultes, A. Andreazza, K. H. Becks, M. Citterio, K. Einsweiler, S. Kersten, P. Kind, S. Latorre, P. Mattig, C. Meroni, and F. Sabatini. The power supply system for the atlas pixel detector. In *Nuclear Science Symposium Conference Record, (c) 2004 IEEE*, volume 3, pages 1711–1715 Vol. 3, 2004.
- [35] T. Henss, A. Andreani, J. Boek, G. Boyd, M. Citterio, K. Einsweiler, S. Kersten, P. Kind, K. Lantzsch, S. Latorre, P. Mättig, C. Meroni, F. Sabatini, and J. Schultes. The hardware of the atlas pixel detector control system. *Journal of Instrumentation*, 2(05):P05006–P05006, 2007.
- [36] M. Imhauser, K. H. Becks, S. Kersten, P. Kind, P. Mattig, and J. Schultes. The control system for the atlas pixel detector. *Nuclear Science, IEEE Transactions on*, 51(3):502–507, 2004.
- [37] C. Bayer, S. Berry, P. Bonneau, M. Bosteels, G. Hallewell, M. Imhäuser, S. Kersten, P. Kind, M. Pimenta dos Santos, and V. Vacek. Studies for a detector control system for the atlas pixel detector. In *7th Workshop on Electronics for LHC Experiments - LEB 2001*, volume CERN/2001/005 - CERN/LHCC/2001/034, page 396. CERN, Sept. 10-14 2001.

- [38] H Boterenbrood and B I Hallgren. The development of embedded local monitor board (elmb). In CERN, editor, *9th Workshop on Electronics for LHC Experiments - LECC 2003*, volume CERN-LHCC-2003-055, page p.331, Amsterdam, The Netherlands, Sep. 29 - Oct. 3 2003.
- [39] S. Schmeling. Common tools for large experiment controls-a common approach for deployment, maintenance, and support. *Nuclear Science, IEEE Transactions on*, 53:970–973, 2006.
- [40] Joachim Schultes, Karl-Heinz Becks, Tobias Flick, Tobias Henß, Martin Imhauser, Susanne Kersten, Peter Kind, Kerstin Lantzsch, Peter Mattig, Kendall Reeves, and Jens Weingarten. Validation studies of the atlas pixel detector control system. *Nucl. Instr. and Meth. A*, 565(1):90–96, 2006.
- [41] Martin Imhauser, Karl-Heinz Becks, Tobias Henß, Susanne Kersten, Peter Mattig, and Joachim Schultes. First experiences with the atlas pixel detector control system at the combined test beam 2004. *Nucl. Instr. and Meth. A*, 565(1):97–101, 2006.