INFN/AExxx, LAL-xxx, SLAC-R-xxx

Super$B$ Detector Technical Design Report

Abstract

This report describes the technical design detector for Super$B$.

Contents

11 Electronics, Trigger, Data Acquisition and Online

Breton/Marconi/Luitz Pages : 7-10

11.1 Open Issues for Pisa Meeting

- Finalize agreements on clock distribution and use of mezzanines
- Review of the WBS for ETD and Online. Make sure subdetectors will update their cost and labor estimates.
- Discussion of the SEU rates and "system architecture for error handling" and recovery from loss-of-(c)lock in FEE
- Review and discussion of the power supply section
- Remaining R&D
- Discuss ETD/Online org structure (for TDR)
- Review of subdetector contributions to ELEX chapter

11.2 Architecture Overview

The Super*B* [1] Electronics, Trigger, Data acquisition and Online system (ETD) comprises the Level-1 trigger, the event data chain, and their support systems. Event data corresponding to accepted Level-1 triggers move from the Front-End Electronics (FEE) through the Read-Out Modules (ROMs), the network event builder, the High Level Trigger (HLT) to a data logging buffer where they are handed over to the offline for further processing and archival. ETD also encompasses the hardware and software components that control and monitor the detector and data acquisition systems, and perform near-real-time data quality monitoring and online calibration.

The system design takes into account the experience from running *BABAR* [2] and building the LHC experiments [3], [4], [5]. To minimize the complexity of the FEEs and the number of data links, the detector side of the system is synchronous, and the readout of all subdetectors is triggered by a fixed-latency first-level trigger. Custom hardware components (e.g. specialized data links) are only used where the requirements cannot be met by off-the-shelf commercialy available components (e.g. Ethernet). Radiation levels are significantly higher than in *BABAR*, making it mandatory to design radiation-tolerant on-detector electronics and links. Fig. 11.1 shows an overview of the trigger and the event data chain:

A first-level hardware trigger processes dedicated data streams of reduced primitives from the sub-detectors to provide trigger decisions to the Fast Control and Timing System (FCTS).

The FCTS is the central bandmaster of the system; it distributes the clock and readout commands to all elements of the architecture, and initiates the readout of the events in the detector FEE. In response to readout requests, the FEE send event fragments to the ROMs which perform a first stage internal event build and send the partially constructed events to the HLT farm where they are combined into complete events, and processed by the HLT which reduces the amount of data to be logged permanently by rejecting uninteresting events.

The trigger, data acquisition and support components of the ETD system are described

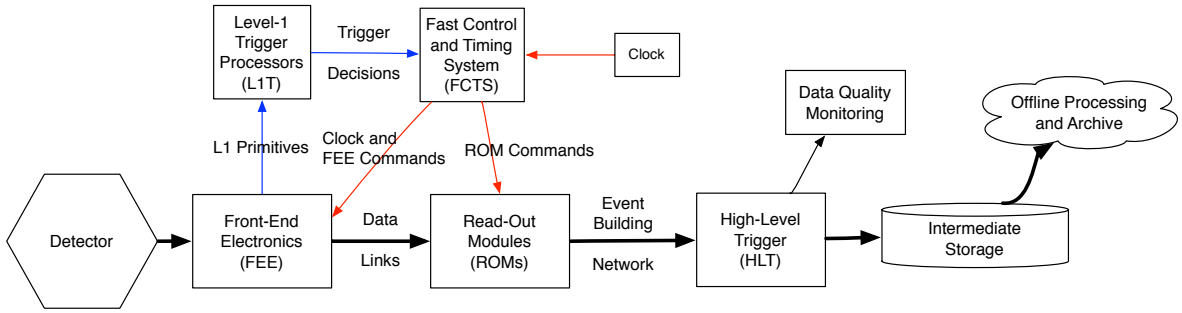

Figure 11.1: Overview of the Trigger and Data Chain

in this chapter, subdetector electronics, power supplies, grounding and shielding, and the cable plant are described in the next chapter.

11.2.1 Trigger Strategy

The *BABAR* and Belle [6] experiments both chose to use "open triggers" that preserved nearly 100% of $B\overline{B}$ events of all topologies, and a very large fraction of $\tau^+\tau^-$ and $c\bar{c}$ events. This choice enabled very broad physics programs at both experiments, albeit at the cost of a large number of events that needed to be logged and reconstructed, since it was so difficult to reliably separate the desired signals from the $q\bar{q}$ ($q = u, d, s$) continuum and from higher-mass two-photon physics at trigger level. The physics program envisioned for Super$B$ requires very high efficiencies for a wide variety of $B\overline{B}$, $\tau^+\tau^-$, and $c\bar{c}$ events, and depends on continuing the same strategy, since few classes of the relevant decays provide the kinds of clear signatures that allow the construction of specific triggers.

All levels of the trigger system are designed to permit the acquisition of prescaled samples of events that can be used to measure the trigger performance.

The trigger system consists of the following components [1]:

Level 1 (L1) Trigger: A synchronous, fully pipelined L1 trigger receives continuous data streams from the detector independently of the event readout and delivers readout decisions to the FCTS with a fixed latency. Like in *BABAR*, the Super$B$ L1 trigger operates on reduced-data streams from the drift chamber and the calorimeter.

High Level Triggers (HLT) — Level 3 (L3) and Level 4 (L4): The L3 trigger is a software filter that runs on a commodity computer farm and bases its decisions on specialized fast reconstruction of complete events. An additional "Level 4" filter may be implemented to reduce the volume of permanently recorded data if needed. Decisions by L4 would depend on a more complete event reconstruction and analysis. If the worst-case per-event performance of the L4 reconstruction algorithms does not meet the near-real-time requirements of L3, it might become necessary to decouple L4 from L3 – hence, its designation as a separate trigger stage.

11.2.2 Trigger Rate and Event Size Estimation

SL + UM

The Super$B$ L1-accept rate design standard of 150 kHz is based on an extrapolation from *BABAR* (see the Super$B$ CDR [1] for more detail). The *BABAR* Level-1 physics configuration produced a trigger of approximately 3 kHz at a luminosity of 10^{34} cm^{-2}sec^{-1}, however changes in background conditions produced large variations in this rate. The *BABAR* DAQ system performed well, with little dead time, up to approx-

[1] While at this time we do not foresee a "Level 2" trigger that acts on partial event information in the data path, the data acquisition system architecture would allow the addition of such a trigger stage at a later time, hence the nomenclature.

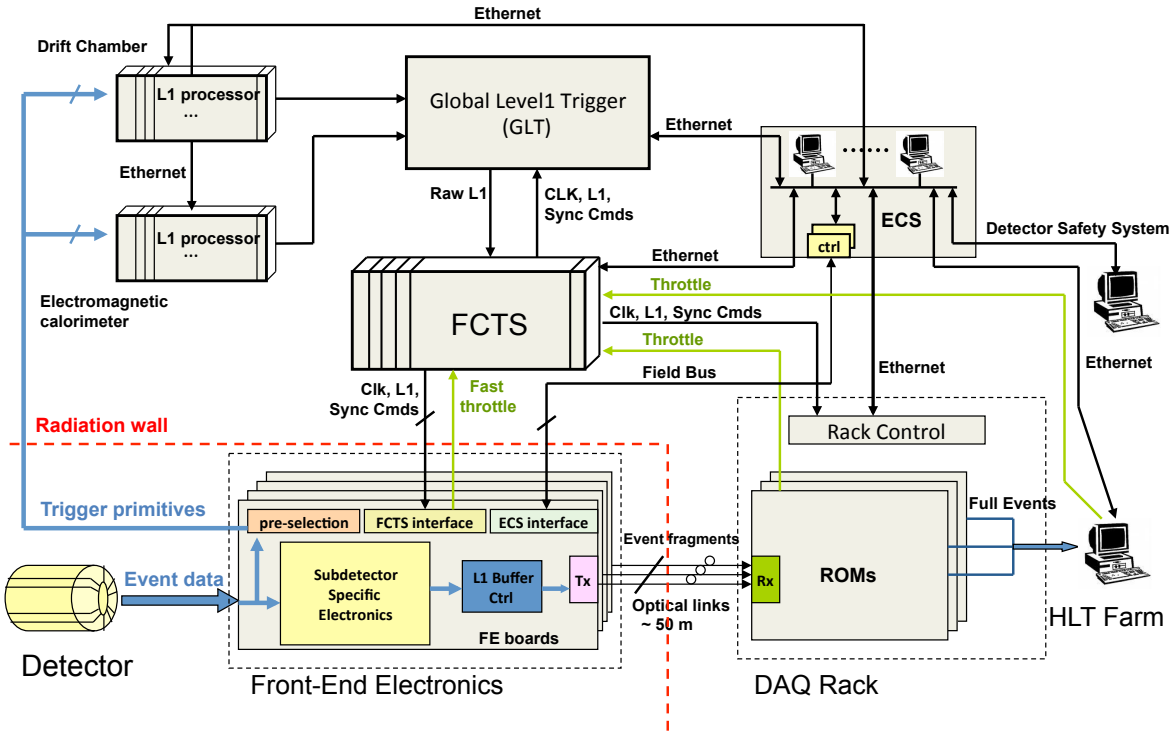

Figure 11.2: Overview of Level-1 Trigger, FCTS, FEE and ROMs.

imately 4.5 kHz. This 50% headroom was very valuable for maintaining stable and efficient operation and will be retained in Super$B$.

In *BABAR* the offline physics filter's output corresponded to a cross-section of approximately 20 nb and included a highly efficient Bhabha veto. We take this as an irreducible baseline for an open hardware trigger design; in fact this is somewhat optimistic since the offline filter used results from full event reconstruction. The accepted cross-section for Bhabhas in Super$B$ is approximately 50 nb (Note: to be checked). At the Super$B$ design luminosity of 10^{36} cm^{-2}sec^{-1} these two components add up to 70 kHz L1-accept rate without backgrounds. Scaling the beam backgrounds from *BABAR* obtain 30 kHz of background-related L1 accepts, resulting in a total of 100 kHz. Without a Bhabha veto at L1 (which is not part of the Super$B$ baseline trigger design) this is the minimum rate the readout system has to handle. By adding a *BABAR*-like reserve of 50% to both accommodate the possibility of higher backgrounds than design (e.g. during machine commissioning), and the possibility that the machine exceeds its initial design luminosity, we set the the Super$B$ design rate at 150 kHz.

The event size estimate still has uncertainties. Raw event sizes (between front-end electronics and ROMs) are understood well enough to determine the number of data links required. However, neither the algorithms for data size reduction (such as zero suppression or feature extraction) that can be safely employed in the HLT or possibly in the ROMs, nor their specific performance for event size reduction are yet known. Based on initial estimates for the SVT

and the calorimeter, we predict an event size of 500 kbyte before the HLT that can be reduced in the HLT to 200 kbyte for permanent logging.

11.2.3 Dead Time and Buffer Queue Depth Considerations

The readout system is required to handle the maximum average rate of 150 kHz, and to absorb the expected instantaneous rates, both without incurring dead time of more than 1% under normal operating conditions at design luminosity[2]. Dead time is generated and managed centrally in the FCTS based on feedback ("fast throttle") from the FEE: The FCTS drops valid L1 trigger requests when at least one FEE indicates that these would exceed its instantaneous or average rate capability. The ROMs and the event builder also provide feedback ("slow throttle") to slow down the trigger if they cannot keep up with the L1 accept rate.

The required ability to handle the maximum average rate determines the overall readout system bandwidth; the instantaneous trigger rate requirement affects the FCTS, the data extraction capabilities of the front-end-electronics, and the depth of the de-randomization buffers.

The minimum time interval between bunch crossings at design luminosity is about 2.1 ns—so short in comparison to detector data collection times that we assume "continuous beams" for the purposes of trigger and FCTS design: The inter-event time distribution is exponential and the instantaneous rate is only limited by the capability of the L1 trigger to separate events in time. Therefore, the burst handling capability of the system (i.e. the derandomization buffer size in the FEEs) is determined by the average L1 trigger rate, minimum time between L1 triggers, and the average data link occupancy between the FEEs and the ROMs. Figure 11.3 shows the result of a simple simulation of the FCTS and the FEE of a typical subdetector: The minimum buffer depth that is required to keep the event loss due to a full derandomization buffer below 1% is plotted as a function of the average data link utilization.

In Super*B*, we will target a link utilization of 90% which will require the derandomizer buffers to be able to hold a minimum of 10 events. It is very important to understand that there is no contingency in the target link utilization — all system-wide contingency is contained in the trigger rate headroom.

Additional derandomizer capacity is be required to absorb triggers generated while the fast throttle signal is propagated from the CFEE to the FCTS.

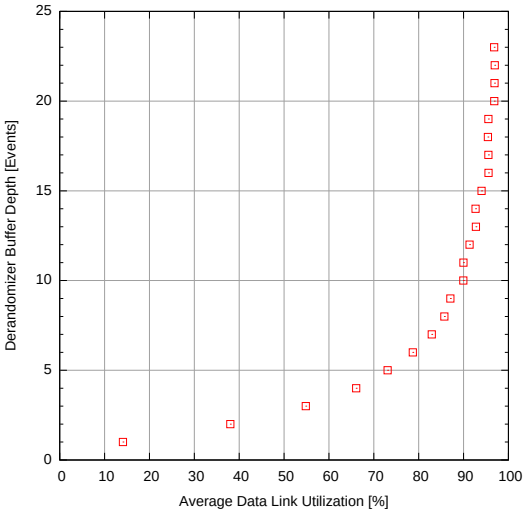

Figure 11.3: Minimum derandomizer buffer depth required to keep the event loss due to "derandomizer full" below 1% as a function of the average data link utilization.

11.3 Electronics in the SuperB Radiation Environment

DB + ?

The high luminosity of the machine and the presence of numerous massive elements close to the interaction region will generate much higher levels of radiation than in *BABAR*, where radiation effects on digital read-out electronics had

[2] The 1% dead time specification does not include events that are lost due to individual sub-detector's intrinsic dead times or the L1 trigger's limitations in separating events that are very close in time.

only been observed after the introduction of FPGAs on the end plate of the drift chamber.

In Super$B$ a large flux of charged particles and photons originating from the interaction point and the beam pipes will cross the detector, and generate large numbers of secondary neutrons as well.

Therefore a common general radiation policy has been put in place at the ETD level.

Long term radiation effects are of two types (ionizing and non-ionizing), while short term effects are linked to instantaneous ionization (Single Events). Radiation levels have been simulated, and Total Ionizing Doses (TID) range from 5 kGy down to a few Gy depending on the electronics location. The neutron fluence is required to estimate the effect of the Non-Ionizing Energy Loss (NIEL) and is still under study.

Shielding of sensitive parts of the detector like electronics is a key point of the design, which has been carefully studied. However, all electronics located on the detector have to be able to handle not only the damages linked to TID and NIEL, but also to present the smallest sensitivity to Single Event Upsets (SEUs), thanks to the intensive use of mitigation techniques like triple modular redundancy (TMR) for the latches and flip-flops and of safety codes (like parity bits) for data stored in memories. All the components used will be validated for their proven capacity to handle the integrated dose and NIEL foreseen. Power supplies are also designed specifically in order to perform safely in the radiative environment. The architecture of the system has been designed in such a way to reduce as much as possible the risk of failure, especially concerning the critical elements linked to experiment control and readout. In case of failure despite all these precautions, misfunctionning will be detected, experiment control system will be immediately warned, and a fast recovery strategy will be deployed in order to limit as much as possible the dead-time due to these radiation effects.

11.4 Trigger and Event Data Chain

The systems in the trigger and event data chain manage event selection, event filtering and end-to-end data flow from the detector to the intermediate buffer.

11.4.1 Choice of Global Clock Frequency

The Super$B$ trigger and data acquisition system requires a global clock that can be distributed to all components that operate sychronously, i.e. FCTS, the L1 trigger and the subdetector FEEs.

To allow features like selective blanking of triggers for a well-defined part of the revolution of the stored beams, it is critically important that the global clock is in well defined relation with the machine RF of 476 MHz and the revolution phase of the stored beams.

In *BABAR* the global clock was generated by dividing the RF by eight, yielding a 59.5 MHz clock that was distributed to all components of the experiment. In addition, a revolution ("fiducial") signal was fed to the FCTS, giving the system knowledge of the revolution phase of the stored beams.

For Super$B$ we will use the same parameters and distribute a 59.5 MHz clock throughout the ETD system.

Question: Would it be still possible to change this to divide by 12 instead and distribute the resulting 39 2/3 MHz clock. This would allow us to use many components and system designs that were developed and qualified for the 40MHz clock of the LHC experiments!

11.4.2 Level-1 Trigger

PB (+ SL?)

The baseline for the L1 trigger is to re-implement the *BABAR* L1 trigger with state-of-the-art technology. It will be a synchronous machine running at 59.5 MHz (or multiples of 59.5 MHz) that processes primitives produced by dedicated electronics located on the front-end boards or other dedicated boards of the respective sub-detector. The raw L1 decisions are sent to the FCTM boards which applies a throttle if necessary and then broadcasts them

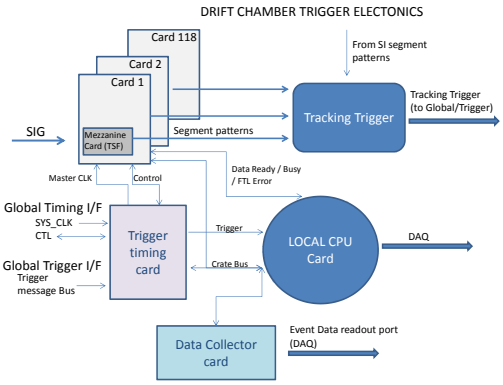

Figure 11.5: DCT Overview

to the FEEs and the ROMs. The standard chosen for the crates would most likely be either ATCA for Physics (Advanced Telecommunications Computing Architecture) for the crates and the backplanes, or a fully custom architecture.

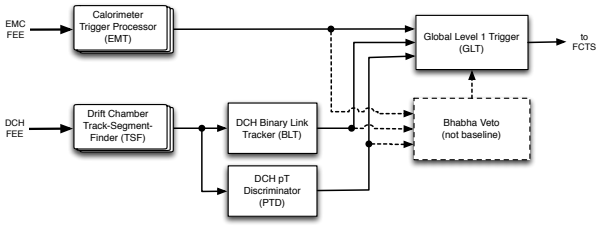

Figure 11.4: Level 1 Trigger Overview

The main elements of the L1 trigger are shown in Fig. 11.4 (see [8] for detailed descriptions of the *BABAR* trigger components):

Drift chamber trigger (DCT): The DCT consists of a track segment finder (TSF) , a binary link tracker (BLT) and a p_t discriminator (PTD). The DCT also extrapolates tracks to the IP; this allows to remove backgrounds that do not originate from the IP.

The layout of the DCT is shown in Fig. 11.5. Track finding is implemented in two stages: In the first stage, the algorithm locally searches for track segments exploiting the data available at the front-end level. The second stage links the track segments and searches for complete tracks.

Local Segment Finding: Two options are under discussion as DCH front-end (FE). In the first option time and charge are measured using FADCs and TDCs. A second option is instead based on a cluster counting technique. Granularity is 64 channels per FE in the first case and 16 channels per FE in the second. In the first case data delivered by the FE will be gathered by a single board on the FE crate and an optical link will connect the crate with the trigger crate. In the other option the DCH electronics has a modularity of 64 channels distributed on 4 contiguous radial planes. This geometry defines the super layer. Table 11.1 shows the super-layer composition for the DCH and the FE cards necessary for its data acquisition. The track segment finder is partially integrated in the DCH front-end.

The TSF will be implemented either at the crate level or at the mezzanine level; the number of optical links and TSFs will be 118 in the case of ADC/TDC-based DCH electronics or equal to the number of DCH front-end crates in the case of cluster counting. With cluster counting, data delivered by the FE will be gathered by a single board in the FE crate and sent to the trigger processor via an optical link.

Each card reads a set of super cell, and each signal will be properly stretched to accommodate at least one drift time. The TSF takes decisions at a programmable rate which can be as high as 59.5 MHz. A possible track segment is shown in Fig. 11.6.

A pattern needs 3 hits to be taken in consideration. The TSF delivers a bit stream to the following DCH stage, whose dimension depends on sampling frequency. This bit stream represents the ϕ of the segment in the super cell. These data are delivered to the trigger crate. To avoid efficiency losses a small number of contiguous channels is collected by all the neighboring cards. It is relevant in this scheme to optimize the sampling frequency as a function of track efficiency and exploited bandwidth.

	Sl1	SL2	SL3	SL4	SL5	SL6	SL7	SL8	SL9	SL10	TOT
Planes	4	4	4	4	4	4	4	4	4	4	
Type	A	A	U	V	U	V	U	V	A	A	
n. wires	736	864	496	560	624	688	752	816	896	960	7392
n. TSF	12	14	8	9	10	11	12	13	14	15	118
n. BCD	1	1	1	1	1	1	1	1	1	1	10

Table 11.1: DCH superlayer (SL) and wire readout

Transmission links: We will use high speed serial links to deliver trigger data to the trigger crate. Therefore signal aggregation delivered by the TSF is made possible, and a single card in the trigger crate has all the data from the pertaining superlayer.

Global Tracking: A simple and efficient tracking algorithm is the Binary Link Track Finder (BLT) developed within the CLEO collaboration. This method starts from superlayer (SL) 2 and moves radially outward combining the TSFs if in the following SL one of the three contiguous TSF is active. The track length varies from the first to the last super layer. It is possible to implement a majority logic function of machine background. Track charge does not affect the algorithm and it can be implemented on programmable logic. DCT primitives are received by the DCH L1 trigger box. In this crate the optical links exploited by the DCH trigger lines are collected in such a way that each board stores the information of a single SL. A bus with a switch topology interconnects the single boards where the SL information is present with the master where all the DC TSFs are present.

Particle Counter: The BLT outputs are used to count different tracks using a multiplicity logic exploiting also isolation criteria. To define a track we can use associative memory based techniques so that in one clock cycle a pattern can be identified. The very same ASIC can in principle be used to compute the transverse momentum and the perigee parameters of the track. Track counter takes in consideration 4 SL long tracks.

Transverse Momentum: Using TSF positions and their ϕ angle it is possible to measure track multiplicity above a pre-defined threshold and their perigee parameter. The data are delivered to the GLT which on the basis of the trigger tables asserts the trigger.

Electromagnetic Calorimeter Trigger (EMT): The EMT processes the trigger output from the calorimeter to find energy clusters. From the trigger point of view the electromagnetic calorimeter in the barrel is composed of 24 (8x3) CsI (tl doped) crystal towers.

The EMT layout is shown in Fig. 11.7.

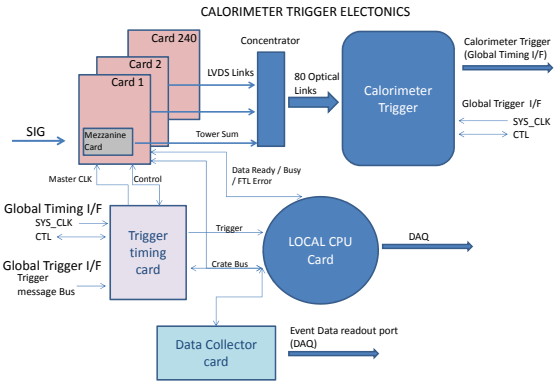

Figure 11.7: EMT Overview

The number of crystals in the barrel is 5760 this implies that the trigger box will handle 240 modules (trigger lines) in the barrel. At the moment we are thinking to divide the end-cap in 5x5 crystal tower this would imply a maximum of 60 towers. Two different strategies are under study. In the first strategy the analogic

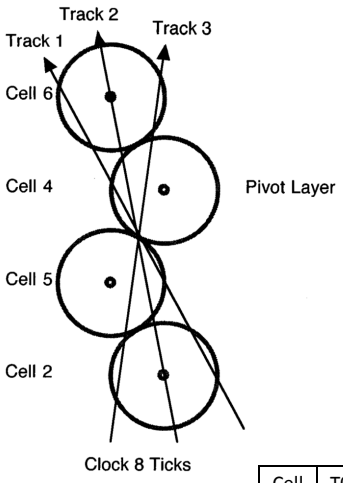

Cell	T0	1	2	3	4	T0	1	2	3	4	T0	1	2	3	4
6	0	0	0	1	2	0	1	2	3	0	0	0	0	1	2
4	0	0	0	1	2	0	0	0	1	2	0	0	1	2	3
5	0	0	0	1	2	0	0	0	1	2	0	0	1	2	3
2	0	0	0	1	2	0	1	2	3	0	0	0	0	1	2
	Track 1					Track 2					Track 3				

Cell	T0	1	2	3	4	5	6	7	8	T0	1	2	3	4	5	6	7	8	T0	1	2	3	4	5	6	7	8
6	0	0	0	0	0	1	2	3	4	0	1	2	3	4	5	6	7	0	0	0	0	0	0	1	2	3	4
4	0	0	0	0	0	1	2	3	4	0	0	0	0	0	1	2	3	4	0	0	0	1	2	3	4	5	6
5	0	0	0	0	0	1	2	3	4	0	0	0	0	0	0	1	2	3	0	0	0	1	2	3	4	5	6
2	0	0	0	0	0	1	2	3	4	0	1	2	3	4	5	6	7	0	0	0	0	0	0	1	2	3	4
	Track 1									Track 2									Track 3								

Figure 11.6: Track segment: *BABAR* vs Super*B* output bitstream. In this example we assumed to double babar samplig frequency.

sum is computed using the information of the same pre-amp used by the front-end. The other strategy foresees to endow each crystal with independent photon detectors (SiPm). The signal delivered by the SiPm is analogically added up at the tower level. The output can be calibrated using a look-up table and the sum can be encoded and delivered exploiting an LVDS link to an electro-optical converter shielded by radiation.

Particle Finders: Global cluster formation The energy can be released on several contiguous towers, therefore, as shown in Fig. 11.8 to assert the trigger signal we'll consider neighboring towers. Towers are scanned and their energy added if a release above a threshold of interest is exceeded in neighbouring towers.

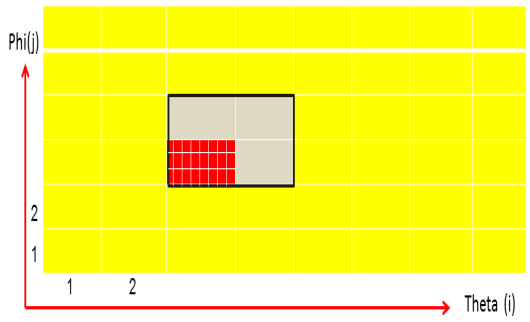

Figure 11.8: EMT trigger tower layout

As in Bbar, there will be 3 separate energy thresholds for clusters in the trigger: M cluster (above a low energy consistent with a minimum ionizing), G cluster (above 500-600 MeV) and cluster (electron Bhabha). The thresholds are all programmable. The Bhabha can be vetoed or pre-scaled. This algorithm will be implemented by the L1 EMT processor.

Particle Counters: Since clusters can be dynamically defined this allows to apply isolation cuts on particle definition.

L1 DCT and EMT trigger processors: A crate based on VME or ATCA technology will host L1 processors. These processors will be based on FPGA technology and have the same optical and electrical interfaces. They will differ by the firmware they run. Track finding and dynamic super-cluster definition will be implemented on the L1 processor backplane.

Global Trigger (GLT): The GLT processor combines the information from DCT and EMT (and possibly other inputs such as a Bhabha veto) and forms a final trigger decision that is sent to the FCTS. The global trigger receives the information of the detector participating to trigger definition through an optical link. The trigger can be asserted considering the information delivered by the DC and the EMC either separately or combined. This module combines the information of the detectors and compares them to pre-defined criteria. Through the use of an FPGA the trigger can be fully programmable and upgradeable.

We will study the applicability of this baseline design at Super$B$ luminosities and backgrounds, and will investigate improvements, such as adding a Bhabha veto to the L1 trigger. We will also study faster sampling of the DCH and the impact of the forward calorimeter design choice.

For the barrel EMC we will need to study how the L1 trigger time resolution can be improved and the trigger jitter can be reduced compared to BaBar. In general, improving the trigger event time precision should allow a reduction in readout window and raw event size. The L1 trigger may also be improved using larger FPGAs (e.g. by implementing tracking or clustering algorithm improvements, or by exploiting better readout granularity in the EMC).

L1 Trigger Latency: The *BABAR* L1 trigger had 12 μs latency. However, since the size, and cost, of the L1 data buffers in the sub-detectors scale directly with trigger latency, it should be substantially reduced, if possible. L1 trigger latencies of the much larger, more complex, ATLAS, CMS and LHCb experiments range between 2 and 4 μs, however these experiments only use fast detectors for triggering. Taking into consideration that the DCH adds an intrinsic dead time of about 1 μs and adding some latency reserve for future upgrades, we are currently estimating a total trigger latency of 6 μs (or less). More detailed engineering studies will be required to validate this estimate.

Monitoring the Trigger: To debug and monitor the trigger, and to provide cluster and track seed information to the higher trigger levels, trigger information supporting the trigger decisions is read out on a per-event basis through the regular readout system. In this respect, the low-level trigger acts like just another subdetector.

11.4.3 Fast Control and Timing System

DC + CB + DB + SL

The Fast Control and Timing System (FCTS) manages all elements linked to clock, trigger, and event readout, and is responsible for partitioning the detector into independent sub-systems for testing and commissioning. Fig. 11.9 shows how the FCTS is connected to the L1 trigger, FEE, ROMs and HLT.

The FCTS will be implemented in a crate where the backplane can be used to distribute all the necessary signals in point-to-point mode. This permits the delivery of very clean synchronous signals to all boards — avoiding the use of external cables. Fig. 11.10 shows the main functional blocks of the FCTS crate. The Fast Control and Timing Module (FCTM, shown in Fig. 11.11) provides the main functions of the FCTS:

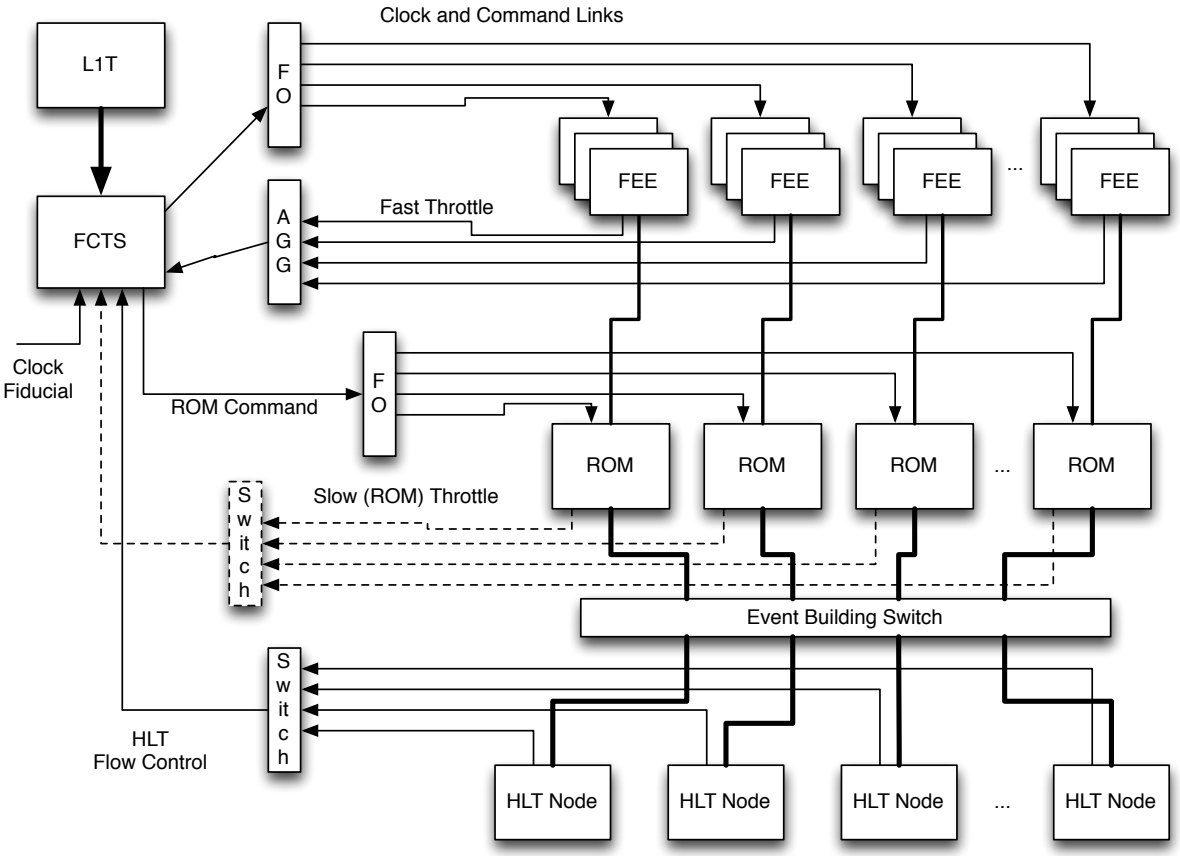

Figure 11.9: Overview of L1T/FCTS/ROM/HLT integration

Clock and Synchronization: The FCTS synchronizes the experiment with the machine and its bunch pattern, distributes the clock throughout the experiment, buffers the clock, and generates synchronous reset commands.

The baseline configuration is to distribute the clock through the FCTS command links[3]. The deserializers will extract the clock from the serialized bistream and report any loss of the clock signal or lock. Possible mechanisms to recover from such a condition are still under investigation and might include the capability of individual FEE to regain synchronization without major interruptions of the data taking.

To globally synchronize the system, all FEE must support the ability to reset all clock dividers, state machines and pipelines at system initialization and reset to ensure that divided versions of the global clock and counters are properly synchronized. This can be achived through a global reset command broadcast by the FCTS. Individual FEEs may be instructed by the ECS to resynchronize.

Trigger Handling and Throttling The FCTS receives the raw L1 trigger decisions, the fast throttle signals from the sub-detector FEE and slow throttle signals from the ROMs, generates readout commands and broadcasts them to the FEEs and the ROMs. The fast throttle is needed becauss during a spike in instantaneous rate or event size, data generated by the FEE can exceed the capacity of the data links. An FEE generates a throttle signal when the free space in the derandomizer buffer falls below a certain threshold (determined by the throttle la-

[3]Should reliability concerns arise during the final design stage, we will consider a separate, dedicated clock distribution network as a backup solution.

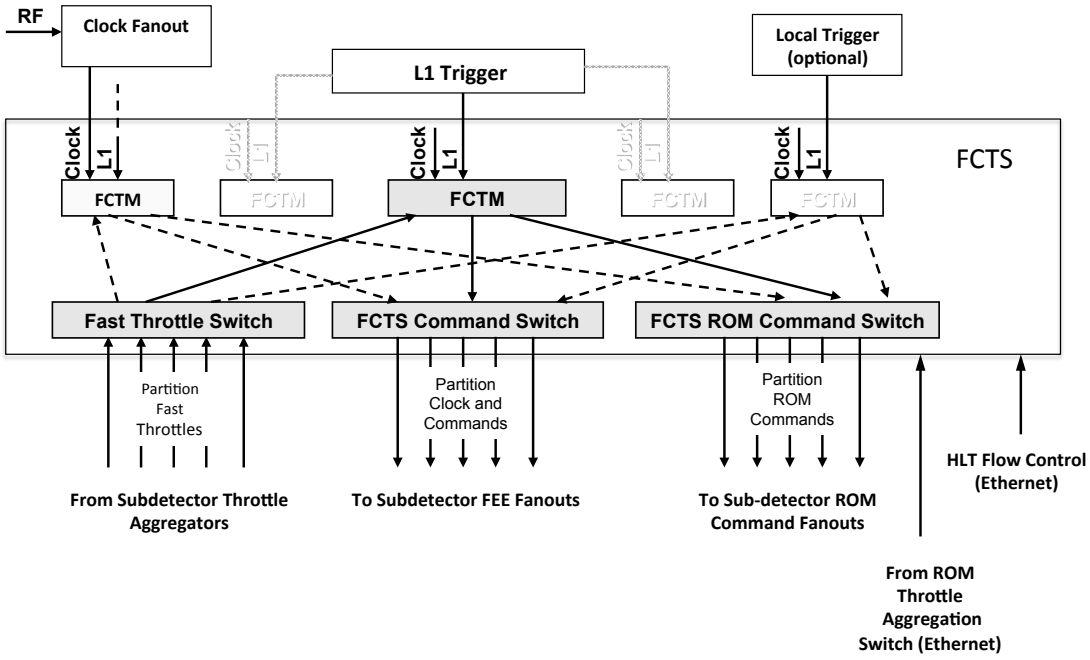

Figure 11.10: The main functional blocks in the FCTS crate

tency), Asserting a throttle signal temporarily inhibits the generation of Level-1 Accept commands in the FCTS and thus stops the readout of new events. Any L1-accepts that are already in-flight when a throttle is asserted will still need to be absorbed by the derandomizer buffer. In addition to being a rate limiting mechanism, the throttle can also be used by the FEEs as an "emergency stop" signal.

As a baseline, the fast throttle will be implemented as a 1-bit level without any framing or encoding. This minimizes the complexity of generating, aggregrating, distributing and monitoring the signals. It also allows to build a simple throttle aggregator device that can easily be standardized across all subsystems.

Every throttle line is monitored (fraction of time a throttle is signal is asserted and a counter of on/off transitions).

Detailed monitoring of the throttle-aggregation devices and detailed sub-detector specific diagnosis of the reasons for a throttle/stop can be performed through the ECS (see below).

Since there is no further local flow control mechanism on the data links, the ROMs cannot assert backpressure through the FEEs but need a separate path to slow down or stop the trigger when they are unable to offload the data through the event builder. With enough buffer space on the data link receiver cards, this throttle ("slow" throttle) does not have stringent latency requirements and can be implemented using Ethernet.

L1-accept Command Format and Distribution The commands broadcast to the FEE have a strict fixed-latency requirement and are as simple and short as possible to minimize the achievable temporal inter-command spacing during transmission and the amount of decod-

ing required in the FEE. We foresee a parity-protected 16-bit command word that in addition to the front end command code and parameters contains an event tag of at least 8 bits that allows the ROM to match the event fragment with its corresponding ROM command word. The FEEs are required to include the FEE command word with each event fragment they send to the ROMs.

The commands distributed to the ROMs are not subject to a fixed-latency requirement. In addition to a copy of the command word sent to the FEEs they contain at least an event timestamp (56 bits), the full trigger word (32 bits) and the HLT destination node (10-12 bits). Additional information such as a per-run event counter or the time since the last trickle injection pulse in the HER or LER might be included in the ROM command word as well. Since latency is variable, the ROM commands can be derandomized before transmission so that the ROM command links only need to sustain the average rate (150kHz) of ROM commands.

FEE and ROM commands are sent in the same order (parallel pipelines). Suitable ways of handling lost or corrupted FEE commands will have to be developed, however this will depend on the detailed failure modes of clock distribution and command links.

Event Management: The FCTS generates unique event identifiers, manages the assignment of events to nodes in the HLT farm and uses a per-node sliding window protocol to load-balance the HLT farm and to stop sending events to unresponsive HLT farm nodes. To do this, the FCTS manages a per-node counter of events it is still allowed to send to this node. A node's counter gets decremented with each event sent to that node. The FCTS determines the next destination node by searching for the next non-zero counter. Every HLT node asynchronously sends generic "event requests" with the number of events it is willing to take — when the FCTS receives such a request from a node it updates the corresponding counter with the value contained in the request. The ROM/HLT/FCTS protocol is described in more detail in the network event builder section below.

Super*B* events are large enough that in an Ethernet-based event builder we will not need to batch events into multi-event-packets (MEPs), but the system does not prclude the addition of an MEP scheme. Such a scheme might be required to send overlapping events to the same HLT node or to accommodate a non-Ethernet event building network that requires transmission of data in significantly larger units.

To generate an MEP, the FCTS would simply send the same HLT destination address for subsequent events which would then be packed into the same MEP (the packing can in fact be determined by the individual ROM). Sending a ROM command with a new HLT destination will then close the current MEP and open a new one.

Handling the event distribution with the help of the FCTS minimizes the intelligence required in the ROMs.

The FCTS also keeps track of all its activity, including an accounting of triggers lost due to FEE throttling or other sources of trigger inhibits.

Calibration and Commissioning: The FCTS can trigger the generation of calibration pulses and flexibly programmable local triggers for calibration and commissioning. To support the EMC source calibration where a meaningful global trigger can not be constructed, the system will be designd support an untriggered readout of individual channels.

The FCTS crate includes as many FCTM boards as required to cover all partitions. One FCTM will be dedicated to the unused subsystems in order to provide them with the clock and the minimum necessary commands.

Three dedicated switches are required in order to be able to partition the system into independent sub-systems or groups of sub-systems: One switch each for FEE clock and commands, ROM commands, and fast throttle feedback from the FEEs.

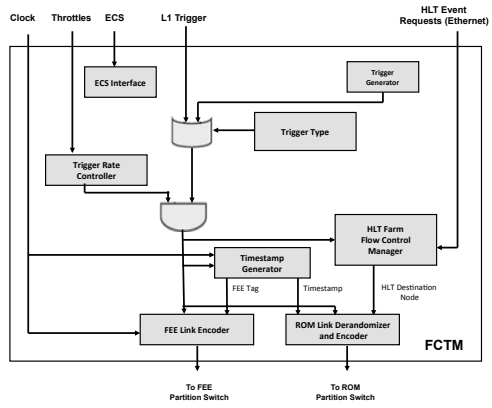

Figure 11.11: Fast Control and Timing Module

11.4.4 Control and Data Links

Requirements High-speed optical links will be used for data transfer, triggering and fast control distribution. The clock recovered from the serial streams will be used to synchronize the whole experiment.

Considerations on event size, average expected trigger rate, topology, and cost of the link components, suggest to choose a per-link throughput in the 1-3 Gbit/s range.

Super$B$ policy explicitly requires the use of commercial off-the-shelf (COTS) components where available, avoiding the design of any application specific integrated circuits (ASIC).

Since the whole ETD is designed to work synchronously with a frequency-divided machine clock, the latency and the recovered clock phase of serial links deployed for trigger and fast control distribution should be fixed and predictable, at a level below 1 unit interval. This requirement does not apply to read-out links, where the latency can be variable.

Moreover, the line coding should be DC-balanced in order to allow AC-coupled transmissions where needed and in such a way not to stress, and consequently not reduce the lifetime of, optical transmitters.

Radiation Tolerance Issues Because the high-speed optical links are a key element of the ETD, their electrical and optical components must withstand the expected radiation levels.

According to preliminary simulations of the luminosity-scaling background component [cite table from Riccardo], as far as the links are concerned, the worst case TID, neutron and hadron (E¿20 MeV) fluence are respectively 7.6 Gy (Si) and $3 \cdot 10^{11} cm^{-2}$and $7.6 \cdot 10^{10} cm^{-2}$ (SAFETY FACTOR INCLUDED?)[4].

An R&D activity has been carried out in order to identify the most reliable devices in terms of radiation tolerance among those compatible with the Super$B$ requirements. This activity included the irradiation testing (with 62-MeV protons) of SERDESes, both standalone (Texas Instruments TLK2711A [10], DS92LV18 [11] and DS92LV2421 -serializer only- [12]), and embedded in SRAM-based FPGAs (Xilinx Virtex-5 [13] and Virtex-6 [14]). The DS92LV18 has also been tested with gamma rays from a ^{60}Co source. For the optical part of the link, a VCSEL-based (Avago AFBR-57R5APZ [15]) transceiver has been tested.

The links implemented with SRAM-based FPGAs have also been protected with triple modular redundancy (TMR) and configuration scrubbing mitigation techniques. The layout of the TMR-protected version has been optimized with placement-hardening algorithms and different layouts have been tested (Fig. 11.12).

Before discussing the irradiation results, it is worth remarking that the optical links are a chain of different devices (SERDES, optical transceiver, fiber) and the failure of any component impacts the whole chain. In general the irradiation with 62-MeV protons generates both total ionizing dose and single event effects.

For the different devices the following dose effects were observed:

- **TLK2711A**: failed persistently (not recoverable by cycling power) after absorbing a dose of $\sim$ 400 Gy (Si).
- **DS92LV18**: did not fail either during testing with protons, or with gamma rays. For a total dose of 2.5 kGy(Si) the only effect

[4] We do not expect that links will be installed in the high-radiation regions that are closer to the beam than the innermost layer of the drift chamber

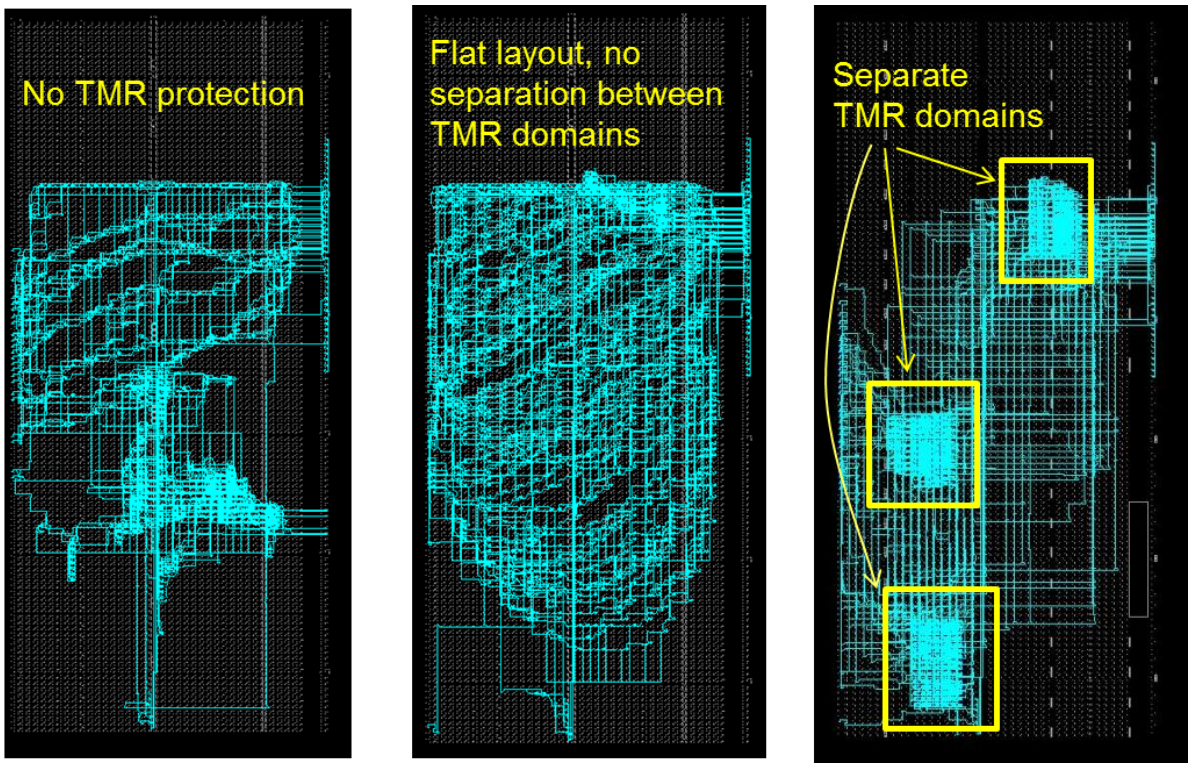

Figure 11.12: Layouts of the serial link implementation in a Xilinx FPGA V5LX50T. Left: link implementation without TMR protection. Center: TMR-protected version with flat layout. Right: TMR-protected version with separation of the TMR domains in distinct areas.

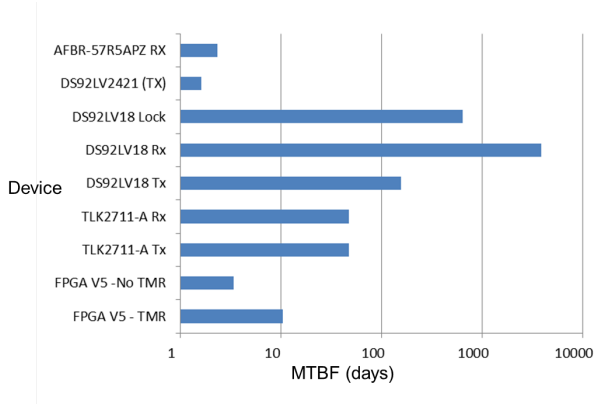

Figure 11.13: Expected failure rates at Super$B$ for a fluence of $3 \cdot 10^{11} cm^{-2}$ 62-MeV protons in one effective year (10^7s).

observed was a moderate ($\sim 0.5\%$) variation in power dissipation.

- **DS92LV2421/22**: did not fail persistently during the testing and did not show any measurable current variation due to TID.
- **Xilinx V5 and V6 FPGAs**: did not exhibit TID-related variations of the power dissipation, however, their power consumption increased gradually during the accumulation of SEUs. Reconfiguration of the devices restored the initial power consumption.
- **Avago AFBR-57R5APZ**: failed persistently at 400 Gy (Si).

Concerning transmission errors related to single event effects, Fig. 11.13 shows the expected

mean time between failures (MTBF) for each tested device, considering a reference fluence of $3 \cdot 10^{11} cm^{-2}$ 62-MeV protons per effective year (10^7 s). In the following, we use the word "failure" to refer to the following:

1. for a standalone SERDESes, to a single bit error either in the transmitter or in the receiver (for the DS92LV18 also to a loss of lock);
2. for a FPGA-embedded SERDES, a persistent failure of the link requiring the reconfiguration of the device.

The separate MTBFs for the transmitter and receiver parts of the TLK2711A and of the DS92LV18 are reported. For the DS92LV18, the rate of losses-of-lock is also shown. The Virtex-6 FPGA MTBF is not shown since it is lower than 10^{-1}days.

These results show that optical receivers are expected to be the main source of errors (we did not find any error in the transmitters during testing). For the reference fluence of $3 \cdot 10^{11} cm^{-2}$62-MeV protons, a MTBF of 2.3 days is expected for errors (either single or burst) in the optical receiver. Among the SERDESes, the most reliable are the DS92LV18 and the GTP embedded in V5 FPGAs, when protected with TMR (we do not consider the TLK2711, since it has a low tolerance to TID). Due to the line protocol of the DS92LV18, on average 10% of the incoming bit errors from the optical receiver will result in a loss-of-lock from the stream, also with the loss of the recovered clock. This means that the effective mean time between two losses of lock for the DS92LV18 will be lowered to nearly 23 days (from the value of 640 days related to the SERDES alone). On the other hand, the FPGA-based links do not lose the lock due to single or (short) bursts errors in the optical receiver and for the TMR-protected versions the failure rates are in the order of 10 days. The results in terms of MTBF suggest that link failures are unavoidable and therefore the whole ETD should be implemented as a fault-tolerant machine. This is further discussed in each ETD component subsection.

However, in the light of the present irradiation results the electrical part of the links should be based either on Virtex-5 FPGAs or on the DS92LV18. For both these devices the fixed-latency operation required by the trigger and fast control links is guaranteed. Although tested at a line rate of 2 Gbit/s, the links in V5 FPGAs can run up to 3.125 Gbit/s (limited by the embedded SERDES throughput) and the coding, including also error detection and correction (EDAC) schemes, can be completely customized by implementing the required logic in the fabric.

The DS92LV18 is limited at 1.32 Gbit/s and uses a proprietary encoding on the line, the payload can only be protected by means of external EDAC components. The line protocol of the device includes a 20-bit symbol consisting of a start bit, an 18-bit payload and a stop bit. When one of these bits is incorrect, the receiver loses the lock to the stream and the recovered clock stops toggling. Since the errors in the optical receiver are uniformly distributed on the parallel symbol, there is a 2/20=0.1 chance of hitting a start or a stop bit and losing the lock at the deserializer level.

Error Detection and Correction Two common solutions are error *correcting* codes (ECC) or the simpler error *detecting* codes (EDC). The ECC solution, for example using a Reed Solomon corrector, appears to be too intensive with regard to the computational effort of both encoding and decoding, resulting in increased hardware complexity, longer link latency and reduced reliability. The expected low error rate will permit a simpler and faster solution in the form of a proper EDC code. This will simply flag the occurrence of the error in the transmitted block, so that the block can just be discarded. As shown below, the expected frequency of this kind of error is extremely low, thus, bit errors will not affect data quality and event reconstruction.

As far as the choice of the error detection code is concerned, both cyclic redundancy check and check-sums have been evaluated. Since CRC

coding has a complexity that is comparable to a full ECC approach, it is not a viable option.

Checksums have a lower error detection ability than CRCs but have the advantage of a lower hardware complexity that meets the latency requirements and allow for a widespread use in the experiment, increasing the overall reliability. A parameter to be evaluated is the overhead in terms of ratio of the number of check-sum bits added to the useful message bits.

Let us suppose to deploy the DS92LV18 SERDES for the optical links. In this case, the 18 bits block managed as a whole by the SERDES will be buffered in groups of 3 blocks to be protected as a single 54 bit super-block. In order to achieve the best possible DC balance, the super-block will then be scrambled, checksum bits will be added and the data will be transmitted by the SERDES and the optical link. On the receiving end, the inverse process will be applied and possible errors will be flagged.

Tools have been developed in order to evaluate error detecting efficiencies of different checksums, such as *Parity*, *Modular Sum*, *One Complement Sum* and *Fletcher Sum*. The choice will be the adoption of Fletcher check-sum which provides detection of all 1 bit errors and two bits burst errors. Less frequent 2 or 3 bits errors are also detected even if not with the same 100% efficiency. Hardware complexity is limited to simple binary adders using the so called one's complement addition. Using the 54 bits block including a 6 bits check-sum, overhead is as low as 12.5%; error detection probability, evaluated by direct simulation mixed with proper analytical approach, provides the following figures for the detection efficiency: 1 bit errors 100%, 2 bits errors 98.7%, 3 bit errors 98.0%, 4 bit errors 98.1%; at the same time complete coverage is provided for burst errors encompassing 2-3-4 bits, an invaluable feature in our case foreseeing these kind of errors. The above figures, for a reference Bit Error Ratio BER=10^{-10} will deliver a probability of undetected error as low as $1.5 \cdot 10^{-19}$. This undetected error rate, in turn, for a reference 2.5 Gbit/s transmission speed, will deliver 8 years average time between undetected errors in a single link; for 1000 links this will deliver 3 days average time between undetected errors in the whole apparatus, or one over 1000 triggered events slightly degraded, which actually will be considered an acceptable degradation in the overall data quality.

Physical Implementation Serial links could be implemented on a plug-in mezzanine board. The main advantages of this approach will be :

- easier maintenance and possibility of upgrading the link performance (data rate, power, jitter) by simply replacing mezzanine cards, this guarantees protection against obsolescence of components too;
- a single design group will be in charge of characterizing performance and testing of the links;
- design parallelism, other ETD components might be designed by other groups independently in the meanwhile;
- decoupling of the link-related issues from the logic-related ones, critical link issues (latency, jitter) will be solved at the mezzanine level;
- design efficiency, the designer will focus on a 'critical but single' board, in fact a single printed circuit board (PCB) will be customized with different components to fit different needs (e.g. jitter cleaners or radhard SerDes for on detector nodes, copper and/or optical lanes where needed).

The implementation of serial links, as separate systems, by means of mezzanine cards is a widespread solution. Example of this trend are the timing trigger and control (TTC) [16] system of the LHC, the S-Link [17]deployed in several CERN experiments, the SODA [18] timing distribution system adopted in the PANDA experiment. In the TTC system, the implementation on a mezzanine allowed to upgrade the receiver with a VCXO-based PLL for improved output jitter ($\sim$ 20 ps rms), compatible with the

requirements of the reference clock of Gigabit/s SERDESes and other jitter sensitive devices.

Due to the limited space in the FCTS crate, the mezzanines for the FEE clock, command and throttle links will be housed in dedicated fanout crates that will also contain the logic to aggregate and monitor the throttle signals.

The ROM PCIe boards have similar space constraints, so a mezzanine solution for the data links will require the mezzanines and circuitry to convert the signals to e.g. LVDS to be housed in external crates.

Backup Solutions In order to optimize the design effort, control and data links should be based on the same components and mezzanine printed circuit boards. Should the tested devices found not be compatible with the radiation tolerance requirements of the experiment, the gigabit optical link (GOL) [19] transmitter could be considered as a backup solution for data links. The radiation-hard GBT [20] transceivers could also be considered, if available within a time scale compatible with the links realization phase.

11.4.5 Common Front-End Electronics

JM + DB

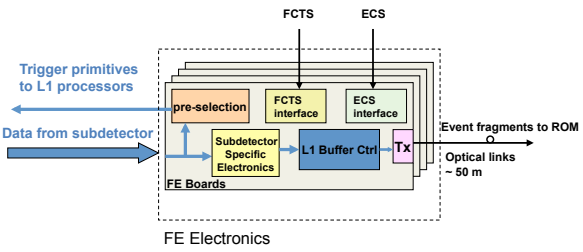

Figure 11.14: Common Front-End Electronics

In our opinion it will be beneficial to implement the different functions required to drive the front-end electronics as independent modular elements. These elements can for example be implemented as mezzanines or as circuits directly mounted on the front-end modules. For instance, as shown in Figure 11.14, one mezzanine can be used for FCTS signals and commands decoding, and one for the ECS management.

Depending on the implementation of the FEE, it can also be useful to decode the FCTS and ECS signals on one mezzanine and to distribute them to the neighbouring boards. This would permit a great reduction in the number of links. Driving the L1 buffers may also be implemented in a dedicated common control circuitry inside a radiation-tolerant FPGA. This circuitry would handle the L1 accept commands and provide the signals necessary to control the data transfer from the latency pipeline into the derandomizer, and finally, the transmission of event fragments to the serializer, as shown in Fig. 11.15 .

The latency buffers can be implemented either in the same FPGA or directly on the carrier boards, and one such single circuit could be able to drive numerous data links in parallel, thus reducing the amount of electronics on the front-end boards.

The control circuit will also handle events that have overlapping readout windows because they were triggered very close in time. The strategy for this is that data shared between events with overlapping readout windows is only read out and transmitted once. Complete events will then be reconstituted in the ROM by duplicating the data from the overlapping region as necessary.

It is also the responsibility of the CFEE to manage the derandomizer buffer and the fast multiplexer that feeds the data link serializer.

Figure 11.14 shows a possible implementation of the L1 buffers, their control electronics and the outputs towards the optical readout links. The control electronics may be located within a dedicated FPGA.

Another important requirement is that all (rad-tolerant) FPGAs in the FE have to be reprogrammable without dismounting a board. While this could be done through decdicated front panel connectors which might be linked to numerous FPGAs, the preferred solution is to program FPGAs through the ECS system without any manual intervention on the detector side.

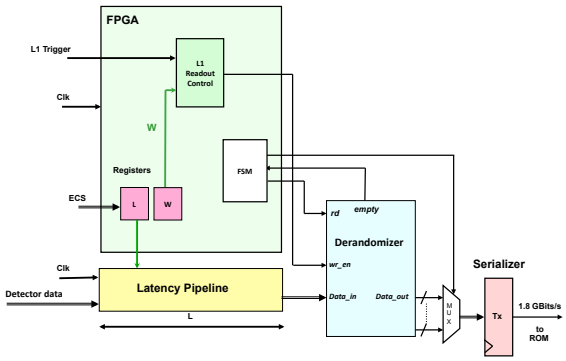

Figure 11.15: Control of latency pipeline and derandomizer in the CFEE

11.4.6 Read-Out Modules

MB + UM

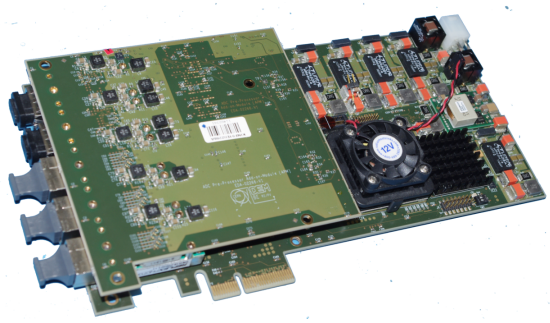

Figure 11.16: Readout Module

The Readout Modules (ROM, Fig. 11.16) are the hardware interfaces between front-end electronics and the software based event builder. ROMs map the synchronous world of FCTS-disciplined hierarchy to the asynchronous domain of the event-building computing farm. Their processing flow starts by receiving event fragments from the sub-detectors' front-end electronics and the corresponding ROM command words from the FCTS. Processing then continues with cross-checking front-end identifiers and absolute time-stamps, buffering them in de-randomizing memories, performing processing (still to be defined) on the fragment data, and eventually injecting the formatted fragment buffers into the event builder and HLT farm. Processing includes duplicating data into that have been partially transmitted because of overlapping readout windows.

Connected to the front-end electronics via optical fibers, they will be located in an easily accessible, no radiation area. The hardware design of the ROMs is identical for all sub-detectors, however, if the need arises, the pecularities of each sub-detector can be addressed by customization of the ROM firmware.

A ROM will have 10 or more optical input channels at a rate of 1 Gbit/s and at least one output channel at 10 Gbit/s. To exploit flexibility and processing power of computer technology, ROM will be manufactured as PCI Express based add-on cards for servers of the event-building farm. This solution will help keeping production costs low (when compared to a field-bus based solution) and opens a new scenario of a mixed hardware-software processing scheme for fragment data; unforeseen changes and upgrades can also be addressed with a low impact on the whole architecture. It is important to note that the space on a PCI express card is rather limited (both for connectors and mezzanines), so the strategy of using mezzanines and the mechanial design have to be reviewed carefully.

A prototype ROM has been manufactured and qualified, in order to assess the suitability of this approach. The prototype ROM is a PCI Express 2.0 based add-on card (Fig. 11.16) based on a Xilinx Virtex6-250T FPGA with 48 high speed serial transceivers. The plug-in module accommodates 3 SNAP12 optical receivers and 2 QSP optical transceivers for a total of 44 receivers and 8 transmitters each operating at max 6.2 Gbit/s. A 8 Mbyte true dual-port ram acts as de-randomizing memory for incoming fragments, while waiting to be streamed to host server memory.

The prototype ROM features a TDC and a ultra-low jitter PLL for the benefit of FPGA (Fig 11.17). A 4x PCIE 2.0 interface with an aggregate bandwidth of 20 Gbit/s per direction allows fragment data to be moved to host memory efficiently. To check worst-case power consumption, the FPGA has been configured with 38 in-

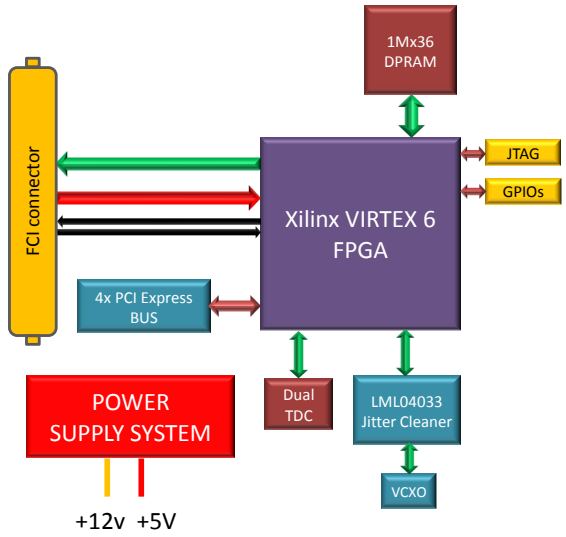

Figure 11.17: ROM block diagram – without optical adapter

put channels at 2.0 Gbit/s, 38 eight FIR filters for data processing, a 10 Gbit/s PCI-express interface, a scatter-gather DMA engine and a timing unit for a total of 92% occupancy of logic resources; consuming approximately 60W, no precaution has been taken for cooling when housed on a low cost 2U Dell server. By means of a custom developed Linux driver we have exercised the data transfer from the ROM dual-port RAM and the host memory; the driver addresses the custom designed scatter-gather DMA engine in the FPGA to efficiently move the fragment data from the ROM directly into user-space buffers; the driver avoids copying data from kernel address space to user address space by handling the map between Unix virtual memory pages and server physical memory regions, tunneling data from ROM dual-port RAM to user-space declared data vectors. Bandwidth tests show that fragment buffers larger than few tenths of Kbytes can be moved at a rate of approximately 950 Mbyte/s in the aforementioned configuration (4x PCIE 1.1), while almost doubling at 4x PCIE 2.0 (with a 256 bytes PCI-E payload). Test results of the prototype ROM outperform requirements in the foreseen deployment and this gives us confidence that this approach can be safely adopted.

11.4.7 Network Event Builder

SL

The ROMs receive event fragments in parallel from the sub-detector front-end-electronics, perform a first stage event-build, store the event fragments in memory buffers, and add the corresponding ROM command received from the FTCS on the dedicated ROM-command links. The resulting partially-built events are then encapsulated in UDP datagrams and sent to the HLT node determined by the destination node field in the ROM command word. Thus, all fragments of an event are sent to the same HLT node which completes the event building process by combining all fragements. Full events are passed to the HLT processes through a queue.

As described above, the FCTS determines the HLT destination node for every event using a round-robin algorithm that is tuned by a simple flow control scheme based on a per-node sliding window maintained by the FCTS. Each HLT node maintains a queue of fully built events that is filled by the event building process and drained by the HLT processes. As long as the event building process is running and the number of events in this queue stays below a high-water mark, the HLT node periodically sends requests for more events to the FCTM that is responsible for the partition. These requests are UDP packets sent over Ethernet; the aggregate rate of the requests can be limited to a fraction of the L1-accept rate, since each request is for multiple events.

If the HLT processes on a node cannot keep up with the incoming events, the node will stop sending requests and after the FCTM has exhausted the node's window of outstanding events, no more events will be directed to the node. Only after the FCTM receives a new event request from the node, it will consider it as a valid event destination again. No events are lost, since the node will queue the outstanding events. This provides a flow-control and load balancing mechanism for the HLT farm. In case of a complete HLT node failure (or the failure of

the event building process) the event loss is less or equal to the number of events last requested.

The event building process is inherently parallel and its rate can be scaled up as needed (up to the bisection bandwidth of the event building network). The baseline technology for the event builder network is standard 10 Gbit/s Ethernet. We will investigate the suitability of end-to-end flow control mechanisms (such as IEEE 802.1Qbb) at the Ethernet layer for avoiding packet loss in the event building network. We will also investigate alternative network technologies and protocols (such as RDMA or Infiniband).

With a L1 trigger rate of 150 kHz and a pre-HLT event size of 500 kbyte, the bandwidth in the network event builder is about 75 Gbyte/s, corresponding to about 750 Gbit/s with network overhead included. To avoid packet loss and to maintain stability, the event building network cannot be operated at an utilization of 100%, so we retain an additional safety factor of ~1.5. Thus, the minimum bisection bandwidth required in the event building network is 1200 Gbit/s. When implemented with 10 Gbit/s Ethernet, this means 120 "source" network interfaces on the ROMs and at least 120 "destination" interfaces on the HLT nodes. Network switches that provide the necessary bandwidth and can host at least $120 + 120 = 240$ 10 Gbit/s Ethernet ports are commercially available at the time of the writing of this document.

11.4.8 High-Level Trigger Farm

SL

The HLT farm needs to provide sufficient aggregate network bandwidth and CPU resources to handle the full Level 1 trigger rate on its input side. The Level 3 trigger algorithms should operate and log data entirely free of event time ordering constraints and be able to take full advantage of modern multi-core CPUs[5]. Extrapolating from *BABAR*, we expect 10 ms core time per event to be more than adequate to implement a software L3 filter, using specialized fast reconstruction algorithms. With such a filter, an output cross-section of 25 nb should be achievable. Using contemporary (as of Spring 2012) hardware and taking into account the CPU overhead for event building, logging and data transfer, a suitable farm could be implemented with approximately 150 nodes.

Level-4 Option To further reduce the amount of permanently stored data, an additional filter stage (L4) could be added that acts only on events accepted by the L3 filter. This L4 stage could be an equivalent (or extension) of the *BABAR* offline physics filter—rejecting events based either on partial or full event reconstruction. If the worst-case behavior of the L4 reconstruction code can be well controlled, it could be run in near real-time as part of, or directly after, the L3 stage. Otherwise, it may be necessary to use deep buffering to decouple the L4 filter from the near real-time performance requirements imposed at the L3 stage. The discussion in the Super*B* CDR [1] about risks and benefits of a L4 filter still applies.

11.4.9 Data Logging

SL + ?

The output of the HLT is logged to disk storage. We assume at least a few Tbyte of usable space per farm node, implemented either as directly attached low-cost disks in a redundant (RAID) configuration, or as a storage system connected through a network or SAN. We do not expect to aggregate data from multiple farm nodes into larger files. Instead, the individual files from the farm nodes will be maintained in the downstream system and the bookkeeping system and data handling procedures will have to deal with missing run contribution files. A switched Gigabit Ethernet network separate from the event builder network is used to transfer data asynchronously to archival storage and/or near-online farms for further processing. It is not yet decided where such facilities will be located, but network connectivity with adequate bandwidth and reliability will need to be provided. Enough local storage must be avail-

[5] The simplest implementation would be to run multiple identical HLT threads or processes which get their input from a single queue of built events.

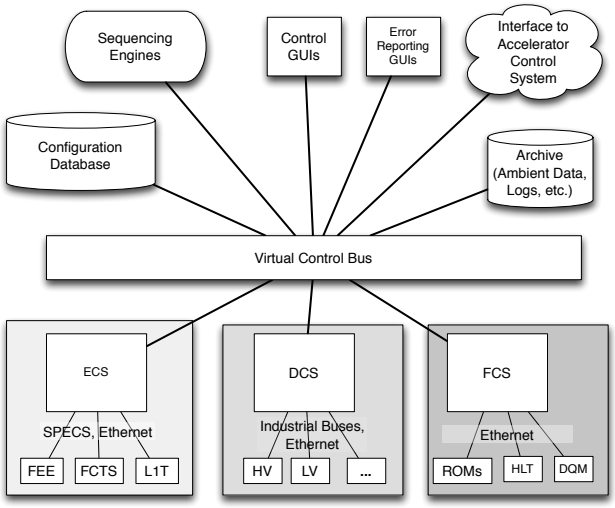

Figure 11.18: The Super*B* unified control system architecture

able to the HLT farm to allow data buffering for the expected periods of link down-time.

Assuming that the HLT accepts a cross-section of about 25 nb leads to an expected event rate of 25 kHz at a luminosity of 10^{36} cm^{-2}sec^{-1}, or a logging data rate of ~5 Gbyte/s.

While the format for the raw data has yet to be determined, many of the basic requirements are clear, such as efficient sequential writing, compact representation of the data, portability, long-term accessibility, and the freedom to tune file sizes to optimize storage system performance.

11.5 System Integration and Error Handling

AA, SC, SL, DC, UM, DB and others

Description of system integration aspects, error-tolerant designs, error detection, fast (downtime-free?) recovery, etc. Participation of multiple sub-systems (Clock?, FCTS, FEE, ROMs, HLT?) is required, so this becomes a separate section.

- evaluation of link failure rate in SuperB assuming 1000 links
- failure IS an option: analysis and recovery strategies
- need for a 'fault-tolerant' DAQ architecture capable to by design to recover from various fault types (mostly SEU-induced) in the FEE and the control and data links.

11.6 Control Systems

A major lesson learned during *BABAR* operations was that achieving high operational efficiency and true "factory-mode" data taking required a high degree of automation and control system integration. The traditional approach of separate control systems for different aspects of the experiment (detector control, run control, farm and logging control) designed and implemented with completely separate tools and very limited capability to communicate with each other, greatly limited the amount of automation and automatic error detection and recovery.

In Super*B*, all routine operations will need to be orchestrated across subsystem boundaries. For example, performing a simple calibration might require high voltages to be ramped to a calibration set-point, the FEE and the FCTS to be configured for calibration, farm nodes to be allocated and configured, the calibration run to be performed, calibration data to be analyzed and after completion the system to be returned to its normal data taking configuration.

For Super*B* we therefore foresee a unified control system that can automatically perform all routine operations. The system comprises the Electronics Control System (ECS), the Detector Control System (DCS), the Farm Control System (FCS), the configuration database, sequencing engines to implement distributed state machines, an archiving and logging system, operator GUIs and the interface to the accelerator control system. All components are connected by a central virtual control bus ("Global Control System").

All operations are driven by the configuration database and executed under the control of one or more sequencing engine.

A large part of the system, including the central virtual control bus, will be implemented using the !CHAOS [9] control system toolkit which is currently under development for the Super*B* accelerator. !CHAOS is a state-of-the-art scalable distributed control system framework that combines high-performance data acquisition and archiving capabilities with a plugin architecture to provide low-level controller interfaces and graphical user interfaces (GUIs).

An overview of the unified control system architecture is shown in Fig 11.18; its major components are described below.

11.6.1 Electronics Control System

SL The Electronics Control System (ECS) controls and monitors the FEE, the FCTS and the Level 1 trigger. Its main responsibilities are:

Configuring the Front-ends: Many front-end parameters must be initialized before the system can work correctly. The number of parameters per channel range from a only a few to large per-channel lookup tables. The ECS may also need to read back parameters from registers in the front-end hardware to check the status or verify that the contents have not changed. For a fast detector configuration and recovery turnaround in factory mode, it is critical to not have bottlenecks either in the ECS itself, or in the ECS' access to the front-end hardware.

Calibration: Calibration runs require extended functionality of the ECS: In a typical calibration run, after loading calibration parameters, event data collected with these parameters is sent through the DAQ system and analyzed. Then the ECS loads the parameters for the next calibration cycle into the front-ends and repeats the operation.

Testing the FEE: The ECS is also used to remotely test all FEE electronics modules using dedicated software. This obviates the need for independent self-test capability for all modules.

Monitoring the Experiment: The ECS continuously monitors FEE boards, FCTS and the L1 Trigger to ensure that they function properly. This might include independent spying on event data to verify data quality, and monitoring operational parameters on the boards (such as voltages, currents, temperatures and error flags). By monitoring these parameters, the ECS also participates in protecting the experiment from a variety of hazards[6].

ECS support must be built into all electronics modules that are to be controlled by the ECS – this includes the FEE.

Programming the FPGA firmware: Where applicable, the ECS can also be used to program or update FPGA configurations. It is highly desirable to be able to perform such operations without physical access the the boards. Depending on the FPGA this may require appropriate programming voltages to be made available on the boards.

The specific requirements that each of the sub-systems makes on ECS bandwidth and functionality must be determined (or at least estimated) as early as possible so that the ECS can be designed to incorporate them. Development of calibration, test, and monitoring routines must be considered an integral part of sub-system development, as it requires detailed knowledge about sub-system internals.

ECS Implementation: The field bus used for the ECS has to be radiation tolerant on the detector side and provide very high reliability. Such a bus has been designed for the LHCb experiment: it is called SPECS (Serial Protocol for Experiment Control System) [7]. It is a bidirectional 10 Mbit/s bus that runs over standard Ethernet Cat5+ cable and provides all possible facilities for ECS (like JTAG (Joint Test Action Group) and I2C (Inter IC)) on a small mezzanine. It could be easily adapted to the Super*B* requirements. Though SPECS was initially based on PCI boards, it is currently being translated to an Ethernet-based system, as part

[6] An independent hardware-based detector safety system, which is part of the DCS, must protect the experiment against equiement damage in case the software-based ECS is not operating correctly.

of an LHCb upgrade, also integrating all the functionalities for the out-of-detector elements. For the electronics located far from the detector, Ethernet will be used for ECS communication. The Super*B* ECS will be implemented using SPECS; an interface to !CHAOS will be developed.

11.6.2 Detector Control System

The Detector Control System (DCS) is responsible for ensuring detector safety, controlling the detector and detector support system, and monitoring and recording detector and environmental conditions. The DCS also provides the primary interface between the accelerator and the detector.

Efficient detector operations in factory mode require high levels of automation and automatic recovery from problems. Here, the DCS plays a key role and a tight integration with the Accelerator Control System (ACS) is hightly desirable. The DCS in conjunction with the ACS manages the accelerator-detector interlocks and beam and detector states and has access to all information from the accelerator and the detector that is needed to fully automate data taking operations.

The DCS-ACS connection is also used to provide the accelerator with beam measurements performed by the detector (such as beam spot positions or bunch-by-bunch luminosities). Operational experience from *BABAR* has shown that mutual access between machine and detector to their respective archived control system data (such as records of background levels, detector currents, trigger rates on the detector side and vacuum pressures, temperatures and stored currents on the machine side) are invaluable for improving the accelerator and detector performance.

The DCS will be implemented with !CHAOS. Due to its distributed nature and modular storage design, !CHAOS will allow us to federate the independent instances of DCS and ACS and provide a unified query interface for data archived by the respective systems.

Low-level components and interlocks responsible for detector safety (Detector Safety System, DSS) will be implemented as simple circuits or with programmable logic controllers (PLCs).

11.6.3 Farm Control System

SL

Processes on the ROMs and on the HLT farm will be started, controlled and monitored by the Farm Control System (FCS) and will be implemented using the !CHAOS framework or a traditional network inter-process communication system such as DIM.

11.7 Other Systems

11.7.1 Data Quality Monitoring System

SL

Event data quality monitoring is based on quantities calculated by the HLT, as well as quantities calculated by a more detailed analysis on a subset of the data. A distributed histogramming system collects the monitoring output histograms from all sources and makes them available to automatic monitoring processes and operator GUIs.

11.7.2 Other Components

SL

Electronic Logbook: A web-based logbook, integrated with all major Online components, allows operators to keep an ongoing log of the experiment's status, activities and changes.

Databases: Online databases such as configuration, conditions, and ambient databases are needed to track, respectively, the intended detector configuration, calibrations, and actual state and time-series information from the DCS.

Configuration Management: The configuration management system defines all hardware and software configuration parameters, and records them in a configuration database.

Performance Monitoring: The performance monitoring system monitors all components of the Online.

Software Release Management: Strict software release management is required, as is a tracking system that records the software version (including any patches) that was running at a given time in any part of the ETD/Online system. Release management must cover FPGAs and other firmware as well as software.

Computing Infrastructure Reliability: The Online computing infrastructure (including the specialized and general-purpose networks, file, database and application servers, operator consoles, and other workstations) must be designed to provide high availability, while being self-contained (sufficiently isolated and provided with firewalls) to minimize external dependencies and downtime.

11.7.3 Software Infrastructure

GM + SL

The data acquisition and online system is basically a distributed system built with commodity hardware components. Substantial manpower will be needed to design the software components—taking a homogeneous approach in both the design and implementation phases. An Online software infrastructure framework will help organize this major undertaking. It should provide basic memory management, communication services, and the executive processes to execute the Online applications. Specific Online applications will make use of these general services to simplify the performance of their functions. Middleware designed specifically for data acquisition exists, and may provide a simple, consistent, and integrated distributed programming environment.

11.8 R&D for Electronics, Trigger and Data Acquisition and Online

All

Note: This is copied from the Whitepaper. Needs updating!

The baseline design presented in this chapter can be implemented with technology and components available at the time time of writing of this document. However, we expect that by the times when we have to freeze various aspects of the design to start construction or purchasing, components that are significantly more performant and/or cost effective will be available. In order to take advantage of these developments, we will need to develop a detailed plan on when we have to finalize the parts of our design.

Data Links: Discuss technology tracking (e.g. LHC Upgrade) here?

Event Builder and HLT Farm: The main R&D topics for the Event Builder and HLT Farm are (1) the applicability of existing tools and frameworks for constructing the event builder; (2) the HLT farm framework; and (3), event building protocols and how they map onto network hardware.

Software Infrastructure: To provide the most efficient use of resources, it is important to investigate how much of the software infrastructure, frameworks and code implementation can be shared with Offline computing. This requires us to determine the level of reliability-engineering required in such a shared approach. We also must develop frameworks to take advantage of multi-core CPUs.

11.9 Organizational Structure of Electronics, Trigger, Data Acquisition and Online

DB + UM + SL

11.10 Conclusions

The architecture of the ETD system for Super*B* is optimized for simplicity and reliability at the lowest possible cost. It builds on substantial in-depth experience with the *BABAR* experiment, as well as more recent developments derived from building and commissioning the

LHC experiments. The proposed system is simple and safe. Trigger and data readout are fully synchronous—allowing them to be easily understood and commissioned. Safety margins are specifically included in all designs to deal with uncertainties in backgrounds and radiation levels. Event readout and event building are centrally supervised by a FCTS system which continuously collects all the information necessary to optimize the trigger rate. The hardware trigger design philosophy is similar to that of *BABAR* but with better efficiency and smaller latency. The event size remains modest.

The Online design philosophy is similar—leveraging existing experience, technology, and toolkits developed by *BABAR*, the LHC experiments, and commercial off-the-shelf computing and networking components—leading to a simple and operationally efficient system to serve the needs of Super*B* factory-mode data taking.

Bibliography

[1] M. Bona *et al.*, *Super*B*: A High-Luminosity Heavy Flavour Factory. Conceptual Design Report*, arXiv:0709.0451v2 [hep-ex], INFN/AE-07/2, SLAC-R-856, LAL 07-15, also available at http://www.pi.infn.it/SuperB/CDR.

[2] B. Aubert *et al.* (*BABAR* Collaboration), *The BABAR Detector*, Nucl. Instrum. Methods Phys. Res., Sect. A **479**, 1 (2002) [arXiv:hep-ex/0105044].

[3] The ATLAS Collaboration, *ATLAS Detector and Physics Performance Technical Design Report*, http://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/TDR/access.html.

[4] The CMS Collaboration, *CMS Detector Technical Design Report*, http://cmsdoc.cern.ch/cms/cpt/tdr/.

[5] The LHCB Collaboration, *LHCB Technical Design Reports*, http://cmsdoc.cern.ch/cms/cpt/tdr/.

[6] The Belle Collaboration, *The Belle Detector*, Nucl. Instrum. Methods Phys. Res., Sect. A **479**, 117 (2002).

[7] The SPECS Web Page, https://lhcb.lal.in2p3.fr/Specs/.

[8] The *BABAR* Trigger Web Pages, http://www.slac.stanford.edu/BFROOT/www/Detector/Trigger/index.html.

[9] L. Catani *et al. Exploring a new paradigm for accelerators and large experimental apparatus control systems*, ICALECPS 2011, http://accelconf.web.cern.ch/accelconf/icalepcs2011/html/clas013.htm

[10] Texas Instruments, *1.6 to 2.7* Gbit/s *Transceiver*, September 2009, http://www.ti.com/lit/ds/symlink/tlk2711a.pdf

[11] Texas Instruments, *DS92LV18 18-Bit Bus LVDS Serializer/Deserializer - 15-66* MHz, Literature Number: SNLS156D, June 2006, http://www.ti.com/lit/ds/symlink/ds92lv18.pdf

[12] Texas Instruments, *DS92LV2421/DS92LV2422 10 to 75* MHz*, 24-bit Channel Link II Serializer and Deserializer*, January 2011, http://www.ti.com/lit/ds/snls321a/snls321a.pdf

[13] Xilinx, *Virtex-5 FPGA RocketIO GTP Transceiver User Guide*, Xilinx UG196, v1.7, 2008, http://www.xilinx.com/support/documentation/user_guides/ug196.pdf

[14] Xilinx, *Virtex-6 FPGA GTX Transceivers User Guide*, Xilinx UG366, v2.6, 2011, http://www.xilinx.com/support/documentation/user_guides/ug366.pdf

[15] Avago Technologies, *AFBR-57R5APZ Data Sheet*, 2007, http://www.avagotech.com/docs/AV02-0881EN

[16] P. Moreira, E. Murer, *TTCrx and QPLL Mezzanine Card (TTCrq)*, 2003, http://bonner-ntserver.rice.edu/cms/TTCrqSpec.pdf

[17] O. Boyle 1 , R. McLaren & Erik van der Bij, *The S-LINK Interface Specification*, 1997, http://vanderby.web.cern.ch/vanderby/s-link/spec/spec/s-link.pdf

[18] I. Konorov; H. Angerer; A. Mann; S. Paul;, *SODA: Time distribution system for the PANDA experiment*, Nuclear Science Symposium Conference Record (NSS/MIC), 2009 IEEE , vol., no., pp.1863-1865, Oct. 24 2009-Nov. 1 2009

[19] P. Moreira et al., *A radiation tolerant gigabit serializer for LHC data transmission*, in Proc. of the 7th Workshop on Electronics for LHC Experiments, Stockholm, Sweden, 10 - 14 Sep 2001, pp.145-149

[20] P. Moreira, J. Troska, *Radiation-Hard Optical Link for Experiments*, PH Faculty Meeting April 25, 2008 CERN, Geneva, Switzerland, http://ph-dep.web.cern.ch/ph-dep/NewsMeetings/FM/FM25Apr08/Troska_Moreira.pdf

12 Subdetector Electronics and Infrastructure

Breton/Marconi/Luitz

12.1 Subsystem-specific Electronics

12.1.1 SVT Electronics

The full details of the SVT electronics has been given in chap. **??**. Here we recall the main features relevant for the data collection, trigger distribution, system programming and monitoring.

In the baseline SVT option, the detector will be equipped with double sided silicon strips (or striplets) in all layers. Custom front-end chips will be used to read the detector and transform as soon as possible the analog information of a particle traversing the detector in a digital one, characterized by position (layer, strip), energy deposit (signal time over threshold) and time. At least two different custom chips will be developed to handle separately the first layers (expecially layer 0), characterized by high occupancy and short strips (good signal/noise) and the external layers (layer 4 and layer 5) where the main concern is the length of the strips (worse signal/noise) and not the occupancy. Both front-end chips will have the same digital readout architecture and will present the same interface to the DAQ chain, although with different settings. Possible upgrades of the SVT internal layers might require to move to a pixellated detector, for which a different custom front-end chip will be designed. The digital architecture of the pixel chips, partially already develloped, will be based on the same general readout architecture so that from the point of view of trigger and DAQ system they will share the same interface.

Common characteristics of the chips will be capability to work both in data-push and data-pull mode, the presence of internal buffers to allow for a trigger latency up to 10 μs, the use of a periodic signal to time-tag the recorded hits, the serialized hit output and the chip programmability via (two) digital lines.

The full SVT data chain will therefore be able to provide and distribute all the signals, clocks, triggers, time-tagging signal to all the front-end chips in a system-wide synchronous way.

A sketch of the full data chain is given in Fig. 12.1. Starting from the detector and going to the ROM boards, the chain contains: a) the front-end chips mounted on an HDI placed immediately at the end of the sensor modules; b) wire connections to a transition card (signals in both directions and power lines); c) a transition card, placed about 50 cm from the end of the sensors, hosting the wire-to-optical conversion; d) a bidirectional optical line running above 1 Gbit/s; e) a receiver programmable board (front-end board: FEB).

Each HDI will host from 5 to 14 FE chips servicing a side of the silicon strip module (a so called ROS: Read-Out Section). All the chips will share the same input lines (currently: reset, clock, fastclock, trigger, time-stamp, registerIn) and at least a registerOut line. Programming of the chips can be done individually by addressing

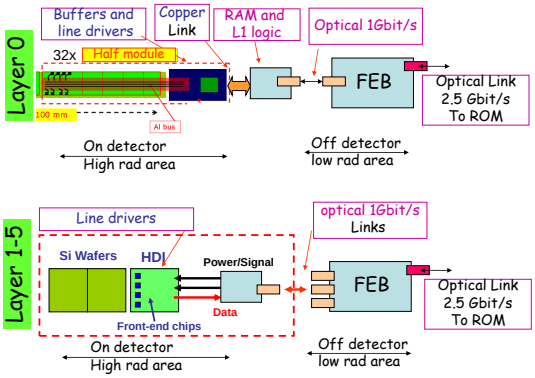

Figure 12.1: SVT Electronics

a sigle chip or via broadcast command sent to all the chips in an HDI. Hits will be serialised on a programmable number of lines (1,2,4 or 6) using the fastclock signal. Each HDI will have a maximum number of 16 output lines running at the fastclock rate.

The role of the transition card is threefold: a) it will distribute the power to the HDI, b) it will receive all the input signals for the front-end chips via the optical line connected to the programmable board and c) it will ship the data to the programmable board for data acquisition. For the inner layers (0-3) there will be a transition card for each HDI. For the outer layers (4-5) it will be possible to group the data of the two HDIs servicing the two sides of the same silicon sensor into the same transition card, reducing in this way the total number of optical links needed.

The received programmable board (FEB) will handle all the comunications with the front-end chips, the FTCS, the ECS and the upper DAQ systems. Each FEB board will be connected to a variable number of transition cards (up to 12 in the current design). Important and critical roles of this board are the clock distribution, the trigger handling and the data collection. The clock-like signals, such as an experiment clock (about 60 MHz), a fast clock (120-180 MHz) and a time-stamping clock (up to 30 MHz), have to be distributed in a system-wide synchronous manner. Special care will be taken to measure at each power-up the latencies of all the serializers and deserializers in the signal and DAQ chain so that the sent signals can be suitably adjusted in phase in order to have a system synchronous at the sensor (or front-end) level. The time-stamping clock will be used to time-tag the hits and can be different for the inner layers, where the high track rate requires short signal shaping times and short daq time windows and outer layers where the long strips and lower track rates allow for a longer shaping times (800 ns - 1 μs) and longer DAQ windows. The estimation of data volumes have been performed assuming a time-stamping of 30 ns period in the inner and outer layers. The acquisition window will be defined in a time window centered around the L1 trigger window and lasting at least 10 time-stamps (300 ns) for the inner layers and 33 time-stamps (990 ns) or more for the outer layers. The trigger request will be sent to the chips via the optical links. The most important function of the board is to collect the data coming out of the front-end chips both for monitoring and for the final daq. The data will be deserialized in the board and the redundant information will be stripped. A possible further data compression can be envisaged in order to reduce the final data volume. Finally the data will be sent-out via an optical link to a ROM module.

Data volumes. As discussed in the SVT chapter, the data rates and volumes are dominated by the background. In the design of the SVT front-end chips and DAQ chain, the latest background Bruno simulations have been considered at the nominal luminosity and a safety factor of 5 has been applied on the simulation results (design inputs). Due to the strong non-uniformity of the particle rate on the sensors, the front-end chip characteristics have been adapted to the peak hit rates, while the data volumes have been extracted from the mean rates for each layer. To evaluate the data rate a 150 kHz trigger rate with 10 μs of maximum latency and 100 ns of time jitter has been considered. An hit size of 16 bits is used as the FE chip output in the calculation that becomes 20 bits during serialization due to the 8b/10b protocol. The bandwidth needed by a layer0 ROS in a data-push configuration is of the order of 20 Gbit/s/ROS. A difficult-to-handle rate that moved us to consider a fully triggered SVT. In table 12.1 for each layer type the mean expected data load is shown.

For events accepted by the L1 trigger, the bandwidth requirement is only 1 Gbit/s and data from each ROS can be transferred on optical links to the front-end boards (FEB) and then to ROMs through the >8 Gbit/s optical readout links.

In total, the SVT electronics requires 18 FEBs and 18 ROMs, 18 optical links at 10 Gbit/s, 172 links at 1 Gbit/s (radiation

Table 12.1: Electronic load on each layer, Readout section and optical link.

Layer	Layer type	chips/ ROS	available channels	Backgnd (MHz/cm^2)	Gbits/trig (per GROS)	FE Boards	Event Size (kB)
0	striplet u	6	768	151	0.99	4	5.3
0	striplet v	6	768	151	0.99	4	5.3
1	strip z	7	896	14.0	0.56	1	2.2
1	strip phi	7	896	16.0	0.64	1	2.6
2	strip z	7	896	9.6	0.54	1	2.2
2	strip phi	7	896	10.3	0.58	1	2.3
3	strip z	10	1280	4.2	0.50	1	2.0
3	strip phi	6	768	3.0	0.36	1	1.5
4a	strip z	5	640	0.28	0.26	1	1.4
4a	strip phi	4	512	0.43	0.38	1	2.0
4b	strip z	5	640	0.28	0.26	1	1.4
4b	strip phi	4	512	0.43	0.40	1	2.1
5a	strip z	5	640	0.15	0.17	1	1.0
5a	strip phi	4	512	0.22	0.24	1	1.5
5b	strip z	5	640	0.15	0.17	1	1.0
5b	strip phi	4	512	0.22	0.24	1	1.5

hard). The average SVT event size is 88 kB, 30% coming only from the layer0.

12.1.2 DCH Electronics

12.1.2.1 Design Goals

The Super*B* Drift Chamber (DCH) Front End Electronics (FEE) is designed to extract and process the about 8000 sense wire signals for tracking and energy loss measurements purpose and to provide informations for Global Trigger generation (trigger primitives).
Two possible scenarios have been foreseen for the dE/dx measurement. The first one is based on the measurement of the sense wires integrated charge (Standard Readout), while the second one is based on electron clusters detection (Sampled Waveforms). A description of the requirements for the two scenarios can be found in the DCH chapter.

12.1.2.2 DCH Front-end system (block diagram)

The DCH FEE chain block diagram is the same for the two options and is shown in Fig.12.2 for the Standard Readout case. The chain is split in two blocks:

- ON DETECTOR electronics: HV distribution and preamplifier boards located on the backward end-plate to preserve sense wire Signal to Noise Ratio.
- OFF DETECTOR electronics: Data Readout and Concentrator Boards located on the top-side of the experiment.

12.1.2.3 Standard Readout - OFF DETECTOR electronics

Front End Boards - Block Diagram Front End Boards (FEB) will host up to 64 channels and will be made of three stages as shown in fig.12.3. The first one receives signals from preamplifiers and generates analog and discriminated outputs. The second stage provides digitization for charge and time measurements and includes the logic for trigger primitives generation as well. Finally, the third stage contains the Latency and Readout Buffers and

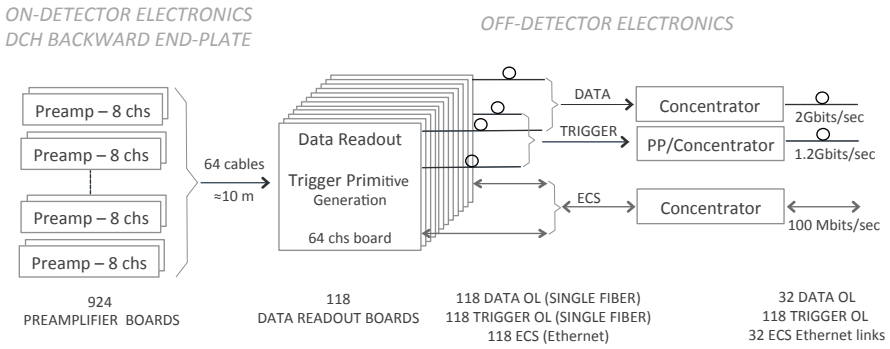

Figure 12.2: DCH front-end block diagram (Standard Readout)

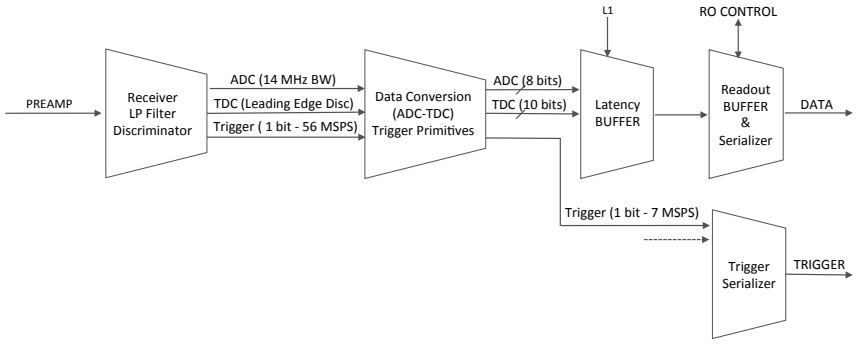

Figure 12.3: Front End Board Block Diagram (Standard Readout)

the dedicated control logic. FEBs will include an ECS section as well (not shown in the block diagram) for parameters setting/sensing and FEE chain test.
Input signal connections will be implemented by means of twisted or mini-coaxial cables, while output connections will depend on the data path. DAQ chain and Trigger chain data path will use optical links while ECS will use copper links (see ETD Control System chapter).

Front-End-Boards - Receiver Section This stage amplify and split the preamplifier output signal to feed an anti-aliasing 14 MHz low-pass filter (charge measurements) and a leading edge discriminator (time measurements).

Front-End-Boards - Digitization Section The 14 MHz filter output signal is routed to an eight bits and (about) 28 MSPS FADC whose outputs feed a section of the system Latency Buffer implemented in a FPGA. Thirty-two FADC output samples (corresponding to about 1.14 μs, enough to span the full signal development) will be readout in presence of a Level 1 (L1) trigger and, eventually, stored in a Readout Buffer.
The comparator output is routed to a FPGA where it is spit in two paths. The first one is sent to a TDC (implemented in the FPGA itself using the oversampling method) for time measurement. The second one, synchronized with the system clock and conveniently stretched to remove redundant informations, is sent to the DCH Trigger Segment Finding (TSF) modules (see ETD Trigger chapter).
The TDC outputs are routed to the second section of the Latency Buffer that, again, will be readout in presence of a L1 trigger signal.
The event data structure will not have a fixed length as L1 triggers spaced less than single

event readout time will extend the time window to include the new event. Nevertheless, the board structure will be also compatible with local Feature EXtraction (FEX) implementation, i.e. the extraction of relevant information from the digitized data. In case of FEX implementation the transferred data stream would have a fixed length; an example of a possible readout data structure is shown shown in table 12.2.
In both cases the first, non zero, FADC output can be used for amplitude correction in discriminated data then minimizing input signal slewing effects.

Table 12.2: FEX based data stream (fixed length structure)

Data stream example
Digitizer Module Address (2 bytes)
Flag (1 byte)
Trigger Tag (1 byte)
Counter (1 byte)
Charge (2 bytes)
Time (2 Bytes)
1st ADC sample (different from baseline) for time walk correction (1 byte)

12.1.2.4 Sampled Waveforms - OFF DETECTOR electronics

The Cluster Counting technique for dE/dx measurement is based on the detection of electron clusters (primary ionization measurement) then high bandwidth devices and high sampling frequency digitizers (at least 1 GSPS) must be used.
Moreover, because data throughput must sustain the Super$B$ foreseen 150 kHz average trigger rate, a very fast processing is required. The constraint, at the state of art of technology, result in a huge power requirement and, as a consequence, in a low FEB modularity. Concerning tracking requirements, if we assume that full efficiency in single electron cluster detection is achieved, Cluster Counting measurement already includes information for tracking purpose (it is just required to store clusters arrival time instead of simply counting them).
Sampled Waveforms DCH FEE chain block diagram is similar to the diagram shown in fig.12.2; but, because of the lower board modularity both the number of crates and boards increase significantly (see table 12.3).

Front End Boards Front End Boards will be based on high sampling rate ($\geq$ 1 GSPS) digitizer, then a limited number of channels can be packaged in a single board, mainly because of power requirements. At present, up to 8 channels working at 1 GSPS have been packaged in a single VME 6U board. Nevertheless, in the next future, board modularity could get to 16 (or 24) with small increase in power requirements.
The circuit structure is very close to the block diagram shown in 12.3. Differences arise in digitizing section as no TDC is required for time measurements and also trigger primitives are generated from FADC outputs.
A sensitive issue concerns the FEX. Because of the large amount of data per channel (about 1 thousand of bytes) we can not transfer raw data to the DAQ then FEX must be implemented in the FEB itself. Thus when an L1 accept is raised all the event samples must be examined to identify clusters. The time required to implement the procedure is still compatible with the average trigger rate foreseen at the nominal Super$B$ luminosity ($\sim$ 150 kHz), but it could be a limit if the luminosity increases.
Another issue concerns the radiation background, as high performances RAM based FPGA must be used in the design.

12.1.2.5 Front End Crates

Each Front End Crate will host up to 16 FEB, a Power Supply board for VFEB, Data Concentrators and, eventually, Trigger Patch Panel o Trigger Concentrator. Custom backplanes will be designed to:

- distribute power and common signals to the FEBs
- allow the use of Interconnection Boards to collect low modularity VFEB cables (fig 12.4)
- route some of the trigger signals to the neighbors boards (see ETD Trigger chap-

Table 12.3: Number of links (Data, ECS, Trig), FEBs and crates for 64 channels Standard Readout (SR) and 16 channels Sampled Waveforms (SW) board modularity

	Mod	Data	ECS	Trig	Boards	Crates
SR	64	32	32	118	118	8
SW	16	32	32	118	462	29

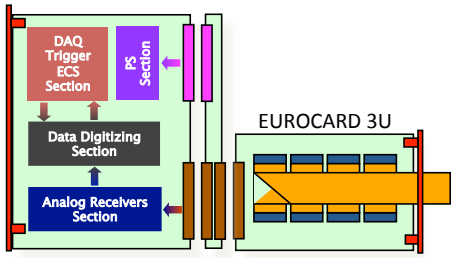

Figure 12.4: VFEB - FEB connections (FEB board - Custom Backplane - Interconnection Board

ter) and (eventually) to distribute common signals to the boards

12.1.2.6 Number of crates and links

Table 12.3 shows the estimation of the number of links, boards and crates required for both DAQ and Trigger FEE chains (each crate is supposed to host up to 16 FEB). As shown in the table the number of Trigger OL do not change despite the different FEB number for Standard Readout or Sampled Waveforms scenarios. This is because the Sampled Waveform option foresees a Concentrator board also for the trigger chain to compensate for FEBs low modulerity. The board will collect Trigger OL coming from several FEBs and will deliver a single OL to the TFS modules.
The estimation has been done assuming:

- 150 kHz L1 trigger rate, 7392 sense wires (subdivided in 10 super-layers)
- 10% chamber occupancy in 1 μs time window
- 48 bytes per channel data transfer
- single link bandwidth is 2 Gbits/sec for DAQ data path and 1.2 Gbits/sec for Trigger data path

12.1.2.7 ECS

Each FEB will host a mezzanine board to manage ECS communication. Besides the control of the board, ECS mezzanine should provide the capability of data readout for debugging purpose. Detail can be found in the ETD ECS chapter.

12.1.2.8 Cabling

Because the large number of channels involved, DCH cable layout must be carefully designed. The main requirement concerns the possibility of replacing, in case of failure, a VFEB without disconnecting too many cables. Thus, signal and HV cables should be routed through the chamber outermost layer to minimize cables overlap. Table 12.4 shows the foreseen number of cables and a rough estimation of cable sizes.

12.1.2.9 Power Requirements

A very preliminary power requirement estimation is shown in table 12.5. The estimation for ON-DETECTOR electronics is based on preamplifier simulation and prototype test, while estimation for OFF-DETECTOR electronics come from the state of art of digitizing board available. Local (VFEB) voltage regulation is supposed to be implemented by means of linear low-drop hard-rad regulator.

Table 12.4: Estimation of the number and dimension of DCH cables (7392 sense wires - VFEB modularity = 8 channels)

	LVPS	HV	Signal (coax)	Signal (twisted)
Quantity	118	32	7392	924
N of cores	16	25	1	16
Cond. area (mm^2)	0.5	0.07		
Overall diam. (mm)	12	12	1.8	6.5

Table 12.5: Power requirement estimation for bot Standard Readout (SR) and Sampled Waveform (SW)

	Channel	Board	Overall
SR VFEB	30 mW		250 W
SW VFEB	150 mW		1.2 kW
SR FEB		40 W	5 kW
SW FEB		40 W	19 kW
SR & CC Data Conc.		30 W	240 W
SW Trig Conc.		30 W	870 W

12.1.3 PID Electronics

The electronics for the FDIRC can be seen as an upgrade of the electronics of the *BABAR* DIRC. The new requirements of the experiment (Trigger rate, background, radiation environment) and FDIRC specific requirements (resolution, number of channels and topology) have led to a similar but new design of the electronics chain.

The electronics will equip the 18,432 channels of the 12 sectors of the FDIRC. The electronics chain is based on a high resolution / high count rate TDC, a time associated charge measurement on 12 bits and an event data packing sending event data frames to the data acquisition system (DAQ). The target performance of the overall electronics chain is a time resolution of 100 ps rms. This chain has to deal with a count rate per channel of 100 kHz, a trigger rate up to 150 KHz and a minimum spacing between triggers of about 50 ns.

The estimate radiation level is expected to be less than 100 rads per year. The use of radiation tolerant components or off the shelves radiation-qualified components is mandatory. However, the expected energy of the particles may make the latch-up effect almost impossible. Thus, the design has to take into account only Single Event Upsets. We selected the Actel family FPGA components for their non-volatile flash technology configuration memories, which are well adapted to radiation environment.

Several architectures have been considered which can be summarized as follows:

- All electronics directly mounted on the FBLOCK.
- All electronics mounted next to the detector and linked to the PMTs by cables.
- A part of it on the detector (the Front-end boards) and the other part, called crate concentrator, situated close to the detector, (this board is in charge of interfacing with the Front-end, reading out event data, packing and sending it to the DAQ.

The first solution has been chosen as baseline for the TDR for two main reasons:

- The cost of the cables (PM to Front-end boards) is estimated to be close to 200 kEuros (1/3 of the price of the overall electronics cost), making this solution too expensive. Moreover, the possible option to have pre-amps on the PMT bases doesn't prevent from having electronics and power supplies on the detector.
- The large amount of data per channel leads to have the L0 derandomizer and buffer on the Front-end boards. The FCTS receiver could be individually located on each Front-end board but the number of cables needed pushes to distribute all the control signals on a backplane. Consequently the board dedicated to receiving and transmitting FCTS signals on the backplane naturally tends to also become the event data concentrator and the link to the DAQ.

The baseline design assumes a 16-channel TDC ASIC offering the required precision of 70 ps rms- embedding an analog pipeline in order to provide an amplitude measurement transmitted with the hit time. Thanks to a 12-bit ADC, the charge measurement will be used for electronics calibration, monitoring and survey purposes. The Front end board FPGA synchronizes the process, associates the time and charge information and finally packs them into a data frame which is sent via the backplane to the FBLOCK control board (FBC). The FBC is in charge of distributing signals coming from the FCTS and ECS, packing the data received from the FE boards to a n-event frame including control bits and transferring it to the DAQ.

12.1.3.1 The TDC chip

A former TDC chip offering the requested performances of resolution has already been designed for the SuperNemo experiment. It provides a time measurement with both a high resolution 200 ps step (70 ps RMS) and a large dynamic range (53 bits). The architecture of this chip is based on the association of Delay Locked Loops (DLLs) with a digital counter, all of these components being synchronized to a 160 MHz external clock.

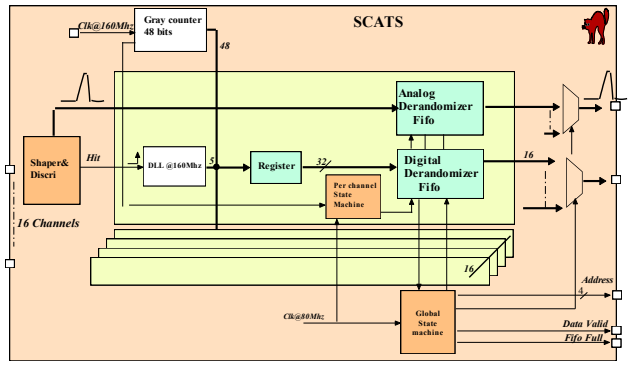

Figure 12.5: Block diagram of the Super*B* FDIRC TDC chip — SCATS —

The Super*B* chip (SCATS, Fig. 12.5) will keep the same philosophy but the high input rate requirement lead to a complete re-design of the readout part, in order to minimize the dead time per channel. Instead of registers and multiplexer which are the bottlenecks of the SuperNemo chip readout, it makes use of an individual FIFO memory per channel in order to derandomize the high frequency bursts of input data

With this architecture, data from the DLLs and the coarse counters are transferred into the FIFO memory within two clock cycles. When the transfer is complete, the channel is automatically reset and ready for the next hit. Simulations of the readout state machine showed an output FIFO data rate capability of 80 MHz. Time ranges for the DLLs and the coarse counter can be easily customized by adjusting the output data format (16, 32, 48 or 64 bits). Therefore, the chip is suitable for various applications with either high count rate and short integration time or low count rate and long integration time.

A FIFO depth of 8 words has been selected after simulation with a exponential distribution model of delta time between hits (mean rate of about 1MHz) applied to inputs. Thus the simulation gives a dead time of approximately 1% with 500 kHz input rate on each channel.

To design this FIFO a full custom RAM has been developed. It permits reducing the size of the chip and consequently its cost. The chip is designed using known and proved mitigation techniques to face single event upset (SEU) issues due to the low-level radiation environment.A first version of the chip without the analog FIFO and the discriminator has been submitted in November 2011 and the test are ... Note: something is missing here.

We plan to submit in 2012 one chip PIF dedicated to the currently missing parts:

- A low walk (approx. 50 ps) discriminator based on a CFD like design.
- A track/peak detector to be able to sample the maximum of the signal.
- An analog pipeline synchronized with the digital FIFO and providing analog output for charge measurement.

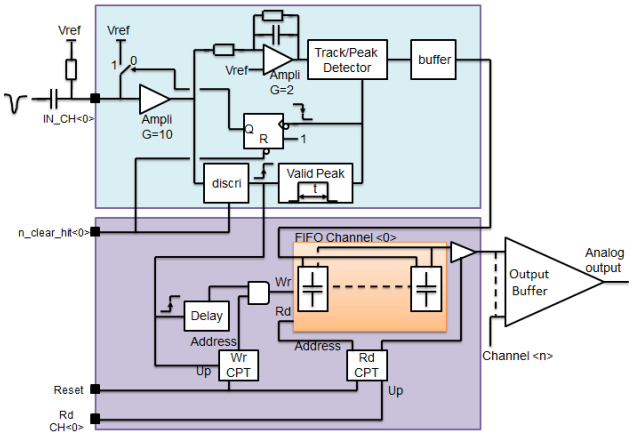

Figure 12.6: PIF One Channel

After testing and validation, it is foreseen to be included in the final version of SCATS taking benefit of sharing the FIFO pointers of the

analog and digital parts. The chip will be assembled and submitted end of 2013.

12.1.3.2 The Front-end Crate

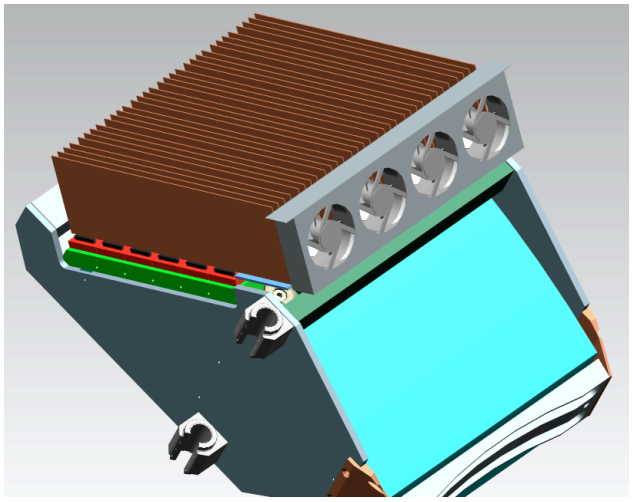

Figure 12.7: The FBLOCK equipped with the boards and fan tray

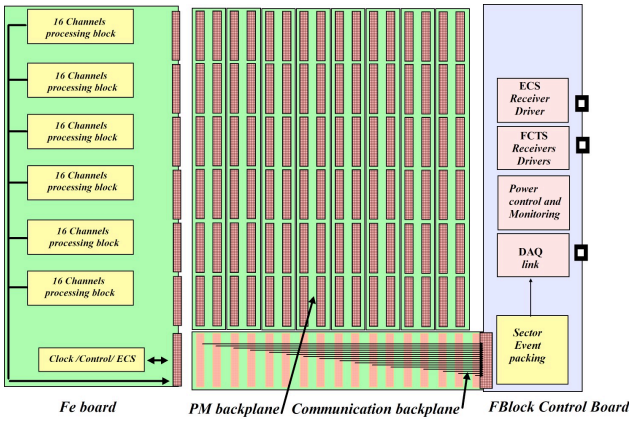

Figure 12.8: The Frontend Crate

The board input will fit the topological distribution of the PM on the FBLOCK. The PMs are arranged as a matrix of 6 in vertical direction by 8 in horizontal direction. Each column of 6 PMs will fit to one FE board. One vertical backplane (PM Backplane) will interface between the 4 connectors of each PM base to one connector of FE board. The PM Backplane is also in charge of distributing the High Voltage, thus avoiding HV cables to pass over the electronics. The FB crate will use as much as possible the elements of a commercial crate, in order to avoid the design of too many specific elements like board guides.

12.1.3.3 The Communication Backplane

Distributes the ECS and FCTS signals from the FBC to the 8 FE boards thanks to point to point LVDS links. Connects each FE board to the FBC for data transfer. A serial protocol will be used between FE board and the FBC in order to reduce the number of wires and consequently ameliorate the reliability. It will also distribute JTAG signal for FPGA board reprogramming and all signals for monitoring and control of the crate.

12.1.3.4 The PMT Backplane

It is an assembly of 8 motherboards, each one corresponding to a column of 6 PMTs. One motherboard receives 2 Fe-board. The 64 channels from 4 connectors per each PMT are merged on the motherboard into two connectors to get into the Front end board to get 16 channels per half PMT, i.e., 6 PMTs correspond to 96 channels per FE-board. It also insures the ground continuity between FE-boards — crate — FBLOCK.

12.1.3.5 Cooling and power supply

The electronics is located on the detector in a place enclosed by the doors. There are 2 major consequences: one is the problem of the cooling which must be carefully studied in terms of reliability and capability and the second is that the location is naturally shielded. Consequently the use of magnetic sensitive components as coils or fan trays is possible.

An estimation of the overall electronics consumption lead to approximately 6 kW, not including the external power supplies. This can be broken down to individual contribution as follows:

- Electronics: 0.325W/channel, 500 W/sector and 6 kW/system.
- HV resistor chain: 0.19W/tube, 9.1W/sector, and 109 W/system.

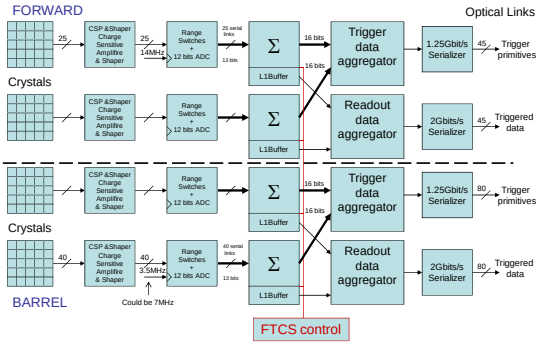

Figure 12.10: EMC Electronics

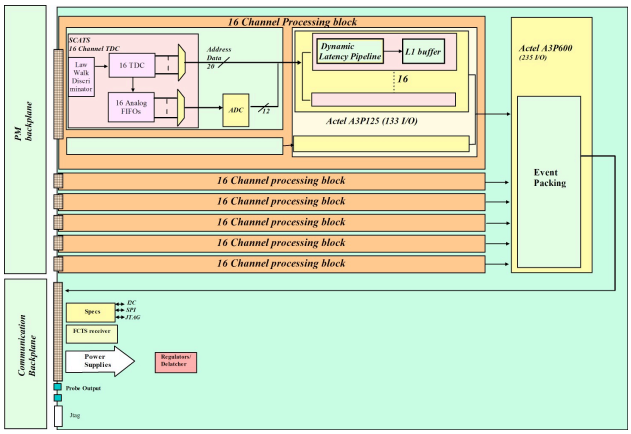

Figure 12.9: PID Front-end Board

The cooling system must be designed in order to maintain the electronics located inside at a constant temperature close to the optimum of 30 degrees. The air inside the volume must be extracted while the dry, clean temperature controlled air will be flowing inside. Each FB crate will have its own fan tray like in a commercial crate. Targeting a difference of 10 degrees between inside and outside temperature drives to a rough estimate value of $300m^3/h$ per crate, $4000m^3/h$ can be considered as the baseline value for the whole detector.

12.1.3.6 The front-end board

One Front-end board is constituted of 6 channel-processing blocks handling the 96 channels. The channel-processing block is constituted by one SCATS chip, one ADC, one Actel FPGA and the associated glue logics.

The FPGA receives event data from the TDC and the converted associated charge from the ADC. From one 16 bit bus of the 16 channels coming from the TDC, it de-serializes to 16 data path where events are keeping in a buffer until they are thrown away if there are too old (relatively to the trigger) or sent upon its reception.

The PGA master receives event data from the 6 channel processing blocks and packs the event.The FE board transfers the event frame in differential LVDS to the FBC via the communication backplane.

12.1.3.7 The crate controller board (FBC)

The FBC is the board which gathers the front-end data, control and monitors the crate. There is one board per crate. The board handles several functionalities:

- Receive the event data from the Front-end boards via the communication backplane, constitutes and event data fram for the DAQ
- Spy data building for monitoring and commissioning purposes
- Distributes the ECS (SPECS) signals to the front end. Distributes the JTAG.
- Deserialize clock and control signals from FCTS.
- Monitor the crate; temperature, power supplies, fans.

12.1.4 EMC Electronics

This is still the version from the Whitepaper!!!

Two options have been considered for the EMC system design—a *BABAR*-like push architecture where all calorimeter data are sent over synchronous optical 1 Gbit/s links to L1 latency buffers residing in the trigger system, or a "triggered" pull architecture where the trigger system receives only sums of crystals (via synchronous 1 Gbit/s links), and only events accepted by the trigger are sent to the ROMs through standard 2 Gbit/s optical links.

The triggered option, shown in Fig. 12.10, requires a much smaller number of links and has been chosen as the baseline implementation. The reasons for this choice and the implications are discussed in more detail below.

To support the activated liquid-source calibration, where no central trigger can be provided, both the barrel and the end-cap readout systems need to support a free running "self-triggered" mode where only samples with an actual pulse are sent to the ROM. Pulse detection may require digital signal processing to suppress noisy channels.

Forward Calorimeter The 4500 crystals are read out with PIN or APD photodiodes. A charge preamplifier translates the charge into voltage and the shaper uses a 100 ns shaping time to provide a pulse with a FWHM of 240 ns.

The shaped signal is amplified with two gains ($\times 1$ and $\times 64$). At the end of the analog chain, an auto-range circuit decides which gain will be digitized by a 12 bit pipeline ADC running at 14 MHz. The 12 bits of the ADC plus one bit for the range thus cover the full scale from 10 MeV to 10 GeV with a resolution better than 1%. A gain is set during calibration using a programmable gain amplifier in order to optimize the scale used during calibration with a neutron-activated liquid-source system providing gamma photons around 6 MeV.

Following the *BABAR* detector design, a push architecture with a full granularity readout scheme was first explored. In this approach, the information from 4 channels is grouped, using copper serial links, reaching an aggregate rate of 0.832 Gbit/s per link to use up most of the synchronous optical link's 1 Gbit/s bandwidth. A total of 1125 links are required. The main advantage of this architecture is the flexibility of the trigger algorithm that can be implemented off-detector using state of the art FPGAs without constraining their radiation resistance. The main drawback is the large cost due to the huge number of links.

The number of links can be reduced by summing channels together on the detector side, and only sending the sums to the trigger. The natural granularity of the forward detector is a module which is composed of 25 crystals. In this case, data coming from 25 crystals is summed together, forming a word of 16 bits. Then the sums coming from 4 modules are aggregated together to produce a payload of 0.896 Gbit/s. In this case, the number of synchronous links toward the trigger is only 45. The same number of links would be sufficient to send the full detector data with a 500 ns trigger window. This architecture limits the trigger granularity, and implies more complex electronics on the detector side, but reduces the number of links by a large factor (from 1125 down to 90). However, it cannot be excluded that a faster chipset will appear on the market which could significantly reduce this implied benefit.

Barrel Calorimeter The EMC barrel reuses the 5760 crystals and PIN diodes from *BABAR*, with, however, the shaping time reduced from 1 μs to 500 ns and the sampling rate doubled from 3.5 MHz to 7MHz. The same considerations about serial links discussed above for the forward EMC apply to the barrel EMC. If full granularity data were pushed synchronously to the trigger, about 520 optical links would be necessary.

The number of synchronous trigger links can be drastically reduced by performing sums of 4×3 cells on the detector side, so that 6 such energy sums could be continuously transmitted through a single optical serial link. This permits a reduction in the number of trigger links so as to match the topology of the calorimeter electronics boxes, which are split into 40 ϕ sectors on both sides of the detector. Therefore, the total number of links would be 80 both for the trigger and the data readout toward the ROMs, including a substantial safety margin (> 1.5).

12.1.5 IFR Electronics

ver 0.5 sep 18 2012, Angelo Cotta Ramusino

The full description of the IFR readout electronics, going from the design constraints determined by the IFR detector's features to the detailes of the adopted baseline design is given in the subdetector chapter. The present subchapter is meant to highlight those features of the IFR electronics description which are most related to the Super$B$'s common ETD and ECS infrastructures. The active layers of the IFR are equipped with modules in which the detector elements (of different widths for the PHY and the Z views in the barrel) are assembled in two orthogonal layers. The overall IFR electronics channel count amounts to 11604 for the barrel section and 9540 for the endcaps, assuming that, for both sections, the IFR is instrumented with 9 active layers.

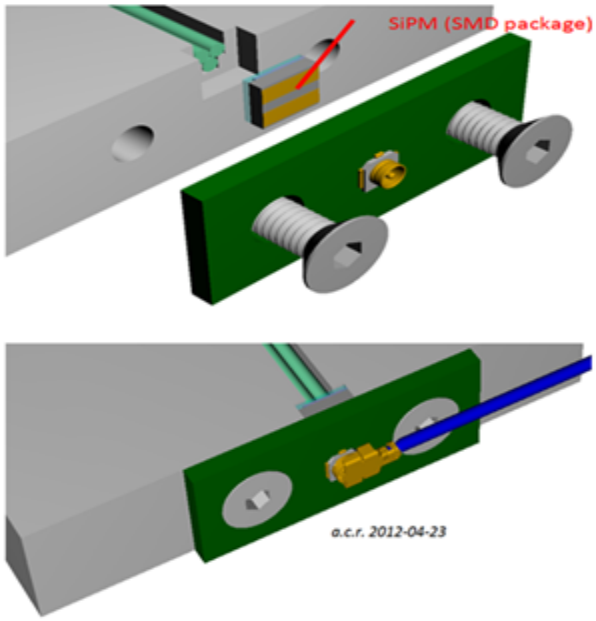

Figure 12.11: Estimation of the event size and data bandwidth for the IFR "Barrel" section

The IFR detector will be read out in "binary mode": the output of each SiPM device will be amplified, shaped and compared to a threshold, the binary status of each comparator's output being the variable to be recorded to reconstruct the particle tracks within the IFR detector. The downstream stages of the IFR electronics are the "digitizers" which sample the comparators' outputs at Super$B$ clock rate and store the samples into local "on-detector" circular memories. These buffers are designed to keep the data for a time interval at least equal to the Super$B$ trigger command latency. The last stage of the "on-detector" readout system is based on Finite State Machines (FSM for short) which extract, from the local latency buffers, the data selected by the FCTS trigger command.

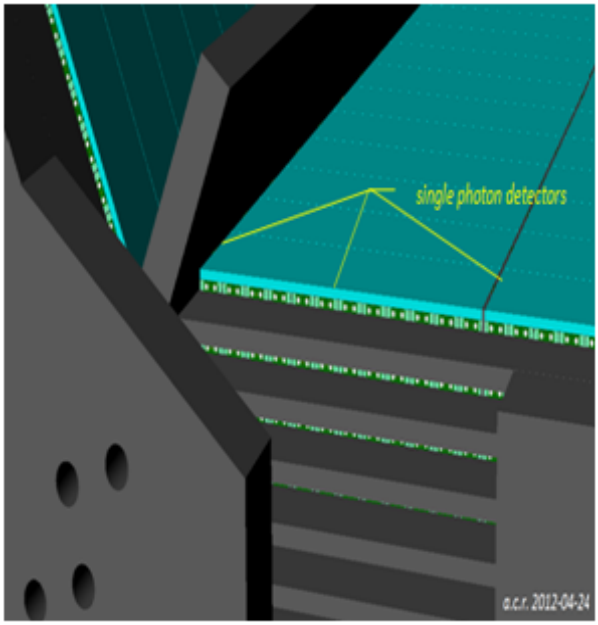

Figure 12.12: Estimation of the event size and data bandwidth for the IFR "Endcap" section

In the baseline version of the IFR electronics design the digitizer blocks introduced above have 32 input channels. The "trigger matched" data output by each digitizer in response to a received trigger pulse is sent, via serial links on copper, to the "data merger" units. These units assemble the trigger matched data from a number of front end digitizers into event data packets of suitable size and forward them to the ROMs via the Super$B$ fast data links. The application of the binary readout scheme to the IFR prototype detector developed in the R&D

phase has allowed the determination of parameters, such as the sampling frequency and the number of samples, which are relevant to the performance of IFR readout system.

Tables in figures 12.11 and 12.12 show the estimated event size and data bandwidth required for the IFR readout at the nominal SuperB trigger rate of 150kHz. The figures presented in the tables take into account a trigger jitter of 100ns. The event size and required bandwidth could be reduced by applying the data reduction scheme proposed by the ETD group: if a new trigger would select data that has already been sent in response to a previous one, the repeating data is not retransmitted and a pointer to the start of the repeating data within the previous packet would be sent instead.

LAYER WIDTH	LAYER	No. Modules per layer	LAYER ENABLE	PHI ASSUMING 50MM BARS	ZETA ASSUMING 106MM BARS
1963	1	6	1	13	17
1987	2	6	1	13	17
2050	3	6	1	13	17
2113	4	6		14	17
2176	5	6		14	17
2240	6	6		14	17
2304	7	6	1	15	17
2367	8	6		15	17
2431	9	8		12	17
2494	10	8		12	17
2569	11	8	1	12	17
2641	12	8		13	17
2712	13	8		13	17
2784	14	8	1	13	17
2879	15	8		14	17
2973	16	8	1	14	17
3068	17	8		15	17
3144	18	8	1	15	15
3296	19	8	1	16	15
NUMBER OF MODULES per sextant:		64		TOTAL PER SEXTANT	1940
TOTAL NUMBER OF MODULES	384			TOTAL CHANNELS PER BARREL	11640

Figure 12.13: Block diagram of the IFR readout ASIC

ENDCAP		
top section	horizontal bars per module:	37
	vertical bars per module:	51
center section	horizontal bars per module:	36
	vertical bars per module:	64
bottom section	horizontal bars per module:	37
	vertical bars per module:	51
AVERAGE channel count per module:		92
NO_OF_MOD_PER_LAYER_EC:		3
horizontal bars per layer		110
vertical bars per layer		166
NUMBER OF LAYER PER DOOR		9
NUMBER OF DOORS IN ENDCAPS		4
TOTAL NUMBER OF MODULES IN ENDCAPS:		108
horizontal bars per door:		990
vertical bars per door:		1494
TOTAL HORIZONTAL BARS:		3960
TOTAL VERTICAL BARS:		5976
TOTAL CHANNELS IN ENDCAPS:		9936

Figure 12.14: The IFR cable conduits for 2 sextants

While a readout system based on "COTS" ("Components Off The Shelf") was designed and successfully exploited for the readout of the IFR prototype, the front end stages of the electronic readout chain for the Super*B* IFR detector, for the reasons illustrated in the IFR subdetector chapter, will be implemented in an ASIC, whose block diagram is shown in figure 12.13.

The IFR scintillation detectors, each equipped with a single-photon counting device (silicon photomultiplier or SiPM for short), are grouped in modules which are inserted in selected gaps of the flux return steel or around it. The signals from all SiPMs of a module are carried by thin coaxial cables which exit the module's aluminum enclosure and are mass-terminated to a high density connector on a carrier PCB (printed circuit board). The average length of the coaxial assemblies connecting the sources (SiPMs) to the first amplification stages is in the order of a few meters because it is convenient to locate the ASICs' carrying boards where they would be accessible for any maintenance eventually needed.

Fig. 12.14 shows the situation for the IFR barrel: in this representation the metal enclosures of the modules are not shown; the closest convenient locations for the ASICs carrying boards are indicated by the yellow callouts.The front end cards are installed inside the 3" x 5" cable conduits as suggested in figure

12.15. While the figure only shows the double shielded, multi differential-pair cables (green jacket) needed to route the front end cards' output serial lines to the data merger units, the cable conduits are also meant to host cables for the distribution of the SiPM bias voltages (one bias voltage common to all SiPMs of a module), of the fast (FCS) and slow (ECS) commands and, finally, of the supply voltages for the front end cards.

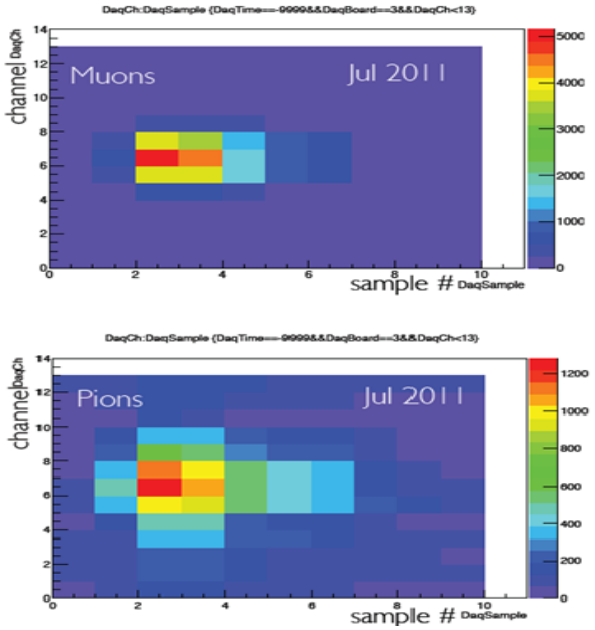

Figure 12.15: Detail of front end cards installed in the IFR cable conduit

Figure 12.16 shows the perspective view of the two doors of an endcap: the yellow callouts indicate the openings in the side-lining steel through which the data, control and power cables for the detector modules will be routed. The signals emerging from these openings are routed to the crates indicated by the red callouts. For the endcap section of the IFR, the front end cards carrying the IFR read out ASICs could be installed directly in the crate and connected to the data merger cards through the crate's backplane interconnections.

The "data merger" crates are interfaced to the Super*B* FCTS (Fast Control and Timing System) and ECS (Experiment Control System) from which they receive and execute the commands controlling the data acquisition and the detector configuration respectively. The data merger units also collect the trigger matched data from the front end cards and merge the input streams into the output packets sent, via the high speed optical data links, to the SuperB ROM (ReadOut Modules).

BARREL		ENDCAP	
Max. channel count per module:	32	Avg. channel count per module	92
Number of MOD32 processing units per module	1	Number of MOD32 processing units per module	3
TOTAL NO. OF MODULES IN THE BARREL SECTION	384	TOTAL NO. OF MODULES IN THE ENDCAPS	108
Total Number of modulo 32 processing units	384	Total Number of modulo 32 processing units	324
Sampling period = 1 / FCTS_clock (ns)	17,86	Sampling period = 1 / FCTS_clock (ns)	17,86
Number of samples in the trigger matching window	10	Number of samples in the trigger matching window	10
BARREL EVENT SIZE (kB)	**15,36**	**ENDCAP EVENT SIZE (kB)**	**12,96**
TRIGGER RATE (kHz)	**150**	**TRIGGER RATE (kHz)**	**150**
TOTAL BANDWIDTH (Gbps) (including 8b/10b overhead)	**23,04**	**TOTAL BANDWIDTH (Gbps)** (including 8b/10b overhead)	**19,44**
Number of data links	24	Number of data links	16
Bandwidth per link (Gbps)	0,96	Bandwidth per link (Gbps)	1,215

Figure 12.16: Perspective view of IFR endcaps

The IFR barrel will be equipped with a total of 6 "data merger" crates while 8 are foreseen for the endcaps. Further details on partitioning, location and construction of the data merger crates are given in the IFR subdetector chapter, which also provides a description of the services needed by the IFR readout in the experimental hall: - SiPM supply voltages: 384 (barrel) + 324 (endcaps) = 708 "HV" channels (one for each group of 32 SiPMs) featuring a programmable output voltage up to 100V @ 25mA. The power dissipated by the SiPMs in the whole IFR under nominal operating conditions is of the order of a few tens of Watts - front end cards' supply voltages: 200mA @3.3V for each unit handling 32 channels; lower supply voltages eventually necessary would be derived from on-board low-drop out (LDO) regulators. The number of

modulo-32 processing units (MOD32) needed to readout the IFR detector is 384 (barrel) + 324 (endcaps) = 708. Assuming that for both the barrel and the endcap case each front end card would carry 2 ASICs, the number of individually programmed and monitored "LV" (3.3V @ 1A) supply channels would be: 192 for the barrel and 162 for the endcaps. - cooling: the temperature of the IFR steel should be controlled to within a few degrees, as it was for BaBar; the design of the barrel cooling system should then take into account the estimated 600W dissipated in total by the front end stages of the IFR barrel and endcap electronics.

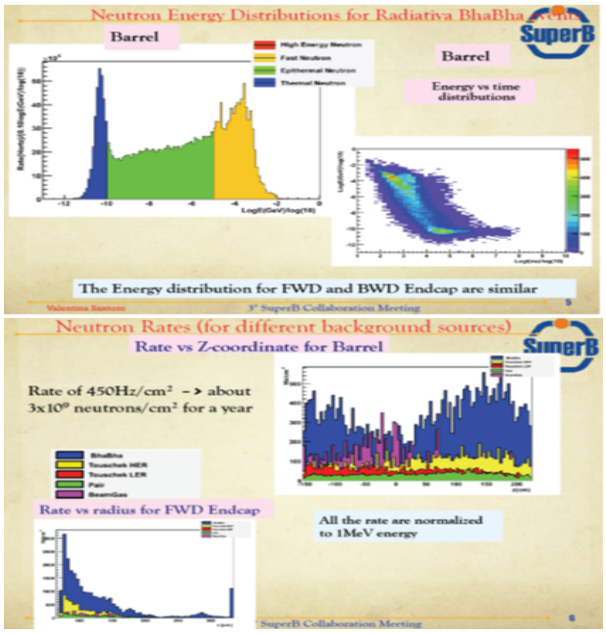

Figure 12.17: Main functions performed by the electronic units in the "data merger" crate

12.2 Electronics Infrastructure

12.2.1 Power supplies, grounding and cabling

Gianluigi Pessina

12.2.1.1 Power Supply to the Front-end:

The voltage supply system is normally composed by a cascade of AC/DC, DC/DC and linear regulators. Depending on the power dissipation and noise requirements some of the above elements can be avoided. The supply system of an accelerator-based experiment has known additional constraints to be satisfied. The large particle fluence and the presence of a strong magnetic field can have an impact on the aging and behavior of the electronic equipment. The first constraint is addressed only by adopting radiation hardened (rad-hard) technology and using suitable layout recipes for the monolithic circuits. This is common to all the devices that sit in the detector area. Magnetic field has generally less impact except for AC/DC and DC/DC converters that need to use inductances and/or transformers, having ferromagnetic cores.

Power Supply outside the detector area: In the following we will describe our solution, able to face the above constraints. The strategy we would adopt is to minimize the number of regulators in the detector area. The distance between the regulators and the front-end can be a few tens of meters. The energy the cable is able to store in its inductive component can be large and attention must be adopted to protect the connected electronic equipment in case of accidental short circuit to ground. Fig. 12.18 shows an example of the recovery from a short circuit of $50\,m$ cable with $4\,mm^2$ section. The short circuit current was limited to $20\,A$. A N-MOS, IPP50CN10NG, with $50\,m\Omega$ ON resistance simulated the short circuit and it breaks down above about $100\,V$ when in open state, that explains the reason of the clipping. The measurement has been taken in the worst condition of no applied load. It is clear from this that an accidental short circuit is very critical in producing possibly destructive damaging.

We have found that most of the cables with 3 or 4 poles and cross-section between $1.5\,mm^2$ and $4\,mm^2$ have an inductance per unit length, L_M, of the order of $0.7\,\mu H/m$, from DC to few hundred of KHz. Calling I_{short} the maximum delivering current available from the power supply and l the cable length, then:

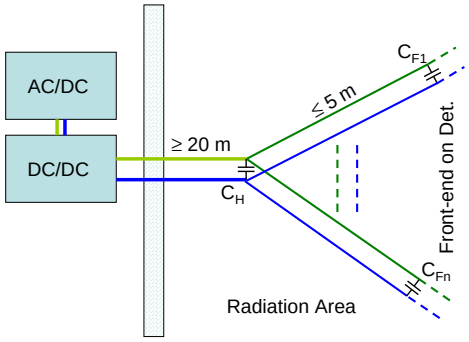

Figure 12.19: Possible layout for a sub-detector for what concerns voltage supply.

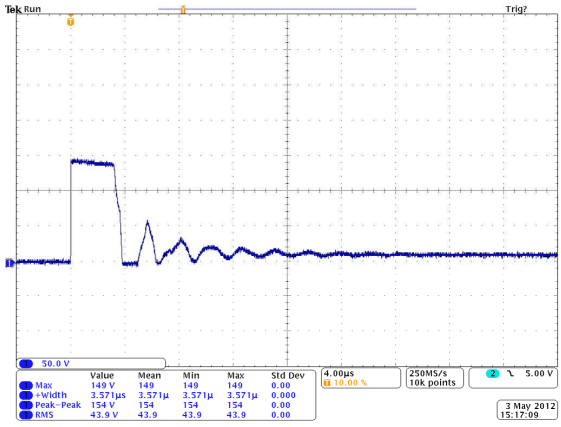

Figure 12.18: Signal at the end of a 50 m cable after a short circuit. The clipping at 150 V is due to the breakdown of the MOS switch used.

$$E_{nergy} = \frac{1}{2} L_M l I_{short}^2 \qquad (12.1)$$

is the available energy driven to the load in case of accidental short circuit. To limit the voltage at a safe level our straightforward solution is to add in parallel to the load a capacitance able to store such released energy. This can be done if we satisfy that:

$$\frac{1}{2} C_{lim} V_{over}^2 = \frac{1}{2} L_M l I_{short}^2 \qquad (12.2)$$

where V_{over} is the maximum voltage that must not be exceeded and C_{lim} is the capacitance whose value must be chosen to satisfy the eq. with the given voltage V_{over}. As an instance, with $50\,m$ of cable length, $I_{short} = 20\,A$ and $C_{lim} = 160\,\mu F$ the maximum over voltage excursion would be less than $9.5\,V$.

Adopting the introduced technique a hub with a distribution to several shorter cables can be implemented as shown in Fig. 12.19. A large value capacitance, C_H, is at the end of the cable that connects the DC/DC regulators from the outside to the inside of the detector area. In our example we continue with considering $50\,m$ of cable length and $C_H = 160\,\mu F$. From this point several shorter cables, or stubs, connect the various parts of the detector or sub-detector front-end. To save space, the sections of these last cables can be smaller since they have each to manage a smaller current. At the end of each of these stubs, $5\,m$ in our present example, a smaller value capacitance, C_{Fx}, $(33\,\mu F)$ is connected. In case of short at the end of a stub all the current flows into it. But as soon as the short is opened capacitance C_H absorbs the energy of the longer cable, while the energy of the shorted stub is managed by the corresponding capacitance C_{Fx}. Fig. 12.20 shows that the signal at the end of the $50\,m$ cable has an over voltage of only about 10 %. The maximum current was $20\,A$ and it can be seen that the baseline before the short is released is about $2\,V$, the dropout generated by the $20\,A$ current across about $0.1\,\Omega$ given by the sum of the stub (as a reference a section of $1\,mm^2$ has an impedance of about $16\,\Omega/Km$) and the ON resistance of the N-MOS. Fig. 12.21 shows the over voltage present at the stub end where the short is generated and released; again the over voltage is contained within about 10 %. As it can be appreciated, the baseline that precedes the release of the short is about $1\,V$, namely $20\,A$ developed across the about $50\,m\Omega$ ON resistance of the N-MOS. In the test setup of the laboratory we implemented 2 stubs and Fig. 12.22 is the signal at the stub end where the short circuit was not present. Again the over voltage is respecting the safety conditions. The principle applies well also if a low regulated voltage is considered and Fig. 12.23 is an example of release from a short

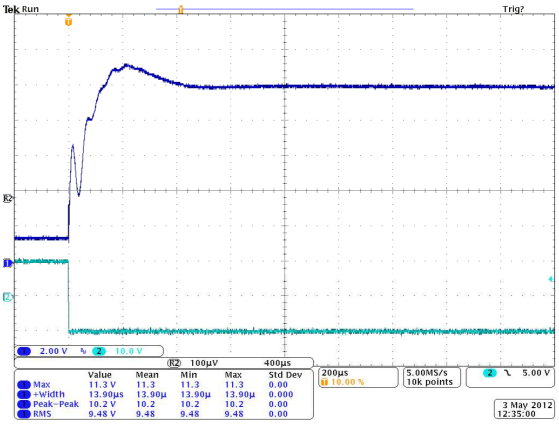

Figure 12.21: Signal at the end of a stub of Fig. 12.19 after a short circuit (blue line). Two stubs were present in the test setup. C_H is $160\,\mu F$, the C_{Fx} are $33\,\mu F$ and the short current is $20\,A$; the supply voltage is $10\,V$. The green line is the voltage driver at the gate of the switched N-MOS.

of more than $25\,A$ on a $5\,m$ cable ($2.5\,mm^2$ of cross-section) loaded with $160\,\mu F$ capacitance.

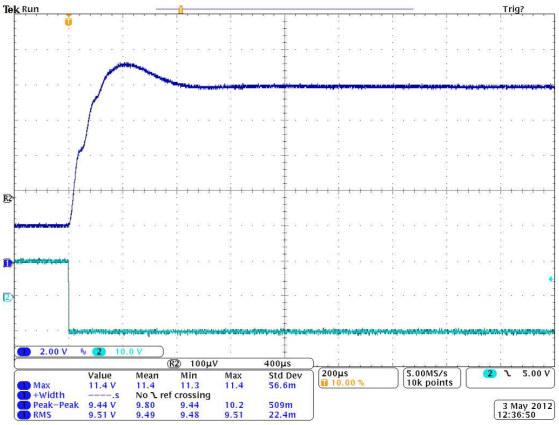

Figure 12.20: Signal at the end of a $50\,m$ cable (blue line) after that a short circuit is generated at one end of a stub. The $50\,m$ cable is loaded with a $160\,\mu F$ capacitor and the short current is $20\,A$; the supply voltage is $10\,V$. Every stub is loaded with $33\,\mu F$. The green line is the voltage driver at the gate of the switched N-MOS.

The suppression capacitance must show a very small series resistance and inductance. Capacitors with plastic dielectric such as Metalized Polypropylene Film satisfy this condition. As an example the $160\,\mu F$ we have adopted for the test has only $2.2\,m\Omega$ of series resistance, but, being big in volume, it shows a series inductance of a few tens of nH. To compensate for this last effect a smaller value (and volume) capacitance ($1\,\mu F$) is put in parallel, able to account for the fast part of the rising signal. Metalized Polypropylene Film capacitances have a range of values limited to a few hundreds of μF. As a consequence, a limited value of current per cable, $10\,A$ to $20\,A$, results in a good compromise. Many commercial regulators, also in the form of the so called bricks and half-bricks layouts, are available on the shelf at low cost. This strategy is particularly usefully for minimizing the dropout along the cable and it is of particular concern when a low voltage is needed.

We cannot forget that an over-voltage can happen due to a possible malfunctioning of the regulator. To reject rapidly and with good precision this effect a stack of fast diodes is a good choice. For instance with a voltage supply of $10\,V$ the series connection of about 20 diodes allows to maintain the safe operating condition provided that they are in contact with a heat sink in case the problem persists for a while. The location of the stack of diodes can be close to the regulator and space occupation would not constitute a problem.

Noise cabling and shielding: The combination of the inductance component of the wires and the suppression capacitance has a twofold utility as it behaves also as a low pass filter. Fig. 12.24 shows the noise at the end of the $50\,m$ cable plus $2 \times 5\,m$ stubs when $160\,\mu F$ plus $2 \times 33\,\mu F$ capacitances load the combination. The applied supply voltage was $10\,V$ and the load $3.3\,\Omega$. In this case a standard commercial regulator has been used. Very low noise DC/DC regulators have been designed [30] and Fig. 12.25 shows the noise performance under the same conditions. Even better performances can be obtained by

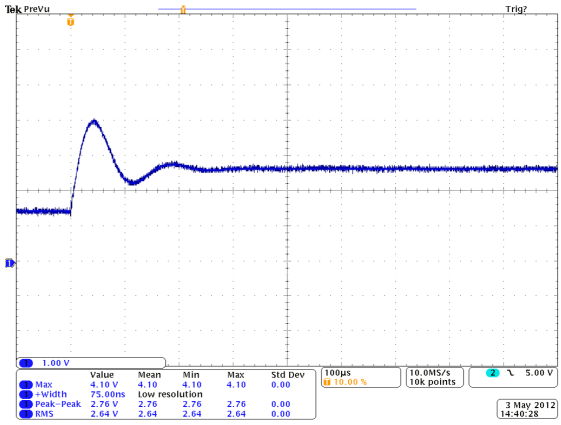

Figure 12.23: Short circuit release on a $5\,m$ cable with $2.5\,mm^2$ cross-section.

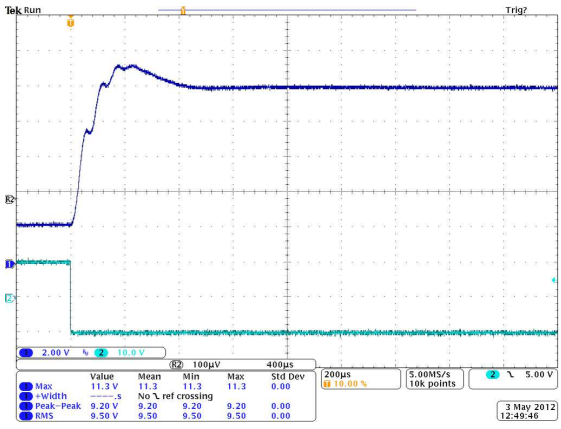

Figure 12.22: Signal at the end of a stub of Fig. 12.19 after a short circuit (blue line) happened at another stub. Two stubs were present in the test setup. C_H is $160\,\mu F$, the C_{Fx} are $33\,\mu F$ and the short current is $20\,A$; the supply voltage is $10\,V$. The green line is the voltage driver at the gate of the switched N-MOS.

cascading the DC/DC to a linear regulator of very good quality [31].

Low noise results are obtained if care is considered on the type of cables adopted. In all the measurements described so far cables used were all armored. This precaution allows to shield the supply voltage from outside disturbances but also to avoid to create disturbances to the outside world. We intend to adopt this kind of layout solution for the final experimental setup. In addition, where needed, we intend to add a double shield by inserting the cables inside a tubular copper mesh.

The connection scheme of Fig. 12.19 is, in a natural way, suitable to route ground. Let's suppose that the ground of every detector or sub-detector to which the cables are routed have their ground isolated. Then, we can route a tinned copper wire (or a copper bar) very close to the power supply cables so as to suppress area sensitive to EMI interferences. Such a routing scheme allows a 'star' connection with only one ground contact node (we remember that AC/DC and DC/DC regulators are floating), that is the standard requirement.

Shielding is considered for whose region were the electric or magnetic field can affect the performances. The shields can be considered for the whole sub-detector or individually on a channel by channel basis. This is particularly true with the effect of magnetic field on those detectors that extend on a large volume, such as photomultiplier tubes (PMTs). Past experience showed that in these cases a local shield implemented with mu-metal around every PMT is essential.

Power Supply in the detector area: We are considering the opportunity to use both DC/DC and linear regulators inside the detector area. Inductances and transformers cannot be based on a ferromagnetic coil. As a consequence they are limited in range of values and the switching speed of the DC/DC must be very large. This is the case for the monolithic DC/DC regulator we are considering [32], developed in $0.35\,\mu m$ CMOS technology based on rad-hard layout and components, and having a switching frequency of the modulator of a few MHz, which allows the use of a coil-free inductance. Based on the same technology a linear regulator is also available [33].

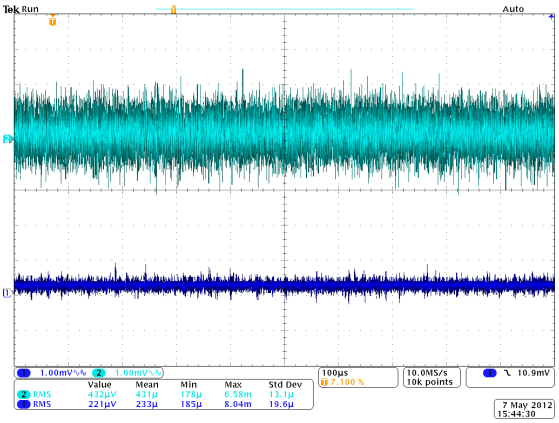

Figure 12.25: Very low noise DC/DC regulator [30]. Measurements condition and setup as for Fig. 12.24

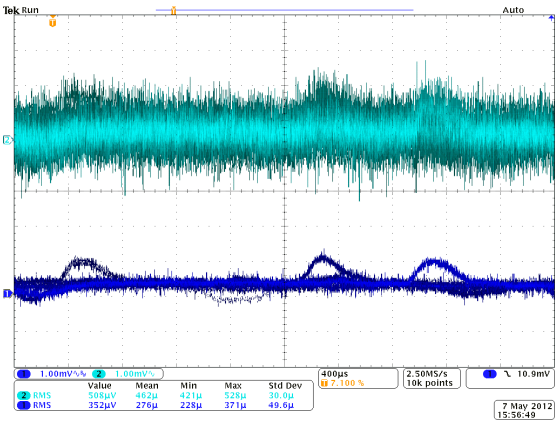

Figure 12.24: Noise after $50\,m$ cable loaded with $160 + 2 \times 33\,\mu F$ capacitance at $10\,V$ and with $3.3\,\Omega$ load. Upper noise is with the full $350\,MHz$ bandwidth of the oscilloscope, Lower noise has the scope bandwidth limited to $20\,MHz$.

12.2.1.2 High Voltage Power Supply to the Detectors:

High voltage power supplies suffer of similar problems as AC/DC and DC/DC regulators. As a consequence these regulators must be located outside the detector, sub-detector area (we do not know about any commercial rad-hard high voltage regulator). The energy released to the load in case of accidental short circuit would be not an issue thanks to the fact the, normally, such regulators have their driving current limited to a few hundred of μA. As an instance, if we load the line with a $1\,nF$ high voltage capacitor and considering $1\,A$ the short circuit current we expect an over-voltage of about $0.2\,V$. Finally, commercial over-voltage protectors based on gas discharge tubes are very efficient and fast.

12.2.2 Grounding and Shielding

This section has been incorporated in the Power Supply section

12.2.3 Cable Plant

This section has been incorporated in the Power Supply section

Bibliography

[1] A. Sharma, F. Sauli, Nucl. Instr. and Meth. A 350 (1994) 470.

[2] C. Avanzini et Al. Nuclear Instruments and Methods in Physics Research A449 (2000). 237-247.

[3] P.R.Burchat and John Hiser Studies of helium gas mixtures in drift chambers SLAC-PUB-5626 SCIPP 91/25 September 1991.

[4] Gianluigi Cibinetto, *Something about the prototype*, SuperB Workshop and Kick Off Meeting, Elba, May 31 2011

[5] Gianluigi Cibinetto, *Beam Test: data for analysis*, I SuperB meeting, London, Sep 15 2011

[6] Valentina Santoro, *IFR Background Report*, III SuperB meeting, LNF-Frascati, Mar 21 2012

[7] Pierre Jarron, CERN 1211 Geneva 23 Switzerland, *Radiation tolerant electronics for the LHC experiments*, Proceedings of the 4th Workshop on Electronics for LHC Experiments, Rome, Sept 1998, p. 7

[8] Federico Faccio, CERN 1211 Geneva 23 Switzerland, *Radiation effects in devices and technologies*, IN2P3 school, May 2005

[9] *ATLAS Policy on Radiation Tolerant Electronics: ATLAS Radiation Tolerance Criteria*, http://atlas.web.cern.ch/Atlas/GROUPS/FRONTEND/WWW/RAD/RadWebPage/ATLASPolicy/Appendix_1.pdf

[10] W. Dabrowski (AGH University of Science and Technology, Krakow, Poland), *Challenges and benefits of designing readout ASICs in advanced technologies*, Nuclear Instruments and Methods in Physics Research A 579 (2007) pp. 821-827

[11] M.Angelone, M.Pillon, R.Faccini, D.Pinci, W.Baldini, R.Calabrese, G.Cibinetto, A.Cotta Ramusino, R.Malaguti, M.Pozzati, *Silicon photo-multiplier radiation hardness tests with a beam controlled neutron source*, Nuclear Instruments and Methods in Physics Research A 623 (2010) pp. 921-926

[12] Flavio Dal Corso (INFN Padova), *Studies of SiPM @ Padova: Neutrons irradiation at LNL 9-13 Oct 2009*, XI SuperB General Meeting , LNF, Frascati, Dec. 1 2009

[13] Enrico Feltresi (INFN Padova), *Neutron Damage on SiPMs studies at LNL*, XV SuperB General Meeting Caltech, Pasadena CA, USA, Dec. 14 2010

[14] Flavio Dal Corso (INFN Padova), *Status of SiPM; studies @ Padova*, XVI SuperB General Meeting , LNF, Frascati, Apr. 4 2011

[15] A. Cotta Ramusino (INFN Ferrara), *SuperB IFR electronics: update*, 2nd SuperB General Meeting , LNF, Frascati, Dec. 15 2011

[16] Orsay MicroElectronics Group Associated (OMEGA) (LAL, Orsay, France), *EASIROC DATASHEET Version: 5*, April 2011

[17] Orsay MicroElectronics Group Associated (OMEGA) (LAL, Orsay, France), *EASIROC SOFTWARE & TEST BOARD USER GUIDE Version: 11*, April 2011

[18] James W. Meadows, Engineering Physics Division, Argonne National Laboratory, Argonne, Illinois, USA, *The $^{9}Be(d, n)$ thick-target neutron spectra for deuteron energies between 2.6 and 7.0 MeV*, Nuclear Instruments and Methods in Physics Research A 324 (1993) pp. 239-246

[19] A. Cotta Ramusino (INFN Ferrara), *SuperB IFR electronics: update on prototype electronics and IFR_DAQ*, XII SuperB Workshop - LAPP Annecy, Mar. 17 2010

[20] E. Grauges et al. [SuperB collaboration], *SuperB progress report – Detector*, arXiv:1007.4241,2010

[21] A. Cotta Ramusino (INFN Ferrara), *SuperB IFR electronics: update*, 3rd SuperB meeting, LNF, Frascati, Mar. 13 2012

[22] R. Mos, J. Barszcz, M. Jastrzb, W. Kucewicz, J. Młynarczyk, E. Raus, M. Sapor, *Front-End electronics for Silicon Photomultiplier detectors*, Electrical Review 11a (2010) pp. 79-83

[23] A.Baschirotto, G.Cocciolo, M.De Matteis, A.Giachero, C.Gotti, M.Maino, G.Pessina, *A FAST AND LOW NOISE CHARGE SENSITIVE PREAMPLIFIER IN 90 nm CMOS TECHNOLOGY*, Journal of Instrumentation, Jinst, V 7 C01003 pp 1-8 2012.

[24] G. Mazza et al. (Dipartimento di Fisica - INFN Torino), *A CMOS 0.13um Silicon Pixel Detector Readout ASIC for the PANDA experiment*, 2012 JINST 7 C02015 doi:10.1088/1748-0221/7/02/C02015

[25] G. Dellacasa et al. (Dipartimento di Fisica - INFN Torino), *A 130 nm ASIC prototype for the NA62 Gigatracker readout*, Proc. of the Int. Workshop on Semiconductor Pixel Detector for Particles and Imaging (Pixel 2010), Nucl. Instr. and Meth. in Physics Research, A 650 (2011) pp. 115-119

[26] A. La Rosa et al., *Ionizing radiation effects on a 64-channel charge measurement ASIC designed in CMOS 0.35 um technology*, Nucl. Instr. and Meth. in Physics Research A 593 (2008) pp. 619-623

[27] PICMG, AdvancedTCA PCI Express PICMG 3.4 Specification, May 2003

[28] M. Baszczyk, P. Dorosz, S. Głb, W. Kucewicz, M. Sapor, *Four channels data acquisition system for silicon photomultipliers*, Elektronika 12 (2011) pp. 28-31

[29] Paolo Carniti, Andrea Giachero, Claudio Gotti, Matteo Maino, Gianluigi Pessina, *An introduction to CLARO: A fast Front-End ASIC for Photomultipliers*, 2nd SuperB Collaboration Meeting, LNF Frascati, Dec 15 2011

[30] C. Arnaboldi, A. Bau, A. Giachero, C. Gotti, M. Maino, A. Passerini, and G.Pessina, *A Very Low Noise AC/DC Power Supply System For Large Arrays Of Cryogenic Detectors*, 2011 IEEE Nuclear Science Symposium Conference Record, pp. 791-793, February 2012.

[31] G. Pessina Rev. Sci. Instr. **70**, 3473, 1999.

[32] S. Michelis, B. Allongue, G. Blanchot, F. Faccio, C. Fuentes, S. Orlandi, S. Saggini, S. Cengarle, F. Ongaro, J. of Instrumentation, JINST, 7, C01072, 2012.

[33] F. Faccio, P. Moreira, A. Marchioro, S. Velitchko, *A radiation-tolerant LDO voltage regulator for HEP applications* , Conference proceeding of the TWEPP-2005, Valencia (Spain).