

Introduction and Overview for FE-I Design Review

K. Einsweiler, LBNL

Recent history: FE-D and FE-H

Summary of changes required/planned for 0.25 μ FE chip

Scope and goals of this review, and status of FE-I

Overview of system specifications and design of module

Overview of FE-I specifications and design at the top level

Recent History: FE-D and FE-H

- Rad-soft prototyping delivered functional chips in 98 (FE-B, FE-C). After successful testing, process of rad-hard conversion began.
- Submitted FE-D1 run, containing FE-D1 front-end chip, DORIC and VDC chips, and prototype MCC-D0 chip (plus test chips). Submission went out in July 99.
- FE-D1 suffered from minor design errors, and very poor yield in two circuit blocks. After considerable investigation, the low yield was related to technology problems (low rate of very leaky NMOS). Foundry never succeeded in isolating the problem, but proposed a series of special corner runs.
- Submitted FE-D2 run in Aug 00, with two versions of FE-D2. In one version, all design errors were fixed, but basic design (including dynamic logic blocks which suffered low yield) was left unchanged. Second version replaced low-yield blocks with static versions, and removed other circuitry (trim DACs) to make room. Run included full MCC-D2 (100mm²) and new DORIC and VDC chips as well.
- Corner runs gave no new information on yield/technology problems. Yield on static chip was better, and some studies are continuing with this chip.
- Began work on FE-H in Dec 99. Chip was almost ready to submit when we received notification of massive cost increases from Honeywell. With wafer cost of 20-30K\$, effort was abandoned before actually building a complete pixel chip.
- The failure of both of our traditional rad-hard vendors left us with 0.25μ approach, based on commercial process and rad-tolerant layout, as the only direction.

Conversion of Designs to 0.25 μ Process

Similarity between IBM and TSMC design rules:

- Implement common design environment where a single design can be streamed out and submitted to either vendor for fabrication.
- Cost of both is similar, but have monthly access with 8-10 week turnaround to TSMC foundry via MOSIS.
- In case of problems with one vendor, we have a back-up. Although TSMC has not been as thoroughly qualified for radiation as IBM, many tests have been performed by FNAL group, and our first results also look encouraging.

Updated CERN/RAL Standard Cell Library:

- Minor layout modifications made to provide compatibility with TSMC rules.
- Low noise mixed-mode design goals suggested separation of digital ground and substrate connection in the cells.
- We have also separated analog ground and substrate connection in individual analog blocks. Two substrate connections are joined at bottom of column, and both are connected to analog ground at the pad frame.
- Prefer to “over-fill” cells in non-routing layers (poly and active), so that problems do not arise later during integration.

Changes required for 0.25 μ FE chip

Analog Front-end:

- One significant constraint for using rad-tolerant layout rules with commercial deep-submicron processes is the limited W/L that can be achieved in the annular NMOS. The limit is $W/L > 2.3$, and for a typical “small, square” geometry, the limit is even more severe.
- This means that all NMOS will have a large g_m . In the FE-D design, long NMOS were critical in the preamp feedback and the discriminator threshold control. The preamp feedback circuit also provided leakage current tolerance.
- This W/L constraint has led to major modifications to the previous FE design.

Digital Readout:

- Concern about effects of SEU has led to the use of a special SEU-tolerant FF design for all configuration registers in the FE chip.
- Have analyzed all state machines to make sure that SEU will cause only transient effects, not long-term changes to chip operation that require re-initialization.

Overall:

- Operate chip supplies at 2.0V or less to maximize lifetime (due to hot carrier and oxide breakdown effects) and operating margins for safe operation.

Changes implemented in 0.25 μ FE chip

Analog Front-end:

- Implement improved threshold control using 8-bit current DACs instead of previous 5-bit voltage DACs.
- Implement two stage design to allow better optimization for timewalk performance.
- Extended present 3-bit TDAC (threshold trim DAC in pixel) to 5 bits, and added a 5-bit FDAC for feedback current trim in pixel. This will allow improved threshold control and trimming of TOT non-uniformities.
- Dropped second analog power supply to simplify services.
- Implemented improved charge injection design. Previous steered current design had problems with risetime on large charge scale.

Digital Readout:

- Implement all static design for both storage cells and logic blocks.
- Implement more synchronous and “top-down” design, and use standard cells and automatic place and route for logic blocks wherever possible.
- Increase timestamps (TSI) from 7 to 8 bits, increasing maximum latency to 6.4 μ s
- Implement column-based TOT calculation to allow digital timewalk corrections.

Present Status and Scope/Goals of this Review

Test Chip program:

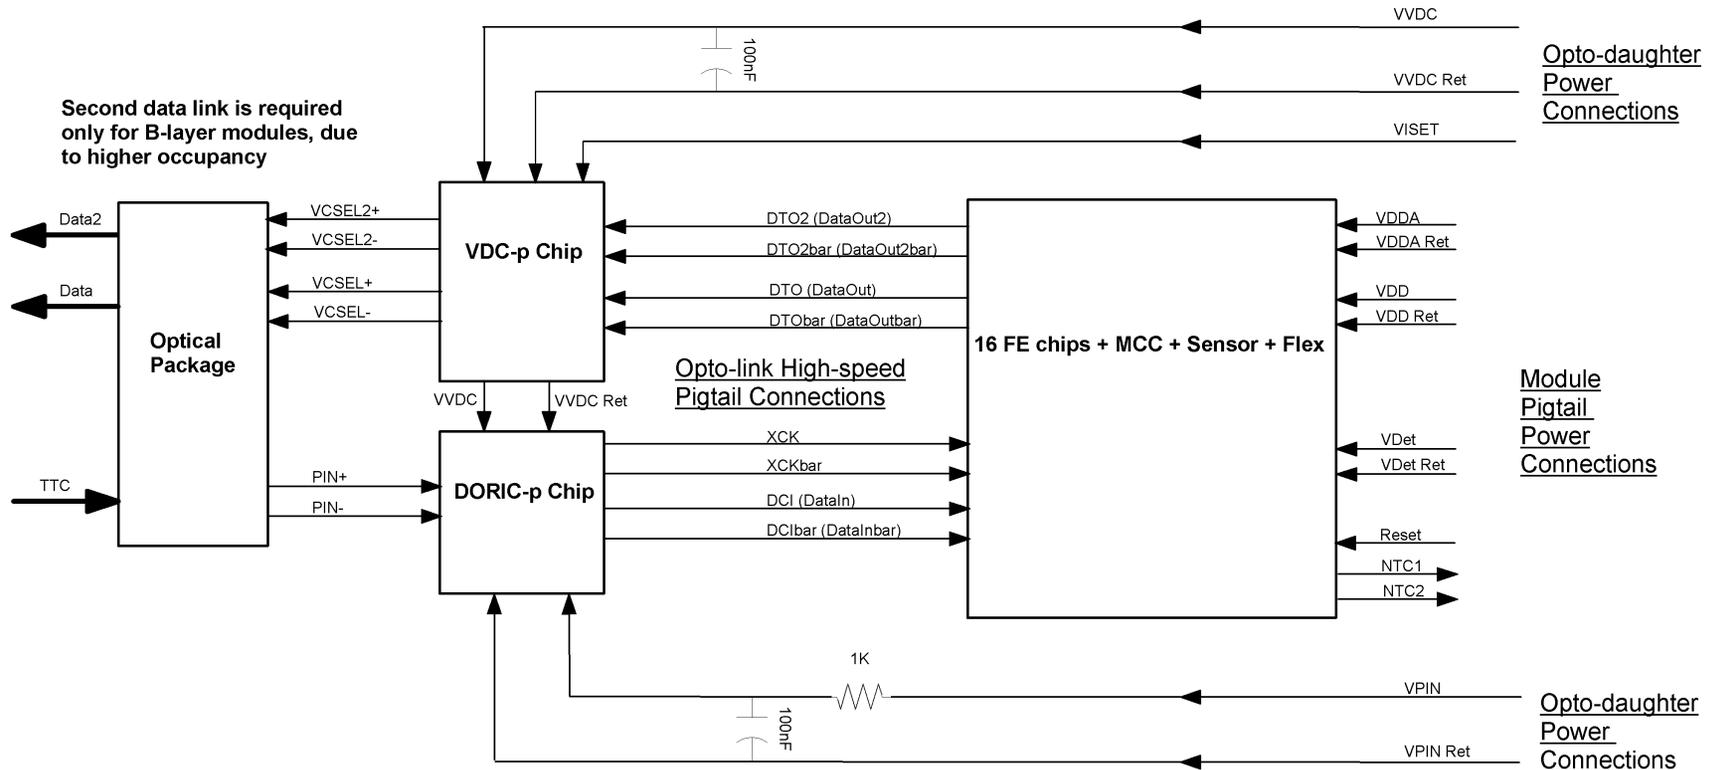
- In order to validate our designs and compare performance to SPICE, we have prepared a Digital Test Chip, submitted to TSMC in Jan 01.
- Due to technical difficulties, we did not complete full verification prior to submission, but we submitted anyway. Chip had many problems, with only one block (Pixel RAM block) which worked. This one block worked as expected, but no further effort was made to characterize this chip.
- Second set of submissions made in end Feb (IBM) and begin March (TSMC). This was our Analog Test chip, containing all basic analog and control blocks (current reference, DACs, Global Register, analog front-end and pixel control logic, LVDS I/O blocks) with approximately final layout.
- Will see first results from TSMC Analog Test Chip later in the review. It has been quite successful, but in the two weeks we have had the chips, we have not managed to fully digest or understand many details of their performance.
- Diced parts from IBM MPW wafers should be available by end of this week. We have our test setup along with us, and plan to try testing some chips next week. A further LBL irradiation is scheduled around June 21, with PS irradiations in July.

FE-I Design and Integration:

- The conversion of our FE-D/FE-H design started seriously in Oct 00. However, the 0.25 μ was a new technology for us, and we have made a large investment in IBM/TSMC compatibility and in our testchip program. This has been extremely useful, but has significantly slowed the integration of the complete pixel chip.
- The design for all critical elements of FE-I has been complete for some time. Many details of interfacing and integration, plus design of some minor blocks still remain. We will present the current status.
- For the analog blocks, it appears that the designs are in good shape, and much of the required performance has been demonstrated in our test chips. Three points of some concern exist at the moment, arising from our analog test chips. They may soon be understood. These are the LVDS driver, the injection chopper, and the large threshold dispersion.
- For the digital blocks, there is only minimal input from the first test chip, in which the pixel SRAM and sense amp readout were tested and worked as expected. There has been considerable simulation with SPICE (ELDO or HSPICE), TimeMill, and Verilog. However, the overall integration of the digital readout has not yet happened, and so there are no top-level simulations to show at this time.
- We would be in better shape for a review in about one month. However, most aspects of the final chip are well-advanced, and we ask for your considered advice on what we have done, what is still to do, and our current problem areas.

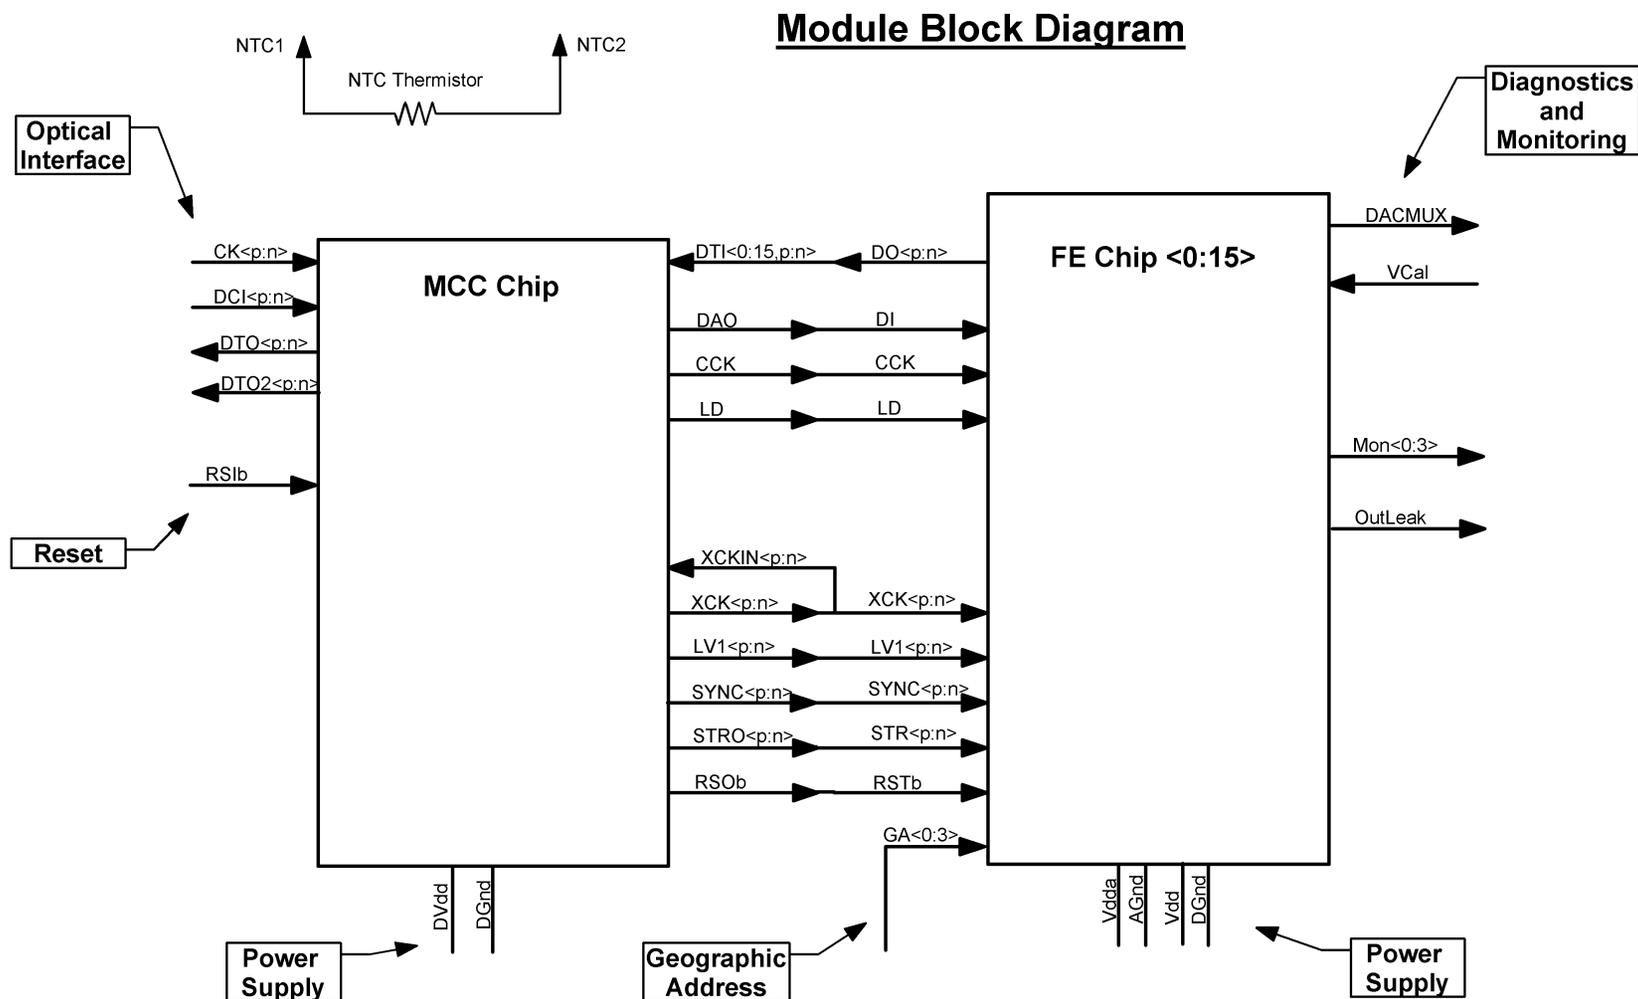
System Design of Pixel Module

Overall system architecture, including interconnections to opto-link and all power supplies:



- Optical package and DORIC/VDC mount on separate opto-card, up to 1m away.
- Module itself uses two LV supplies (analog and digital) and one HV bias supply.
- Communication between module and opto-card uses 3mA LVDS I/O

Block diagram of module itself:



- Two chip design, including a single controller and event-builder chip (MCC), and 16 front-end chips bumped to a single silicon substrate.
- Flex hybrid is used to provide interconnections above.

Features:

- Basic interface to the outside uses a 3-wire protocol (SerialIn, SerialOut, XCK), which maps onto the SCT opto-link protocol
- 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-to-point signals use 0.5 mA drivers (FE chips only), external or bussed signals use 3.5 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).
- There are no remaining analog signals between MCC and FE at this time. All FE chips have internal current references and adjustment DACs to control analog operating points, as well as calibration.
- 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.

Requirements Summary

Power budget (actually current budget for 0.25 μ):

- Present design of services and cooling based on typical and worst-case current and power analyses for DMILL chips. The FE-I should fit the same total current budget for the VDDA (analog) and VDD (digital) supplies.
- The present budgets for the digital supply are 25mA typical, 40mA worst-case. For the analog supply, they are 60mA (21 μ A/pixel) typical and 80mA (28 μ A/pixel) worst-case. The design of the low-mass power distribution system is very challenging, and we want to try hard to stay safely within these budgets. With the reduced voltage of the 0.25 μ process, the total power and cooling are much less difficult issues than they were in the past.

Geometry:

- The active die area for the FE chip is 7.2 x 10.8 mm, of which 7.2 x 8.0mm 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 50 μ pitch has been chosen as reasonably achievable. In the long direction, adequate resolution is obtained with a dimension of 300 μ - 400 μ . The present prototyping program has frozen the length at 400 μ .

Basic 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 $50\mu \times 400\mu$ for the pixel size, with pixels arranged in 18 columns of 160 pixels per column.
- The geometry was mirrored between columns, so that inputs for pixels in column 0 and 17 are on the outside. All other inputs are paired. This gives us 9 column pairs, with a common digital readout in the center, and analog cells on the sides.
- 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 active 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 $7.2\text{mm} \times 10.8\text{mm}$.
- An I/O pad structure of 48 pads, each consisting of a $100\mu \times 200\mu$ wire-bond pad, and a group of 4 bump-bond pads for MCM-D applications, was frozen.
- For final modules, only the central 30 bond pads are available for connections due to mechanical envelopes. Other 18 pads are available for diagnostics.

Feature List for FE-I

Design is largely a conversion of FE-D and FE-H designs.

Analog Front-end:

- 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 a differential second amplifier stage, DC-coupled to the preamp. The reference level (V_{Replica}) is generated in the feedback block, and should match the DC offset of the preamp with no input. The threshold control is applied using two currents to modify the offsets on the inputs to the second amplifier stage, allowing a large range for threshold control.
- The two-stage amplifier is followed by a differential discriminator which provides the digital output sent to the control logic.
- The control logic provides a 5-bit threshold trim capability in each pixel, plus a 5-bit feedback current trim capability for tuning the TOT response. There are four control bits, including Kill (shut down preamp), Mask (block entry of hit into readout logic), HitBus (enable output to global FastOR) and Select (enable injection of charge for testing). The HitBus bit also controls the summing of a current proportional to the feedback plus leakage current in the preamp, allowing monitoring of the feedback current, and of the leakage current from the sensor.

- A global FastOR net is created using all pixels enabled for this type of readout, and provides a self-trigger and diagnostic capability.
- All critical bias currents and voltages on the chip are controlled by internal DACs. There are 11 8-bit DACs for the analog front-end, and an additional DAC for the charge injection. The current DACs are referenced to an internal CMOS current reference, and the DAC values are loaded from the Global Register, and controlled via the Command Register.

Digital Readout:

- It uses an 8-bit Grey-coded 40 MHz differential “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 operates at transfer rates up to 20 MHz in order to meet our requirements. Differential 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 6.4 μ s in this chip). Forty buffers are available for each column pair. The coincidence with the L1 trigger is performed in this buffer. Hits from rejected crossings are immediately cleared.

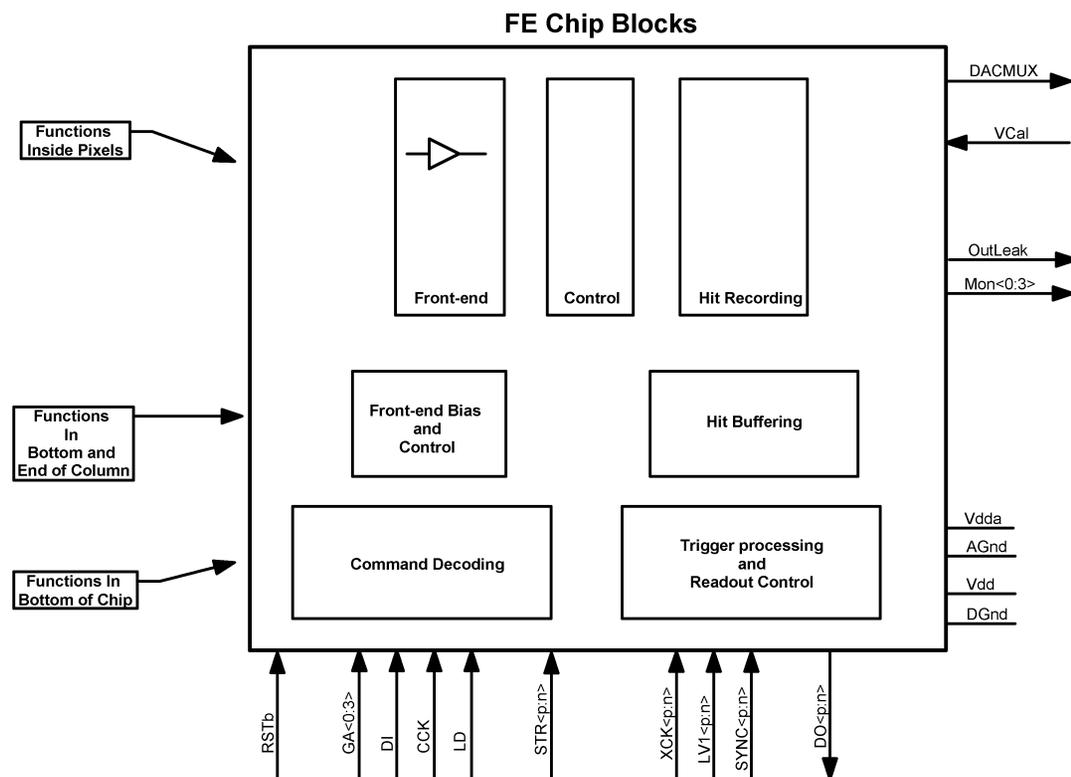
- A readout sequencer stores information on up to 16 events pending readout. As soon as the output serial link is empty, transmission of a new pending event begins. Essentially, sending a L1 trigger corresponds to making a request for the all hits associated with the corresponding beam crossing, which are then pushed off the FE chip to the MCC.

Control Logic:

- Global control of the chip is implemented using a simple command protocol. A Load signal controls whether input bits are interpreted as address and control, or as data. There is a 20-bit Command Register. Individual bits in this register implement specific commands (e.g. ClockGlobal, WriteGlobal, ReadGlobal).
- A Global Register, consisting of about 166 bits, controls Latency, DAC values, enabled columns, clock speeds, and several other parameters. This register is implemented as a shift register and a shadow latch with full readback capability. The shadow latch is SEU-tolerant since it contains critical configuration information.
- A Pixel Register which snakes through the active array provides access to the 14 control bits in the pixel (Select, Mask, HitBus, Kill, FDAC<0:4>, TDAC<0:4>). Readback capability is supported by transferring FF information back into the long shift register for readout. The 14 latches in each pixel are SEU-tolerant.
- Each chip on a module is geographically addressed, and its identity is controlled by external wire-bonds to avoid confusion. A broadcast mode is also supported.

FE-I Block Diagram:

Basic FE block diagram, expanded from module diagram:

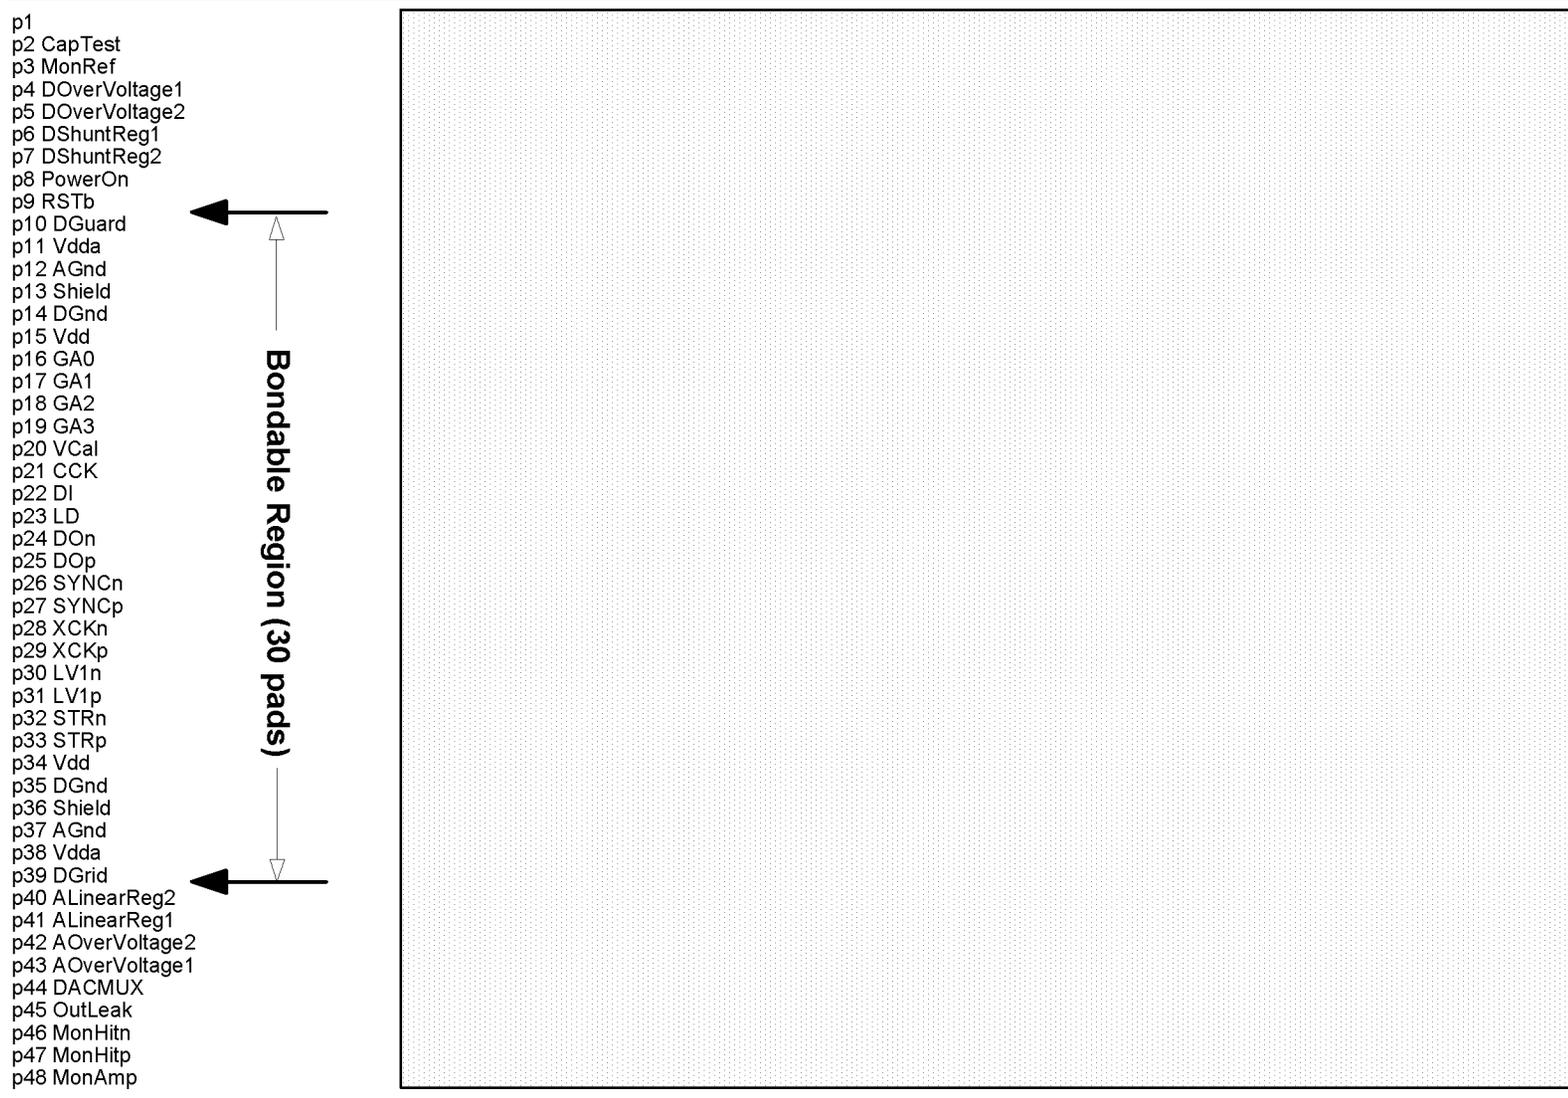


- Within the pixel, there is the front-end (preamp/discriminator), the control block, and the readout block.
- Just below the active pixel matrix is the biasing and control for the front-end blocks, and the data formatting and buffering for the readout blocks.
- Finally, there is the overall readout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

- 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. An analog voltage input (VCal) provides external calibration. Monitoring pins include FastOR, DACs and test pixel output.

FE-I Pinout and Geometry

Sketch of pin assignments and overall geometry of die:



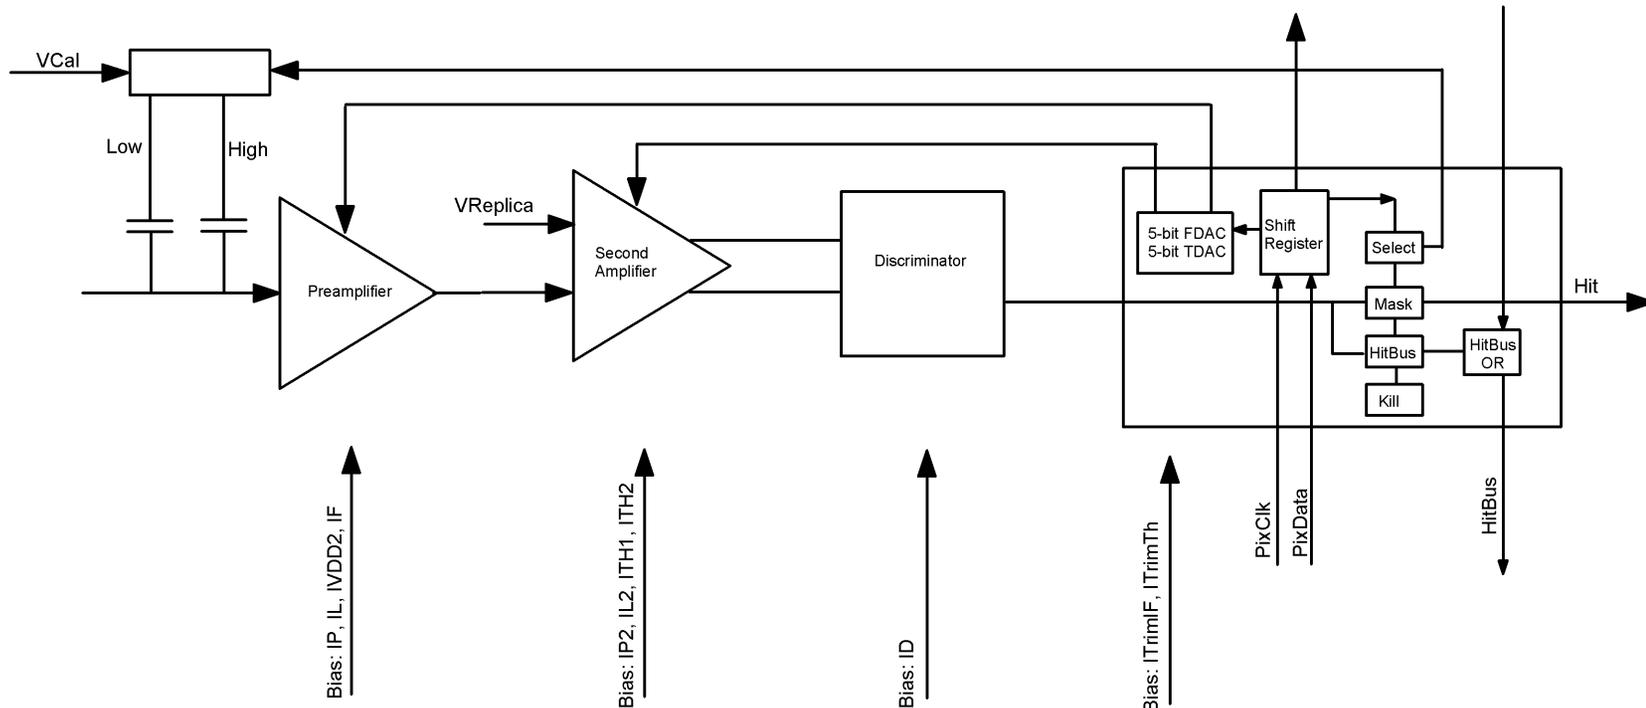
Front-end, biasing and control

Summary of the requirements:

- A nominal capacitive load of 400 fF is expected, roughly half to ground (parasitic) and half to the nearest neighbors (inter-pixel). Good performance should still be obtained with loads above 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 280 μ silicon, and the B-layer should be 200 μ .
- The expected signal after the lifetime dose of 10^{15} n-equiv/cm² is about 6Ke with 200V bias, and about 10Ke with 600V bias. We 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 4Ke (this would be 2.5Ke for the lower bias voltage). This requirement is now defined using a maximum timewalk relative to some large reference charge of 20ns, in order to allow for additional timing dispersion between channels on a chip and module, as well as between modules. This could be achieved by for example setting a 3Ke 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.

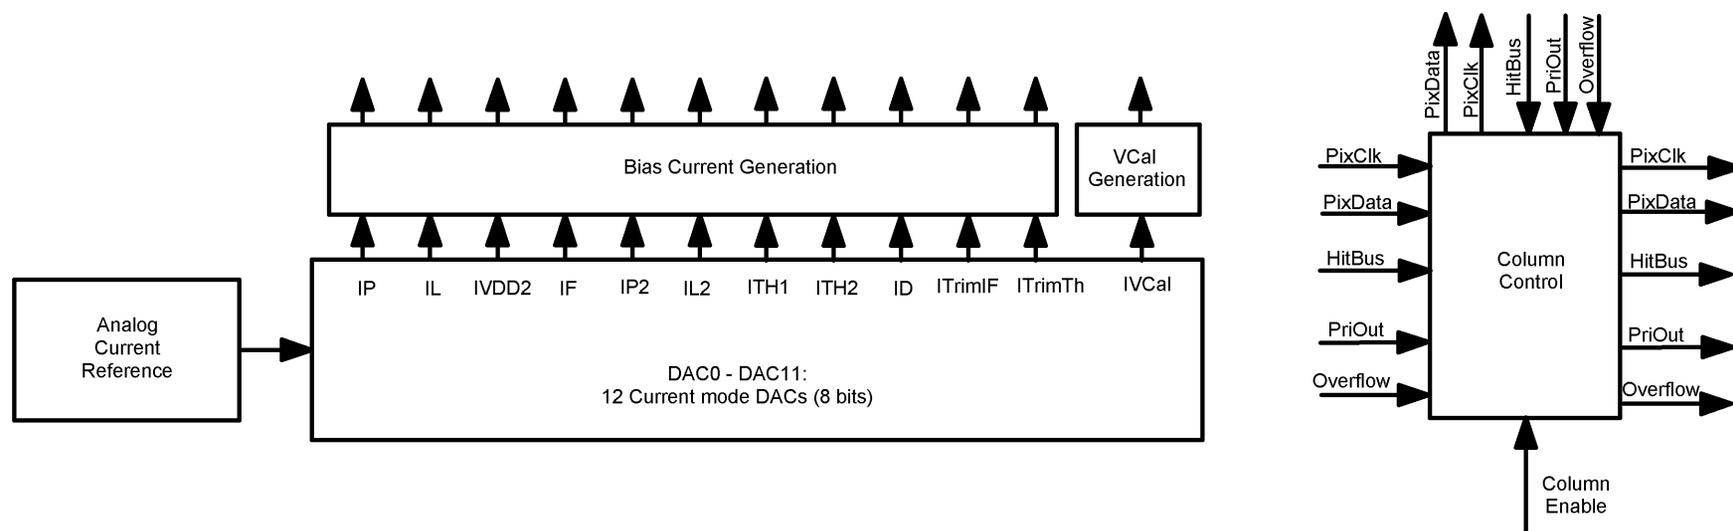
- Noise should be less than 300e and threshold dispersion less than about 200e, leading to an overall threshold “variation” of less than 400e.
- Leakage current tolerance should be at least 50nA per pixel, without significant changes in operating performance, and independently achieved for each pixel.
- 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 $2\mu\text{s}$ is required for the outer layers, and $0.5\mu\text{s}$ 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 at least 0 - 6Ke is needed.
- A two calibration injection capacitors of roughly 2.5fF and 25fF should be included in each pixel, giving a low range for threshold and noise measurements and a high range for charge calibration, timewalk, and crosstalk measurements.
- We do not know whether real diode input protection is needed. In all present chips, no explicit input protection is provided, and no identifiable problems have been observed.

FE and Control Blocks:



- Preamp has roughly 15fF DC feedback design, and 15ns risetime. Input transistor is a PMOS with W/L of 25/0.6 μ , and operates at about 8 μ A bias.
- Feedback circuit generates reference voltage for a differential second stage (VReplica). Feedback current is 3nA for a 1 μ s return to baseline and 20Ke input.
- Threshold control operates by modifying the offsets at the input of the second amplifier.
- Discriminator is DC-coupled and differential. It uses a bias current of about 5 μ A.

FE Biassing and Control Blocks:



- A reference circuit is used to supply a $4\mu\text{A}$ reference current to the current mode DACs ($1\mu\text{A}/\text{bit}$, but two LSB generated by simple mirrors).
- The current mode DACs are a rad-tolerant 8-bit design with good linearity.
- The biassing circuits are located directly on top of the DACs, and mirror the current down so that 64 DAC counts provides the nominal bias current. The biases are distributed as V_{gs} voltages, with matching mirrors in each pixel.
- 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, bypassing the sparse scan readout of a column pair when transferring an event out of the FE chip, and bypassing the buffer overflow generation from the End of Column buffers.

Digital readout

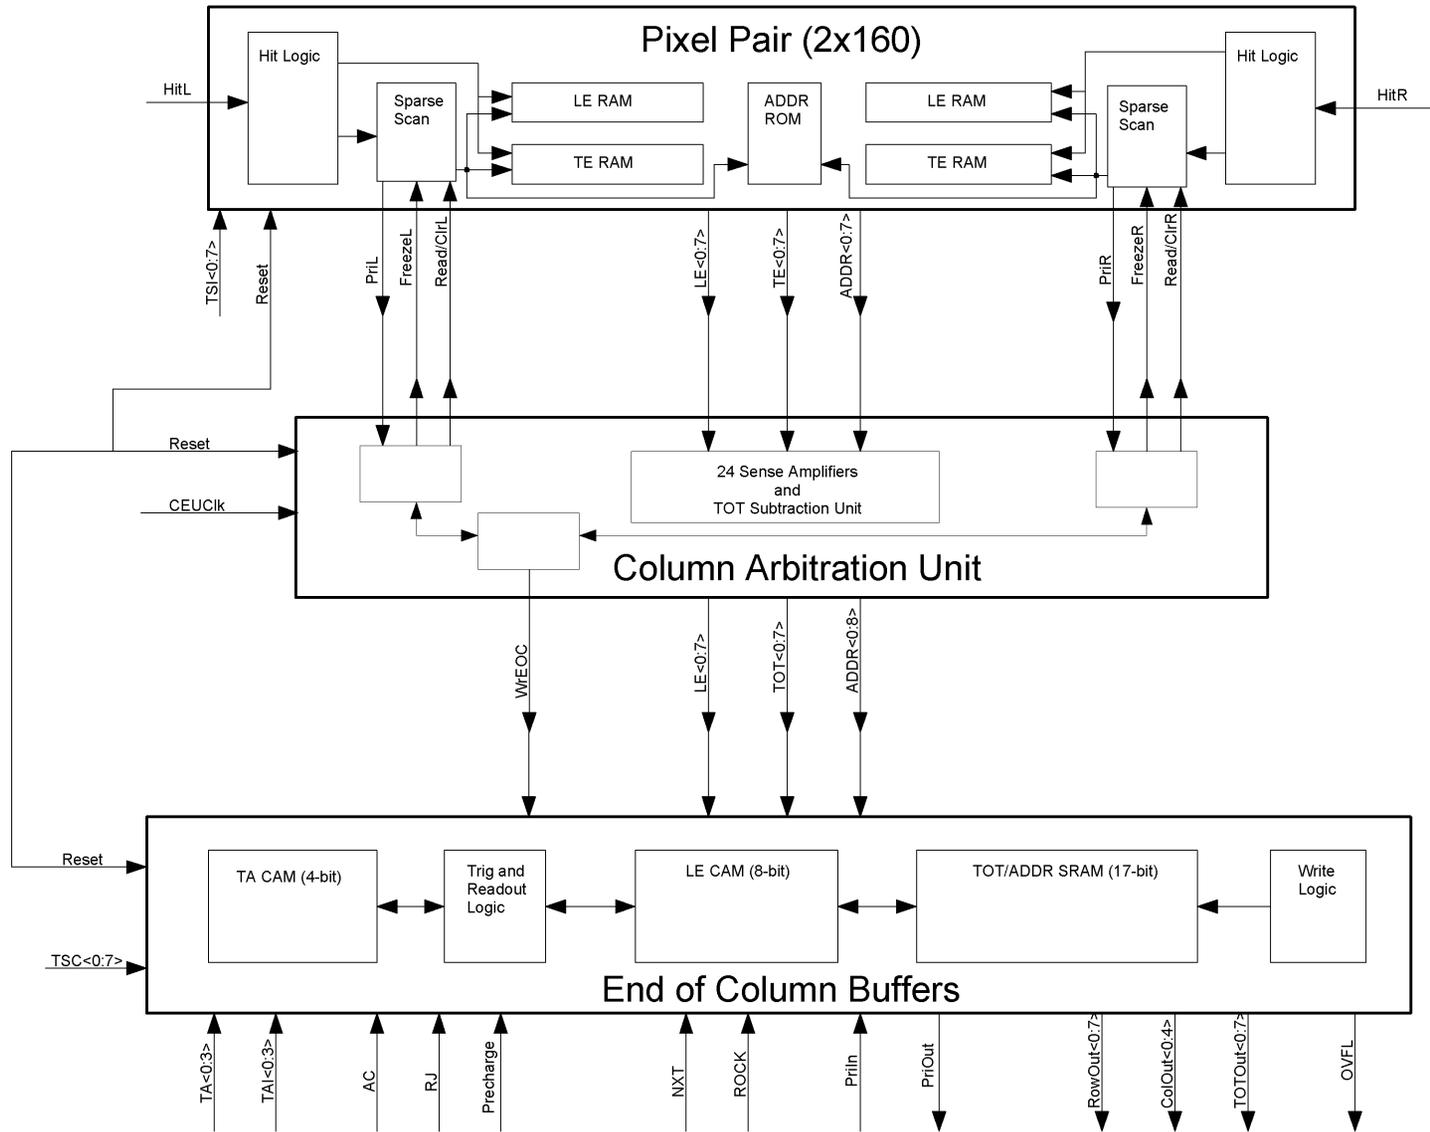
Summary of the requirements:

- Make a unique association of each hit pixel with a 40 MHz beam crossing.
- Store hits in pixel array for L1 latency period, which can extend up to $3.2\mu\text{s}$.
- Make a modest TOT measurement by counting time differences between leading and trailing edges in 40 MHz units.
- Simulations for the current architecture exist, driven by the full GEANT simulation of ATLAS. This suggests that the current architecture needs to operate with a 20 MHz column clock rate and have 25 buffers per column pair in order to provide safe operation of the outer layers. The B-layer requirements are more stringent, and require something like 40 buffers. Further study is underway for this layer.
- There is only a single 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.

Specifications:

- Clock duty cycle specified to be between 40% and 60% (high phase lies between 10ns and 15ns, or nominal +/- 2.5ns).

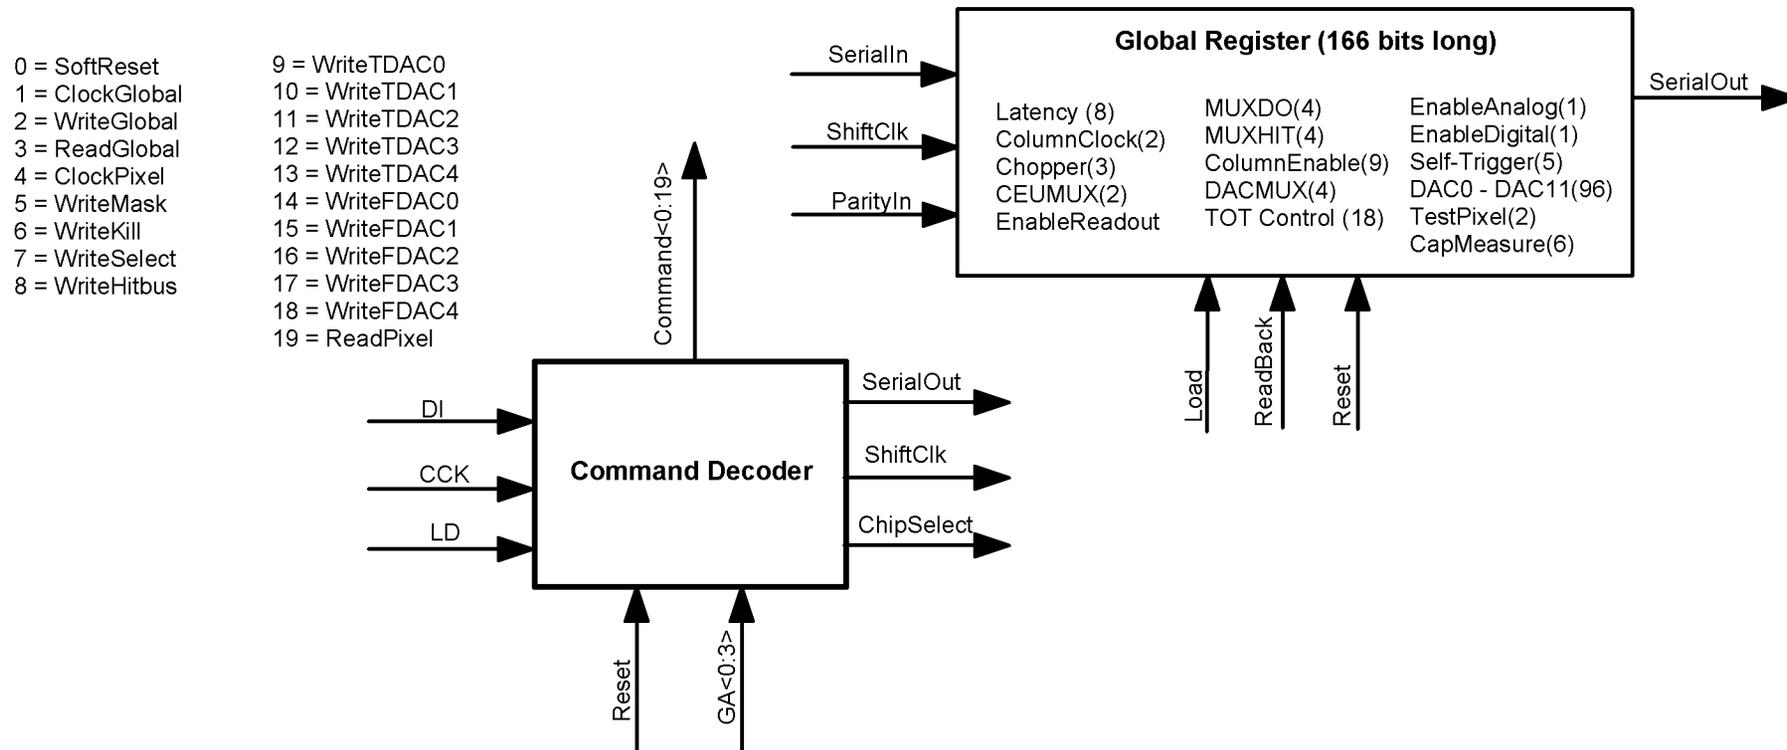
Block diagram of the basic column-pair readout:



Summary of basic steps in readout of pixel data:

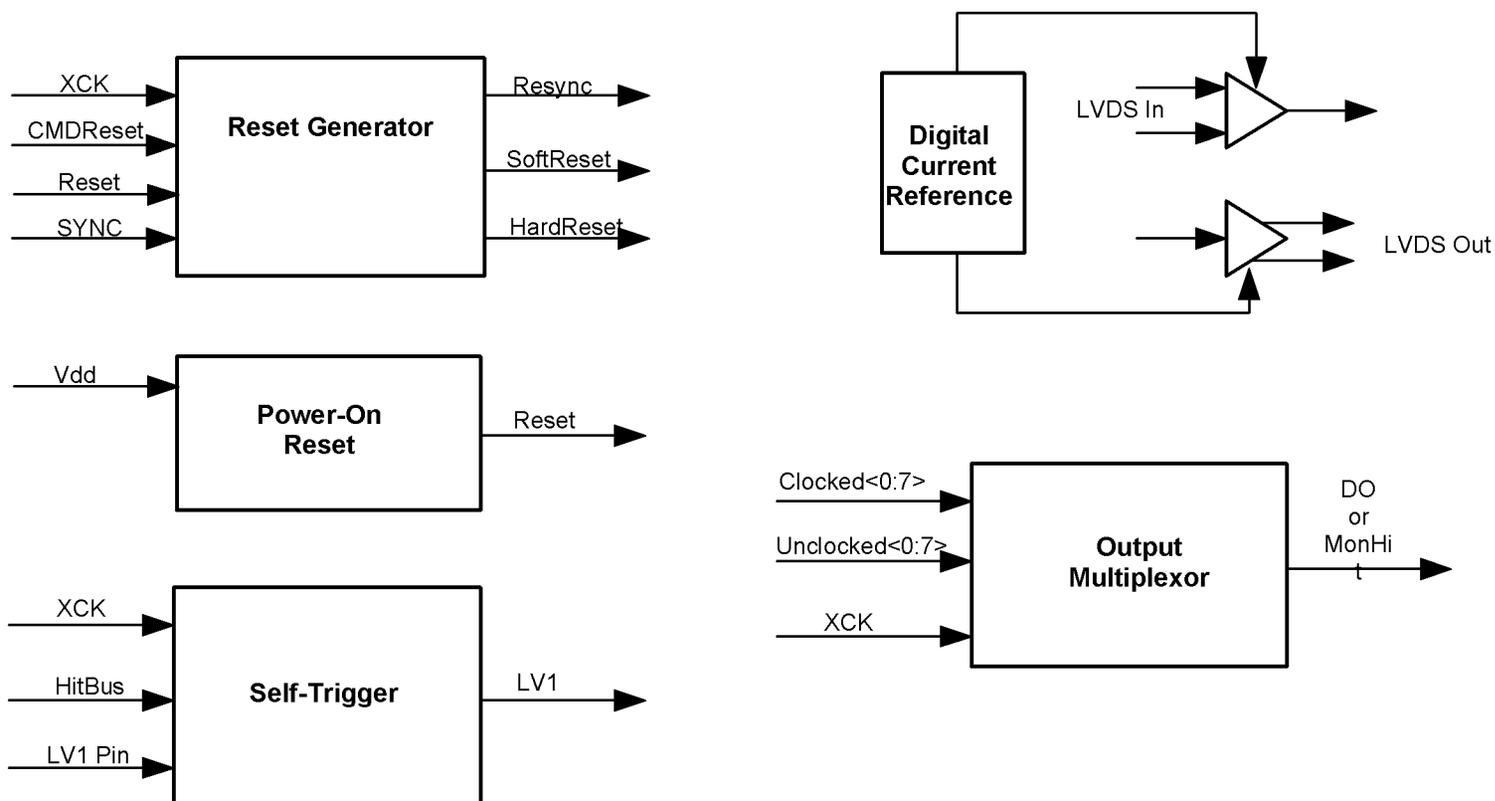
- Transfer hit information (LE and TE timestamp, plus pixel row address) into an EOC buffer. This operation begins when data is complete (after discriminator trailing edge). The transfer of hits from a column pair is synchronized by the CEU in the bottom of column, which operates at a speed of 5, 10, or 20MHz.
- Hit information is formatted by the CEU. Formatting includes TOT calculation (TE-LE subtraction) in all cases. Optionally, a digital threshold may be applied to TOT, a timewalk correction may be applied (write hit twice if below correction threshold, once with LE and once with LE-1), or both. These operation are pipelined to minimize deadtime, but EOC writes cannot occur faster than 20MHz.
- Hit information is written to the EOC buffer, and waits there for a corresponding LVL1 trigger. If a trigger arrives at the correct time (checked using LE timestamp of hit), the data is flagged as belonging to a particular 4-bit trigger number. Otherwise it is reset and the buffer is freed.
- Once the chip has received LVL1 triggers, the trigger FIFO will no longer be empty. This initiates a readout sequence in which the EOC buffers are scanned for the presence of hits belonging to a particular trigger number. If hits are found, they are transmitted to the serializer. After all hits for a given trigger number have been sent, an End-of-Event word is appended to the data stream.
- All of these operations occur concurrently and without deadtime, with all column pairs operating independently and in parallel.

Block diagram of Command Decoder and Global Register:



- Simple command protocol, based on a 5+20-bit command field, after which LD goes high and associated data may be transmitted. This supports 20 different, independent commands.
- Global Register controls overall operation of FE chip. Because of the critical importance of its bits, the actual latched values are stored in special SEU-tolerant FF. First measurements with 0.25 μ prototypes indicate the SEU flip rate for these bits is less than one flip per bit during the lifetime of ATLAS.

Block diagrams of remaining blocks in the periphery:



- There is a reset generator that either uses the external RESET pin, or the width of the SYNC input in XCK clocks, to generate three internal 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 zero.

- There is a Power-on Reset which makes sure the digital part of the chip is powered up into the Reset state. This clears the Global Register, and in particular, clears the EnableReadout bit, blocking XCK distribution to digital readout circuitry. The result will be very low digital power consumption, suitable for performing basic module connectivity tests without requiring operation of liquid cooling system. Basic connectivity checks do not require analog power.
- 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 produce output data once it has been armed by a previous LVL1 trigger, and 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 pin (DO), and for the MONHIT monitoring output pin.
- The LVDS driver/receiver circuits use a second internal current reference to define drive current and receiver bias. The common mode voltage is referenced to the Vdd supply using resistors. Output drive is 0.5mA for low power. Circuits are very stable under process and power supply variations, and easily meet specifications at 2V supply voltage.

Summary

- Will see details of our test chip program, and designs, layout, and simulation results from the various blocks in the complete FE-I chip.
- Some aspects of the test chip results are confusing to us, and we hope to understand more in the next few weeks. We will evaluate the IBM test chip, and do more irradiation work prior to submitting the engineering run.
- Much work on integration and verification remains to be done. We will discuss how we intend to do this, but it is still a work in progress.
- We welcome your comments and advice on everything we will present during the next two days.
- Unlike previous pixel reviews, we would also very much like a written report of your findings. If there are critical issues that we have not properly addressed in this review, we expect you to ask for responses.