System Electronics for the ATLAS Upgraded Strip Detector

Release 2.0 – 05 Feb 2013

T. Affolder[a], F. Anghinolfi[b], A. Clark[f], W. Dabrowski[c], J. Dewitt[d], S. Diez Cornell[k], N. Dressdant[e], V. Fadeyev[d], P. Farthouat[b], D. Ferrere[f], A. Greenall[a], A.A. Grillo[d], J. Kaplon[b], M. Key-Charriere[g], D. La Marra[f], E. Lipeles[e], D. Lynn[h], M. Newcomer[e], F. Pereira[b], P. Phillips[g], E. Spencer[d], K. Swientek[c], M. Warren[i], A. Weidberg[j]

[a] *University of Liverpool*
Liverpool L69 3BX, United Kingdom
[b] *CERN,*
1211 Genève 23, Switzerland
[c] *AGH University of science and technologie*
Krakow, Poland
[d] *University of California, Santa-Cruz,*
1156 High Street, Santa Cruz, Ca 95064, USA
[e] *University of Pennsylvannia*
Department of Physics & Astronomy, 209 South 33rd Street, Philadelphia, PA 19104-6396, USA
[f] *University of Geneva*
DPNC, 24 Quai Ernest-Ansermet, 1211 Genève 4, Switzerland
[g] *Rutherford Appleton Laboratory*
Harwell Oxford, Didcot, OX11 0QX, United Kingdom
[h] *Brookhaven National Laboratory*
P.O. Box 5000, Upton, NY 11973-5000, USA
[i] *University College of London*
Department of physics and astronomy, Gower Street, London, WC1E 6BT, United Kingdom
[j] *University of Oxford*
Physics department, Clarendon Laboratory, Parks Road, OX1 3PU, Oxford, United Kingdom
[k] *Lawrence Berkeley National Laboratory*
1 Cyclotron Road, Berkeley, CA 94720, USA

ATL-UPGRADE-PUB-2013-011
15 February 2013

E-mail: Alexander Grillo <agrillo@ucsc.edu>

ABSTRACT: This document describes the architecture of the readout electronics of the ATLAS upgraded silicon strips detector and the design of the front-end electronics system and ASICs.

KEYWORDS:

Release History

Rel. No.	Date	Pages	Description of Changes
1.0	30-Jun-12		Initial Release
2.0	05-Feb-13	5-7	Figure 1 and Table 1 changed to agree with LOI layout.
		9	Power numbers changed to agree with LOI layout.
		10-11	Clarified buffering scheme.
		12	Clarified Data Compression Logic hit reporting.
		13	Expanded description of data packets.
		25-26	Refined description of action at L0A signal.
		26	Clarified output bandwidth options for HCC.
		29	Corrected bandwidth requirement per stave-side.
		32-33	Added section 3.5 on Redundancy.
		36,44	Corrected numbers in Tables 18, 19 & 20.
		44	Corrected services line counts for LOI layout.

Contents

List of Figures

List of Tables

1. Introduction

The basic concept of the front-end system of the Silicon Strip Detector in the Atlas Detector upgraded for the HL-LHC is being elaborated and proposed. The readout electronics of this new detector is based on front-end chips (ABC130), Hybrid Controller chips (HCC) and End of Stave Controller chips (EOSC). This document defines the basic functionality of the front-end system and of the different ASICs.

For purposes of this document, a “link” refers to either an optical transmission line (i.e. a single fibre) or an electrical transmission line using LVDS complementary specifications (i.e. a single pair of electrical conductors). “Data” will refer to the information flow heading off detector, i.e. ABC130 to HCC to EOSC to off-detector. “Trigger, Timing and Control” (TTC) will refer to information flow heading to the detector even though some commands will include data such as register settings, i.e. to EOSC to HCC to ABC130. "DCS" refers to the Detector Slow Control, which includes the local monitoring, the interlock devices, and the related signal transmissions. A “module” refers to the readout unit associated with a single sensor wafer. A “hybrid” refers to an interconnect circuit that has multiple ABC130s. There are one or two hybrids per module.

2. Overview of the readout system

2.1 Detector Lay-out

This section gives a short overview of the organisation of the detector in order to determine the maximum size of modules and staves in terms of channels and later the maximum data rate needed for the readout and the number of TTC links needed.

Figure 1 shows the basic features of the LOI layout [1] used for the basis of the readout organization of the strip barrel and end-cap designs. The layout may evolve between now and the completion of the Technical Design Report if simulations point to possible performance improvements. The general aspects of the readout architecture described here will still be applicable as the layout evolves, however, some details (e.g. exact number of bits in certain fields) may require change. Such changes to layout may also affect the power and material budgets, which will need to be included in the analysis of the resulting changes in performance.

The barrel is organized into staves. While mechanically each stave typically runs from some +z coordinate to a symmetrical –z coordinate, for purposes of the readout architecture, we will use the term “stave” to refer to the half of that full stave since the readout will be collected at both extreme ends. Likewise, the present layout calls for double sided staves, so care must be taken when discussing attributes such as channel count, power, etc. to be clear about a single side or both sides of a stave. The strip modules can host up to 20 256-channel ABC130 chips per side and a stave can gather up to 13 modules of 10 ABC130 each (long strips) or 13 modules of 20 ABC130 each (short strips) per side. These numbers are per side and must be doubled for a double-sided stave.

The end-cap is organized into disks and each disk into wedges often referred to as petals. For purposes of the readout architecture, these petals are equivalent to the barrel staves and each petal will gather data from all its modules at its outer radius as the equivalent of end of stave.

In the following, the worst case in terms of number of channels will be considered (i.e. only double-sided staves will be considered).

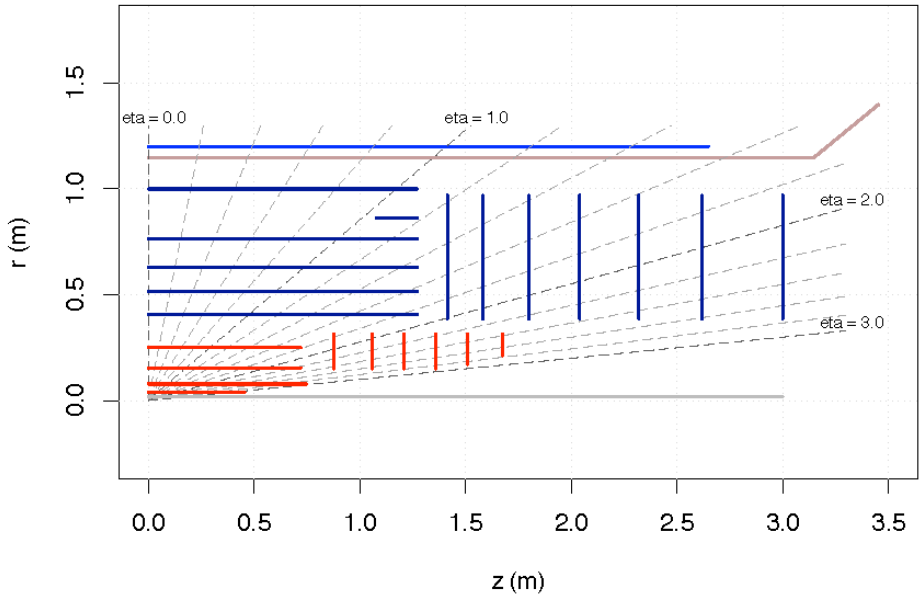

Figure 1: LOI Layout – one quadrant. There are 4 barrel layers and 6 disks of pixels, 3 barrel layers of short strips, 2 full barrel layers and one stub layer of long strips and 7 disks of strips per end.

2.2 Tracker parameters and environmental conditions

This section summarises the main parameters of the tracker from a system point of view (number of channels, total power, etc.) as well as the environmental conditions (radiation level, working temperature).

2.2.1 Number of channels for the strips detector

Assuming double-sided modules are used everywhere, and using the layout shown in Figure 1, the total number of modules of 256-channel ABC130s and the total number of channels are given in Table 1. These numbers are subject to change with the layout but the order of magnitude is certainly correct.

2.2.2 Working temperature

For a safe long-term operation of the tracker it is necessary that the sensors and electronics have good contact to the cooling pipe at -35°C. This requires a system for thermal monitoring with high granularity. Some of the signals in that system will be used for hardware interlock. Due to various heat sources within the thermal enclosure the average ambient temperature will be several degrees above this, but the exact amount is difficult to estimate now.

2.2.3 Radiation levels

The upgraded detector will run until a 3000 fb^{-1} integrated luminosity will be obtained. The detector and its electronics have to be designed for twice as much, i.e. 6000 fb^{-1}. The TID will be about a factor 10 higher than that of the current detector and is given in Table 2.

The amount of highly energetic neutrons, very damaging for the sensors, has to be reduced. This will be done by adding some moderating material that will have the effect of decreasing the neutron energy. Figure 2 gives the expected fluence in the tracker volume. At the level of

the pixel detector, more than 10^{15} n.cm^{-2} are expected while for the strip detector it is in the range $10^{14} – 10^{15}$ n.cm^{-2}.

Barrel	**Layer**	**Type of Strip**	**F Segmentaion per Stave**	**Number of Single Sided Modules in Z per Half Barrel**	**Number of 256-channel ABC130 per Module**	**Number of ABC130 per Half Stave**
	0	**Short Strips**	**28**	**2x13**	**20**	**2x260**
	1	**Short Strips**	**36**	**2x13**	**20**	**2x260**
	2	**Short Strips**	**44**	**2x13**	**20**	**2x260**
	3	**Long Strips**	**56**	**2x13**	**10**	**2x130**
	Stub	**Long Strips**	**64**	**2x2**	**10**	**2x20**
	4	**Long Strips**	**72**	**2x13**	**10**	**2x130**
	Total number of staves for the Barrel					**300**
	Total number of ABC130 ICs for the Entire Barrel					**184'000**
End-cap	**Disk**	**F Segmentation per Petal**	**Ring**	**Hybrid**	**Number of 256-channel ABC130 per Hybrid**	**Number of ABC130 per Petal-Ring**
	0-6	**32**	**0**	**H0**	**8**	
	All Disks Identical			**H1**	**9**	
				H2	**9**	
				H3	**10**	**36**
			1	**H4**	**11**	
				H5	**11**	**22**
			2	**H6**	**12**	**12**
			3	**H7**	**7**	**14**
			4	**H8**	**8**	**16**
			5	**H9**	**9**	**18**
	Total Number of ABC130 ICs per Petal					**236**
	Total Number of ABC130 ICs for One Disk					**7'552**
	Total Number of ABC130 ICs for One End-cap					**52'864**
	Total Number of ABC130 ICs for Entire End-cap					**105'728**
Total Number of ABC130 ICs for the Barrel and End-caps						**289'728**
Total Number of Channels of Readout						**74'170'368**

Table 1: Number of modules and channels for the strip detector assuming double-sided modules everywhere and the layout of Figure 1 but subject to change for other layouts.

Radius in cm	Dose in kGy
5.05	4500
12.25	1000
29.9	300
43.9	200
51.4	100
108	40

Table 2: Total ionising dose in kGy for 3000 fb-1 integrated luminosity at various radii.

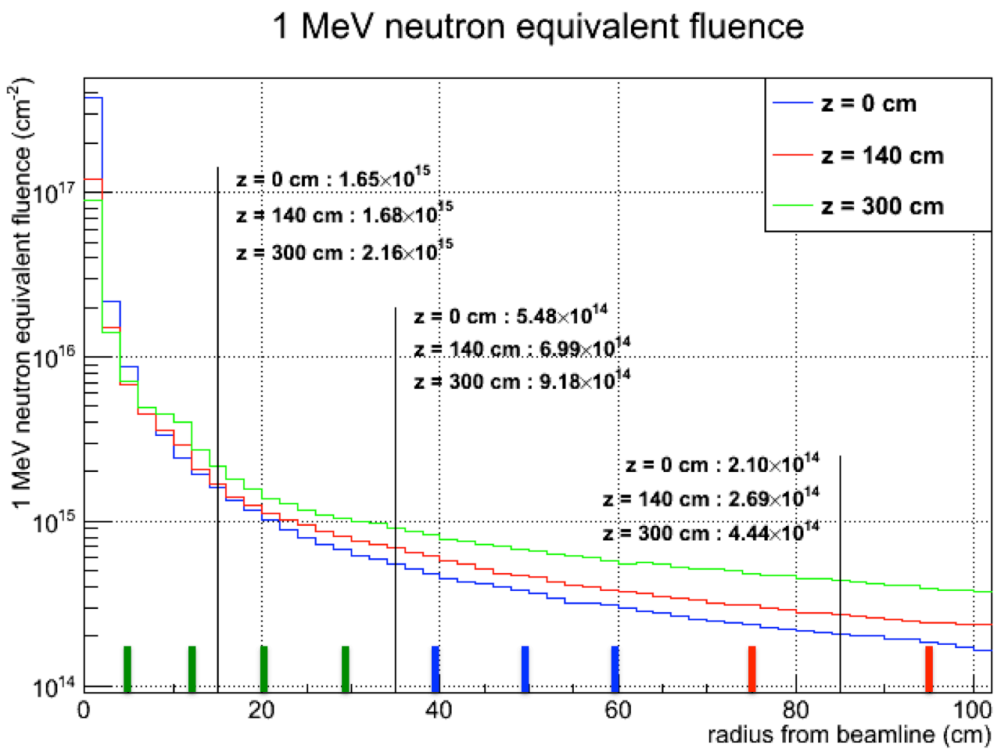

Figure 2: 1 MeV neutron equivalent fluence for 3000 fb-1 integrated luminosity. Vertical bars show the positions of the pixel layers (green), short strips (blue) and long strips (red) for the original straw-man layout. [1]

2.3 Readout units for the strip detector

The detector is organised in staves with up to 26 modules counting both sides (a module being a 10x10 cm^2 sensor wafer and its front-end electronics). Each stave is made of two layers in order to have a double-sided detector. These two layers will be treated separately from the readout point of view. It is assumed that the data of a single-sided stave will be collected by the end of stave (or service) controllers (EOSC) located at each end of the staves; each EOSC will readout up to 13 modules. Each module on the stave has either one (long strips) or two (short strips) readout hybrids. In both cases, the readout unit is the readout hybrid. The readout hybrid will host 10 ABC130 for both short strips and long strips and a so-called hybrid controller chip (HCC) that will gather the data of the hybrid and transmit them to the EOSC. The short-strip hybrid will readout two rows of 2.5-cm long strips and the long-strip hybrid will readout two rows of 5-cm long strips. In both cases the two rows of strips are interleave connected into the input channels of the ABC130 chips. It could in addition contain specific power components (Serial Power Protection or DC-DC converters).

Figure 3 is a sketch of a single-sided stave of short strips.

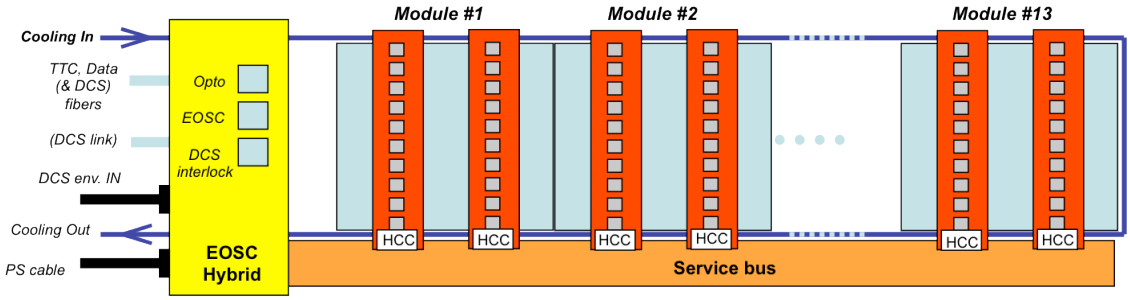

Figure 3: Sketch of a single-sided stave of short strips. Up to 13 modules are readout by 26 hybrids of ten 256-channel ABC130s each. An end of stave controller gathers the data from all module hybrids and interfaces to the off-detector electronics. Each hybrid houses a hybrid controller that interfaces the ABC130s to the EOSC. The DCS data can either be sent together with the readout data or be sent on a dedicated link.

2.4 Rough estimate of the total power in the strip detector

Recent detailed estimations have shown that the maximum power consumption per channel should be of the order of 0.7 mW. Assuming a pessimistic 1 mW per channel for the strip ABC130 and a 1.2 V working Vdd, one needs to feed about 210 mA per 256-channel ABC130. Hence, one has to feed a total current (for the barrel and both end-caps) of order 61 kA to the front-end. The power dissipated in the front-end will depend on the powering scheme but will in any case include some inefficiency because of the embedded regulators or converters. 80% efficiency would lead to about 93 kW dissipated in the tracker volume, while 70% efficiency would lead to about 106 kW. These numbers show the extreme attention one has to pay to the design of the different power components so that their efficiency is maintained as high as possible.

The current SCT and TRT detectors are fed with about 12 kA through the available cables. Assuming that the amount of services cannot be increased (not because it is planned to reuse existing service but because the volume available will not dramatically change), the powering scheme to be used must limit the amount of current to be fed at that level. That is about $1/5^{th}$ of the current needed by the front-end electronics. Either at least a factor of 5 reduction by DC-DC conversion or a serial powering scheme of at least 5 modules could be used.

2.5 Interface to the trigger DAQ

The requirement document [2] remains valid although some parameters could change because of the difficulties that could be encountered by the level-1 trigger system to maintain a maximum 200 kHz L1A rate.

2.5.1 Level-1 Accept rate

This document assumes that the level-1 (L1A) rate will not exceed 200 kHz. This is doubled from the original 100 kHz specification when the Upgrade R&D first started. Simulations run so far indicate that the architecture described in this document will handle the resulting data rates. Some discussion includes descriptions of alternate solutions should more bandwidth be required. For example, see section 3.4 covering end of stave electronics. These solutions require either doubling the number of links or exploiting more advanced IC technologies. The former would represent increases in material and power; the latter is certainly possible as our R&D work continues.

2.5.2 Possible tracking as part of the Level-1 trigger

In case that the calorimeter and muon level-1 trigger electronics cannot efficiently select events and that some information from the tracker is needed to either improve the cuts on muon momentum or to improve the electron isolation process, two main roads are possible:

1. Add specific trigger layers to produce a self-seeded track trigger;
2. Use the higher rate Level-0 trigger (L0) coupled with region of interest information to extract a limited amount of track data in time to be used by the Level-1 trigger logic thus reducing its rate to a manageable level. This scheme is called Regional Readout Request (R3).

The second solution is certainly the less difficult one to be implemented, as it only requires some extra buffer management in the front-end electronics and some reasonable increase of bandwidth for the data readout.

Although at the time this document is written it is not clear whether or not such a "track trigger" feature is necessary, it is deemed safer to specify and design the system assuming the R3 scheme is used. In addition some extra circuitry will be provided in the first ABC130 iteration to allow possible testing of the self-seeded track trigger. See section 3.1.3.

2.5.3 Regional Readout Request (R3) scheme

It is assumed that a very first level trigger called L0 is generated at a maximum rate of 500 kHz together with region of interest (RoI) addresses. It is assumed that there are up to 4 RoI's per L0 and that the total of the ROI's for one L0 covers about 10% of the detector. The L0 signal must be received within a fixed, but programmable to a maximum 6.4 μs, latency (currently 6.4 μs) after the event of interest such that the data for the event is at the end of the pipeline (named L0_Buffer). Upon receipt of the L0 signal, event data along with the associated BCID are transferred into a Random Access Memory, the second stage buffer named the R3L1_Buffer, at an address specified by a local 8 bit L0 counter (L0ID). A matching counter is maintained by the central trigger system, the value of which is sent as part of RoI and L1A readout requests to select the event address in the R3L1_Buffer. These memory access reads are non destructive so that data may be requested multiple times. Occasionally an RoI will overlap with a L1A selection. Those ABC130s that are part of a RoI will send the associated data to the hybrid controller (which will send them to the end of stave controller, etc.) so that the off-detector triggering process can continue. After some time (undefined yet but less than 256μs) the level-1 trigger electronics will issue a decision. There is no need for having a fixed latency between an event occurrence and the availability of RoI or L1A signals. The L1A rate is expected to be lower than 200kHz. The L1A signal together with the associated L0ID is sent to the front-end. The ABC130s will extract the corresponding data and send them to the readout system. The data in the second stage buffer, which can hold 256 events, will be overwritten when the buffer is full, the oldest L0ID first. For a L0 trigger rate of 500 kHz, the buffer will hold data for up to roughly 500 μs, but the exact maximum will depend upon the L0 rate. It is also assumed that the ROI and L1A signals will arrive with time ordered L0IDs. Therefore, the ABC130s will also delete from their second stage buffer the data corresponding to preceding L0IDs. Figure 4 gives a rough sketch of the R3 Scheme and it is discussed in more detail in section 3.1.

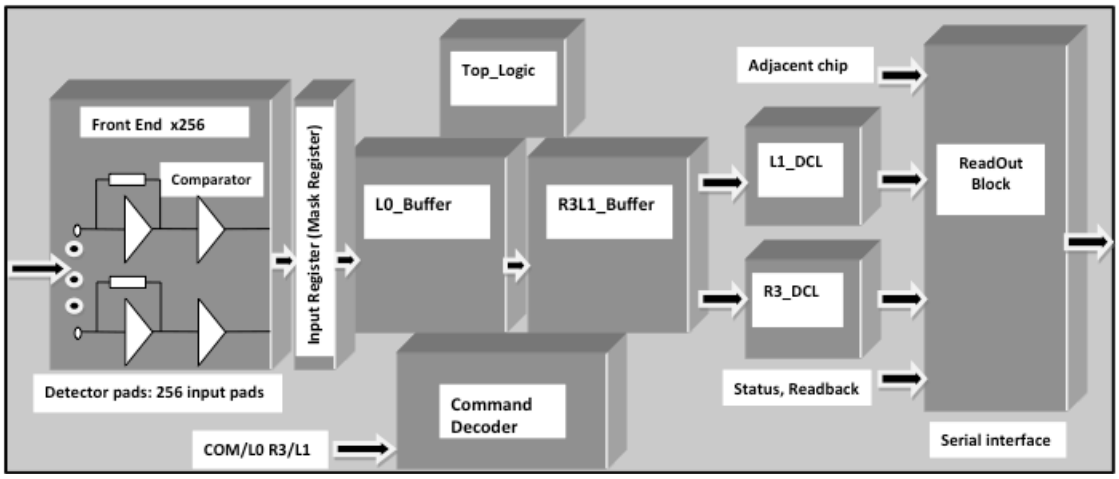

Figure 4: Block diagram of the ABC130 readout part with the R3 capability. All the ABC130s receive L0A and extract the corresponding data from the pipeline (L0_Buffer). These data are stored together with a L0ID in the R3L1_Buffer. Only those in the region of interest receive R3 signal and send immediately the corresponding data. L1A comes later together with the L0ID to which it refers. This L0ID is used to select data in the R3L1_Buffer and send them. The TTC system requires higher bandwidth for the R3 Scheme as some additional information is to be sent together with L0A and L1A.

2.6 Policy for fighting SEU

The amount of energetic hadrons (more than 20 MeV) susceptible of generating SEU is very high and hence SEUs will appear everywhere. The following policy will be adopted:

- The static registers holding thresholds, masks, etc. will be implemented with a triple redundant logic;
- The "physics data" themselves will not be protected, as a SEU acts as a small excess of noise and because the data do not stay for a long time in the ABC130s;
- Level-1 identifier (L1ID) and Bunch Crossing identifier (BCID) will not be protected as they stay a very short time in the front-end. In addition an error is very easily detected in the off-detector electronics and the policy of periodic resets will be maintained (There is a Bunch counter reset every 90 μs and an Event counter reset at a relatively high frequency [order of Hz] which can be used to reconfigure the front-end.);
- Special care will have to be taken for the transmission to the front-end of the trigger and control (TTC), as the receiving PIN diodes are very sensitive to SEU. Specific measurements have shown that SEU can generate multi-bit errors in high-speed links [3]. Adapted error detection and correction must be implemented.

3. Data readout chain

As shown in Figure 3, the readout chain consists of readout hybrids housing the ABC130 and the HCC ASICs, a service bus to the end of stave control board, the latter housing an EOSC and a DCS ASIC as well as the opto-electrical devices interfacing to two optical fibres. In addition, in case a self-seeding track trigger is to be implemented, a concentrator ASIC would be needed. The following sections will describe the different ASICs and their interfaces.

3.1 Front-end ASIC: ABC130

This chip provides all functions required for processing the signals from 256 strips of a silicon strip detector employing the binary readout architecture. The simplified block diagram of the hookup of the chip is shown in Figure 4. The main functional blocks are: front-end, input register, pipeline, derandomizing buffer, data compression logic blocks, command decoder, readout logic, threshold & calibration control, power regulation.

The architecture chosen for the ABC130 allows a multi-trigger data flow control retaining the Beam Crossing synchronous pipeline transfer signal (L0 here) from previous ABC versions and a new asynchronous Regional Readout Request (R3 here) and a second level asynchronous data readout intended for a global readout (L1 here).

The ABC130 contains 256 analogue preamplifier-shapers followed by discriminators with individual threshold trimming capabilities. The shaping time should allow 0.75 fC pulse detection separated by 75 ns. The binary outputs of the discriminators (data) are sampled at the bunch crossing clocking rate (BC) and stored for 6.4 μs in the "pipeline" or L0_Buffer memory bank. At the reception of a L0 signal, the data in the memory that were stored at some fixed latency time before the L0 signal are extracted from the pipeline and transferred to the derandomizing buffer (R3L1_Buffer) and stored. The data corresponding to 3 consecutive time slots are transferred for each L0 and form an event. There is enough room in the R3L1_Buffer for 256 events. After 256 events are stored, new events overwrite the R3L1_Buffer content.

At reception of a "R3" signal (R3s_L1 output from the HCC), which carries a L0 identifier, the event that has the same L0 number is extracted from the R3L1_Buffer and processed through the R3 Data Compression Logic block that performs zero suppression and cluster identification. The same happens at the reception of a "L1" signal (R3s_L1 output from the HCC), but the event with the correct L0 number is processed through the L1 Data Compression Logic block.

The R3 and L1 Data Compression Logic (DCL) differ in their algorithms to detect clusters and perform cluster identification, however, both algorithms will apply the same hit criteria as the present ATLAS-SCT chip, the ABCD. That is, a hit channel can be defined by the time sequence X1X, 01X, or XXX where 0 is no-hit, 1 is hit and X is don't care and the three bits represent the discriminator result for three consecutive beam crossings centered on the trigger BC. Which criteria to use is set by the configuration register. The information extracted from the 2 DCLs is different, for example the R3 unit is limited to identify 4 clusters at maximum and produces one data packet. The L1 unit is not limited in the number of clusters and number of packets produced.

The information extracted from the DCL contains a channel number to identify each cluster and bits that represent the cluster shape. The data are stored in local FIFOs to be transferred at the proper time to the Readout block.

The Readout block contains a packet builder and a fast serializer. The packet builder gets the data from either:

- The adjacent chip (external)
- The R3_DCL (internal)
- The L1_DCL (internal)
- The internal registers (internal)

The priority of data is set according to this order. The data from an adjacent chip is already formatted and is simply transferred to the serializer. The data from the R3_DCL, L1_DCL or from internal registers are formatted in packets of 60 bits as shown in Table 3. As

soon as there is some formatted data packet ready, it is sent out through the fast serializer running at the RCLK clock rate (nominal 160 MHz).

There are four types of data packets, each having the same header format but slightly different payload formats as shown in Table 5 through Table 8. The R3 packet can contain up to four clusters and only reports the address of the first channel of each cluster. Only the single beam crossing of the L0 trigger is considered for this cluster finding. This abbreviated data with only one packet per R3 and at most four clusters is used to minimize transmission time for data to be included in the L1 trigger decision.

The 1BC-L1 packet can contain up to three clusters with the address of the first hit channel in the cluster and a bit pattern showing hit or no-hit in the three adjacent channels. Clusters of more than 4 channels are reported as multiple clusters and multiple packets per event are transmitted if necessary.

The 3BC-L1 packet contains only one cluster of four consecutive channels with three beam crossings per channel reported (the trigger BC plus one BC before and one after the trigger). Sufficient packets will be sent to report all found clusters. This packet format, selectable by the configuration register, is intended for commissioning or timing studies.

The payload of the Internal Register packet is just the 32-bit contents of the register.

26 bits Header	34 bits Payload
60 bits packet size	

Table 3: ABC130 Packet Format

Start Bit	ChipID	TYP	L0ID	BCID
1	5	4	8	8

Table 4: ABC130 Packet Header

Address	Address	Address	Address	OvF	Stop Bit
8	8	8	8	1	1

Table 5: ABC130 Packet Payload, R3 packet

Address	Hit	Address	Hit	Address	Hit	Stop Bit
8	3	8	3	8	3	1

Table 6: ABC130 Packet Payload, 1BC-L1 Packet

Address	Hit	Hit	Hit	Hit	TBD	End Bit
8	3	3	3	3	13	1

Table 7: ABC130 Packet Payload, 3BC-L1 Packet

Data	TBD	End Bit
32	1	1

Table 8: ABC130 Packet Payload, 32 bits register packet.

The time interval between transmissions of packets out of a single chip is variable and depends on the arrival time of the R3 or L1 signals and on the presence of packets in adjacent chips. A "Xoff" mechanism prevents the adjacent chip to transmit its packet to the current chip if the latter is not ready to accept it (internal FIFO Full).

There are two sets of bi-directional signals for communicating with adjacent chips, XOFF1 and DATA_1 for communication with the chip to one side and XOFF2 and DATA_2 for communication with the chip on the other side. This is to allow data to flow in either direction through the chain of ABC130s on a hybrid. All of the data can flow in one direction or the other or the ABC130s can be split into two chains or loops on the hybrid. The latter case can be used to eliminate a malfunctioning chip from the readout by breaking the chain at that point.

The ABC130 ASICs contains addressable registers to configure:

- The analogue front-end (bias and threshold)
- The channels calibration
- The channels mask
- The prioritization of packets and other configurations of the digital functions
- Select which set of bi-directional lines to use for data readout
- The DCS local features

All the registers can be read-back, also during data taking: the registers data packets are treated with the lowest priority. One specific register can be set so that it is read with the highest priority: this register may contain the critical error flags.

3.1.1 Control Signals

The control signals are given in Table 9.

Name	Type	Description
In1- In256	Analog	Preamplifier Inputs from silicon strip sensors
RCLK	LVDS	(80 or) 160 MHz Clock input primarily intended for Data
BC	LVDS	Beam Crossing Clock at 40MHz
L0_CMD	LVDS	L0 Synchronous Trigger with BC falling edge, CMD with BC rising edge
R3s_L1	LVDS	R3L0ID with BC falling edge, L1L0ID with BC rising edge

Table 9: Control signals

RCLK - The data clock determines the rate at which data is clocked out of the chip. It is expected that this will be either 80 or 160 MHz. This clock will be delayed by an arbitrary amount with respect to the incoming stave clock to accommodate a hybrid wide phase shift to align sensor signals with the BC.

BC - The 40 MHz LHC beam crossing clock.

L0_CMD - This 80 Mb/s signal is comprised of two time-multiplexed 40 Mb/s components:

L0 - A one bit BC synchronous signal signaling that data in the pipeline has been identified as an event to store into the R3L1_Buffer. The resultant action is to transfer the data from the pipeline into the R3L1_Buffer with a BC tag and a L0 tag generated internally.

CMD - The input for the Command Decoder that will be processing commands. Commands can be received independently of the status of the triggering system and during physics data taking. The normal length of a command is 53 bits.

R3s_L1 - This 80 Mb/s signal is comprised of two time-multiplexed 40 Mb/s components:

R3L0ID - The Regional Readout Request (R3L0ID) signal is delivered on one phase of the BC clock and has two fields and a length of 11 bits: 3 bits act as start bit pattern (101), followed by 8 bits containing the L0ID identifier of the event to retrieve in the L1Buffer.

L1L0ID - The second trigger level (L1L0ID) signal is delivered on one phase of the BC clock and has two fields and a length of 11 bits: 3 bits act as start bit pattern (110), followed by 8 bits containing the L0ID identifier of the event to retrieve in the R3L1_Buffer.

3.1.2 Bi-Directional Signals

Table 10 lists the bi-directional signals of the ABC130.

Name	Type	Description
XOFF1	LVDS	Chip to Chip Data Packet Flow control
XOFF2	LVDS	Chip to Chip Data Packet Flow control
DATA_1	LVDS	Chip to Chip Data Packet Input/Output
DATA_2	LVDS	Chip to Chip Data Packet Input/Output

Table 10: Bi-Directional Data Communication Signals

XOFF1/XOFF2 - When true, the chip is not ready to receive data from the adjacent chip in the serial chain. XOFF will be controlled by a priority encoder set by the chip position number as well as data in the queue in the local ASIC. Can be set as input or output.

DATA_1/DATA_2 - Serialized Data passing from chip to chip. Can be set as input or output.

3.1.3 Fast Cluster block for self-seeded track tracker

The Fast Cluster block is a prototyping block included on the ABC130 that, when enabled, will provide prompt, beam clock synchronous, cluster position data to an external device that will correlate clusters between tracking layers and select high PT coincidences to send to the Trigger processor. The fast cluster algorithm assumes that the ABC130 is bonded with the two rows of strips interleaved as described in section 2.3. Strip data is provided from the input of the ABC130 pipeline after the mask register. This allows noisy channels to be masked off to avoid having a cluster registered at most or all beam crossings from improperly functional channels. Data from these banks is processed independently in parallel.

Its block diagram is shown on Figure 5 and its main feature are:

- Parallel processing of cluster finding from each 128 strip bank
- One dedicated High Speed Serialized output per bank of 128 strips (2 per ABC130)
- Rejection of Low momentum tracks by limiting the number of hit strips in a cluster to 2
- Half Strip Precision achieved by listing how many strips were hit (1 or 2)
- Fixed Delay from BC - Serialized data will be sent with a fixed delay from cluster's BC
  - 5 BC delay with a 160 MHz Serializing Clock
  - 2 BC, dead timeless, delay with a 640 MHz clock
- Power (when enabled) with Drivers
  - Logic + Serializer 4.5 mW + 9 mW for two Drivers @160 MHz
  - Logic + Serializer 5.5 mW + 12 mW for two Drivers + 3 mW for receiver @ 640 MHz

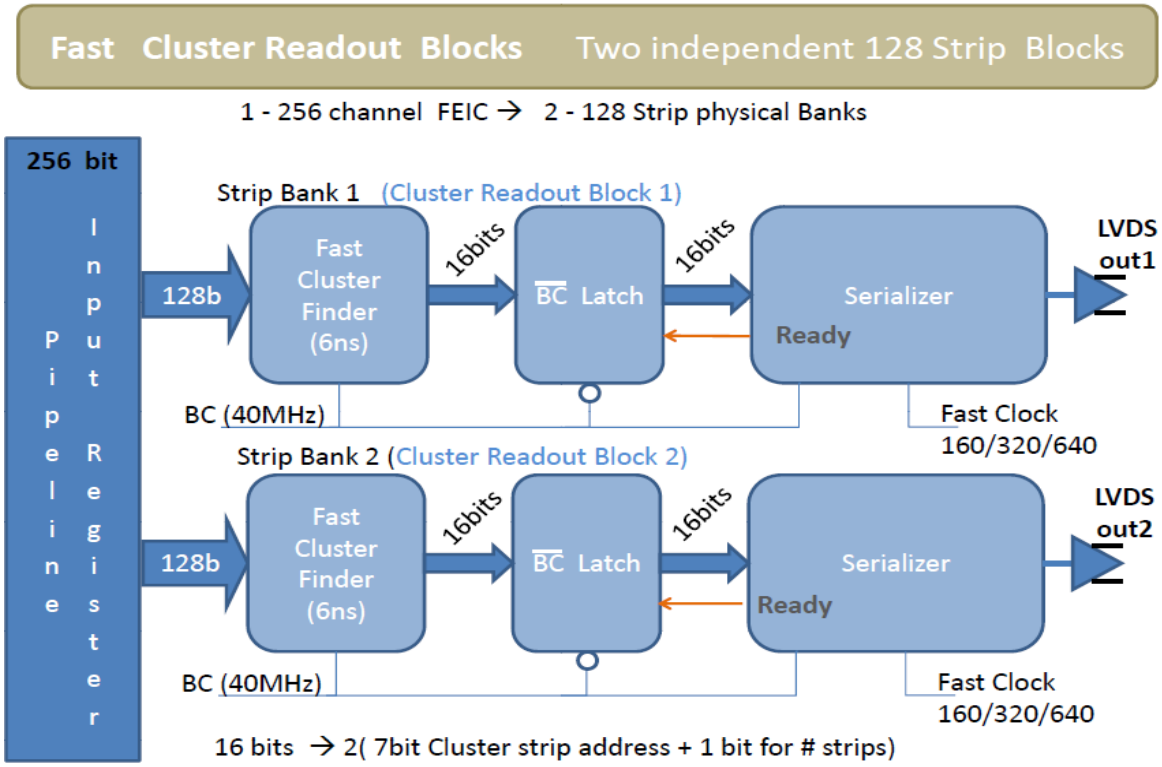

Figure 5: Block diagram of the fast cluster blocks in ABC130

Fresh strip data is loaded into the pipeline at the rising edge of the BC. This data (odd channels to one cluster finder and even channels to the other) is evaluated using combinatorial logic in 6 ns. The Cluster Finder output consists of two 8-bit words for each bank, where the 8-bit word represents a 7-bit address plus one bit to indicate the number of hit strips (1 or 2). Given the 80 µm strip pitch for the upgraded strip detector this will effectively yield a phi resolution of 40 µm. In the event that fewer than two clusters are found, the output 8 bits are set to an illegal combination, all 1's. The BC bar latch will record the Cluster Finder data if the serializer is ready. The serializer will start on the next rising edge of the BC clock. When serialized at 640Mbps, the 16 bits of cluster data is output to the correlator every BC period within a 2 BC period delay and no cluster finder data will be lost. Figure 6 shows a timing diagram of this process. In the event of more than two qualifying clusters in a single BC, the

two clusters nearest the edges of the bank (lowest and highest addresses) are sent out and clusters between these are lost.

A special framing mode is provided when a bit is set by the command decoder. In this mode the serializer outputs send 8 1's followed by 8 zeros to establish the data frame according to the local clocks. The correlator will need a corresponding mode to establish a proper read delay for each incoming channel.

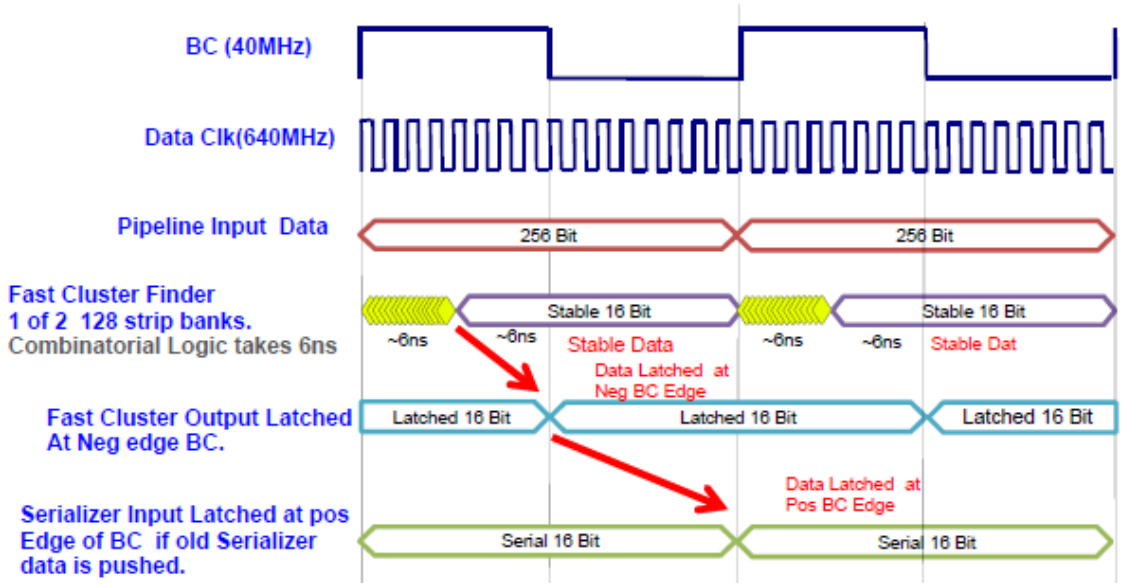

Figure 6: Cluster finding Logic Data synchronization. Data Latched at the serializer input (bottom trace) is immediately sent out serially.

3.2 Hybrid controller ASIC: HCC

The HCC is the interface between the stave service bus (stave side) and the front-end ASICs on the strip detector hybrids (hybrid side). Figure 3 depicts its position in the readout architecture. Each HCC in the inner barrel will service one half of a sensor for inner layer short strips, 2560 strips with ten 256-channel FEIC's (ABC130), and one whole sensor or module for outer layer sensors with long strips.

Five, hard wired, hybrid address bits (A0-4) uniquely identify each HCC on the stave; these may be set by bonding or by an Efuse. The latter requires two additional ASIC mask steps. All active control and data lines use a form of Low Voltage Differential Signalling (LVDS like). The stave side control lines are HCC input signals generated by the EOSC (GBT Elinks) and bussed to all stave resident HCC's (up to 26). Stave side control signals are received, interpreted and redistributed as required to the hybrid by the HCC. The hybrid side control signals are bussed to all ABC130s and each has a single differential pair termination. The HCC differential signal drivers will be current programmable in 8 steps to allow the circuit to be adapted to load and receiver conditions in various implementations. Two independent control signals are time multiplexed (phase encoded) onto each differential I/O pair at 80 Mbps such that each control signal has a 40 Mbps update rate. The HCC interprets stave side commands downloaded on the L0_CMD input and responds to a subset of these for its setup and response to monitoring requests. Unrecognized commands are sent through to the hybrid side. The HCC monitors the stave side (40 MHz) BC Stave signal generated and sent on a

single differential pair to all HCC chips on the stave by the GBT at the EOSC. The HCC will implement a radiation tolerant PLL that will synthesize binary weighted frequencies up to 640 MHz for use in providing programmable phase hybrid clocks for data and the ABC130 pipeline clocking in time with the arrival of interaction particles at the location of the hybrid sensor. The HCC will provide both a Beam Clock (BC) and a separate hybrid level data readout clock (DRC) with a common phased offset remotely programmed into the HCC. The HCC clocks may be turned off to operate the FEIC's in a power reduction mode. Hard reset is made possible by holding the R3_L1 in the true state for 16 or more BCs. Each HCC sends data back to the GBT receiver via a private data path. Two drivers in the HCC are connected in parallel to one hybrid specific data line on the stave (26 data pairs/stave). The current drive for each driver is selectable in 8 steps in the range 0 – 3 mA. For low latency operation the data rate can be as high as 320 Mbps corresponding to receiving data from both hybrid data loops at 160 Mbps. When the hybrids are serially powered the operating voltages between the EOSC and the HCC will vary by position, a difference of up to 30 V may be realized depending on the length of the serial chain. AC coupled stave side communications will be necessary. The AC coupled HCC receivers will use internal hysteresis to ensure memory of the last state on the AC coupled side. This idea has been successfully prototyped in the SPI chip, a Serial Power Control ASIC [4]. The HCC data drivers sending data to the GBT will also need AC coupling capacitors. Since the GBT has no provision for hysteresis, area balanced communication techniques will be required to prevent common mode drift of the AC coupled GBT receivers. Either scrambling or the more constrained encoding such as 8b/10b will be used. It is foreseen to make the implementation of the coding technique selectable using a bit in a status register. The HCC can report errors/reduce power (interlock) and by request output monitor data.

3.2.1 Description of the HCC I/Os

Figure 7 shows a block diagram of the HCC with the different I/Os.

Stave Side HCC Inputs and Outputs:
The signal connections between the HCC and the stave bus are listed in Table 11.

Name	Description	Speed
BC_Stave	40 MHz GBT Clock	40 MHz
L0_CMD	Time multiplexed beam synchronous L0 and CMD, each @ 40 Mbps. CMD uses inverted logic.	80 Mbps
R3_L1	Time multiplexed Regional Readout Request (R3) and second Level Trigger (L1), each @ 40 Mbps. L1 uses inverted logic.	80Mbps
HARD RESET	Holding R3_L1 True for 16BC initiates a HARD RESET.	na
Data_I and Data_II	Stave side point to point output DATA from each HCC to the EOSC. Data I and Data II are separately programmable output drivers that may be individually enabled.	80, 160 or 320 Mbps
Address Bits (A0–4)	Hard wired static address bits to uniquely identify each hybrid on the Stave bus. These may be fuse programmable.	na

Table 11: Stave Side HCC Inputs and Outputs

Hybrid Side Inputs and Outputs:
All the signals are shared among all ABC130 chips on the hybrid and are listed in Table 12. All signals are sent with the same clock phase as the BC by the HCC. Control signals will be encoded on both edges of the 40 MHz local clock.

Name	Description	Speed
BC	Beam crossing clock phase shifted by programming the HCC, used for BCID counting / Pipeline clocking / Phase marker for L0_CMD and R3s_L1 signals to ABC130s	40 MHz
L0_CMD	Independent CMD and L0 streams each at 40 Mbps	80 Mbps
R3s_L1	Independent L1 and R3s streams each at 40 Mbps	80 Mbps
DRC	Clock generated with an HCC programmed phase	80 or 160 MHz

Table 12: Hybrid Side HCC Inputs and Outputs

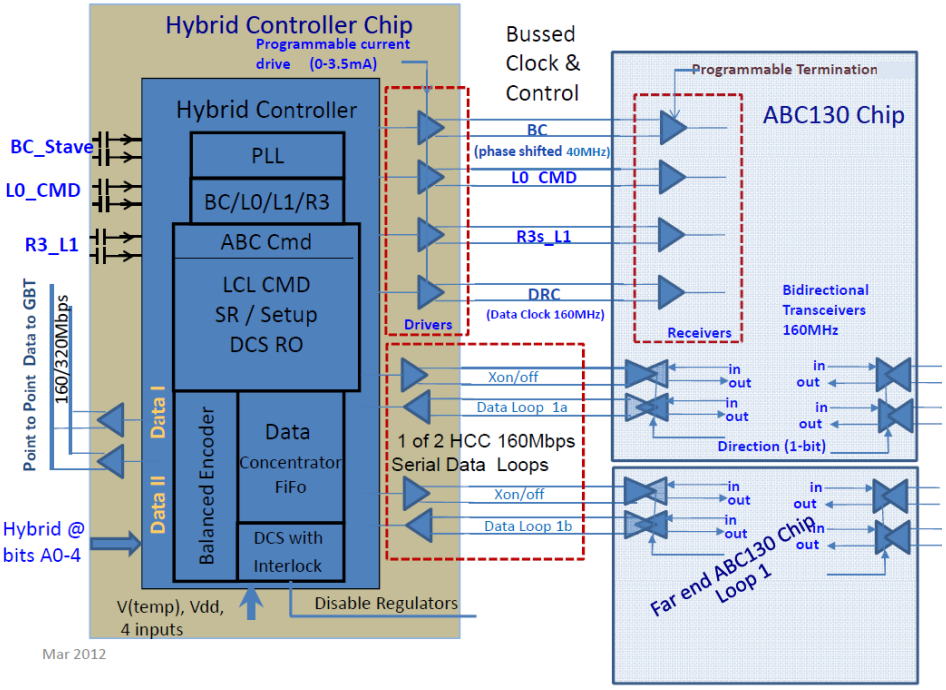

Figure 7: Block diagram of the Hybrid Controller Chip (HCC) and related I/O lines.

3.2.2 GBT Generated Signals

The input signals to the HCC, which are generated by the GBT at the EOS, are listed in Table 13. Note that the GBT generated signals L0_CMD and R3_L1 are each made up of two multiplexed signals, L0 with CMD and R3 with L1 with one issued on the rising edge of the BC_Stave clock and the other on the falling edge. The second portions of each, CMD and L1, will be issued with inverted logic in order to maximize the number of transitions during idle periods so that the AC coupled logic will preserve the proper logic levels. When no L0_CMD or no R3_L1 signals are being sent their inputs will follow the BC clock. The HCC will eliminate the inversion prior to sending CMD to the hybrid side.

Name	Description	Speed or size
BC_Stave	BC clock. Base clock for all HCC and ABC130 operations. A PLL in the HCC will provide a 640 MHz internal clock to be used for phasing the clock going to the Hybrid in steps of 1.6 ns.	40 MHz
L0_CMD	L0 (40 Mbps): Issued on the rising edge of BC_Stave. L0 is a beam synchronous signal sent a fixed number of BCs following the event of interest such that the data from that event are at the end of the ABC130 pipeline when the L0 is issued. L0, sent to all ABC130s in parallel, is used by the ABC130 logic to signal the transfer of data from the pipeline to the current L0ID memory address and increment the L0ID memory address counter. CMD (40 Mbps): Issued on the falling edge of BC_Stave with inverted logic includes commands for the ABC130 and HCC.	80 Mbps
R3_L1	R3: (40 Mbps): the Regional Readout Request (R3) sequence is transmitted on the rising edge of BC_Stave to prompt the readout of data stored in R3L1_Buffer from a specific region of the detector. The R3 transmission contains 4 fields including the module address and memory address of the selected hybrid data in the R3L1_Buffers. R3 is expected within ~50 μs following the event of interest and before that data is requested for L1, second level, readout. All the modules of a stave are represented by a bit map in the second field of the R3 bit stream. The HCC compares bits A0-A4 of its numerical module address with the corresponding bit position in the serialized R3 module address field. If the corresponding bit in the serial stream is high, a R3s (R3 short) sequence is sent to the ABC130s in the hybrid. (R3s retains only the L0ID address for the ten ABC130s on the selected hybrid and thus is shortened to 12 bits). It is possible for a single module to respond to requests for data from contiguous L0ID's. L1 (40 Mbps): Issued on falling edge of BC_Stave with inverted logic requests full detector readout. L1 is appropriately phase shifted and retransmitted to the ABC130 without intervention. L1 consists of 3 fields: header, L0ID location and trailer	80 Mbps
R3 fields (27 bits total)		
	Header	2 bits
	Module Address Bits (One bit for each module)	15 bits
	L0_ID	8 bit L0ID
	Trailer	2bits
L1 fields (12 bits total)		
	Header	2 bits
	L0_ID	8-bit L0ID
	Trailer	2bits
Hard Reset	R3_L1 input true for 16 BC's will initiate a Hard Reset.	
Idle	Low state while no L1 or R3 data is being transmitted.	

Table 13: GBT Generated Signals

3.2.3 Hybrid Side Common Signals

The phases of all the signals listed in Table 14 are adjusted for a common delay.

Name	Description	Speed
BC	40 MHz beam synchronous clock. Phase adjusted to account for the time of flight from the interaction point to the local sensor positions to ensure a common BC for event data. BC also serves to provide the phase information for time multiplexed control signals. The BC may be turned off to reduce power dissipation in the hybrid	40 MHz
DRC	Data Rate Clock. A selectable 80 or 160MHz clock used to clock serial data transmission between ABC130's and the HCC	80 or 160 MHz
L0_CMD	Phase shifted version of the L0_CMD signal on the stave side. HCC commands may be removed from this stream although it may not be necessary	80 Mbps
R3s_L1	R3s is a shortened version of the stave side signal R3. The HCC will strip off the module address field (12 bits) once it recognizes a match with its physical address. The total length of R3s is 12 bits. It is phase adjusted and sent out to the ABC130's. L1 is phase adjusted also and sent to the ABC130's	80 Mbps

Table 14: Hybrid Side HCC Common Signals

3.2.4 Hybrid Side Data collection

Data received from the FEIC's must be sorted into a single output stream and routed to the selected encoder as shown in Figure 8. Priority encoding is used to ensure that data collected from the various hybrid side streams is gathered at uniform rates,

Data from the HCC monitor and control and status registers must be interleaved with Hybrid FEIC data.

A send ID mode may be invoked to allow proper identification of the hybrid data stream all the way through the DAQ chain. This may be sufficient to allow the hybrid address bits to be dropped from the output during normal data taking.

3.2.5 Stave side Data Out

The HCC houses two data drivers that are available on the stave side for data transfer. Either or both may be programmed to output data to the Stave. It is expected that their outputs will be coupled in parallel to a single stave side data pair of lines to minimize mass. The Stave side drivers will set a common mode output via internal resistors at ~0.6V. These drivers will transmit data at 160 or 320 Mbps.

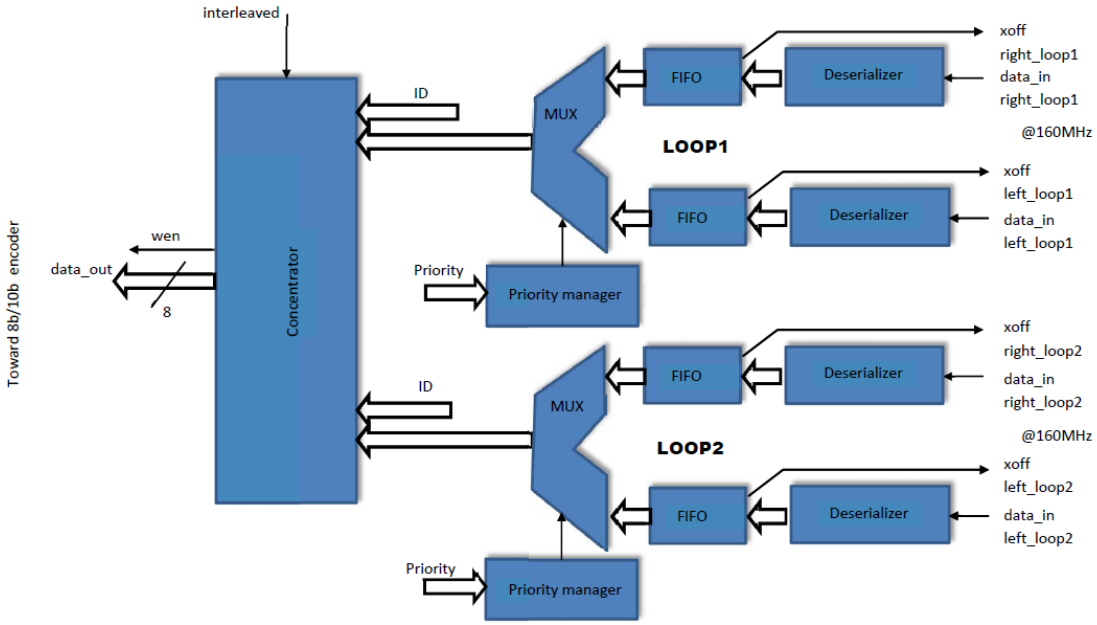

Figure 8: Block diagram of the HCC data concentrator. Two loops of ABC130 ASICs on the hybrid may be programmed to send data to the HCC through any of 4 input ports. The diagram above indicates the function of the Data Concentrator block to organize the hybrid ABC130 data for transmission out to the stave side driver.

3.2.6 Autonomous Monitoring

The HCC will contain a voltage based analogue monitoring block with 1 mV sensitivity from 0 to 1.1 V (10 bit range and accuracy) for two internal and four external values. Each monitored quantity will have a programmable upper and lower limit that will be downloaded through the HCC command decoder. An out of limit value will set a monitor flag and have the possibility to reduce the power in the hybrid by turning off the BC clock or disabling the regulators on the ABC130s with an external logic level. The operation of these interlock functions will be programmed in the HCC control register through the command decoder. Operationally the monitor block will use a counter driven linear ramp generator as a reference to compare with monitored voltages. When the reference voltage exceeds the sensed value, the current step counter value will be compared with the upper and lower limit and be recorded in the monitor registers. If enabled an interlock will occur at that time. The update rate will be every ~10 ms. The Monitor will be compatible with both external NTC temperature measurement and internal diode based temperature monitoring. It is envisioned to provide measurement of the internal bandgap as well to relate the ramp counter to a reference voltage.

3.3 Data correlator chip for self-seeded track trigger

The correlator chip is a notional concept being detailed at the behavioural Verilog level as a partner ASIC to receive the fast cluster outputs from the ABC130. Its basic function is to provide a low mass platform for intra-layer coincidence identification of high PT tracks in the barrel Silicon Strip Detector. Due to the relatively low probability of accidental coincidences at the beam crossing rate and the low occurrence rate of high PT tracks it is expected that any inner and outer barrel layer bank of 128 strips will have a relatively low coincidence identification rate. For this reason the correlator chip is envisioned to work loosely at the hybrid

granularity level although there is no fundamental reason that one correlator chip must readout all chips from one hybrid exactly. Figure 9 shows the concept. Fast Cluster serialized inputs derived from 20 banks of 128 contiguous strips from the inner side of a barrel stave layer are compared with the 20 pairs of Fast Cluster inputs from the outer side of a barrel stave layer radially aligned with the inner layer strips.

Data from the Fast Cluster finder on each ABC130 will be sent at 640 Mbps, providing its two cluster outputs (16 bits each) at the beam crossing rate. The correlator chip will collect this data simultaneously from all 20 front-end chips it is connected to. Ignoring edge effects between clusters, each inner and outer bank of 128 strips will have 4 possible combinations of potentially interesting coincident ionizations (High PT) because there are up to two clusters identified in each bank. These combinations will be tested simultaneously in a very simple manner. Each of the cluster pairs will be combined to form an address in a memory loaded with values at the beginning of the run. Since the number of possible combinations (216) is far greater than the number of interesting high PT coincident clusters, each interesting memory location will be loaded with a predetermined unique number or tag. Assuming each inner bank strip has 16 or fewer possible interesting outer bank high PT coincidence positions, only 2048 unique tags are needed for each matched bank pair. This number requires transmission of only 11 bits.

If the address of the coincident bank is sent in the remaining 5 bits, then the correlator chip will require only one output at 640 Mbps to report coincident clusters between layers for a full hybrid. Coincident clusters for four correlator chips (2 modules or 10240 strips) could be combined in an aggregator MUX to feed a single fiber that transmits its beam synchronous coincidence data out to a trigger processor in USA15. Using this technique a seeded track trigger would require 3 fibers per (half, 12 module) stave.

Given the high transmission rate from different drivers on different front-end chips and different hybrids, a special framing mode would be necessary to allow setting individual delays of the receiver clocks on the Correlator ASIC.

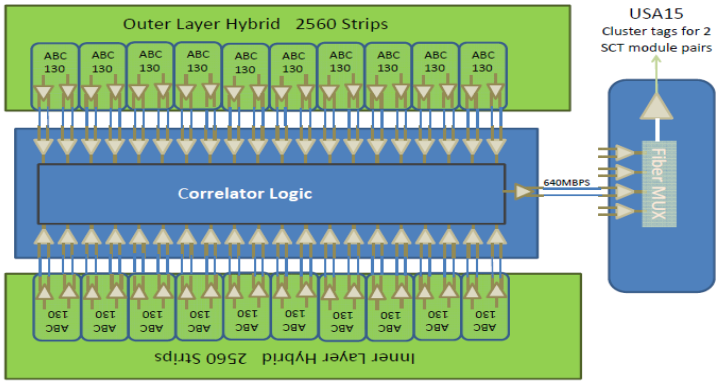

Figure 9: Notional schematic of the hookup of the Correlator chip to an inner and outer side of a barrel stave. One correlator chip would process coincidences for one Hybrid. A total of 4 Correlator chips (representing high PT triggers distilled from 10240 strips) would send beam synchronous data to an aggregator MUX for transmission up to USA15.

3.4 End of stave electronics

The end of stave module contains an End of Stave Controller, a DCS dedicated ASIC and the interface to two fibres, one transmitting data from the detector to the off-detector electronics and one receiving information from the off-detector electronics. The GBT [5] together with the Versatile Link components [6] are very good candidates for these purposes as they allow transmission in both directions of the TTC, data and control information and as they are designed to be used in high radiation environment.

3.4.1 Using the GBT

The GBT uses the concept of e-links to connect to the specific front-end readout ASICs. In the stave case, these e-links would be used to connect up to 26 HCCs of the readout hybrids to the GBT.

Each e-link consists of one clock line, one TXDATA line ("up-link" to readout the detector) and one RXDATA line ("down-link" to transmit trigger and commands) and has the following characteristics:

- Different e-links can run at different speed from 80Mbps to 160Mbps;
- Inside an e-link, RXDATA and TXDATA can run at different speeds;
- Unused clocks and RXDATA and TXDATA lines can be disabled.

It is envisaged to use up to 26 TXDATA lines at 160 Mbps to readout a short strip stave in the final system. The present GBT in development has a bandwidth limit of 3.2 Gbps payload if run in normal mode with forward error correction (FEC) enabled. However, it is possible to run the GBT in Wide Mode with FEC disabled, which then allows a payload bandwidth of 4.48 Gbps. In fact, in Wide Mode it is possible to configure the GBT to run up to 28 input lines at 160 Mbps utilizing the full 4.48 Gbps bandwidth and accommodating the full half-stave number of modules. There may be a more advanced version developed in time for the final system that increases the bandwidth with FEC. Early prototyping may run the TXDATA lines at 80 Mbps. The final system will either run in Wide Mode, with a more advanced GBT or with two GBTs. See the discussion of the needs for error correction in Section 3.4.2. If it becomes necessary to run the TXDATA lines at 320 Mbps (See "low latency option in Section 3.2.), a faster GBT or two GBTs will be necessary.

The vast majority of the e-links will only use TXDATA (i.e. the associated clock and RXDATA lines will be switched off and not routed on the stave service bus); a few clock and RXDATA lines will be used to transmit the TTC information from the GBT to the HCC.

TTC information needed

Assuming the Regional Readout Request scheme is to be used, we need the following TTC signals and information (See section 3.2 for a more complete description of these signals including how some are multiplexed together.):

- BC signal. The 40 MHz beam crossing signal, also used as the system clock for each hybrid;
- L0A signal. Can run up to 1 MHz (L0-rate) but expected to be no more than 500 kHz. It must be transmitted with fixed latency and it must be possible to issue it at any BC without dead time. When an L0 occurs, the corresponding data plus the L0ID are stored in R3L1_Buffer waiting for a R3ID or L1A. See section 2.5.3. If the module is part of a region of interest (RoI), a subsequent R3ID signal will be received, or if the event is selected by the level 1 trigger, a L1A signal will be received. In either case, the data selected by its L0ID will be transferred to the appropriate data compression logic block

(R3_DCL or L1_DCL) for clustering and compression and then sent to the readout block (See figure 4.);
- R3ID. The address of the readout hybrids that are part of RoI's and hence must send their data immediately. This information will be available some time after L0A;
- L1A. Trigger decision after the tracker RoI data have been analysed. It can run at 40–200 kHz (L1A-rate) and must be accompanied with a L0ID so that the data from the corresponding event can be retrieved and sent out. L1A and L0ID can be sent asynchronously without any (major) latency constraints;
- Commands (fast and slow) such as bunch counter reset (BCR), event counter reset (ECR) or registers access. It is important to note that the fast commands such as BCR could occur at the same time as L0 or L1A and that synchronicity is important.

TTC information to the GBT

The off-detector to GBT downlink is running at 3.2 Gbps, meaning that with each BC the GBT receives an 80-bit word (D0–D79). When the GBT uses e-links, these data bits are allocated to e-link RXDATA in the following way:

- if e-links are running at 80 Mbps, D0–D1 are allocated to e-link #1, … D78–D79 are allocated at e-link #40;
- if e-links are running at 160 Mbps, D0–D3 are allocated to e-link #1, … D76–D79 are allocated at e-link #20.

Several transmission tests have been done on stave bus prototypes at 40, 80 and 160 Mbps. In order to be safe, we should limit the bandwidth to 80 Mbps meaning that in order to get 4 bits per BC a HCC must be connected to 2 RXDATA lines of 2 e-links. Multidrop lines with up to 26 loads have been demonstrated and, hence, 2 e-link RXDATA lines would be used to transmit the TTC info from the GBT to the HCCs (and 26 e-link TXDATA are used for transmitting readout data from the HCCs to the GBT).

DC balanced codes

Scrambling is probably not an option for a 2-bit data transmitted at 80 Mbps but very simple ways of balancing the lines can be envisaged. L0A and Commands are sharing one 80 Mbps line while L1A and R3ID are sharing another 80Mbps line. If L0A and L1A use the opposite polarity than Commands and R3ID for their idle states, this should result in a reasonably balanced transmission.

Summary

The connection scheme to the GBT is summarised in Figure 10. The first version of the GBT will only allow 26 e-links at 160 Mbps with FEC disabled or at 80 Mbps with FEC enabled. Full stave prototypes will likely run with 26 e-links at 80Mbps and the need for FEC will be tested. A faster GBT may be available by the time of construction. It is hence important that the ABC130 and the HCC readout can run at either 80 or 160Mbps where 80Mbps will be used for prototype testing with the present GBT. Operating the HCC data output at 320Mbps is a further upgrade option.

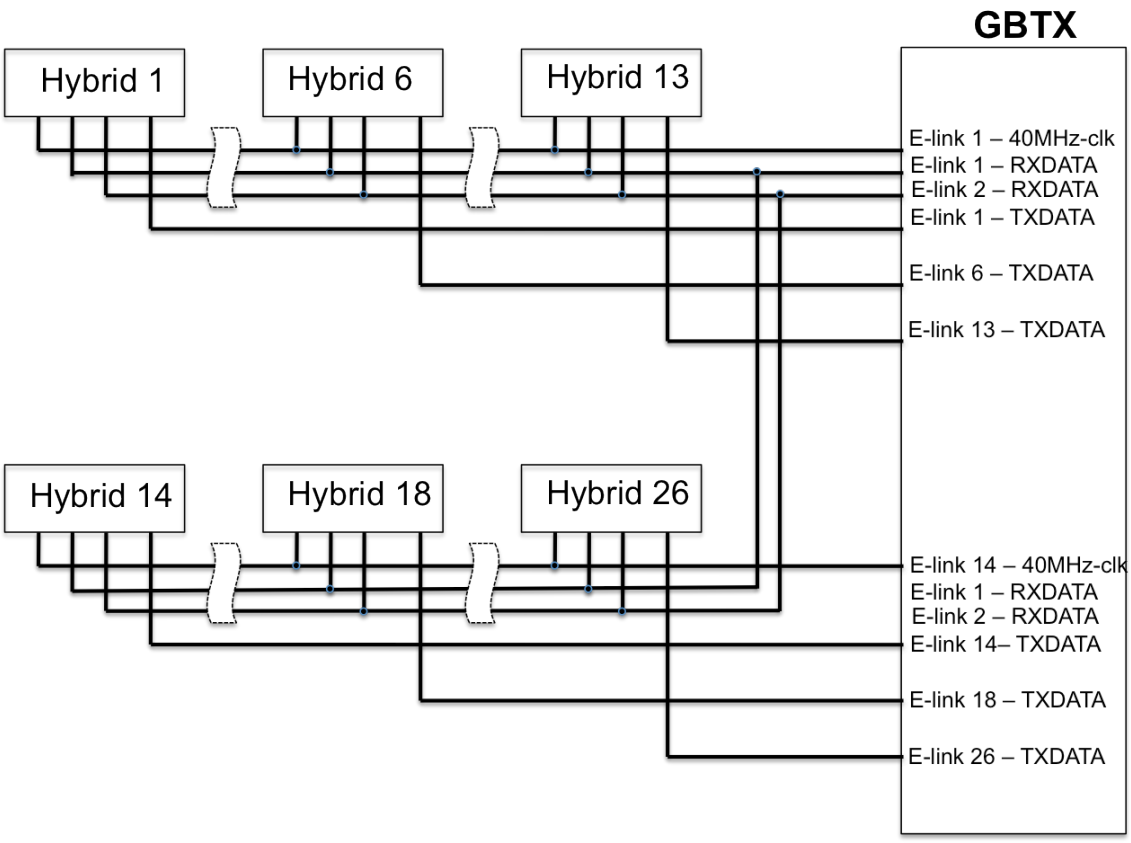

Figure 10: View of the GBTx and 26 readout hybrids. 26 e-links TXDATA are used to readout the 26 hybrids. Two e-links 40 MHz clocks are used to feed 2 sets of 13 readout hybrids. 2 e-links RXDATA are used to transmit 4 bits per BC to the 26 hybrids.

3.4.2 The Versatile Link

In the current generation of LHC experiments, different optical links have been developed for the four experiments. Even within ATLAS, there is a wide diversity of systems used for the various sub-systems, as well as differences between data and TTC links for the same sub-system. There have been problems with several of the optical links used in ATLAS [7] particularly with those based on non-commercial systems. For the upgrade it would, therefore, be better to put more resources into developing one common and very reliable system. This is of particular concern for the inner tracker as access to the on-detector optoelectronics, once the detector is installed, will either be impossible or at best extremely difficult. The use of optoelectronic components in the tracker impose some constraints that are not relevant for commercial components:

1. All components must be sufficiently radiation tolerant to withstand the fluences and ionizing doses corresponding to an accumulated luminosity of 3000 fb^{-1};
2. The system must be sufficiently robust in the presence of Single Event Upsets (SEU) caused by the radiation field;
3. The components must be of as low mass as possible and use preferentially low Z material;
4. Magnetic materials should be avoided.

There are common challenges faced by ATLAS and CMS for the optical links. Therefore, for HL-LHC groups from ATLAS and CMS are collaborating on a joint project to develop a Versatile Link (VL), which could be used in all the sub-systems required for ATLAS. Another advantage of this common approach is that it will minimise the development cost, compared to that of developing many solutions. This is very important because, although the full costs of the development stages are difficult to quantify, they are significant compared to the production costs.

The strategy adopted by the VL project is to make minimal modifications to existing commercial components, thus benefiting from the extensive reliability testing of large-scale manufacturers as well as the extent to which these components are "qualified by the customer". Such a strategy has been successfully adopted by CMS for their existing optical links. For the off-detector system, there are no additional constraints, so fully commercial components should be used.

The planned versatile link optical transmission system will be summarized here but a more complete description including technology evaluation data can be found in a companion document *Optical Links for the ITK* [8].

Bandwidth Requirements

In the present plans the output data rate from the HCC is 160 Mbits/s. For the barrel short strips with a half-length of ~1.2 m, there would be up to 13 modules per side. Allowing for two hybrids and HCCs per module, the resulting data rate at the End of Stave (EOS) would be up to 4.16 Gbits/s. Note that for the case of the RoI track trigger, there would be no additional bandwidth requirements and the trigger information would be interleaved with the normal data readout. The self-seeded track trigger would require additional bandwidth but those requirements have not yet been evaluated. The data rates required for the TTC links are lower but it will be convenient to use the same speed for the TTC links as for the data links. The current straw-man layout for the strip detector requires 1672 bi-directional links. If similar fibre cables are used as in the current SCT and Pixel detectors (which contain up to 96 fibres), this would require about 20 fibre cables of diameter 10.5 mm. If it were decided to use full redundancy for the optical links, this would require doubling the number of optical links.

System Aspects

The VL is based on bi-directional optical links. The increased number of channels and hit occupancy for the upgrade tracker compared to the existing tracker will require an order of magnitude higher total bandwidth. It is not feasible from the cost or material point of view to increase the number of optical links by that amount, therefore, the bandwidth of each link should be greatly increased from the very low values used in the existing SCT & Pixel detectors (40 to 80Mbits/s). The most economical strategy is to use the highest practical bandwidth. This is limited not by the fibres but by the speed of custom radiation hard electronics. The VL being developed is based on the SFP+ standard, which operates at a speed of 4.8 Gbits/s so there should be fewer links required than the existing ATLAS tracker. There are several advantages and disadvantages of operating at 850 nm (multi-mode) compared to 1310 nm (single-mode). (See. reference 7.) Therefore, the VL project is developing systems for both wavelengths. Which version will eventually be used by ATLAS will be decided before the production phase based on performance and cost considerations. The VL is planned to be used in conjunction with the GBT chip set [5] as illustrated in Figure 11.

As described in Section 3.4.1 above, the GBT acts as a multiplexer taking data from several modules and creating a high-speed serial signal. Environmental monitoring data can also be fed into the data stream through the GBT. The serial data stream communicates with a custom radiation tolerant laser driver ASIC (GBLD [8]), which sends the data over fibre via a laser. The laser could be a VCSEL for the 850 nm option or an Edge Emitting Laser (EEL) for 1310 nm. However, if suitable VCSELs are developed in industry for SM operation at 1310 nm, this could be an attractive option. For the TTC links the optical signal is received by a photodiode and followed by a radiation tolerant Trans-Impedance Amplifier (TIA). For the TTC links the GBT acts as a de-multiplexer to send common TTC data to different modules.

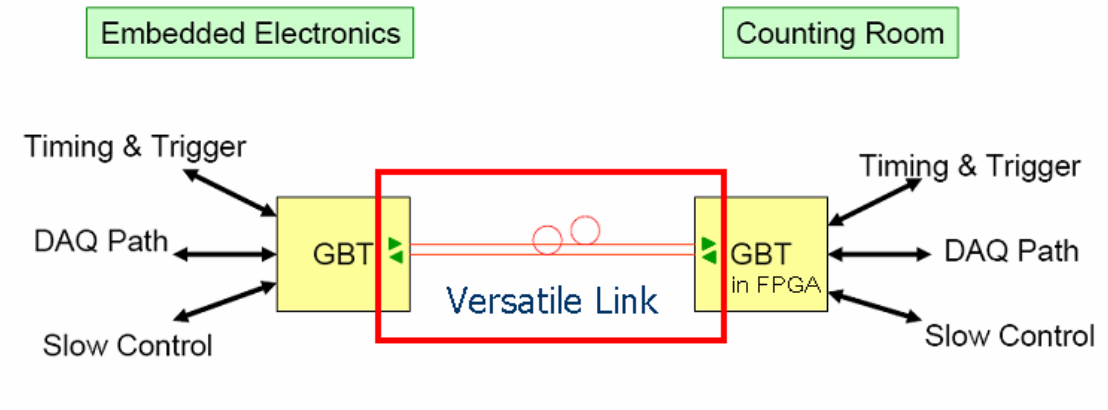

Figure 11: GBT and VL system

The VL is protocol agnostic, apart from requiring DC balance. For the data links the GBT wraps the data packets (payload) into a header, adds data for error correction, and uses data scrambling to ensure DC balance. Data Scrambling is an attractive way of achieving DC balance as it requires no additional overhead in data rates (as opposed to other methods to achieve DC balance such as 8b10b). The current error correction scheme allows for 80 bits of payload for each 128 bits, thus providing 3.2 Gbits/s of useful data. This provides multi-bit error correction which is required for the TTC links since there will be significant SEU rates in the *p-i-n* diodes. However, a lighter error correction scheme could be used for the data links (for which no significant SEU rates are expected) to provide higher data rates (e.g. 4.16 Gbits/s as required for the short strips).

The simplest architecture for the optical links is Point-to-Point. However, the VL is compatible with more sophisticated architectures used in Passive Optical Networks (PONs) as shown in Figure 12. One variant of a PON that is being considered for application to the TTC links is the use of optical couplers. This would allow the optical data from one TTC link to be split and sent to several VL receivers. This option has the advantage of reducing the number of channels of transmitters required in USA15 as well as reducing the number of fibres. Note that this option would require increasing the width of the address field in the GBT, which could be achieved in a second iteration of the chip.

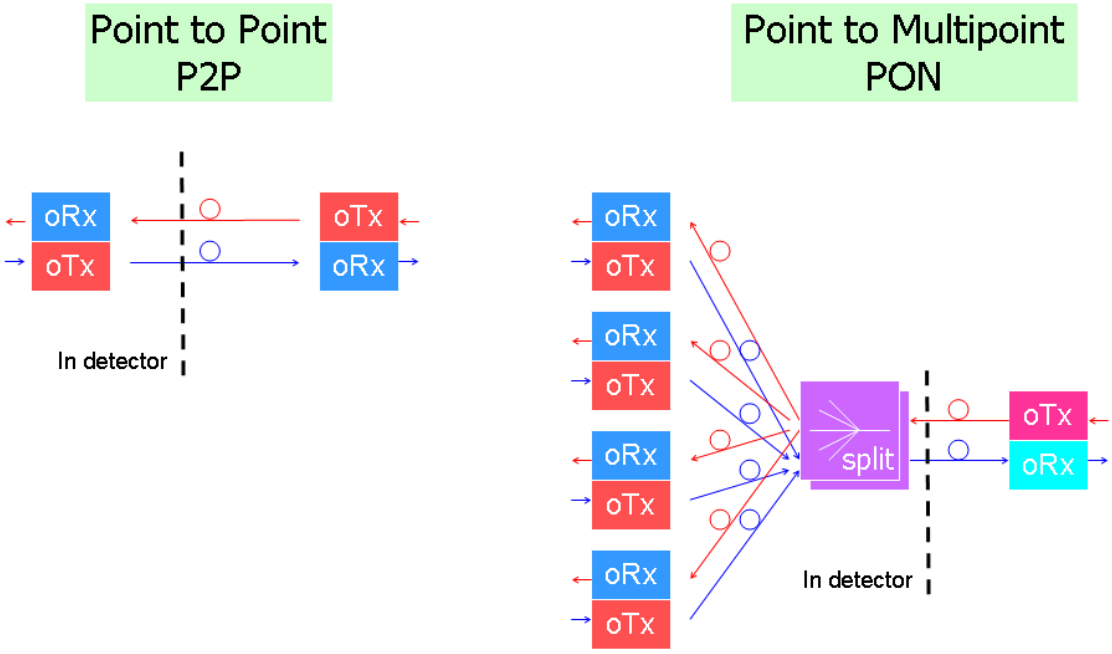

Figure 12: VL P2P and PON architectures. OTx (ORx) refers to optical transmitters and receivers. In the Inner Tracker, OTx and ORx will be combined into a transceiver, the VL TRx.

The optical transmission system can be broken down into on-detector components, passive components and off-detector optoelectronics. For the first two classes of components, radiation tolerance, reliability, power and material are critical considerations.

On-detector Components

The basic building block of the on-detector component is the Versatile Transceiver, VTRx, which consists of a transmitter (VTx) and receiver (VRx). The Tx consists of a laser driver custom radiation-hard ASIC (GBLD [9]) and the laser with an LC fibre connector. The Rx consists of a *p-i-n* diode coupled to an LC fibre connector and a Trans Impedance Amplifier (TIA). The package is based on the commercial SFP+ with the minimal set of modifications required to make it compatible with operation inside an LHC detector. These are detailed in other documents [8,9]. The GBLD is being specifically designed to be compatible with either VCSELs or EELs. This implies that it can operate with low output drive currents required for the VCSELs as well as the higher currents required for EELs. The drive current can be adjusted to accommodate the expected increase in threshold current with radiation damage. The TIA will provide an effective digital output. A photo of a prototype package is shown in Figure 13.

The custom ASICs in the VTRx are being fabricated in the IBM 130 nm technology and are expected to be sufficiently radiation hard as are the other on-detector readout ASICs. An R&D programme to select suitably radiation hard lasers and *p-i-n* diodes has identified candidate devices. Radiation tests, which included the expected annealing processes over the lifetime of the detector, have shown that with small increases in drive current of the GBLD the candidate VCSELs and EELs will survive the expected radiation levels. Likewise, InGaAs *p-i-n* diodes have been shown to operate with good responsivity after expected lifetime fluence provided that the TIA is designed to cope with the measured increase in leakage current. Data errors due to Single Event Upset (SEU) are also a concern especially for the VRx since, at the expected bandwidth, one SEU could corrupt several consecutive bits. SEU rates were measured

in pion test beams at PSI including the effects of multiple bit errors. The error correction implemented in the GBT was demonstrated to be fully efficient at correcting the observed errors. No evidence for SEUs was found with the VTx and it may be possible to omit use of error correction for transmitted data in order to increase the useful bandwidth but this decision awaits measurement of SEU rates for the GBTx itself.

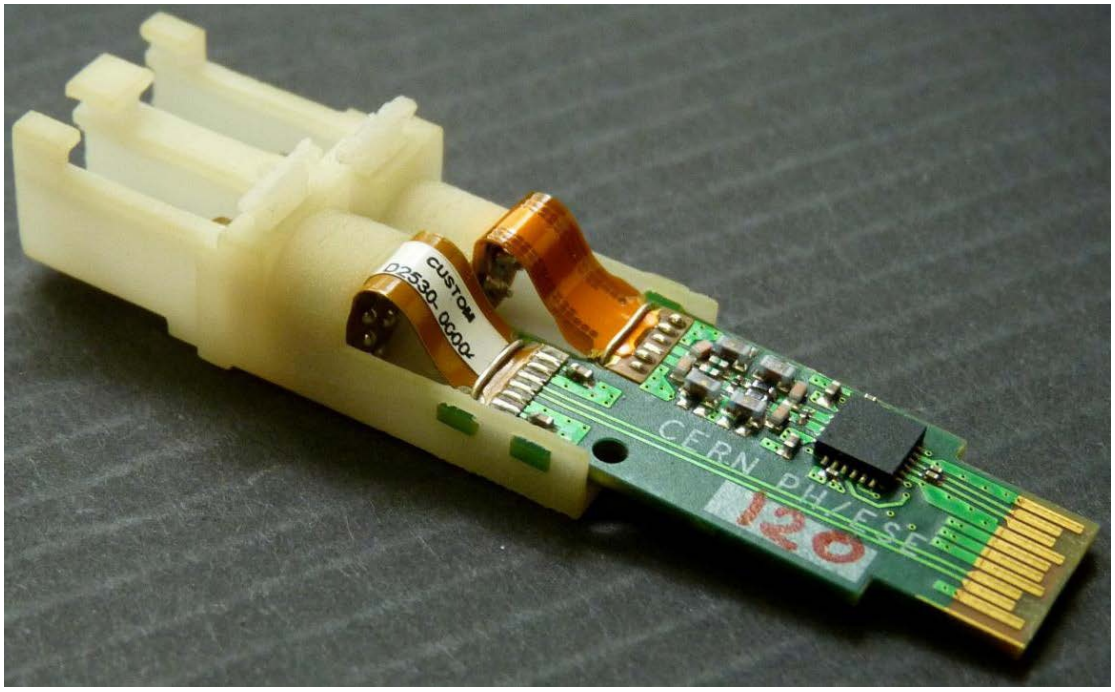

Figure 13: VTRx prototype showing the plastic part housing the OSAs and the LC fibre connector latches as well as the PCB with a laser driver. The TIA is integrated into the ROSA package in order to minimise stray capacitance.

Passive Components

The passive components include fibres, optical connectors, fibre cables and jackets. Many commercial optical fibres are not sufficiently radiation hard for application in the HL-LHC environment. Radiation hardness of fibres has been studied for many applications, including the current generation of LHC experiments, however, the radiation hardness needs to be verified for the higher doses expected at HL-LHC. In the upgraded tracker, the fibres will be exposed to radiation at low temperatures, around T~ -25°C. Previous results in the literature (and confirmed by our tests) show that the effect of radiation damage on fibres at low temperature is greatly enhanced because the beneficial thermal annealing is suppressed. Five fibres, both single mode (SM) and multimode (MM), were found to have acceptable radiation induced attenuation (RIA) at temperatures down to -30°C as shown in Table 15. The mechanical strength of the fibre core and the acrylate buffer as well as that of the cables and jackets needs to be evaluated. This is still in progress. Both glass based and silica based optical couplers have been tested with the silica based devices showing acceptable and significantly reduced damage compared to the glass based devices.

Fibre	RIA (dB)
Corning SMF-28e (SM)	0.961
Draka Elite SRH-SMF (SM)	0.0643
Corning Clearcurve OM4 (MM)	0.580
Draka Elite SRH-MMF (MM)	0.143
Producer X (SM)	0.277

Table 15: Expected RIA for fibres from the tracker to USA15. The data from Producer X was taken at higher dose rates in 2010. For the Corning SMF-28e, separate data was used to evaluate the RIA inside the tracker (cold) and outside (warm). For the other fibres only the cold fibre RIA data was available. In addition no corrections for dose rate effects were made. Therefore, these values should be considered as conservative upper limits.

Off-detector Optoelectonics

The off-detector optoelectronics do not have any of the constraints of the on-detector components, therefore, commercial components are being evaluated to identify candidate devices that can satisfy the remaining links requirements, for example, compatibility with Tx power and Rx sensitivity on-detector, and overall time jitter. Several commercial solutions are available and options are being evaluated to maximize the channel density, for example VCSEL and *p-i-n* arrays.

Power Budget

It is important to minimize the electrical power for the on-detector components. The electrical power budget of the VTRx is given in Table 16.

Item	Minimum Power (mW)	Maximum Power (mW)
GBTx		1800
GBLD	217	735
GBTIA		120

Table 16 Electrical power consumption for the components of one VTRx. The GBTx is included for completeness. The power consumption for the GBLD includes the power dissipation in the lasers. The minimum power for GBLD refers to the case of a single channel, operating at minimum bias and modulation current, and the maximum power refers to the case of two channels being operated at maximum bias and modulation currents. The maximum power, therefore, covers the worst case for the power using EELs after radiation damage.

3.5 Redundancy

There is some limited redundancy in the readout chain but only at the hybrid level. On the hybrid, each ABC130 has bi-directional lines to allow transmission of data to the chip on either side thus forming two possible data loops to the HCC. As discussed in section 3.1, if one ABC130 fails, data from chips on either side of the failed one can still be sent to the HCC. Also, the HCC has two output drivers connected to the same data line on the stave service bus. Only one driver will be used but if it should fail, the other output driver can still make the HCC and the hybrid functional provided that the failed output driver does not short the output lines.

At the present time there are no plans to include two independent HCCs on each hybrid or two independent data lines on the stave service bus. This is due to lack of space and desires to minimize material in the tracker volume. Likewise, there is no redundancy of the functionality on the EOSC. As pointed out in section 3.4.2, duplicating the EOSC functionality, especially the GBT and Versatile Link chip set would significantly increase the material. This does leave several possible single points of failure compared to the presently installed ATLAS SCT.

4. Power distribution

This section covers the low voltage and high voltage power distribution as well as the needed cable plant for different options. Only the barrel section will be considered here since the detailed design of the end-cap section (number of ABC130s per module and number of modules per petal or disk) has not been finalized.

The low voltage power section details the two powering schemes (DC-DC converters and serial power) in terms of power needed and overall efficiency for different options (different voltages for the analogue and digital parts, use of on-chip additional DC-DC converters). The first version of the two ASICs will use exclusively 1.2 V for analogue and digital power. The ABC130 will include separate externally enabled Low Drop Out (LDO) regulators for the analog and digital power domains. Their use will be optional. No provision will be included on this first version of the 130 nm chip set for on-chip DC-DC power conversion. The option of reducing the digital power rail to 0.9 V will require more study especially as it affects performance of the standard cells and SEU susceptibility. For this reason, the total power numbers listed in section 4.1.2 are the best estimate for the first versions of the ASICs.

4.1 Low voltage distribution

The powering scheme is not yet defined and will not be defined before specific R&D work is done to validate the different options (serial powering or DC-DC converters) at the component and system levels. As a consequence, the readout architecture has to be kept compatible with these different options and in particular with the serial powering scheme. In this scheme different parts of the system are at different potentials and this has some implications on the data transmission, which will have to be AC coupled.

It is assumed that the ABC130s of one hybrid will NOT be serially powered (i.e. they will all be at the same potential) while modules on a stave could be serially powered. Hence only the links between the hybrid controllers and the end of stave controllers are affected.

Independent of the power scheme used, powering only part of the system has to be feasible. Several system issues will have to be looked at. Some are related to DCS (e.g. where can we measure and control the currents and voltages), others are related to actions to be taken after a failure is observed (e.g. in case one module of a stave loses contact with the cooling, switching off the ABC130s and HCC of this module).

The current 128-channel ABCnext using 0.25 μm CMOS technology needs 40 mA for the analogue functions and 90 mA for the digital ones (using 2.5 V Vdd). Detailed estimates of the analogue power consumption of an ABC130 in 0.13 μm [10] have shown that the analogue part of the circuitry would need 160 μW per channel for the short strips and 300 μW per channel for the long strips assuming a 1.2 V Vdda, i.e. 34 mA per 256-channel ABC130 for the short strips and 64 mA per 256-channel ABC130 for long strips. Other analysis [11] has shown that one could expect 92 mW per 256-channel ABC130 (360 μW per channel) at 0.9 V Vddd, i.e. 102 mA per 256-channel ABC130 for the digital part.

Based on such current consumptions (34 mA and 102 mA for the analogue and digital parts), and assuming a 1.2 V Vdda and a 0.9 V Vddd working voltages, each readout hybrid of a short strips single-sided module would consume 1.33 W (520 µW per channel) and require at most 340 mA for the analogue part and 1.02 A for the digital one (1.342 A in total). Similarly, a readout hybrid of a long strips single-sided module would consume 0.85 W (660 µW per channel) and would require 320 mA for the analogue part and 510 mA for the digital part (0.830 A in total). The power consumption of the hybrid controller, estimated to be 180 mW (200 mA @ 0.9 V) is to be added. Table 17 summarises the power requirements per readout hybrid and stave for the strip detector.

	Hybrid		***Stave***	
	Power [W]	Current [A] (1.2 V Vdda & 0.9 V Vddd)	Power [W]	Current [A] (1.2 V Vdda & 0.9 V Vddd)
Short strips	1.51	1.55 0.35 for analogue 1.2 for digital	39.2	40.3 9.1 for analogue 31.2 for digital
Long strips	1.03	1.03 0.32 for analogue 0.71 for digital	13.4	13.4 4.2 for analogue 9.2 for digital

Table 17: Power consumption per readout hybrid and per stave for short and long strips. Estimated current for a 0.13 µm ABC130 with 1.2 V Vdda and a 0.9 V Vddd.

In this chapter we will not consider the protection schemes but only the power dissipated for different options.

4.1.1 Power model of a short strip single sided stave

A single sided stave for short strips consists of up to 26 readout hybrids, each of them containing 10 256-channel ABC130s and one hybrid controller, and 1 end of stave controller board.

The following assumptions are made for the different components' power consumption:

- 256-channel ABC130 (short strips case):
  - Analogue 34 mA @ 1.2 V per chip
  - Digital 102 mA @ 0.9 V per chip
- Hybrid controller:
  - Digital 200 mA @ 0.9 V (2 times the ABC130)
- End of stave controller board:
  - Opto devices 500 mW (200 mA) at 2.5 V
  - End of stave controller: 1.3 W (1080 mA) at 1.2 V (pessimistic guess)
  - Slow control: 50 mA at 1.2 V

Hence from the power point of view, a readout hybrid can be modeled as a device requiring 340 mA at 1.2 V and 1220 mA at 0.9 V; the end of stave controller board can be modeled as a device requiring 200 mA at 2.5 V and 1130 mA at 1.2 V as shown in Figure 14.

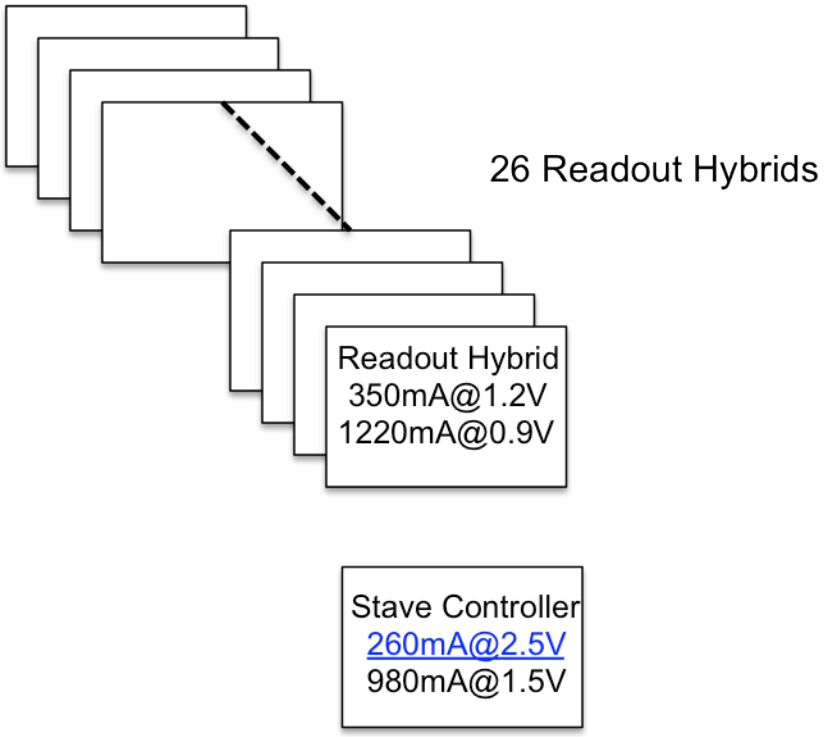

Figure 14: Power model of a short strip single sided stave.

This model has been used to assess the performance of a DC-DC powering scheme and of a serial powering scheme [12]. However, some SEU tests have shown that the SEU sensitivity of 130 nm CMOS technology when running at lower voltage than 1.2 V is high. Hence, it has been decided to set both Vddd and Vdda to 1.2V and to have Vdda and Vddd delivered by internal LDO regulators. Only this scheme is being described in the following section.

4.1.2 Effect of a 1.2V Vdd and of additional LDOs in the ABC130

The power model becomes:

- 256-channel ABC130 (short strips case):
- Analogue 34 mA @ 1.2V per chip (unchanged)
- Digital 102 mA @ 0.9 V per chip becomes 136 mA @ 1.2 V
- Hybrid controller:
- Digital 200 mA @ 1.2 V (2 times the ABC130)
- End of stave controller board:
- Opto devices 500 mW (200 mA) at 2.5 V
- End of stave controller: 1.3 W (1080 mA) at 1.2 V (pessimistic guess)
- Slow control: 50 mA at 1.2 V

Hence, from the power point of view, a readout hybrid can be modelled as a device requiring 1900 mA at 1.2 V; the end of stave controller board can be modelled as a device requiring 200 mA at 2.5 V and 1130 mA at 1.2 V. The total useful power of a single sided short strips stave is then (up to 26 readout hybrids and one end of stave board) about 61 W.

LDOs in the ABC130 will deliver 1.2 V from a 1.3 V input. In the case of a DC-DC solution, this 1.3 V would be delivered from a 12V-1.3V DC-DC for which an 80% yield is assumed.

DC-DC converters case

Table 18 summarises the current and power needs for single-sided short strip stave when using DC-DC converters.

End of stave controller board		
Needed current at 1.2 V	1130.0	mA
Needed current at the input of the 12 V-1.2 V DC-DC	161.4	mA
Needed current at 2.5 V for the opto-devices	200.0	mA
Needed current at the input of the 12 V-2.5 V DC-DC	49	mA
Needed current at the 12 V input	210.5	mA
Readout Hybrid		
Needed current at 1.3 V	1900	mA
Needed current at the input of the 12 V-1.3 V DC-DC	257	mA
Total current for the single sided stave at 12 V	**6.9**	**A**
Total dissipated power on the single sided stave	**83**	**W**
Useful power for a single sided stave	**61**	**W**
Efficiency	**74**	**%**

Table 18: Details of the currents at different stages, assuming 80% efficiency for the 12 V–1.3 V DC-DC converters and a 1.2 V Vddd and Vdda delivered by LDOs in ABC130.

Serial power case

Table 19 summarises the current and power needs for single-sided short strip stave when using serial power.

End of Stave board		
Needed current at 1.2 V	1130	mA
Needed current at the input of the 2.5 V-1.2 V DC-DC	614	mA
Needed current at 2.5 V for the opto-devices	200	mA
Total needed current at 2.5 V	816	mA
Power dissipated in the end of stave board	2.0	W
Readout Hybrid. Shunt regulator delivers 1.3V		
Needed current	1900	mA
Needed current in the shunt regulator	2235	mA
Total input current to the stave	**3**	**A**
Total dissipated power on the single sided stave	**78**	**W**
Useful power for a single sided stave	**61**	**W**
Efficiency	**79**	**%**

Table 19: Details of the currents at different stages for serial powering, assuming 85% efficiency for the main shunt regulator and a 1.2 V Vddd and Vdda.

Summary

Comparing the above numbers with the numbers given in [12] shows that the effect of using a single 1.3 V supply and 1.2 V LDOs to deliver the Vddd and Vdda of the ABC130 induces an approximate 50% increase in power dissipated on the stave from about 50 W to about 70–80 W, compared to an optimised power distribution.

4.1.3 DC-DC physical implementation

A development of radiation hard and magnetic field tolerant DC-DC converters is going on [13]. It can provide up to 3 A with an output voltage in the range 1.2 – 5 V from a 10 – 12 V input voltage. The converter includes over-current, over-voltage and over-temperature protection circuitry as well as the capability to turn it ON and OFF. The feasibility of using an additional switched capacitors converter has also been studied [14].

One converter is to be used for the end of stave board and one for each module. A single 10 – 12 V power line per single-sided stave is needed to power all the converters of this single-sided stave. Tests done with a prototype version on a 4-module stavelet prototype [15] (See Figure 15.) have shown that using these converters does not degrade the noise performance of the detector and does not introduce any pathological behaviour.

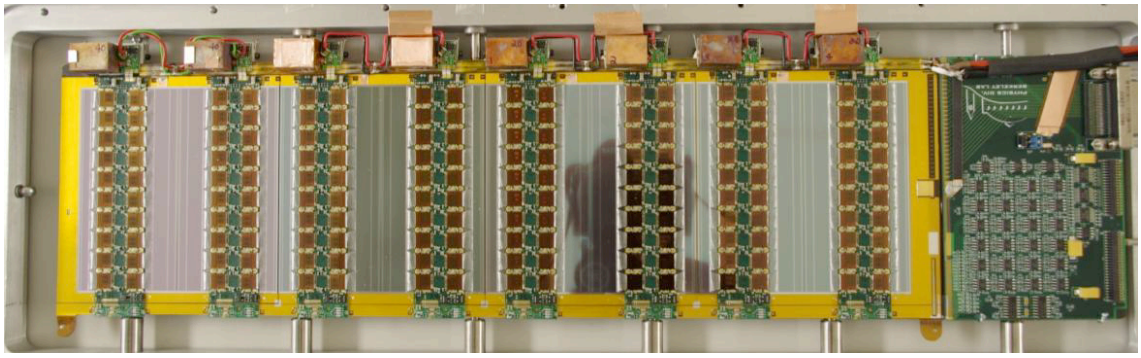

Figure 15: The DC-DC powered stavelet showing four modules and the 8 DC-DC converters powering the 8 readout hybrids.

A point of concern with the DC-DC converters is related to their amount of material as they use several components (ASIC, capacitors and inductor) in addition to a shield. The latter is needed because of the use of air-core inductors (necessary to withstand the high magnetic field in the ATLAS Inner Detector). Recent work [16] has shown that one can reduce the material contribution of all the DC-DC converters of a stave to about 0.032% X0.

4.1.4 Serial power physical implementation and Serial Power and Protection ASIC

The Serial Power and Protection (SPP) ASIC is intended to provide independently powered regulation and protection with remotely addressable operational control for a serially powered hybrid or module hereafter referred to (mostly) as a serially powered module. Power and digital control signaling is provided by a single line on each stave for up to 12 SPP chips to minimize the overhead required for support. The SPP is an integral part of the serial powering approach and will play an important role in demonstrating its robustness. The regulated module voltage may be preset to a wide range of voltages making it compatible with present and future front-end ASICs.

The specific SPP operational goals are:

- Regulation of a serially powered module to a pre determined operating voltage (.9V to 2.7V);
- Autonomous over voltage detection and shut down of a serially powered module;
- Remotely controlled module powering modes: Regulate or Shut down;
- Testing of the implementation of on chip current shunting shut down transistors capable of handling up to 1.6A with less than 120 mV drop across the module. This function provides a failsafe in the event that the module has a catastrophic bonding failure. (Requires the SPP to be mounted on the stave bus tape).

SPP chips on a stave derive their operating current through a resistor attached to an independent power bus that will be included in the stave tape. Pulse width modulated logic signals comprised of an SPP specific address and the operational mode desired for the addressed module will be encoded directly onto the "Global Power" line.

A prototype SPP has been produced. It employs very robust dual gate transistors that can handle higher voltages and currents than the native 130 nm CMOS devices. This preliminary version does not use the special ELT layout techniques required to prevent operational shifts in the threshold due to ionizing radiation. Ionizing radiation studies on a prototype ASIC that contained the analogue circuitry of the SPP chip have demonstrated a highly robust behaviour as shown in Figure 16. The chip was powered via an external power supply set at a high voltage through a series resistor to allow the SPP shunt regulator to set its voltage using the internal bandgap. The lack of significant variation in its self-regulated voltage indicates a robust operation over the detector expected lifetime. Additional studies using protons need to be performed to understand single event effects and combined ionization and bulk damage effects. These will be performed on SPP Version 0. If SEU or other effects are observed in the digital electronics it will be necessary to re-work the digital logic to implement some form of mitigation. This should be relatively easy given the small number of gates in the circuit.

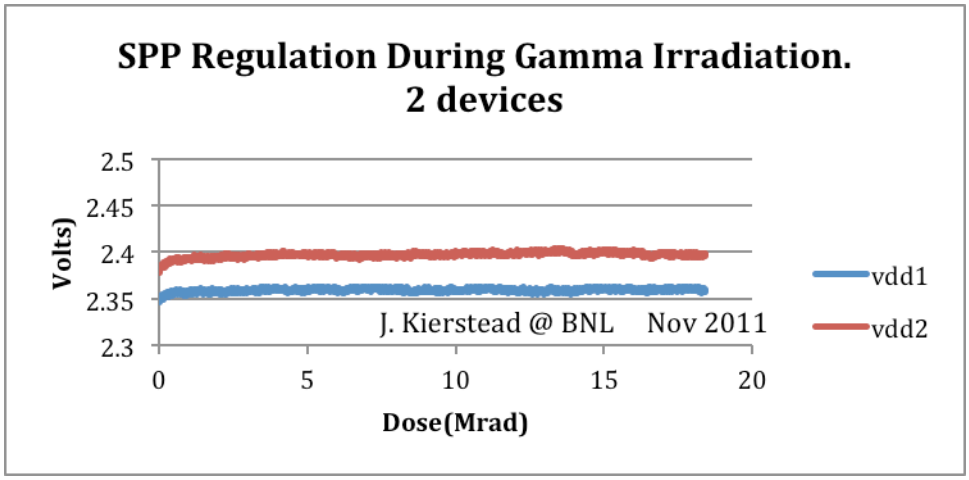

Figure 16: Measured variation of the SPP internal shunt regulator Vdd up to 18 MRad. This measurement suggests that the use of the CMOS8RF dual thickness gate transistors is sufficiently robust for analogue use in the upgraded strip detector. Tests with protons need to be performed to complete the study.

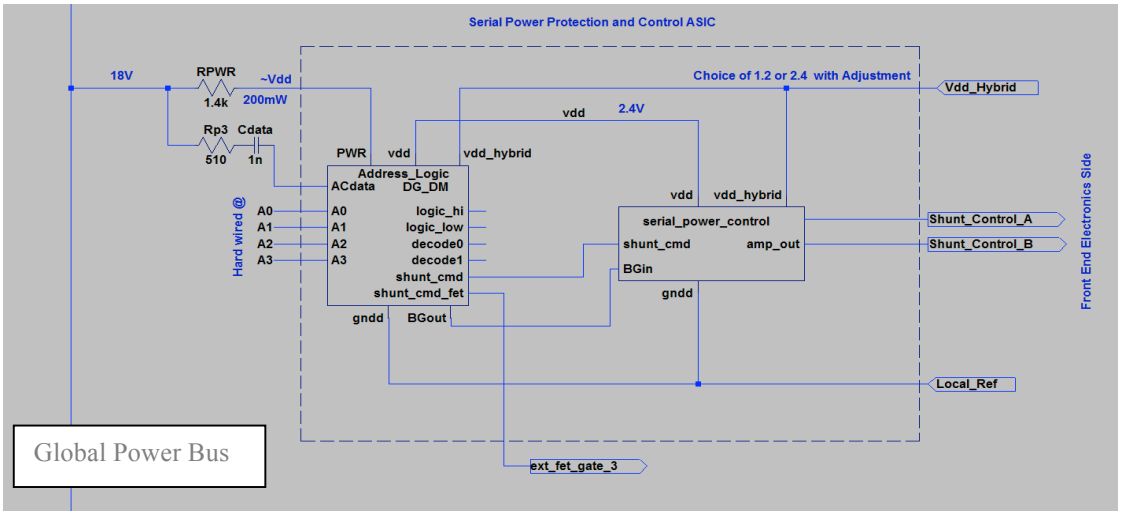

Figure 17: Block Level Schematic of the SPP ASIC. It derives its operating current and pulse width modulated remote commands to regulate or power down a module (or Hybrid) from the Global Power Bus. Its reference potential is that of the module (hybrid) it controls, labelled as Local_Ref above. It monitors the voltage on Vdd_Hybrid and uses the redundant shunt control lines that go to shunt mirrors on separate ABC chips to regulate the module (hybrid) voltage at preset values.

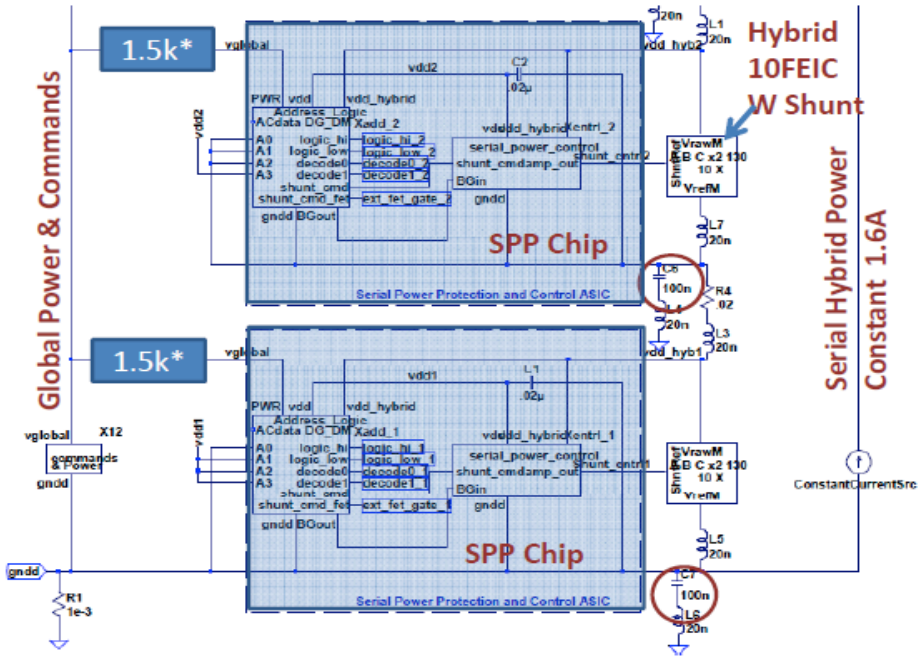

Figure 18: The schematic above demonstrates the hookup of two SPP ASICs on a serially powered stave. The small blocks on the right side are behavioural models of hybrids (each with 10 ABC130 chips) being controlled by the SPP. This kind of modelling was used for extensive simulations of the behaviour of a stack of twelve serially powered hybrids.

Figure 17 shows a schematic of the top level hierarchical blocks in the SPP. The block on the left is responsible for the on chip shunt regulation and command interpretation. The block

on the right is responsible for regulation of the hybrid voltage. The SPP derives the current for its operation through an external resistor attached to a "global power bus" common to all SPP ASICs. Its on board shunt regulator allows it to adjust the current draw across the series resistor to regulate its operational voltage to ~2.4 V. The SPP can also be addressed through the "global power bus" using pulse width modulation. Figure 18 shows the schematic connection of several SPP's in series with their hybrid loads on the right. The hybrid model implements a relatively accurate analogue, (linear) and digital (quadratic) variation of power with voltage. Realistic values for stray inductance, capacitance and resistance are used throughout the model. Figure 19 shows the results of an example simulation where pulse width modulated communication over the Global Voltage bus was used to turn a hybrid off and then turn it back on to its normal regulation mode. The "ext_fet_gate" signal is intended to drive the gate of a large FET to clamp the voltage across the module, shorting it to reduce power in the module while allowing the serial power current to continue on to other modules in the serial string. Its 2.4 V output provides plenty of gate voltage to maintain a low on resistance in the FET.

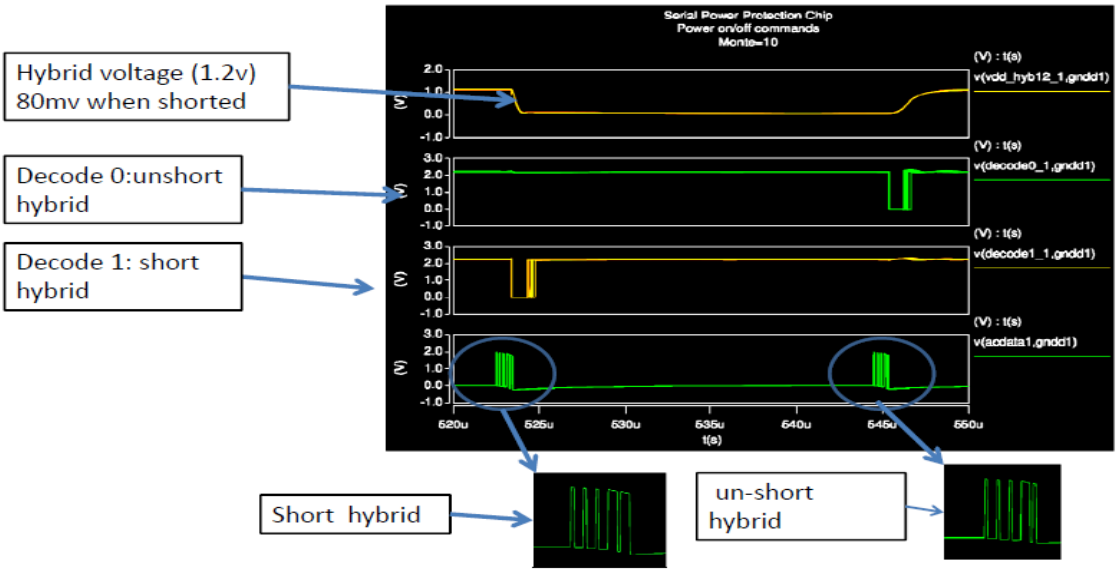

Figure 19: Results of simulation of 12 hybrids demonstrating the shorting of a hybrid using external pulse width modulation.

A prototype of the analogue blocks has been produced and tested. They have been shown to provide excellent control at the expected preset voltages. Furthermore a test shunt for inclusion in the ABC130 chip has been shown to work both in regulation and in shut down mode. Figure 20 shows two of the SPP prototyping boards in use on a module. Note that the SPP chip and a filter capacitor is all that is required for serial powering. The test board is large enough to provide easy access to multiple monitoring points and separates various chip level functions for independent testing. The plot in the lower left shows the minor overshoot associated with turning on the serial power current for the ABC130 based module.

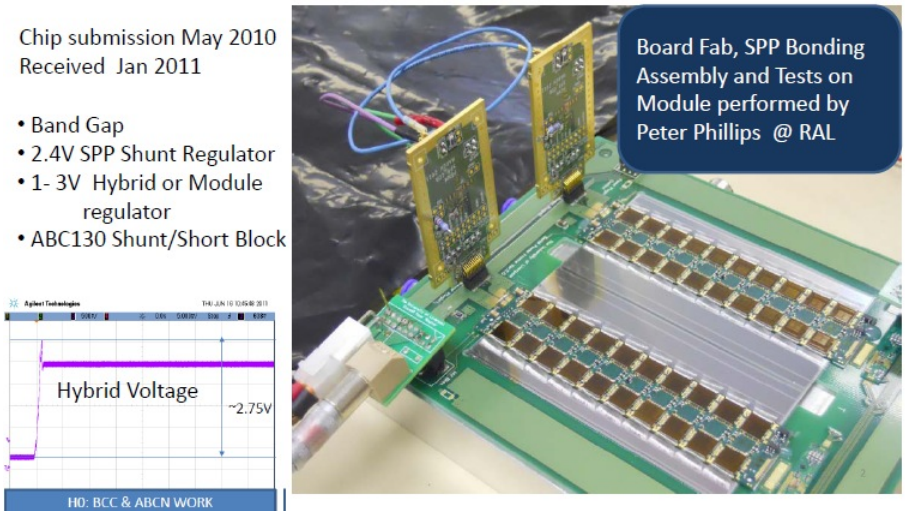

Figure 20: Two SPP analogue only prototype boards powering a module. The plot on the left shows the hybrid voltage during power up. The rapid turn on of the serial power current results in a 250 mV overvoltage followed by proper regulation at 2.5 V. Note that in normal operation a 2 mm X 2 mm SPP die, a resistor and a few capacitors are the only components required for Serial Powering.

SPP Version 0 Submission

The layout of the SPP version 0 die is shown in Figure 21. 40 SPP dice have been fabricated and will be tested with modules and against radiation. The next version of the SPP ASIC will include remote variation of module voltage and temperature based interlock.

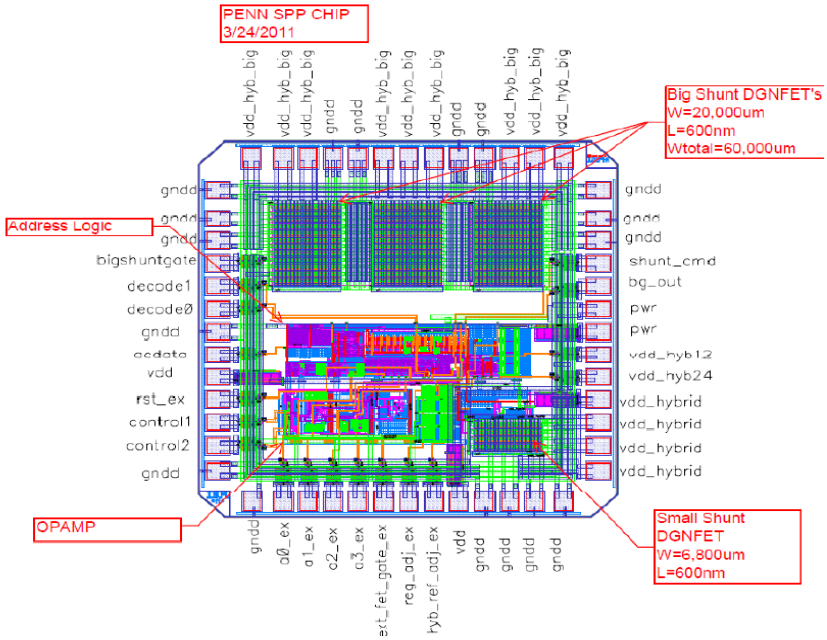

Figure 21: The SPP Version 0.

4.1.5 Off-detector low voltage bulk supply

The off-detector low voltage bulk supplies are of different types, depending on the powering scheme used.

In the case of a DC-DC converter scheme, the bulk supplies are standard LV power supplies able to deliver 12–15 V with 8–10 A capability (assuming one LV bulk supply per stave side). Devices similar to those currently used in ATLAS would be adequate. These devices are two-fold: relatively high voltage AC-DC converters located in USA15 deliver 48 V or 280 V to radiation tolerant DC-DC converters located in UX15. This scheme limits the current, the Ohmic losses and hence the amount of cables needed between USA15 and the cavern.

In the case of serial powering, the LV bulk supplies have to act as a current source. The Serial Powering Current source performs important protection and monitoring functions. During power-up it will be important to monitor the slope dV/dI. When all readout hybrids reach their regulating voltage, additional current will have very little effect on the voltage at the source. This change in slope will give the system feedback as to the proper operational current as the system ages. The monitored voltage will also give confirmation that a hybrid has been shut down, set to a low impedance to eliminate its power dissipation. Powering up one hybrid at a time and monitoring dV/dI will allow the monitoring system to show if one hybrid requires additional current or if it is not regulating at the proper voltage. A prototype device is already available, including over-voltage protection, isolated USB interface and programmable PID coefficients for system tuning.

4.2 High voltage distribution

The options for HV distribution are described for the barrel stave. The same options are also valid for the end-cap stave. Each sensor requires a detector bias as large as -500 V. In the baseline barrel stave design sufficient real estate exists to bias each detector individually as shown in Figure 22 (a). The 12 detectors shown represent a ½ stave (either top or bottom) which could be increased to 13. This solution is the most robust in that a failure by a single sensor will not affect the operation of any other stave sensors. Additionally, detector currents are easily monitored at the power supply. Negatives to this approach include the large number of cables and power supplies needed to supply a stave. The stave layout must ensure that two adjacent detectors can withstand a potential difference of 500 V when one detector is disabled and its neighbour is fully biased.

Re-use of the current SCT HV cables or other constraints may not permit a solution as in Figure 22 (a) due to the large number of cables that are needed. It may be possible and necessary to supply a ½ stave with just a few HV connections as in Figure 22 (b). In this case the failure of a single sensor may disable several sensors. Individual currents may not be monitored at the HV supply. However, it is possible to design current monitoring circuitry either on-hybrid or at the End-of-Stave to measure the individual currents and to read out their values via DCS.

The lowest mass option is to supply a ½ stave with a single HV connection as in Figure 22 (c). This option, however, is unattractive from a reliability standpoint.

A variant of this latter option is shown in Figure 22 (d) in which a single HV connection supplies all sensors. However, each sensor is connected through a HV switch under DCS control that permits a malfunctioning sensor to be isolated so that it does not affect the operation of the others. At present no HV switch that is sufficiently radiation hard and operates in such a

large magnetic field is known. HV MOSFETs and likely BJTs cannot operate to such radiation levels. However, two technologies are currently being investigated that may permit the HV distribution topology shown in Figure 22 (d). These technologies, Gallium Nitride and Silicon Carbide, as well as specific design details are presented in reference [17].

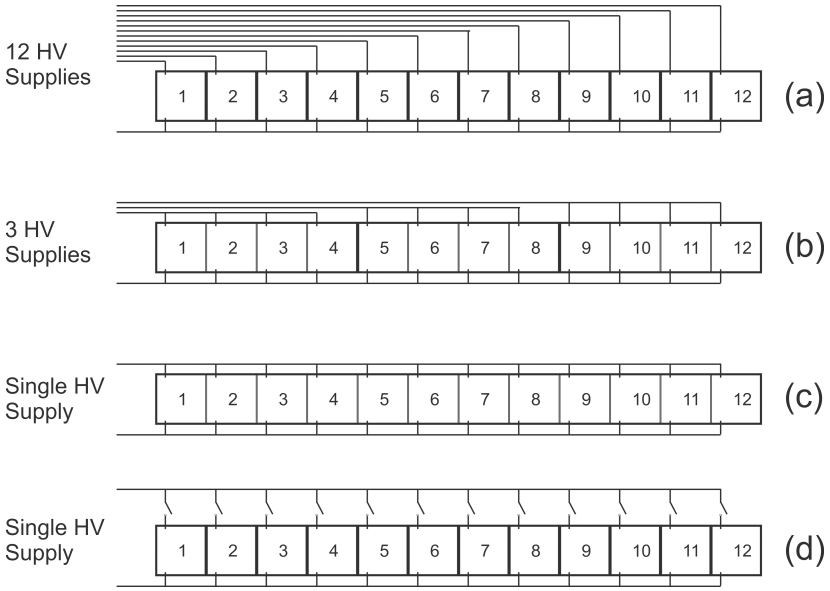

Figure 22: Possible HV distribution schemes for barrel staves.

4.3 Cable plant for power distribution

This section describes the needed cable plant for the high voltage and low voltage distribution for different options:

- Individual or grouped sensor bias for the high voltage;
- Serial power or DC-DC converters scheme for the low voltage.

4.3.1 High voltage cable plant

The number of needed "lines" is given. A "line" could be either a coaxial cable or a twisted-pair cable.

In the case sensors are individually biased, a total of 16,960 high voltage "lines" are needed. The lines must be rated 500 V, 3.1 mA.

In the case a single "line" for four sensors (no multiplexing) is used, 4240 "lines" are needed. These lines must be rated 500 V, 8 mA.

In the case a single HV "line" with multiplexing can be used for a single-sided stave (i.e. 12 or 13 modules), only 1584 "lines" are needed. These lines must be rated 500 V, 25 mA.

4.3.2 Low voltage cable plant

The content of this section is extremely premature and can only give an order of magnitude of the volume of services. The question of reusing existing cables is also not addressed. A power "line" consists in a pair of cables of equal cross section, one for the power feed and one for the power return.

It is assumed common 1.2 V Vddd and Vdda are used (as described in Section 4.1.2). It is also assumed that, in the case of serial power, all the readout hybrids of a single-sided stave or petal are serially powered and are using the same serial power feed, and that a separate power feed is used for the end of stave board. In the case of a DC-DC scheme, it is assumed that a single power "line" is needed per single-sided stave or petal.

As an example, this leads to the following needs per single-sided short strips stave:

- One "line" 12 V and return at 6.4 A needed in case DC-DC converters are used;
- One "line" 31 V and return at 2.3 A and one "line" 2.5 V and return at 0.9 A needed in case of serial power.

Based on these assumptions, the total number of power "lines" is listed in Table 20.

		Serial Power			**DC-DC Converters**		
		Number of Power lines	Voltage [V]	Current [A]	Number of Power lines	Voltage [V]	Current [A]
Barrel	Short strips	432	34	2.2	432	12	6.9
	Long strips	512	17	2.2	512	12	3.5
	Stub	256	2.6	2.2	256	12	0.7
	End of Stave	1200	2.5	0.8	-	-	-
End-caps	Strips	896	17	2.6	896	12	3.3
	End of Stave	896	2.5	0.8	-	-	-
Total number of lines		***4192***			***2096***		

Table 20: Number of power "lines" with their voltage and current for the two powering options.

4.3.3 Power distribution cable plant summary

Table 21 summarises the cross-section of copper needed per side of the strips detector. The following dimensions are assumed for each type of cable:

- LV cables: assuming a 2-V drop total and a 100-m length, 2-mm^2 of copper per lead and per ampere is needed;
- Bias cables: 1.05-mm diameter.

	LV serial power	**LV DC-DC**	**HV individual**	**HV 4 sensors**	**HV stave**
Barrel	864 * 4.4 mm^2 1536 * 3 mm^2 2400 * 2 mm^2	862 * 12.8 mm^2 1024 * 4.4 mm^2 512 * 1 mm^2	11840 * 0.87 mm^2	2960 * 0.87 mm^2	1200 * 0.87 mm^2
End-caps	1792 * 5 mm^2 1792 * 2 mm^2	1792 * 4.8 mm^2	8960 * 0.87 mm^2	2240 * 0.87 mm^2	896 * 0.87 mm^2
Total	***258 cm^2***	***247 cm^2***	***181 cm^2***	***46 cm^2***	***19 cm^2***

Table 21: Total cross-section of copper needed for different schemes of HV and LV powering.

5. DCS & Interlocks

The detector operation is very demanding in term of monitoring, control and safety. The status of the powering, the cooling and the environment are essential information to survey and control during the detector lifetime. The construction of the large silicon strip Upgrade requires that the DCS scheme takes into account the new powering, cooling and readout system with an architecture that is compatible for the whole inner detector [18]. The DCS system will be partly coupled to the readout chain providing wide flexibility for detector diagnostic. Therefore, the ASICs designs will implement some DCS features, which will be integrated into the readout chain.

In addition an autonomous interlock systems will have to guaranty safe operation to protect the detector at various stages against over overheating scenarios either in case of a cooling failure or in case of excessive local heating. The requirement for the whole Inner Detector (ID) upgrade that services volume not increase while the number of channels increases by possibly an order of magnitude compared to the present ID forces the DCS architecture to limit the number of communication paths as much as possible. Furthermore, two powering schemes are being considered for the strip detector, a serial powering scheme or local DC.DC conversion. These options have different features and will require designing some specificity for each in the DCS architecture.

5.1 The DCS architecture

All the use cases, from construction quality assurance, ID commissioning and finally ID operation must be considered. Various electrical test conditions must also be considered, for instance: reduced power without cooling for electrical integrity; full power with warm cooling for quality assurance; and full power with the nominal cooling for commissioning and operation. The system should be designed such that the operational risks are limited and that the survey parameters offer good capability to diagnose alarms.

The DCS architecture should consider the use cases, the detector layout, the powering system, the readout architecture and the service constraints between the detector and the counting room.

The detector layout drives the location of temperature sensors on the modules along the local support and the locations for monitoring other parameters like the local power supplies. Some DCS parameters will be digitized locally by ADCs, which will be integrated in ASICs located on the hybrid (both ABC130 and HCC) and on the EOSC (GBT). This data will be stripped out of the data path by the off-detector GBT-SCA and passed to the Global Strip DCS as shown in Figure 23. Even if the layout and geometry will be different for the barrel and the end-cap, the architecture will remain identical. The number of electrical lines between the cavern and the counting room is limited to the existing service that is now used by the inner detector. This limiting factor, in addition to the requirement that the material budget within the ID volume should be minimized, has implications on the number of lines assigned for powering and for DCS. Therefore, it is proposed to readout most of the DCS information via the optical fiber data line using a low and limited bandwidth. The survey or diagnostic data will be accessible only when powering the various readout stages. A few DCS electrical lines per stave/super-module or petal will have to be used for functions such as the cooling loop interlock and the environmental monitoring.

As described in Figure 23 the proposed strip DCS architecture integrates the information from the detector substructure and the detector environment as well as the power supplies

information. The direct DCS readings either from sensors or power supply monitoring will be done using up to four twisted pair cables linking the EOS electrical side to the embedded ADCs and interlock system. One of these lines will be dedicated for the power supply interlock based on the cooling loop temperature. Redundancy for this power supply interlock will be provided by sensors on both sides of the double-sided detector structure linked to each EOS side. The DCS crate with embedded ADCs will digitize the information from the cooling loop temperature, the environmental temperature and the LV and HV crate power supply. The readout and data handling will be done at a higher software level allowing for communication with the DCS crate, the GBT-SCA decoder, the DAQ via the DDC (DAQ DCS Communication) and the configuration database. The global strip DCS Graphical User Interface (GUI) will centralize all the information coming from the various DCS sources.

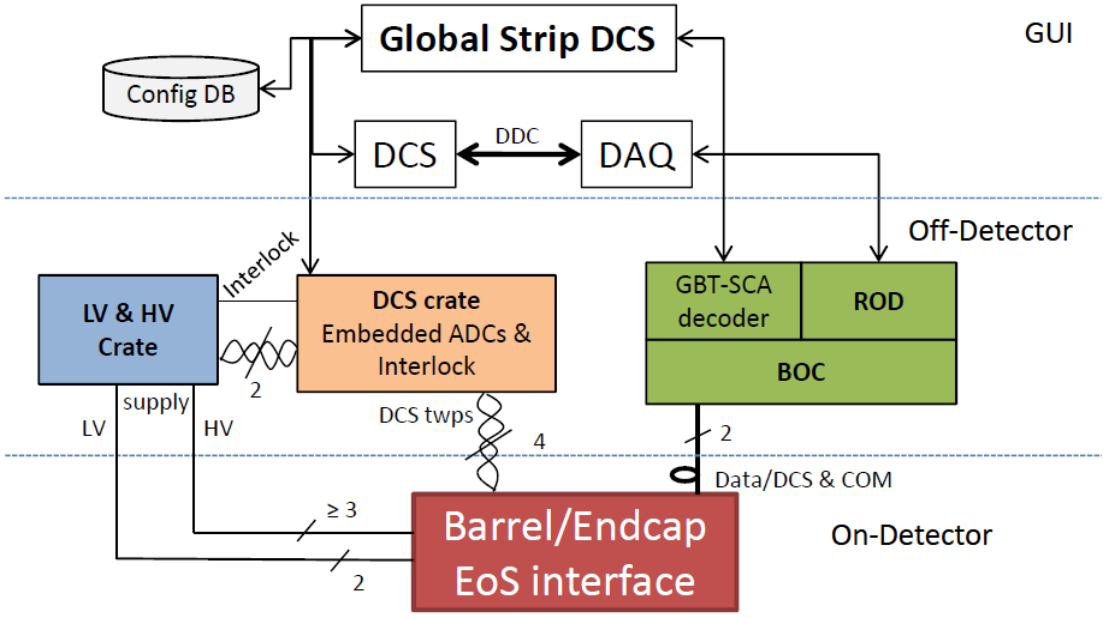

Figure 23: Proposed strip DCS architecture connecting the On-detector, the Off-detector and the software user interfaces. The DCS cables are maximum 4 twisted pairs per electrical substructure side.

5.2 The detector safety and interlocks

The estimated nominal power dissipation of the ID Upgrade (pixels plus strips) is 80 kW rising to possibly 180 kW after irradiation and including safety factors [19] meaning that the cooling is a key aspect of the detector operation. Given the heat density and the silicon sensor runaway risk, the module design requires that the thermal path be efficiently coupled to the cooling tube. For detector safety three main risks are therefore identified:

- Cooling failure: many reasons could be anticipated like the main cooling plant failure, a local leak, or even a pipe blockage.
- A broken thermal path like delamination of a component along the thermal path, which may happen either due to aging or to induced stress into glue interfaces.
- Thermal runaway: given that the silicon bulk current significantly increases with radiation damage (around 6 orders of magnitude), if the temperature is not below a safe

limit, the silicon sensor will make an uncontrolled and damaging current and temperature excursion.

The proposed safety action as currently implemented for the existing SCT detector [20] and concerning the first cooling risk listed above is to cut the LV and HV at the power supply crate only for the problematic loop. This action will require highly reliable equipment, which will be independent of the powering and the data acquisition stage. In addition each cooling loop will have a redundant temperature sensor for the interlock and the global detector survey.

The safety action for the other failure risks, either a broken thermal path or silicon thermal runaway, will be to locally disable the LV and HV power supply along the detector substructure. Only the problematic module will be affected by this action while the neighboring modules will not be affected and remain operational. As soon as the LV and HV power will be disabled the local temperature will decrease. The operator will then have to diagnose carefully the problematic module to identify the source of the problem. The temperature information will be based on NTC (Negative Temperature Coefficient) thermistors that will be read out and surveyed for the interlock logic by either the HCC or the GBT-SCA as described on Figure 24. The interlock trip level will be based on a programmable temperature set by a DAC. The interlock function could be disabled by the operator in some use cases and especially if a failure is identified. The redundancy of the module interlock will be ensured by the second hybrid of a module side. The use of this type of interlock system requires that the ASIC (HCC or GBT-SCA) be powered independently of the FE. The interlock block will have to be SEU protected and all the registers that may affect the decision level will have to be triplicated.

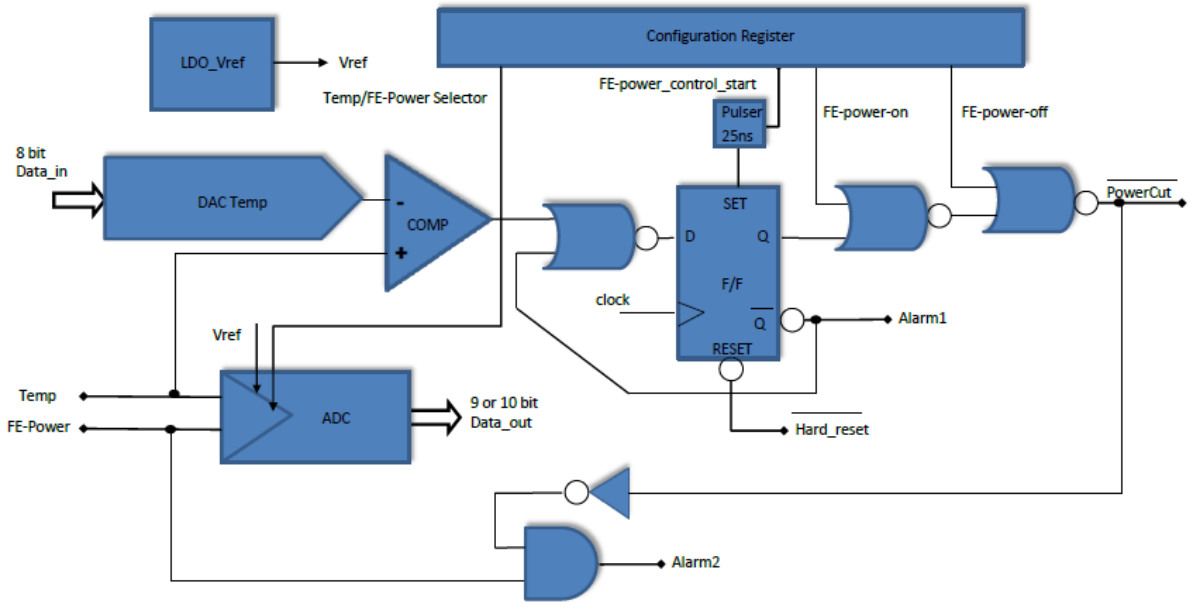

Figure 24: Proposal of a DCS block diagram to be integrated into the HCC for temperature monitoring and FE power control

5.3 Control

While the safety interlock system is defined to protect the detector against failures, the detector control and commands are functionalities that are provided to the operators allowing switching ON and OFF global or local powering stages. The powering units are mainly divided into two

groups: the on-detector integrated devices and the off-detector power cards and crates. Most of the on-detector power will be consumed by the FE readout chips for which the number is defined by the detector granularity. When all or even a single module is powered, it is compulsory that the cooling is up and running. The other powered items of the readout chain along the module substructure are the HCC (2 per hybrid side), the GBT chip and possibly the opto-transceiver components at the EOS. The EOS components and HCC could be powered with warm cooling or no cooling for a limited time while the FE chips require full cooling.

For the off-detector, even if the modularity of the power units are not defined yet it is anticipated that a LV card or channel will provide the power for one side of the substructure. The powering sequence is to first supply the EOS components allowing for optical communication and access to some DCS information. Then, once the opto-communication is operational, the other chips can be powered with first the HCC to allow accessing some register and DCS communication. The very last parts to be powered are the FE chips either via the serial powering chain or via the DC-DC converter stages. The operator will enable/disable or shunt the powering chain depending of the use cases.

For the HV providing separate lines to each module will require a large volume of cables. Therefore, a multiplexing system with HV switches is being considered to be installed either on the EOS card or elsewhere inside or outside the detector volume. (See Section 4.2 above.) Such a system is under R&D, and if it is qualified for operation in a magnetic field and high radiation environment, it would allow the same modularity for the LV and the HV. The operation of the HV switches will have to be controlled either via the GBT slow control chip or via another command line, which will have to be defined. In the powering sequence, the HV is usually the last part to switch ON before physics or calibration runs. In debug mode or when doing I-V characterization, the operator could operate with FE unpowered or even with limited or no cooling.

5.4 The monitoring stages

The monitoring of the power supply, the temperature and the relative humidity will be provided by three groups of data acquisition blocks: the direct environmental information via embedded ADCs like ELMB [21], the temperature reading via the GBT SCA and the DCS and diagnostics via the FE and HCC ICs to be decoded with the readout system.

The DCS crate is not defined yet in terms of hardware and modularity, but it is proposed to locate it in the counting room allowing better accessibility for hardware maintenance during the detector operation. As compared to the current SCT detector, the cable length will be of the order of 100 m instead of 30 m, which will require using shielded twisted pairs to minimize the parasitic noise. Up to 4 twisted pairs are proposed to be used per EOS electrical side: 1 for the cooling and interlock, 1 for the LV sense and 2 for local or environmental monitoring. The DCS crate will monitor the cooling and the LV and HV power supply data as well as holding the interlock logic.

Figure 25 illustrates the DCS information that will be encoded into the optical fiber with the data link. It is foreseen that the FE and the HCC chips will have some ADC channels allowing survey of the temperature and the power supply state. The information will be treated as registers for which a priority level could be set for an alarm flag to be transmitted with the data. The GBT-SCA may monitor additional parameters like temperature or powering state along the detector substructure. The SCA chip data will have as part of the 5 gigabit optical

fiber link a dedicated 40 Mbps bandwidth that will not be embedded and coupled with the FE and HCC data stream.

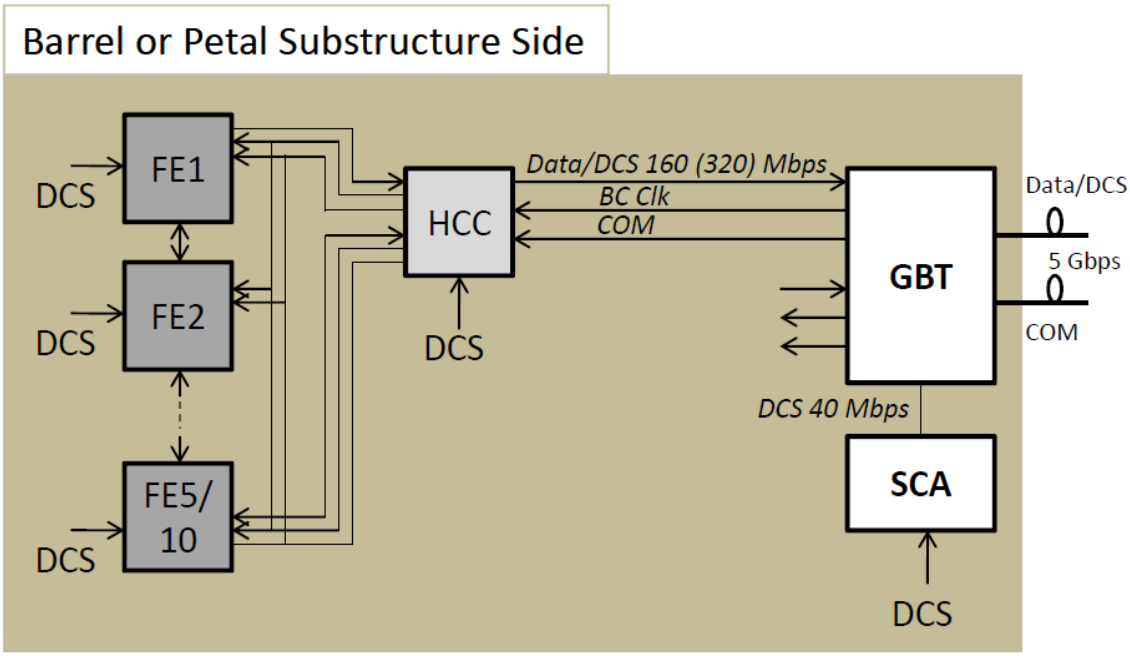

Figure 25: DCS inputs integrated into the readout chain

6. Development plans

A first version of ABC130 and HCC will be produced in the course of 2013 with the objective of validating some concepts, assessing the radiation hardness of the device and studying the power consumption. This version will be used on stave prototypes and the two options for power distribution will be extensively tested.

The final version of the chip will only be designed when key parameters of the upgraded system will be defined. The length of the long strips will not have any effect at the chip level. The need for a level-1 track trigger and the selection of the way of doing it (L0/L1 scheme or self-seeded scheme) will have a major impact on the chip. Similarly the expected L0 and L1 rates will affect the ABC130, the HCC and the EOSC. We now assume an L0 rate of 500 kHz, an L1 rate of 200 kHz and an L0-ROI coverage of 10% of the detector. Simulations to date indicate that these rates can be accommodated with the data transmission bandwidths specified, i.e. 160 Mbps out from each HCC, and $\leq$ 4.48 Gbps out from each EOSC. Further simulation work is needed to confirm that these rates are sufficient. If they prove not to be or if higher trigger rates are envisioned, a complete rethinking of the readout stages of the ABC130 and the HCC will be required as well as a faster version of the EOSC chip set (e.g. a 65 nm version of the GBT) or a duplication of them. The effect on the power would be dramatic. The effect on the overall schedule is difficult to predict but will clearly endanger an already over-optimistic schedule.

References

[1] Layout Requirements and Options for a new Inner Tracker for the ATLAS Upgrade, N. Hessey & J. Tseng, EDMS ATL-P-EP-0001.

[2] Trigger & DAQ Interfaces with Front-end Systems: Requirement Document Version 2.0, ATLAS Internal Note DAQ-NO-103.

[3] Single-Event Upsets in Photodiodes for Multi-Gb/s Data Transmission, J. Troska et al., Proceedings of TWEPP08 page 161.

[4] The SPi chip as an integrated power management device for serial powering of future HEP experiments, M. Trimpl et al, in Proceedings of Vertex 2009, PoS(VERTEX 2009)030.

[5] GBTX specifications, P. Moreira etal., See https://espace.cern.ch/GBT-Project/GBTX/Specifications/Forms/AllItems.aspx.

[6] The versatile link, a common project for super-LHC, L. Amaral et al. 2009, JINST 4 P12003.

[7] VCSEL reliability in ATLAS and development of robust arrays, A R Weidberg etal., 2012 JINST 7 C01098.

[8] Optical Links for the ITK, A. Weidberg, 2012.

[9] GBLD, https://espace.cern.ch/GBT-Project/GBLD/_layouts/viewlsts.aspx?BaseType=1.

[10] The ABCN front-end chip for the ATLAS inner detector upgrade, J. Kaplon, Proceedings of TWEPP08 page 116.

[11] Presentation of Mitch Newcomer during the Amsterdam Atlas tracker upgrade workshop. See http://indico.cern.ch/sessionDisplay.py?sessionId=8&slotId=0&confId=32084#2008-11-04.

[12] Architecture of the Readout Electronics for the ATLAS upgraded silicon strips detector. Interim report of the Strip Readout Working Group, https://edms.cern.ch/document/1070291/1.

[13] Development of custom radiation-tolerant DCDC converter ASICs, F Faccio et al, 2010 JINST 5 C11016.

[14] An integrated DC-DC step-up charge pump and step-down converter in 130 nm technology, M. Bochenek, Proceedings of TWEPP09 page 579.

[15] ATLAS strip tracker stavelets, P W Phillips, 2012 JINST 7 C02028.

[16] DC-DC converters with reduced mass for trackers at the HL-LHC, G. Blanchot et al., 2011_JINST_6_C11035.

[17] Possible Approaches to HV Distribution to Atlas Strip Staves, D. Lynn, https://indico.cern.ch/getFile.py/access?contribId=218&sessionId=46&resId=1&materialId=slides&confId=108365.

[18] D. Ferrère, S.Kersten, Nuclear Instruments and Methods in Physics Research A 636 (2011) S164–S167.

[19] G. Vihhauser, Atlas Upgrade ID cooling system requirements, ATU-SYS-ES-0004, 2009.

[20] The detector control system of the ATLAS SemiConductor Tracker during macro-assembly and integration, A Abdesselam et al, 2008 JINST 3 P02007.

[21] The common infrastructure control of the ATLAS experiment, H.J.Burckhart, et al., Proceedings TWEPP08 page 428.