Overview of On-Detector Pixel Electronics

K. Einsweiler, LBNL

Overview of system specifications and design of module

Summary of organization, previous submissions, radhard issues, and next steps towards production

Overview of FE-D specifications and design at the top level (no schematics - not a design review...)

System Design of Pixel Module

First requirements analysis and discussions led to module system design in 96:

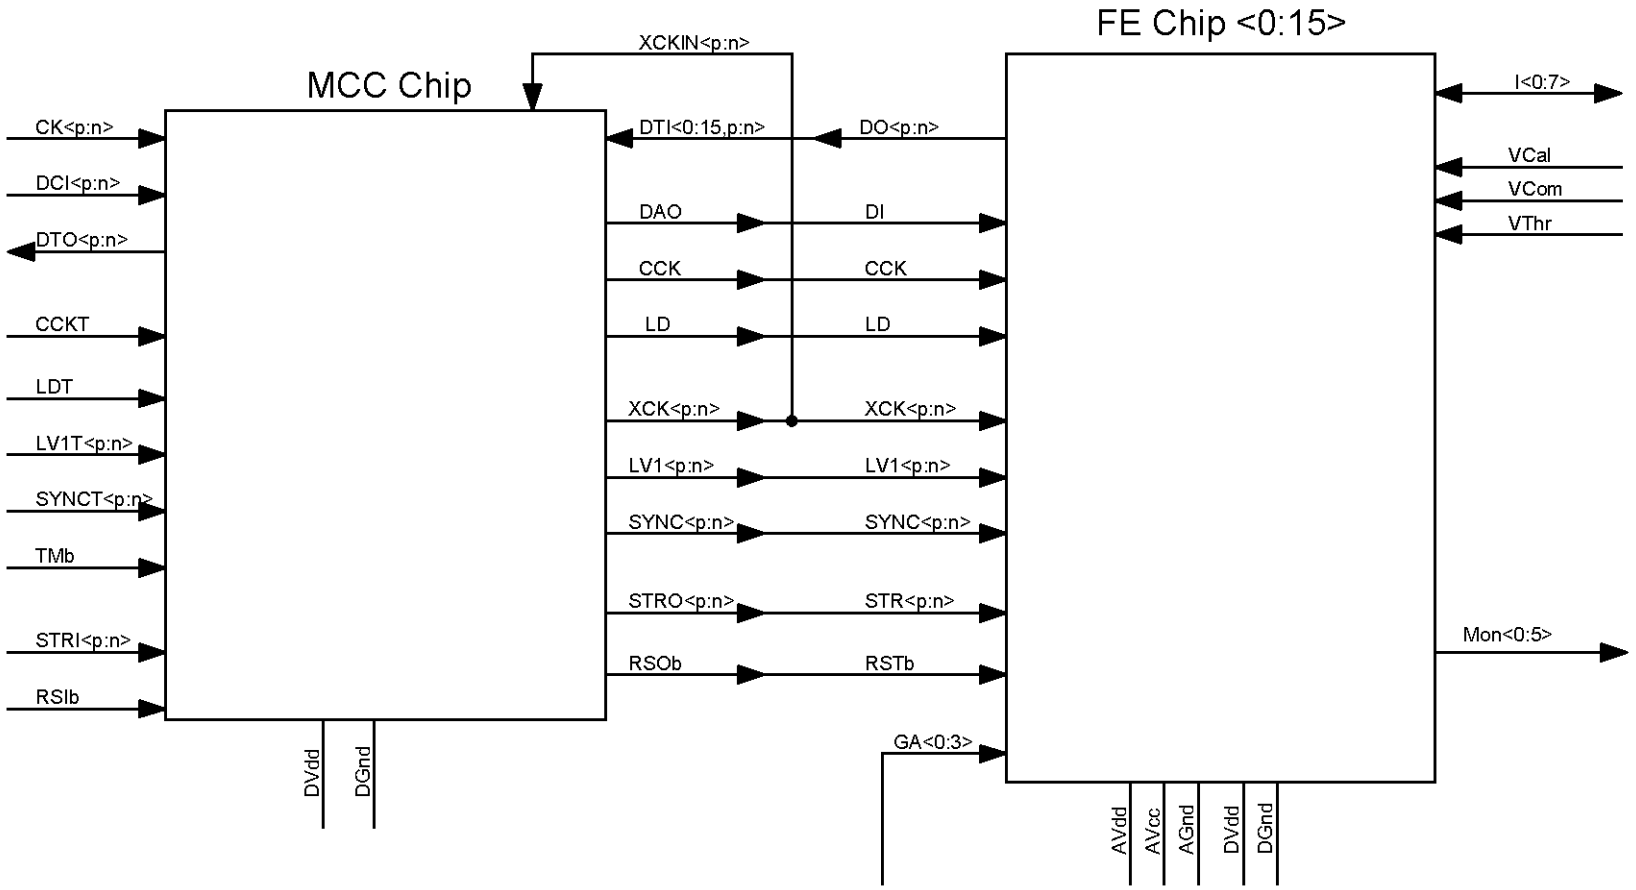

- Two chip design, including a single controller and event-builder chip (MCC) and 16 front-end chips bumped to a single silicon substrate.

Features:

- Basic interface to the outside uses a 3-wire protocol (SerialIn, SerialOut, XCK), which maps onto the ATLAS SCT opto-link protocol
- Provide a “transparent mode” interface through MCC to FE chips for simplified testing.
- Basic interconnections between FE and MCC use bussed signals. Slow control will not operate when recording events, so it uses full-swing CMOS. All fast signals use low-swing differential “LVDS-like” signaling. Point-point signals use 0.5 mA drivers (FE chips only), external or bussed signals use 3 mA drivers.
- To provide enhanced speed and robust module design, the serial output lines are connected from the FE to the MCC in a star topology (16 parallel inputs on MCC).
- Only analog signal between MCC and FE at this time is VCal calibration signal. All FE chips have internal current references and adjustment DACs to control analog operating points.
- Architecture is “data-push” style: each crossing for which LV1 accept is present causes all FE chips to autonomously transmit back hit information for the given crossing. LV1 signal may remain set for many contiguous crossings to allow readout of longer time intervals. MCC merges such events together.
- Synchronization signal available to ensure FE chips label LV1 properly.
- System uses two analog supplies to reduce power dissipation.

Complete Module Design:

- Has evolved to be a complete and independent assembly with only DC connections on Cu/Al and all AC connections on fiber optics:

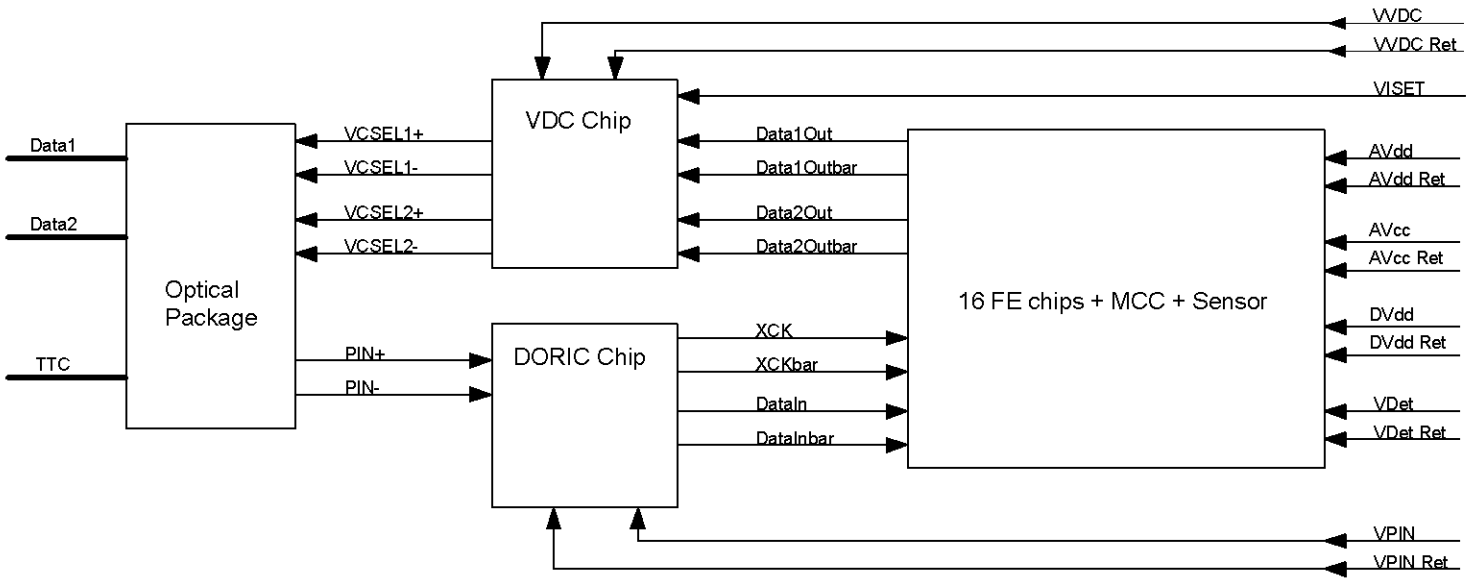

- Opto-interface includes two additional chips:

  → VDC provides current drive for VCSELs, roughly 1 mA quiescent current and 10-20 mA for full on state. The large current provides some contingency even after 10^{15} irradiation of baseline VCSELs (Mitel).

  → DORIC provides PIN preamp, DLL clock recovery and bi-phase mark decoding of command/ control data stream

- Power connections for six supplies, nominally 4V, 3.5V, 1.75V, 3.5V, -600V, and 5V, all brought in on low mass power cables.

Requirements Summary

Power budget:

- Total for module is based on 0.6W/cm^2 power density, or 5.5W total power, including all power dissipation in services within the pixel tracker volume:

  $\rightarrow$ Each FE chip is allowed roughly 250 mW (150mW analog, 100mW digital). The analog budget is further specified to be 40μW/pixel for 300μ pixels.

  $\rightarrow$ MCC is allowed roughly 300 mW on the common digital supply.

  $\rightarrow$ The leakage budget is 500mW (1 mA at 500V), the opto-link is 300mW, and the power distribution is 500mW (about 10% ΔV within pixel volume).

Geometry:

- The active die area for the FE chip is 7.2 x 10.8 mm, of which 7.2 x 8.0 mm is sensitive area for particle detection. The sensitive area of the FE chips must extend to the edge of the die along 3 sides, with all additional logic and I/O concentrated on the remaining side.
- Physics studies indicate that the pixels should be as narrow as possible in one dimension, and a 50μ pitch has been chosen as reasonably achievable. In the long direction, adequate resolution is obtained with a dimension of 300μ - 400μ. The present prototyping program has frozen the length at 400μ. The final production chips will shrink this as far as possible.

Radiation Hardness:

- The outer layers of the detector must survive for the expected 10 year LHC lifetime, defined in ATLAS as 3 years at 10^{33} and 7 years at 10^{34}. This lifetime dose corresponds to roughly 10^{15} 1 MeV n equivalent/cm^2, or 30 MRad, and is almost entirely due to pions.
- The B-layer is expected to be changed regularly, since its lifetime dose would be in excess of 100 MRad, and so we assume a lifetime dose similar to the outer layers for the first version. We are pursuing R&D for a 100 MRad version. Depending on the success of this effort, the initial version may use more aggressive technology.

Institutional Roles for On-Detector Electronics

- FE chip is a collaboration between Bonn, CPPM (Marseille), and LBL. In the present DMILL effort, the rough areas of responsibility are: CPPM does analog front-end, LBL does digital readout, Bonn does overall integration and designs many miscellaneous blocks. All groups follow overall design issues.
- MCC is primarily Genova with technical support from Bonn (CAD help) and LBL (full custom layout support, Honeywell standard cell support).
- Opto-electronics is being developed by ATLAS SCT at RAL. However, they have made several choices that are less-than-optimal for pixels, as the SCT is designing for lower radiation doses, larger modules, and larger power budgets.

  → Design is entirely NPN in a commerical radsoft BiCMOS process (AMS 0.8μ). We are concerned about radiation hardness and power consumption (design biases bipolars hard for improved radiation resistance).

  → Signal transmission requirements of SCT emphasize high-powered drivers, not needed for pixels.

  → We would like an all CMOS design for possible integration into MCC die.
- For these reasons, we have started a pixel-specific effort on opto-electronics, with OSU and Siegen participating. Goal is a rad-hard CMOS implementation of the chips in prototype form later this year.
- Opto-packaging is another problem area, as SCT currently totally reliant on one vendor (GEC-Marconi in UK). For this reason, we have strongly encouraged OSU to explore cheaper, lower mass, pixel options for opto-packaging.

Demonstrator Program

- In 1997, we agreed on overall design specifications for the FE and MCC chips necessary to implement this module design in a prototype form. We decided to pursue two prototypes for the FE chip. This was based partly on history and partly on the goal of submitting designs to two rad-hard vendors.
- One was called FE-A, and was designed for submission to AMS 0.8µ BiCMOS. This process was viewed as a prototype vehicle for DMILL. The chip was submitted in Oct. 97, and testing began in Jan. 98. A second, purely CMOS version, FE-C, was submitted in May 98. This chip has 880K transistors.
- The second was referred to as FE-B, and was designed for submission to HP 0.8µ CMOS. This process was viewed as a prototype vehicle for Honeywell. This chip was submitted in Feb. 98, and testing began in Apr. 98. This chip has about 850K transistors.
- A DMILL prototype matrix (no EOC, simple readout) called MAREBO was also submitted in Jul 97 and tested in Jan 98, to verify the FE design in DMILL.
- The MCC was submitted in May 98, along with the FE-C, and returned in late summer. FE-C chips were tested and bonded for evaluation in Sept. testbeams.
- All of these chips contain minor errors, but in all cases their functionality was quite close to the submission goals. Extensive lab testing and testbeam studies have been carried out on all chips. Excellent performance has been achieved.

Some Features of the FE-B Chip

- This chip was entirely designed at LBL. Most lab results and test beam results have been obtained from these chips.
- Improved version of the readout logic is included in new DMILL chip (FE-D), so that will be described in more detail later in this talk.
- The front-end design is quite different than that of FE-D:
- The FE-B design uses a dual-threshold discriminator in order to separate the very low threshold required to achieve the timing requirements from the higher threshold of interest in reading out the hits.
- In particular, if one achieves adequate overdrive performance, with the charge required to fire the discriminator within 25ns being lower than the threshold separation, all hits passing the second threshold will be “in-time”, and all out of time hits are eliminated.
- This design also controls the speed of the back-end more carefully since simulations show that the time structure of real pulses, and of cross-talk pulses induced by neighboring pixels via the inter-pixel capacitance, is different. Some suppression of cross-talk is possible by slowing the upper discriminator stage.
- The design attempts to minimize dispersion of the threshold difference.
- The FE-B implementation of this design shows excellent cross-talk suppression, but has worse than expected timewalk (see talk of John R).

Summary of Demonstrator Results

- FE chips have worked largely as expected (or better in some cases). High yield of assemblies has been achieved, along with routine operation in the lab and testbeam at 2Ke thresholds.
- Excellent threshold control (100-150e dispersion), noise performance (100-150e), timewalk (1-2Ke overdrive), and crosstalk (less than 5%) have been achieved.
- Operation with irradiated detectors and leakage currents up to 50 nA per pixel have been shown in testbeam, including operation of assemblies at 600V bias and 98% efficiency.
- Operation of 16-chip multi-chip modules has been shown, with some degradations in performance, but what appears to be acceptable operation. Much evaluation work with larger numbers of modules remains to prove we have control of this level of the system.
- Operation of multi-chip modules using realistic interconnects (Flex or MCM-D) and MCC chips in full MCC mode has also been demonstrated.
- Basic proof of principle has been provided. If we are able to achieve this same level of performance with rad-hard electronics, and after irradiation, we can build a high quality pixel detector for ATLAS.

Steps Towards Production

- After rad-soft demonstrator prototypes of FE and MCC chips, next step is radiation hard versions of everything.
- For FE, have agreed to first make a common DMILL design using the demonstrator geometry, but including several new features to bring it closer to a production chip. This chip is referred to as FE-D. After this, the same design (roughly) will be transferred to HSOI, and is referred to as FE-H. The extra capabilities of HSOI will allow us to consider several performance improvements, which may be essential for the B-layer.
- We anticipate that after these two rad-hard designs are fully evaluated, we will make a vendor choice for pixels, and work on a pre-production version of the chip. The definition of this chip would be that, barring significant errors, it could be used for ATLAS. At the present time, this design should be the same for the outer pixel layers and the B-layer, however further experience may alter this.
- After evaluation of the pre-production FE chip, the final production run would begin.
- For MCC, it is more difficult to pursue two vendor options since good standard cell libraries are required. A first rad-hard prototype (not complete) will be submitted soon, to be followed by a complete DMILL design. There are already indications that implementing the required functionality in DMILL within the allowed die area will be very difficult, so an HSOI version may follow.

Two Vendor Rad-Hard Prototyping

- ATLAS Inner Detector insists all major subsystems should prototype in two vendors, in order to minimize risks and costs.
- We strongly support this, and are working with TEMIC DMILL, a 2-metal 1-poly 0.8µ BiCMOS process, and Honeywell Enhanced SOI, a 3-4-metal 1-poly 0.8µ CMOS process. Our experience so far indicates that the HSOI process is capable of far denser layout than DMILL.
- CPPM has made several prototype chips already in DMILL, although so far only one used the TEMIC process, and that was a beta run with many problems. It appears that the process will work for our 25 MRad application. The vendor has supplied corner models, both pre- and post- rad, and these provide the basis for analyzing the design margin prior to submission.
- The HSOI process is much newer. At LBL we have irradiated PM bars and characterized the performance of the production SOI and the enhanced SOI process after irradiation. We find the enhanced process (with improved oxide treatment to minimize threshold shifts) has excellent performance after 50 MRad.
- We have submitted prototype FE chips, using an improved FE-B design, in a multi-project production SOI run in Nov 98. We have also submitted detailed transistor and parasitic extraction arrays. We intend to characterize both before and after irradiation in order to also understand the analog performance issues.

- The present FE-D submission includes a common design, which is relatively low risk since the analog portion has already been prototyped in the same process.
- This chip is built using the 400µ demonstrator geometry and it is clear that we cannot shrink this geometry at all, unless we are willing to give up something significant like TOT charge measurement. It is also unlikely that we could include more than about 25 buffers in the EOC (FE-D has 20 buffers like FE-B).
- The FE-H submission is likely to include two versions of the basic pixel array (all of these submissions use a reticle which is large enough to include 2 pixel chips, plus some test chips). This should lead to a better optimized FE design in HSOI.
- The first HSOI prototype chips indicate that we can build a 300µ pixel using this process. Furthermore, it is also likely that we can use at least 30 buffers, and include some minor upgrades in the readout architecture.
- Finally, there are also cost considerations. Since the size of the pixel chip is fixed by geometry, the use of a denser layout process does not increase the die per wafer. The DMILL process benefits from the lower overheads of running it on a large shared fab line (about 200K wafers per year), whereas Honeywell maintains their own specialty line (less than 20K wafers per year). The wafer costs are therefore different at the 25% level. However, it is possible that yield or design issues could overcome this built in advantage for DMILL.

Demonstrator FE Chip Geometry

- Agreement on pixel size was struck in Sept 96, in order to allow compatible, parallel detector and electronics development.
- The geometry adopted was 50μ x 400μ for the pixel size, with pixels arranged in 18 columns of 160 pixels per column.
- The geometry was mirrored between columns, so that the inputs for pixels in column 0 and 17 are on the outside, and all other columns are paired.
- The input pad geometry in the inner column pairs is then a double row of 50μ pitch pads. The metal pad is specified to be 20μ square, with a 12μ opening in the passivation for the bump-bonding.
- The cut die size must not extend beyond 100μ from the edge of the sensitive area on three sides of the die. Hence, nothing outside of the pixel circuitry is allowed on three sides of the chip, to allow module construction.
- The bottom of the chip (all peripheral logic and I/O pads) are allowed 2800μ, making the total active die region 7.2mm x 10.8mm.
- An I/O pad structure of 48 pads, each consisting of a 75μ x 75μ wire-bond pad, and a group of 4 bump-bond pads, was frozen.

Overview of FE-D

Common design effort, based on combining different aspects of FE-A/B/C chips:

- Front-end design is basically that of FE-C and MAREBO, but is purely CMOS. Although there is a BiCMOS version of this FE (2 bipolars used for discriminator input) which has slightly better timing performance than the purely CMOS version, this cell appears to be too large for FE-D.
- The FE uses a DC-feedback preamp design which provides excellent leakage current tolerance, close to constant-current return to baseline for TOT, and very stable operation with different shaping times.
- It is followed by an AC-coupling stage, and a fast discriminator. A divider network provides fine input baseline (threshold) control and provides the resistance for the high pass filter AC-coupling stage.
- The control logic provides a 3-bit threshold trim capability in each pixel, plus individual mask and calibration inject control. A global FastOR net is created using all pixels enabled for readout, and provides a self-trigger capability.
- All critical bias currents and voltages on the chip are controlled by internal DACs (6 current mode, 2 voltage mode), which are referenced to an internal V_{BE} reference, and controlled via the chip command decoder.

- Readout design is an upgraded version of the readout architecture of FE-B.
- It uses a 7-bit Grey-coded 40 MHz “timestamp” bus as a timing reference throughout the active matrix. All pixels measure their leading and trailing edge timing by asynchronously latching this reference in RAMs.
- Hits (address plus LE/TE timing) are transferred from the pixels as soon as the trailing edge occurs, using a shared bus structure in the pixel column pair. This bus should operate at transfer rates up to 20 MHz in order to meet our requirements. Low swing signal transmission and sense amplifiers are used to achieve this.
- Significant buffering is provided in the end of column region for hit storage during the L1 latency (up to 3.2μs in ATLAS). Twenty buffers are available for each column pair. The coincidence with the L1 trigger is performed in this buffer by comparing the LE timestamp with a reference timestamp delayed by the L1 latency. Accepted hits are labeled by a 4-bit trigger number. Hits from rejected crossings are immediately cleared.
- A readout sequencer stores information on up to 16 events pending readout. As soon as the output serial link is empty, transmission of a new pending event begins. Essentially, sending a L1 trigger corresponds to making a request for all hits associated with the corresponding beam crossing, which are then pushed off the FE chip to the MCC. Adjacent beam crossings can be read out to provide the ability to study timing issues in more detail.

- Global control of the chip is implemented using a simple command protocol. A global register controls Latency, DAC values, enabled columns, clock speeds, and several other parameters.
- A pixel register which snakes through the active array provides access to the 5 control lines in the pixel (Select, Mask, Trim<0:2>).
- Each chip on a module is geographically addressed, and its identity is controlled by external wire-bonds to avoid confusion.

FE-D Block Diagram:

Basic FE block diagram, expanded from module diagram:

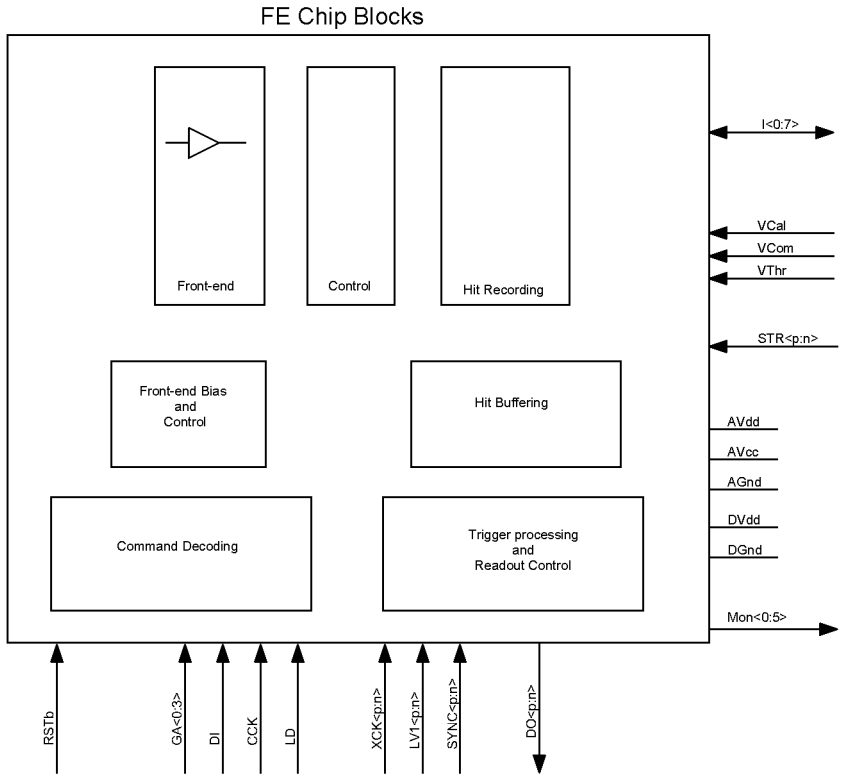

- Within the pixel, there is the front-end (preamp/ discriminator), the control block, and the readout block.
- Just below the active pixel matrix is the biassing and control for the front-end blocks, and the buffering for the readout blocks
- Finally, there is the overall readout control and the command decoding.

3

- Basic Digital I/O shown on bottom: 4 CMOS inputs for control (RSTb, DI, CCK, LD), and 4 fast, differential I/O's for timing and readout (XCK, LV1, SYNC, DO).
- Calibration and monitoring are shown on the right. A fast, differential strobe (STR) supplies calibration timing, and an analog voltage input (VCal) supplies calibration input. Dedicated monitoring pins include FastOR and test pixel output.

FE-D Pinout and Geometry

Sketch of pin assignments and overall geometry of die:

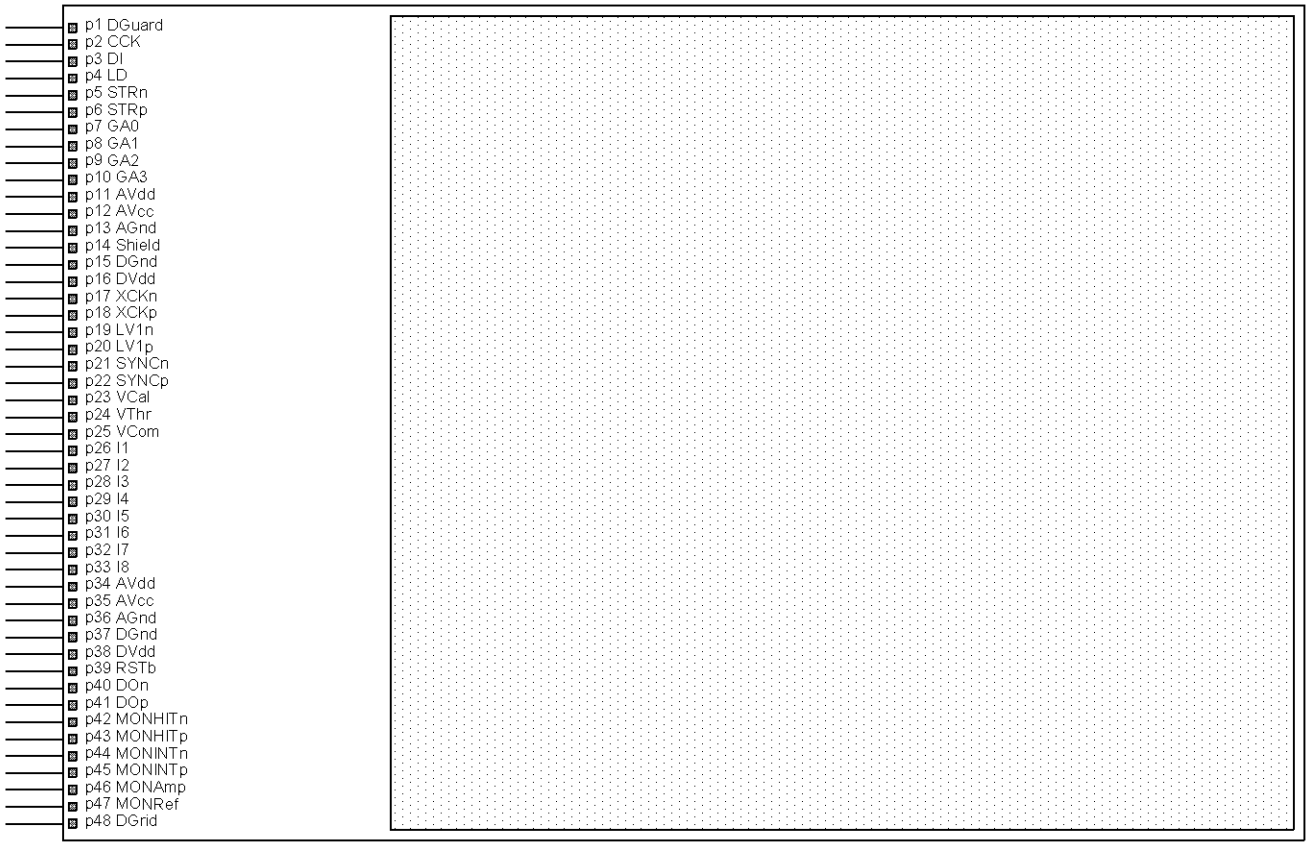

Front-end, biassing and control

Summary of the requirements:

- A nominal capacitive load of 200 fF is expected, roughly half to ground (parasitic) and half to the nearest neighbors (inter-pixel). Good performance should still be obtained with loads of 400-500fF. The n^+ on n-bulk detectors provide negative signals.
- Pixels are oriented to maximize signal and efficiency (minimize charge sharing).
- The outer layers should be 250μ silicon, and the B-layer should be 200μ.
- The expected signal after the lifetime dose of 10^{15} n-equiv/cm^2 is about 6Ke with 200V bias, and about 10Ke with 600V bias. We now propose to operate at the higher bias at the end of the detector lifetime, and have real prototype experience to show that this works well.
- This leads to an in-time threshold requirement of about 3Ke (this would be 2Ke for the lower bias voltage). This requirement has often been defined using a maximum timewalk relative to some large reference charge of 25ns, however, we believe 20ns is the right target for the complete FE chip. This could be achieved by for example setting a 2Ke threshold, and having the required overdrive for a timewalk of 20ns be less than 1Ke. This is the most challenging requirement for our front-end.

- Noise should be less than 300e and threshold dispersion less than about 200e, leading to an overall threshold “variation” of less than 400e.
- Leakage current tolerance should be at least 50nA per pixel, without significant changes in operating performance, and independently achieved for each pixel.
- Noise occupancy should be less than 10^{-6} hits/crossing/pixel.
- Crosstalk between neighboring pixels should be less than 5-10%, where this is defined as the ratio between the threshold and the charge which must be injected into a pixel to fire its neighbors.
- A double pulse resolution of 2μs is required for the outer layers, and 0.5μs for the B-layer, in order to achieve our total deadtime requirements.
- It is required to provide binary readout of each pixel, but a modest analog resolution (4-5 bits) is very desirable if it can be achieved without a large impact on the other performance specifications.
- A threshold range of 0 - 6Ke is needed.
- A calibration injection capacitor of 10fF should be included in each pixel.
- We do not know whether real diode input protection is needed. In all present chips, no explicit input protection is provided, and no major problems have been observed.

FE and Control Blocks:

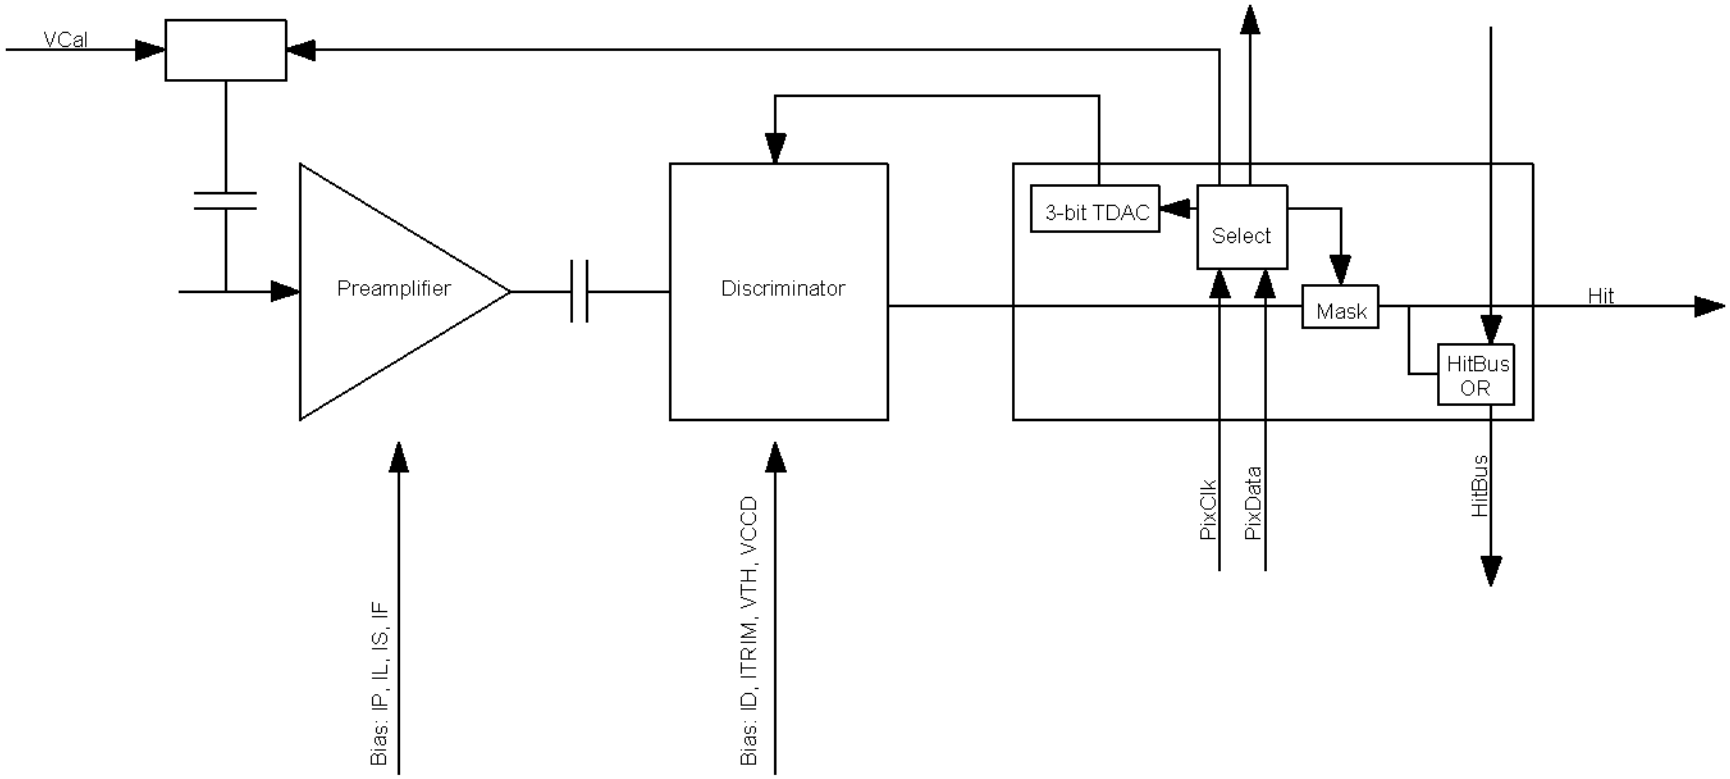

- Preamp has roughly 3fF DC feedback design, 10fF injection cap, and 40ns risetime. Input transistor operates at about 9µA bias and 1.5V. Cascode and follower operate at 3V and about 1.5µA each.
- Discriminator is AC-coupled and uses a bias current of about 6µA, plus two tunable voltage supplies (VTH and VCCD) which control the threshold and AC-coupling time constant.

FE Biassing and Control Blocks:

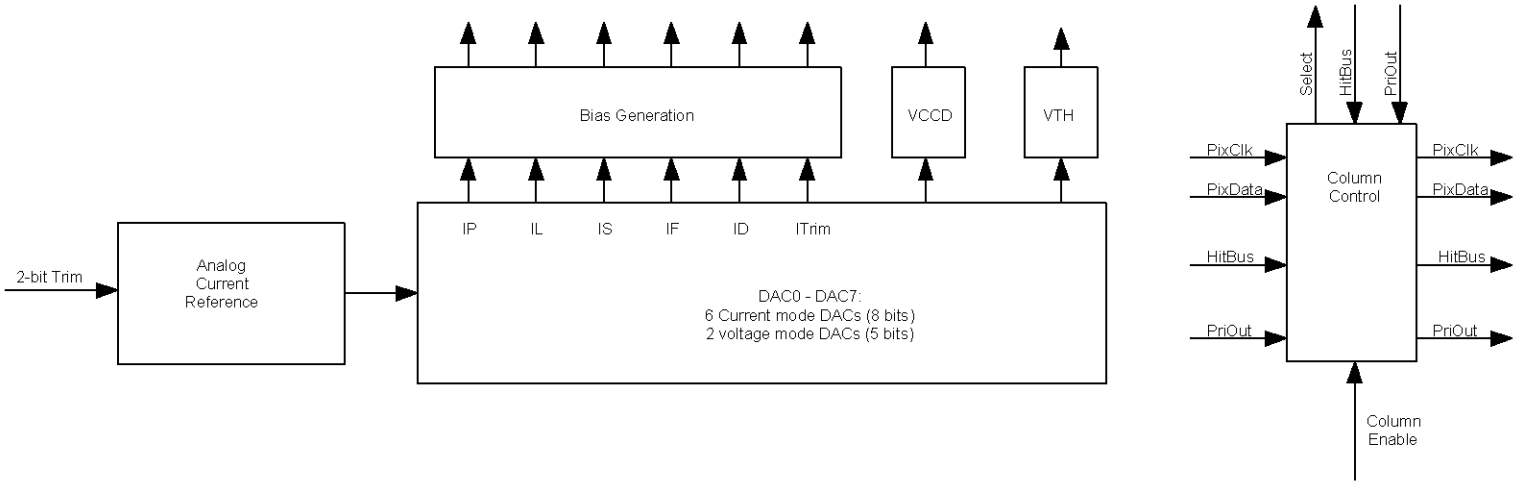

- A V_{BE} reference is used to supply a 16µA reference current to the current mode DACs. It has 2-bit trim capability to ensure roughly +/-1µA output variation for parameter variations within DMILL corners (this is a convenience).
- The current mode DACs are a rad-tolerant 8-bit design with good linearity.
- The two tunable voltage supplies use resistor chains to control output amplifiers which supply the necessary voltages.
- A single column enable bit controls the major operations of a column pair. It allows bypassing a column pair in the pixel shift register chain, bypassing a column pair in the HitBus FastOR net, and bypassing the sparse scan readout of a column pair when transfering an event out of the FE chip.

Digital readout

Summary of the requirements:

- Make a unique association of each hit pixel with a 40 MHz beam crossing.
- Store hits in pixel array for L1 latency period, which is expected to be about 110 crossings, but could increase to 130. Hit losses from front-end deadtime plus any losses during digital readout should be less than 2% at design luminosity and at the maximum L1 rate of 100 KHz.
- Make a modest TOT measurement by counting time differences between leading and trailing edges in 40 MHz units.
- Simulations for the readout architecture exist, driven by the full GEANT simulation of ATLAS. The pixel occupancies are 1-2x10^{-4} hits/crossing for the outer layers and 5-6.5x10^{-4} for the B-layer. The latter corresponds to 20% column pair occupancy. These values assume 300μ long pixels.
- These studies suggest that the current architecture needs a 20 MHz column clock and 20 buffers per column pair to meet specifications for the outer layers. The B-layer requires at least 30 buffers, and may need other improvements. Further study is ongoing for this layer.

- There is only one error condition which occurs, namely overflow of the EOC buffers. In the case where the EOC buffer block in a given column pair overflows, hits are lost until a free buffer exists, and the error condition is stretched to cover a full L1 latency (covering all possible events which could have lost hits due to this condition). The error status is then transmitted in the EOE word.
- An additional mechanism for hit loss arises when hits take longer than the L1 latency to arrive in the EOC buffers. In this case, the hits remain there for an extra L1 latency, and cause potential background for other triggers. For 20 MHz column clock operation, this occurs for less than 1% of the hits, and is negligible.
- Assembly of 16-chip multi-chip modules requires excellent internal testing capability to ensure that at the wafer-probing stage, known good die are good at better than 99.5% confidence. This requires ability to inject known patterns to exercise internal storage (JTAG approach not practical) using digital injection.
- Design should minimize sensitivity to single-point failures disabling a full chip.

Specifications:

- Clock duty cycle is specified to be between 40% and 60% (high phase lies between 10ns and 15ns, or nominal +/- 2.5ns). This specification is exceeded by the present DLL-based DORIC clock regeneration scheme.
- All other timing signals (L1 and SYNC) are validated by the leading edge of the XCK signal, and hence their timing is not critical.

Block diagram of the basic column-pair readout:

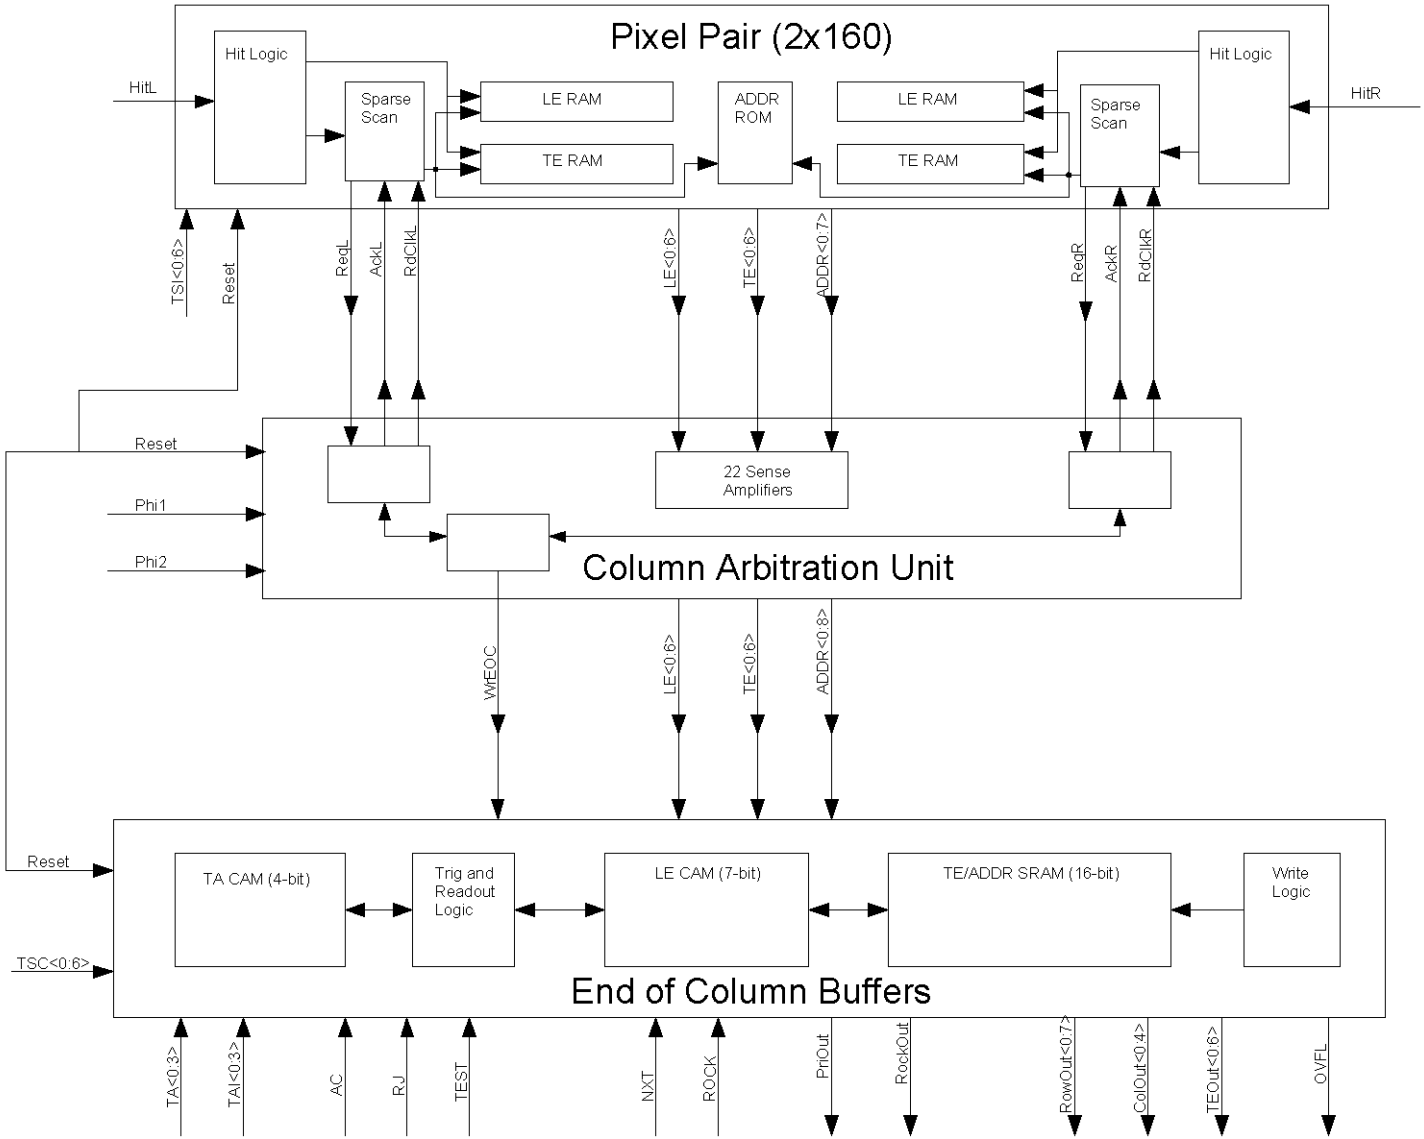

3

Block diagram of the Readout Control:

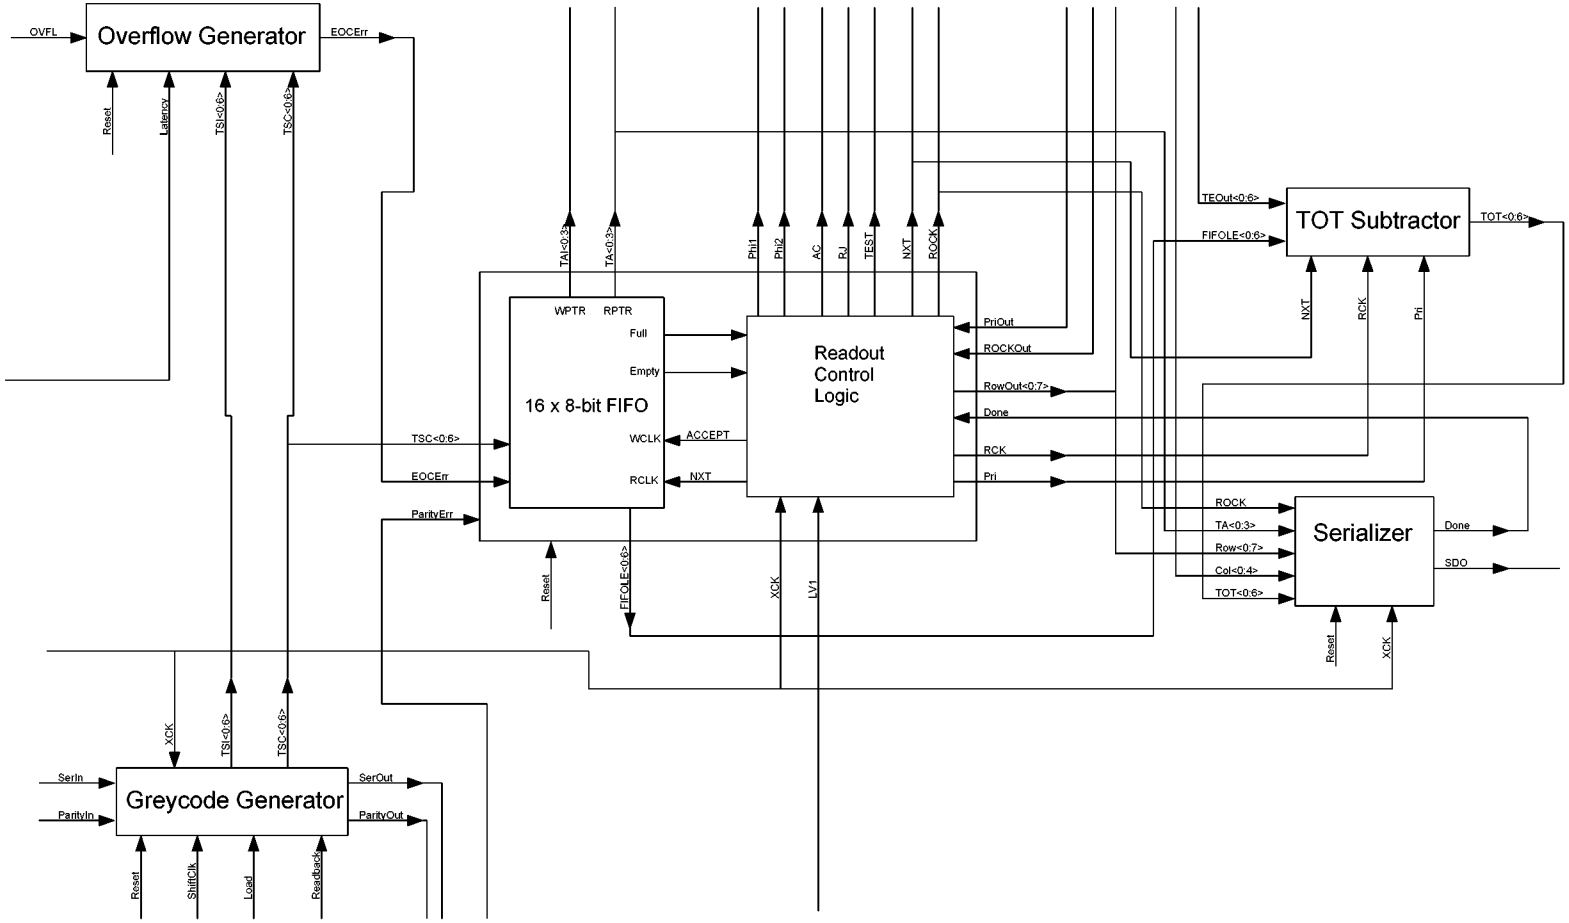

3

Block diagram of Command Decoder and Global Register:

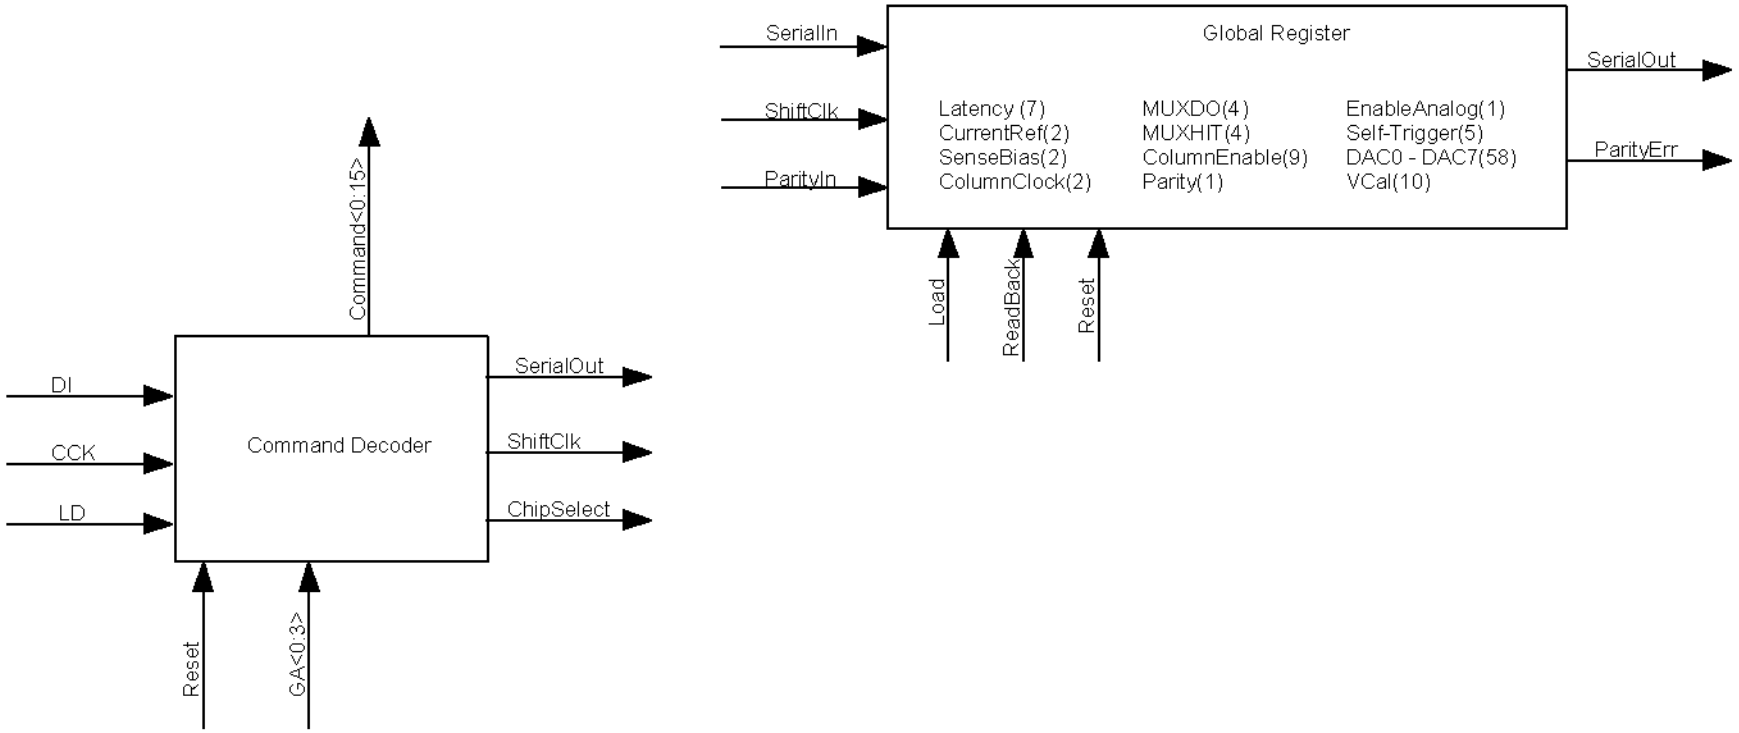

- Simple command protocol, based on a 21-bit command field, after which LD goes high and associated data may be transmitted. This supports 16 different, independent commands.
- Global Register controls overall operation of FE chip. Because of the critical importance of its bits, it does a parity calculation (chain of XORs) of the contents, and generates a Parity Error if the parity is not what is expected (due to SEU). This parity error is transmitted as part of the EOE word to flag such errors.

Block diagrams of remaining blocks in the periphery:

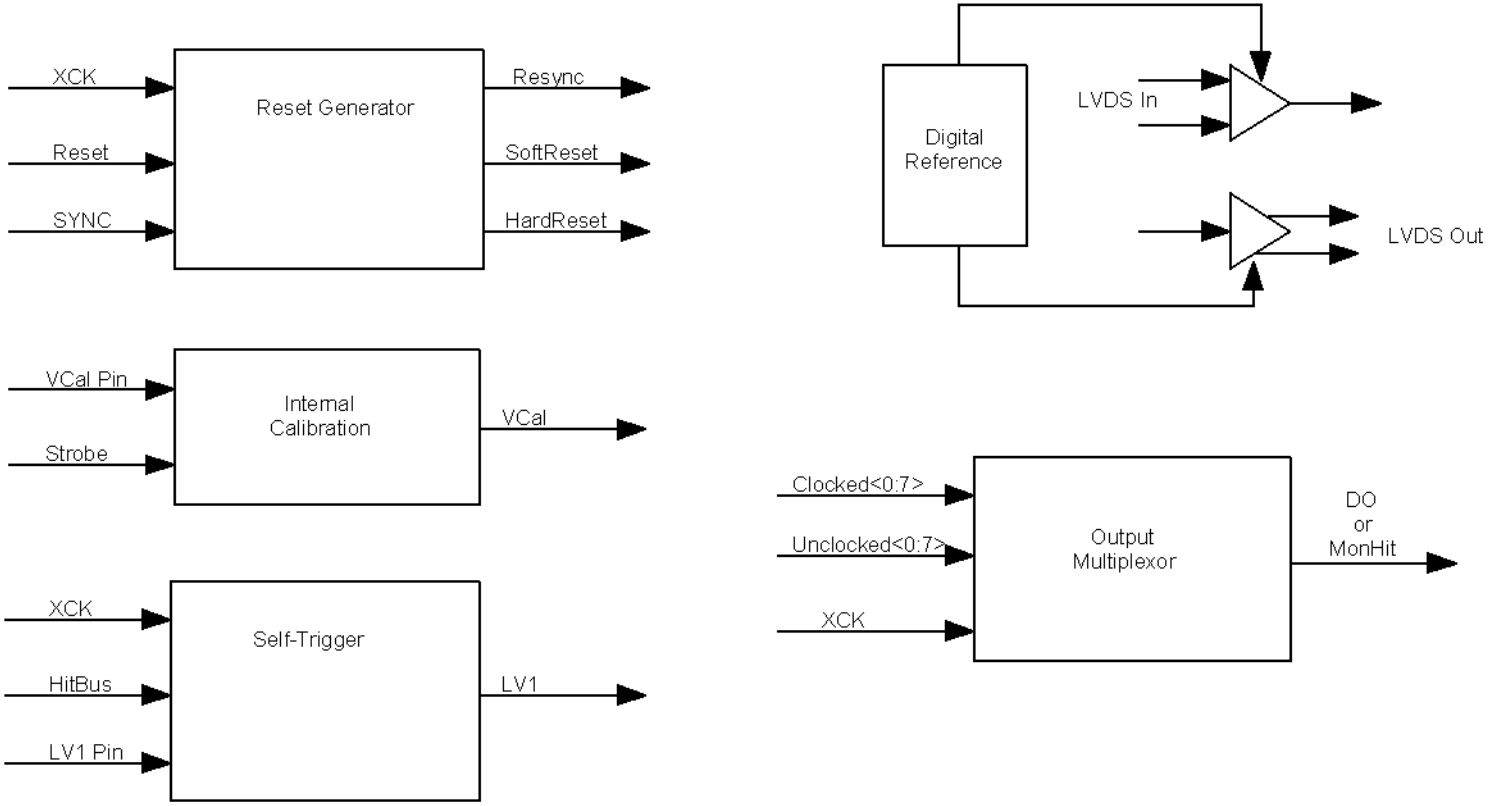

- The previous chips use a compensated FET switch to chop an incoming voltage and inject the corresponding charge into the pixel injection capacitors. A new circuit is under study to perform this task, using a steered current across a resistor. Simulations look promising.
- There is a reset generator that either uses the external RESET pin, or the width of the SYNC input in XCK clocks, to generate internal reset signals.

- There are three basic reset signals. The Resync makes sure that all FE on a module are using the same trigger number (it resets the trigger FIFO). The SoftReset puts the chip into the “empty” state for data, but does not alter any configuration information. Finally, the HardReset also resets all configuration information to default states (typically zero for most register bits).
- There is a self-trigger generator, which either passes the input LV1 through to the trigger processing circuits, or uses the internal FastOR signal to generate its own LV1 signal after a programmable latency. This allows the chip to be used in self-trigger mode with a source, and it will simply produce output data whenever it sees a signal on the internal FastOR.
- There is a dual 8-fold output multiplexor which selects which internal signal or data stream is transmitted off the chip through the serial output. The first eight inputs are synchronized with XCK, while the second eight are not. An identical multiplexor circuit is used both for the standard serial output pins, and for the MONHIT output pins.
- The LVDS driver/receiver circuits use a second internal current reference to define the drive current. The common mode voltage is referenced to the DVdd supply using a resistor. While not quite conforming to the LVDS specification, the design works well enough for local pixel module communication. Have performed corner simulations of the performance of a driver/receiver pair with realistic loading. The most critical signal (XCK) was found to be within specifications.

Summary

- In the last two years, the demonstrator effort has led to realistic prototypes of FE chips and MCC chips, which are the basic building blocks of the ATLAS modules.
- These assemblies have largely achieved our performance requirements. Further minor improvements are needed, and much greater experience with modules needs to be acquired.
- The work on the remaining electrical aspects of the modules, namely the opto-links and the power distribution over low-mass cables, will be prototyped this year.
- The next major step forward is the prototyping of the FE chips in rad-hard processes. If we can achieve the same level of performance as with the rad-soft chips, we will be almost ready to build ATLAS. We should have well-tested and understood rad-hard modules by the middle of 2000 (see schedule talk later).
- Many technical challenges remain in the on-detector electronics, particularly at the level of systems integration and production. However there are no indications of major roadblocks that will prevent us from building the ATLAS pixel detector.
- Further attention will continue to be placed on stretching the performance of all existing components to satisfy the longer term needs of the B-layer, which must cope with about 3 times the occupancy, and a corresponding increase in total dose, over the lifetime of ATLAS.