

Summary of FE-I2 Measurements

K. Einsweiler, LBNL

Summary of register testing of FE-I2 (Bonn + LBL):

- Problems in Command and Global Registers

Summary of additional FE-I2 testing in LBL:

- Basic tests of new digital features in FE-I2 (HitParity, Self-Trigger, Hitbus Scaler)
- First studies of analog performance to investigate new features (threshold control, bias compensation, and auto-tuning).
- This section of the talk is almost entirely the work of John Richardson.

Next Steps to “Fix FE-I2”:

- Six wafers from ATLASPIX2 run are held at back-end processing. It is possible to make metal mask changes and then complete processing on these wafers.
- In parallel, work on real FE-I3 with improved timing analysis and layout for Digital_Bottom block.

Summary of FE-I2 Testing

Initial Testing in Bonn:

- John, Aldo, and myself joined Tobias, Markus, Joern, and Ivan in Bonn and began testing the first wafer on May 22.
- We immediately encountered the problem that the Global Register test always failed under normal operating conditions. After many hours of confusion, during which time the only signs of life were the output from the digital current reference and the change in digital current when the 40MHz clock was applied, we discovered that the control logic would work if the digital supply voltage was lowered to 1.6V.
- Further testing revealed that most of the chip worked fine at the lower supply voltage, but something went wrong with the Global Register when the digital supply voltage was raised. Many different chips were tested on two wafers, all with essentially identical results. This indicated that the yield was high, but there was a systematic problem.
- Subsequently, the Bonn group, with help from Aldo, probed a complete wafer. The analysis of the data has been shown by Aldo, and the yield at low VDD was more than 90%. The VDD threshold for failure was quite well-defined, and was at about 1.85V.
- Finally, measurements of the digital regulator built into FE-I2 showed that the lowest (default) output voltage setting was just barely too high to operate FE-I2.

Implications:

- The only way to use FE-I2 for module construction is to operate it with VDD less than about 1.7V. If the digital regulator in each chip would have provided such an output voltage, we could have guaranteed that all chips in a module would work, and operated the MCC at a higher VDD. The major feature of FE-I2 that does not work at this low VDD is the highest frequency column clock readout, which is only required for high luminosity operation in the innermost modules.
- In the absence of this work-around, we would have to operate the full module at about 1.7V (and many chips would see closer to 1.6V). The present FE-I1 modules do not work at 40MHz at below 1.8V, despite the observation that the individual chips should work fine at this supply voltage. Therefore, it seems rather unlikely that complete modules will work at VDD=1.7V, particularly over the wide range of temperatures and radiation doses which are required for ATLAS.
- This means that there is essentially no possibility to use these chips for production modules. However, for prototyping purposes, we can modify the FlexV5 to separate the VDD for the MCC and the VDD for the FE chips, and potentially make working modules that can be irradiated and placed in test beams. This should be attempted as quickly as possible, and would allow us to evaluate the core functionality of the FE-I2 chip in the full module system environment.

Fault Diagnosis for FE-I2

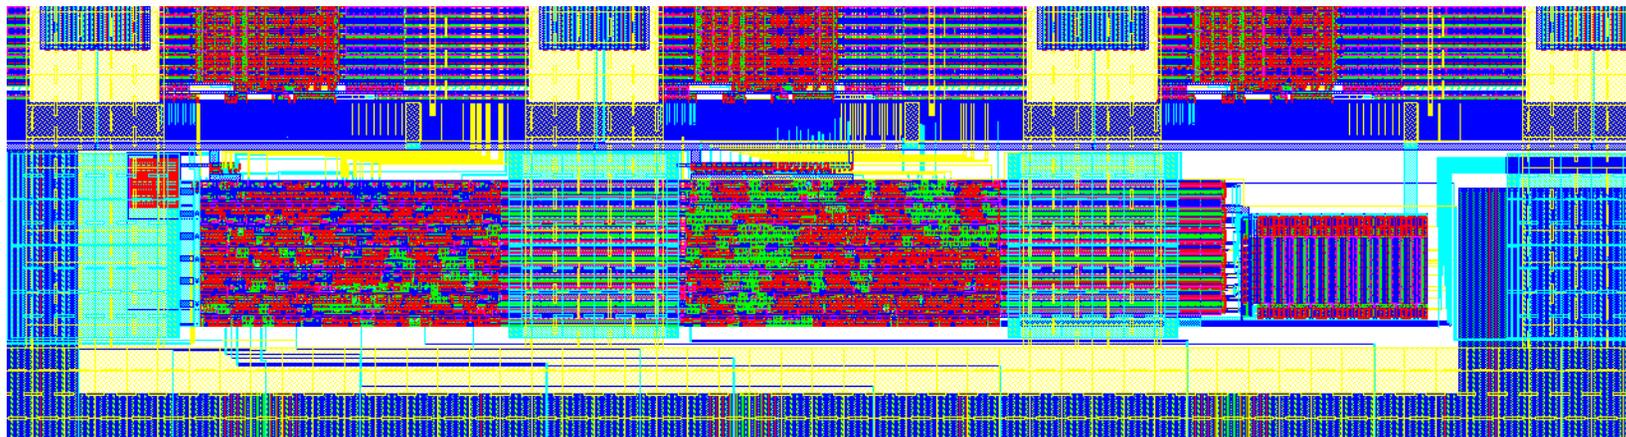
- Many further measurements have been performed to isolate the problem in FE-I2.
- Before plunging into details, describe basic control interface for FE-I2:

FE-I2 Control Interface:

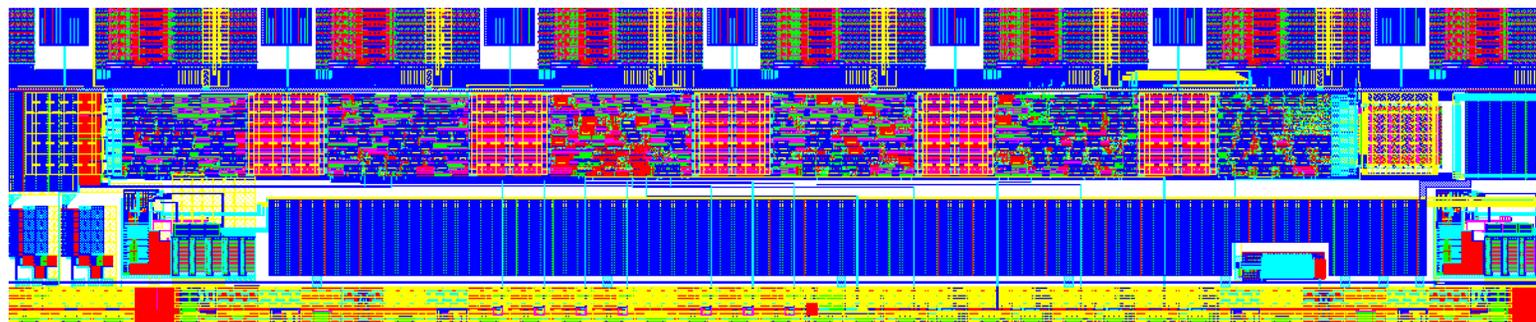
- Most of the control and global readout functions for FE-I (both 1 and 2) are implemented in a large place and route block at the bottom of the chip, named Digital_Bottom. Tools such as Synopsys and Silicon Ensemble are used to compile high-level Verilog descriptions of the functions into circuit blocks using a standard cell library.
- First key block is the Command Decoder (MainCtrl). This block accepts the CCK, DI, and LD signals that communicate the basic configuration commands to the FE chip, and acts upon them. It consists of a 29-bit register, of which 24 bits are for commands to load, write, or read various internal configuration latches, and 5 bits are for addressing information to select a chip within a module. The 24 command bits are latched, whereas the 5 bits of address information are evaluated on the LD high transition to generate an internal chip select signal.
- For FE-I2, this block has been SEU-hardened by implementing two copies of the basic shift register, and using triple-redundancy for the 24 command latches. This has significantly increased the size of this circuit block in FE-I2.

- Second major block is the Global Register. This is divided into 3 pieces internally. The first is a 30-bit ConfigReg located in Digital_Bottom. The second piece is an 8-bit shift-register plus scaler for counting HitBus transitions for hot pixel finding, also located in Digital_Bottom. The third piece is the remaining 193 bits which are implemented in a full-custom bit-slice design that is spread throughout the bottom of the chip, close to the items it must control.
- For FE-I2, this block has been hardened by using triple-redundant latches. In addition, the 30-bit ConfigReg uses special SEU-hardened DFF for the shift register, in order to allow the measurement of SEU cross-sections for these cells.
- In order to monitor the internal signals of the FE chip, there are two, almost identical, 2x8-bit MUXes. One drives the MonHit pin, normally used for the Hitbus output, and one drives the DO pin, used for global register, pixel register, or event data. Using these MUXes, it is also possible to monitor the outputs of the two redundant shift registers in MainCtrl, and to monitor an intermediate point in the Global Register after the 38-bits of place and route register and before the 193 bits of full custom register. These monitoring outputs have been our main tools for analyzing the failures in the FE-I2 control section.

Layout Comparison of Digital_Bottom in FE-I1 and FE-I2:



In FE-I1, there are two regions of place and route circuitry, and a third region containin the trigger FIFO and its control logic.



In FE-I2, there are six regions of place and route circuitry, plus the trigger FIFO. This results in spreading circuit blocks across the regions shown here.

Measurements:

- Since everything appears to work correctly at 1.6V, the basic diagnostic technique is to set up a configuration, including the MUXes, at this low supply voltage. Then a particular operation is performed at a higher VDD, and the chip is returned to the lower voltage to interpret the results.

Command Shift Register:

- For the Command Register, this testing can be done fairly simply, because as long as the address used is different from the address expected by the chip under test, it will not respond to the command in any way, and the data is just transparently clocked in and out of the chip. However, it is still possible to see the shift register outputs on the MUXes. In this way, one can see that there are faults in the command shift registers themselves when operated at higher VDD.
- Different bit patterns have been used to explore the pattern sensitivity of the failure. It is observed that, even at the maximum reasonable VDD of 2.8V, only particular patterns of failures occur. In particular, isolated 1's are lost (become zeros) with 100% probability, whereas for a pair of 1's separated by zeros, only a single 1 is lost. In addition, there are no lost 0's (becoming 1's).
- These failure modes appear consistent with a single-point defect which can overwrite a 1 if it is followed by a 0 in the shift register. Such a problem could occur if the arrival time of the clock at bit $n+1$ minus the arrival time at bit n is quite a bit larger (about 1ns) than the transit time for the data from one bit to the next.

- This can occur in a place and route design if the clock routing is poor. Traditionally, shift registers with single-phase clocking distribute the clock in the opposite direction to the data in order to guarantee that such a race condition cannot occur. This approach has been used in the full-custom part of the Global Shift Register.
- Example of pattern clocked through Command Shift Registers at VDD=1.6V

Trace #3 is CCK

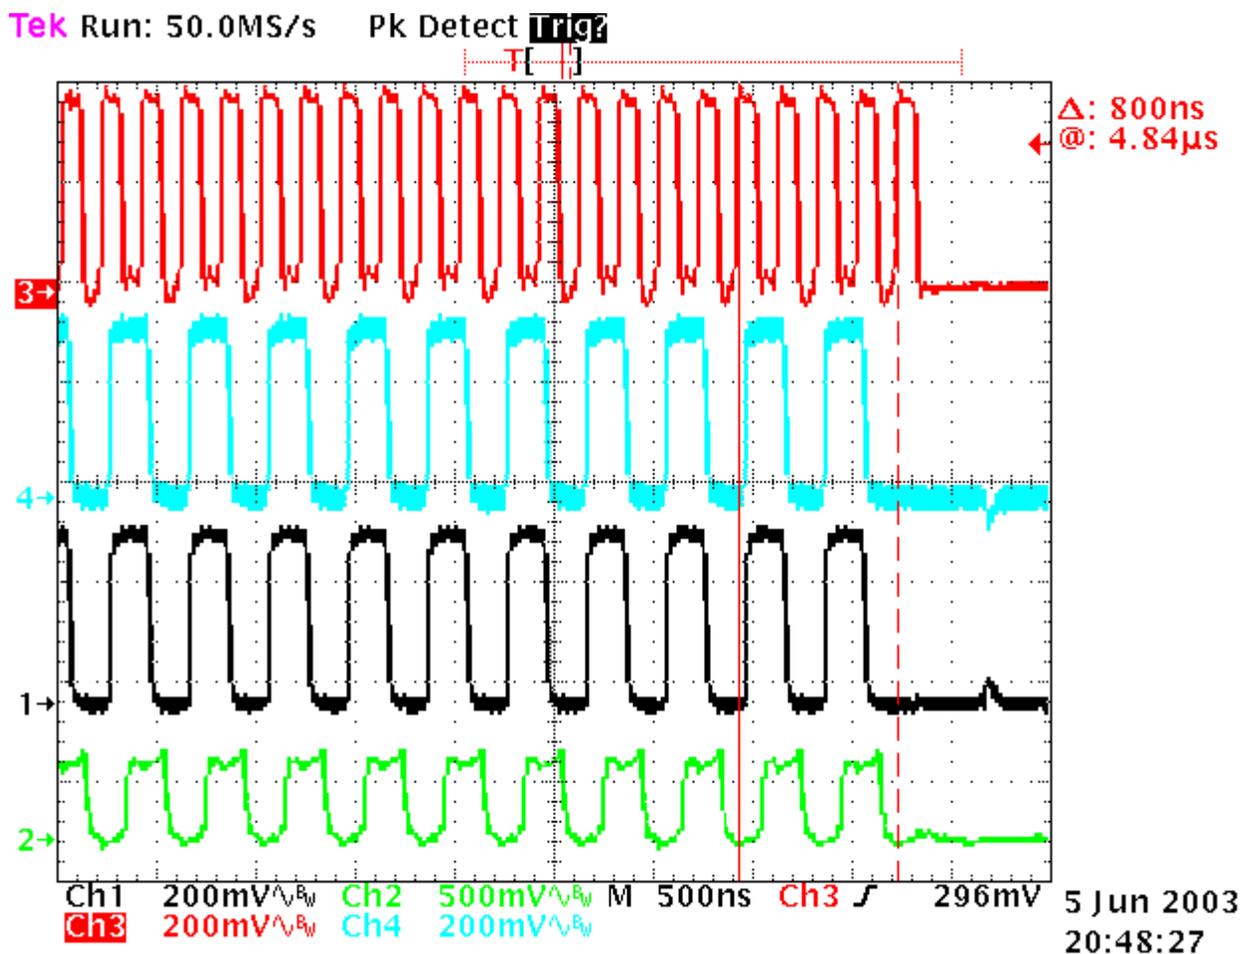
Trace #2 is DI

Trace #1 is MonMux (Serout2)

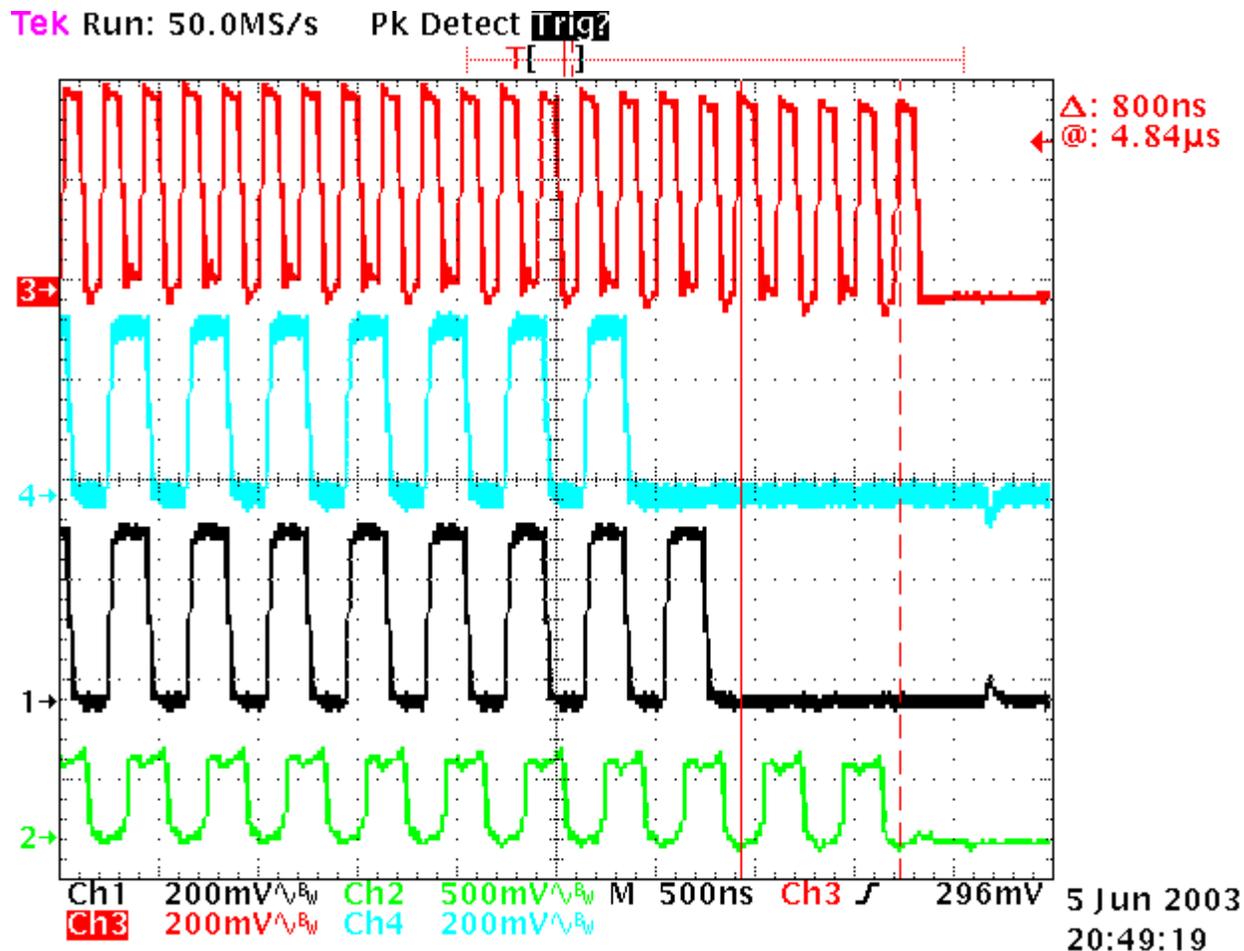
Trace #4 is DOMUX (Serout1)

A pattern of '5's was sent with alternating 0's and 1's.

No data has been lost in the shift registers.



- After pattern was loaded at VDD=1.6V, it is clocked in again at 2.5V. Output data was therefore shifted in at 1.6V and shifted out at 2.5V.



Trace #3 is CCK

Trace #2 is DI

Trace #1 is MonMux (Serout2)

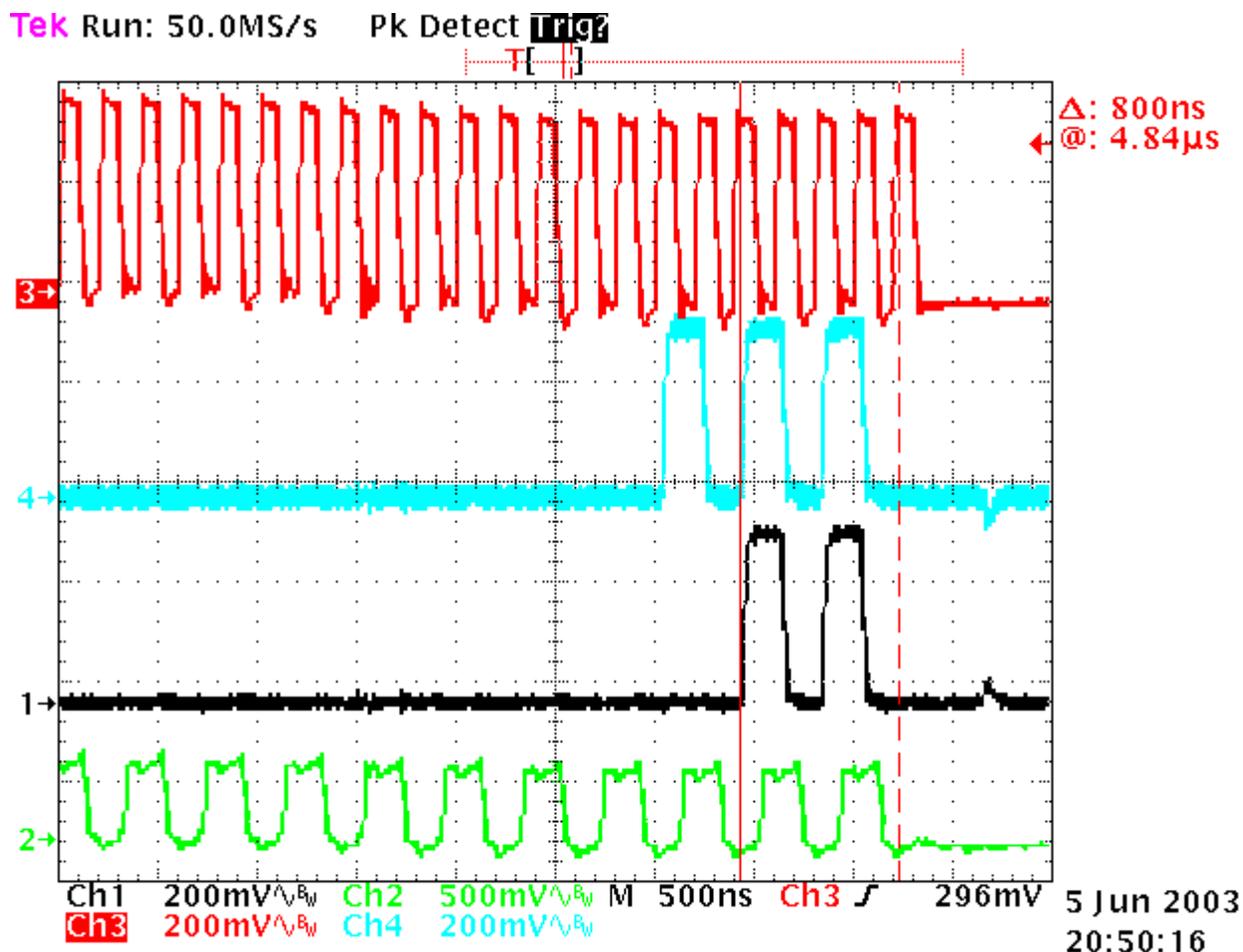
Trace #4 is DOMUX (Serout1)

A pattern of '5's was sent with alternating 0's and 1's.

No data has been lost in the shift registers.

- Observe loss of bits, beginning with 6th bit in Register #1, and 4th bit in Register #2. Clocking in pattern of 'A's indicates losses start in 5th bit. Conclude there is a single point loss between bit 6 and 7 in Register #1 and 5 and 6 in Register #2.

- After pattern was loaded at VDD=1.6V, clocked in again at 2.5V, it is then clocked in again at 1.6V. Output data was therefore shifted in at 2.5V and shifted out at 1.6V.



Trace #3 is CCK

Trace #2 is DI

Trace #1 is MonMux
(Serout2)

Trace #4 is DOMUX
(Serout1)

A pattern of '5's was sent with alternating 0's and 1's.

No data has been lost in the shift registers.

- Confirm previous single-point defect model, since all bits written at 2.5V after the defect are lost. Bits written before the defect are OK, and are eventually shifted out.

Global Shift Register:

- For the Global Shift Register, similar testing strategies can be used. A similar pattern dependence is observed, with loss of isolated 1's, and the loss of a single 1 after a string of zeros, but no losses of zeros.
- Example of start of pattern clocked through Global Shift Register at VDD=1.6V

Trace #3 is CCK

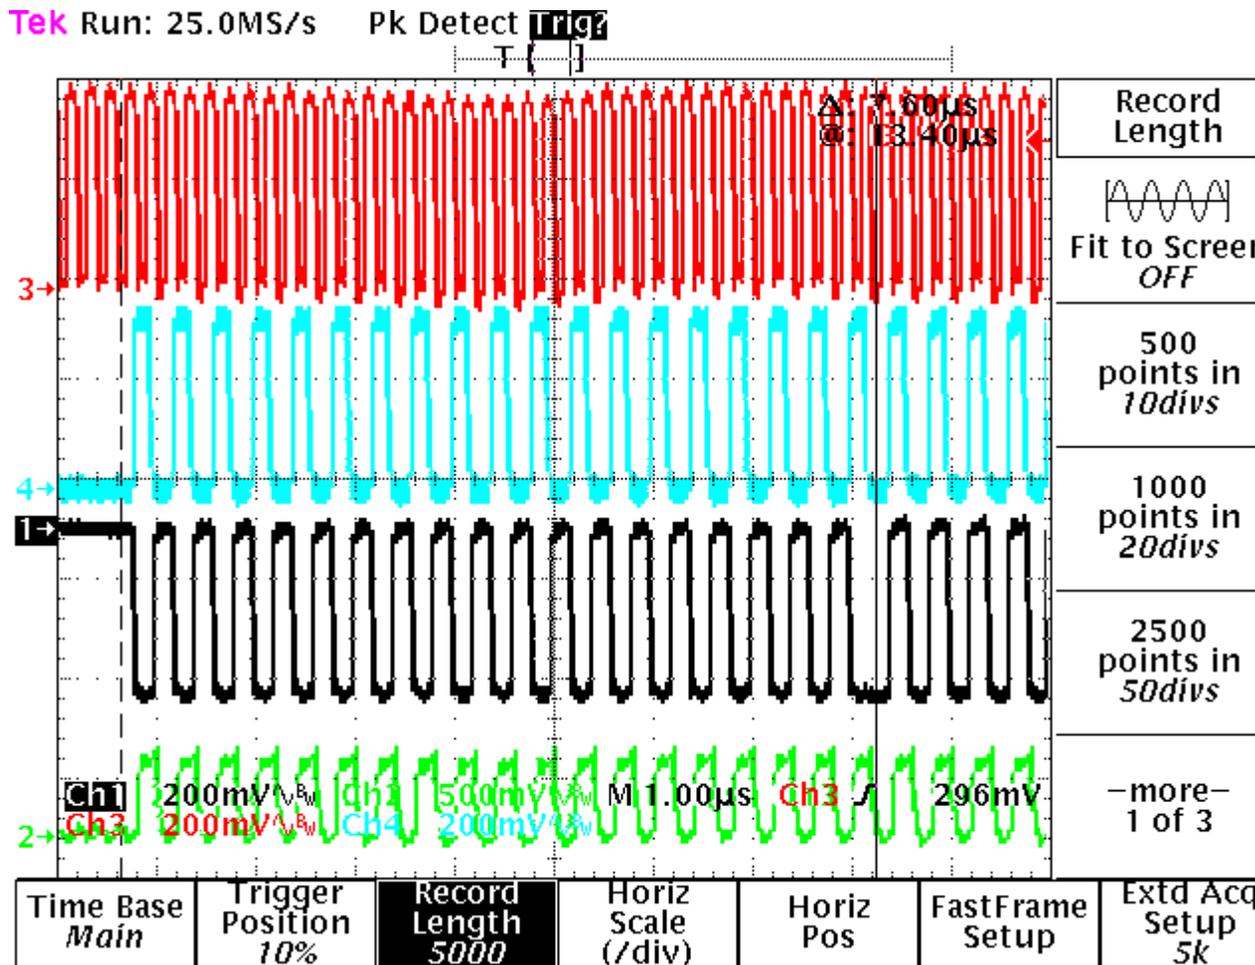
Trace #2 is DI

Trace #1 is register output after first 38 bits.

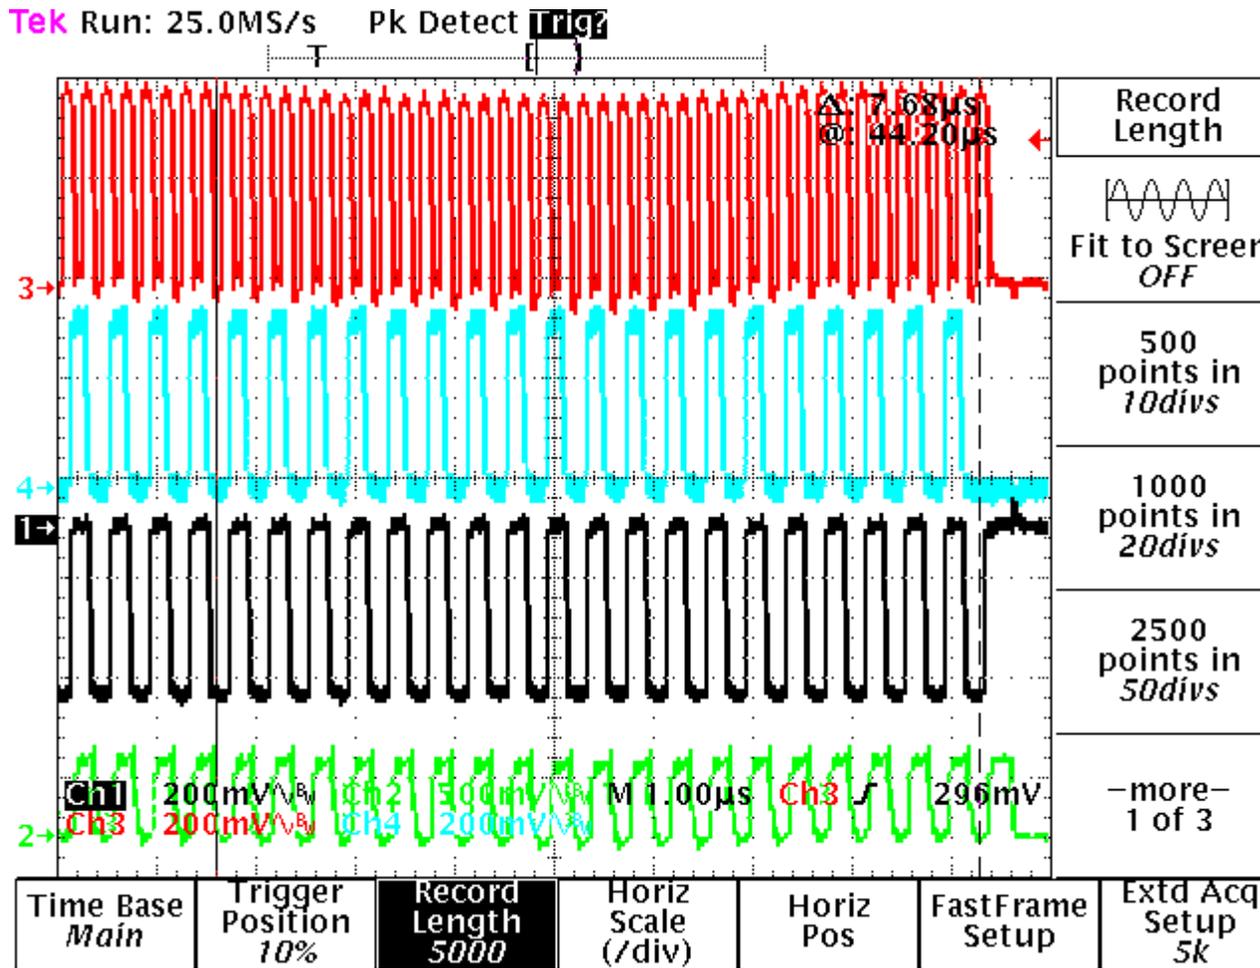
Trace #4 is register output after full 231 bits.

A pattern of '5's was sent with alternating 0's and 1's.

Extra 0 observed as 37th bit in Trace #1 is due to fact that register length is odd and pattern is even.



- Example of end of pattern clocked through Global Shift Register at VDD=1.6V



Trace #3 is CCK

Trace #2 is DI

Trace #1 is register output after first 38 bits.

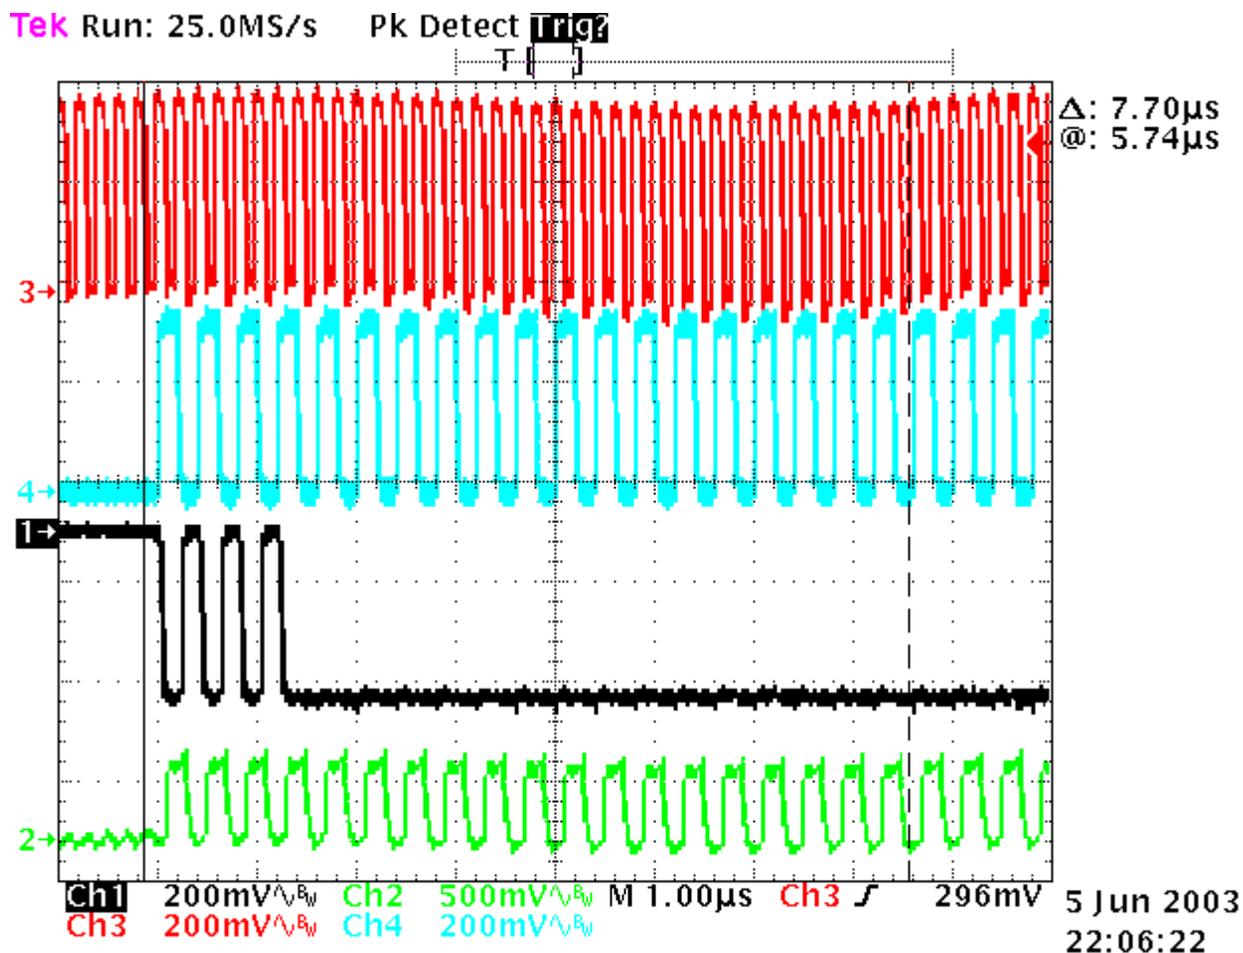
Trace #4 is register output after full 231 bits.

A pattern of '5's was sent with alternating 0's and 1's.

No data loss is observed.

-

- After pattern clocked in at 1.6V, it is clocked again at 2.5V, and start of string is displayed. Output data was shifted in at 1.6V and shifted out at 2.5V.



Trace #3 is CCK

Trace #2 is DI

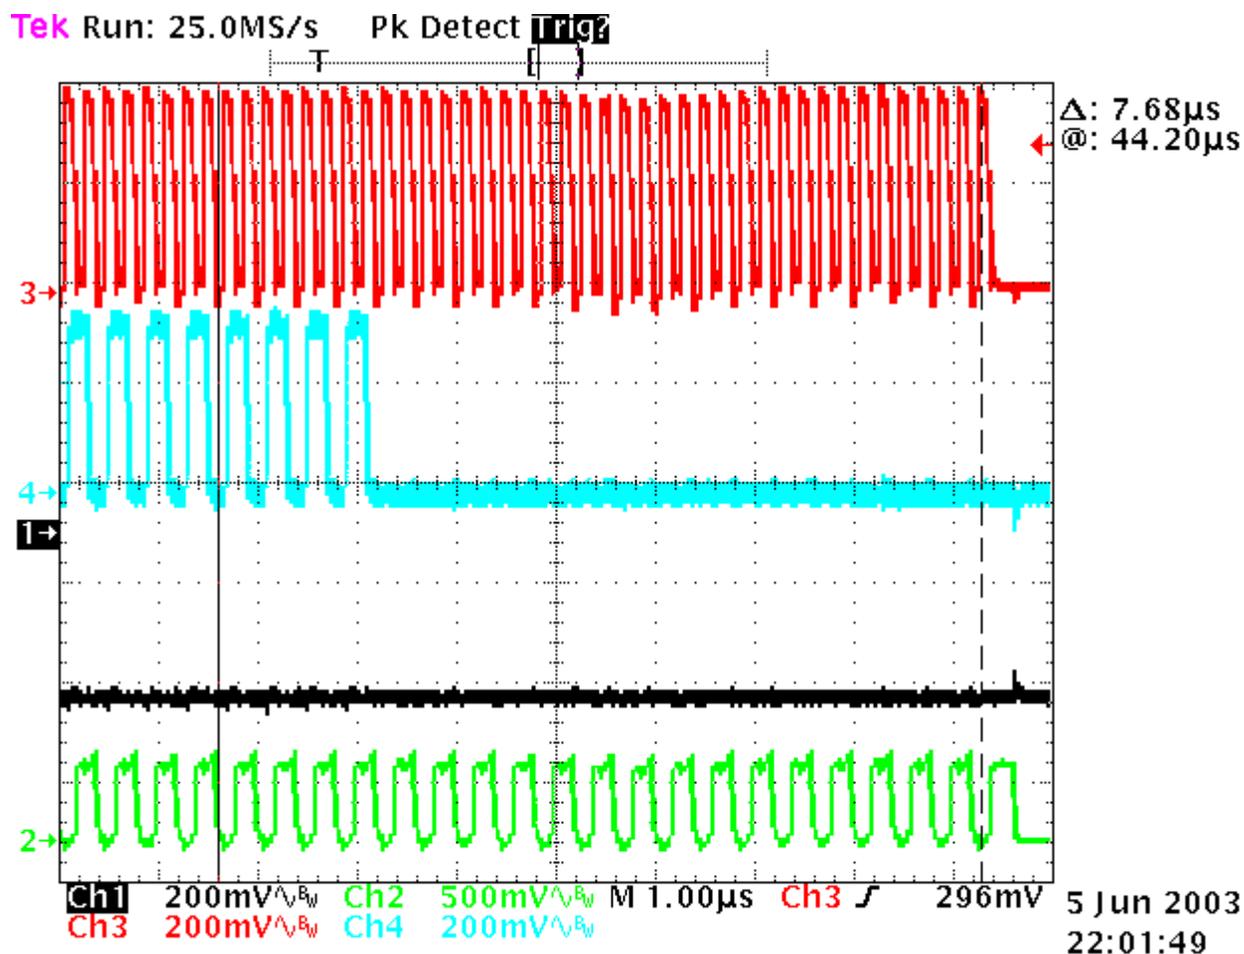
Trace #1 is register output after first 38 bits.

Trace #4 is register output after full 231 bits.

A pattern of '5's was sent with alternating 0's and 1's.

- Observe that only the first 8 bits from the 38-bit place and route part survive. These were the 8 bits which were loaded into the Hitbus Scaler at 1.6V, and have now been clocked out at 2.5V. This suggests a single point failure between blocks.

- After pattern clocked in at 1.6V, it is clocked again at 2.5V, and end of string is displayed.



Trace #3 is CCK

Trace #2 is DI

Trace #1 is register output after first 38 bits.

Trace #4 is register output after full 231 bits.

A pattern of '5's was sent with alternating 0's and 1's.

- Observe the loss of the last 30 bits, with 8 bits surviving from the place and route block. This confirms that the single-point defect appears to be between the exit from the first 30-bit ConfigReg and the entrance of the 8-bit Hitbus Scaler.

- Lowering VDD to 1.6V and clocking again confirms the same failure pattern.
- In the test above, bit 30 is lost and bit 32 survives. When testing with a pattern of 'A's, bit 31 survives and bit 29 is lost, confirming that the defect location is between bits 30 and 31.

Summary:

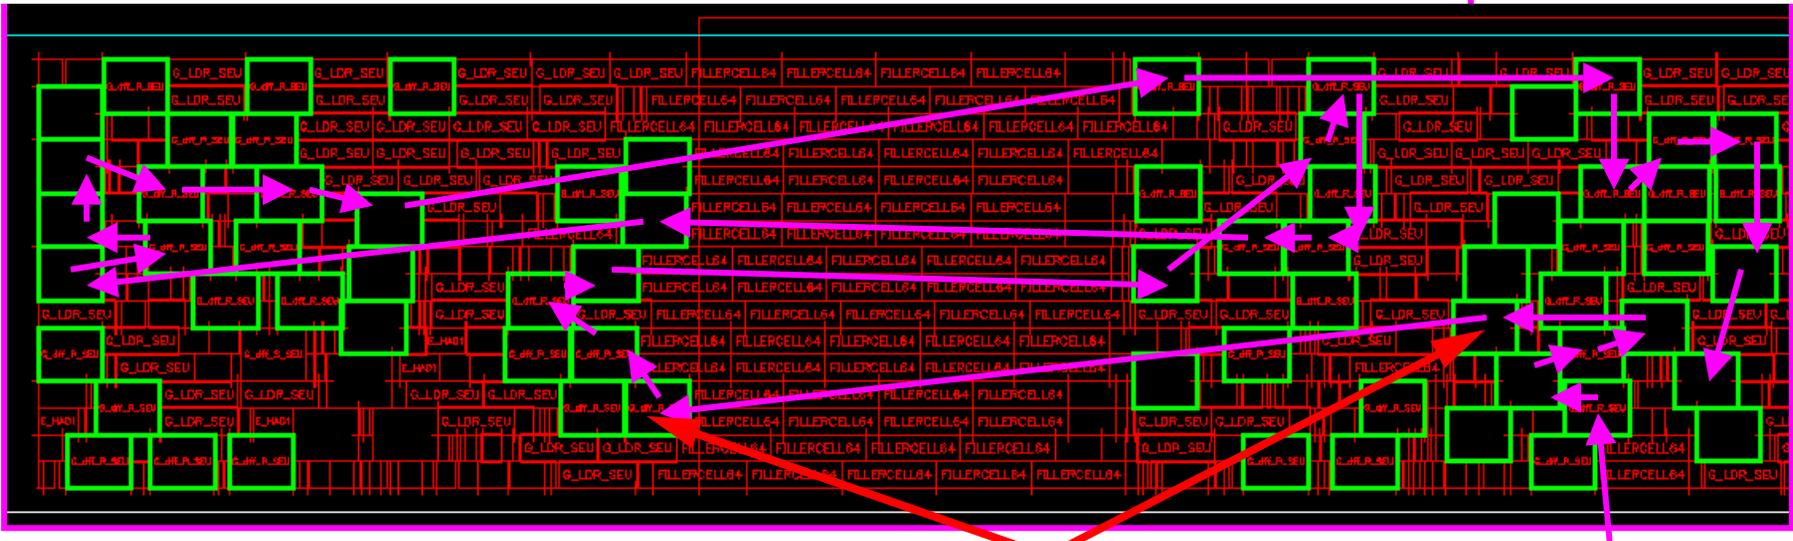
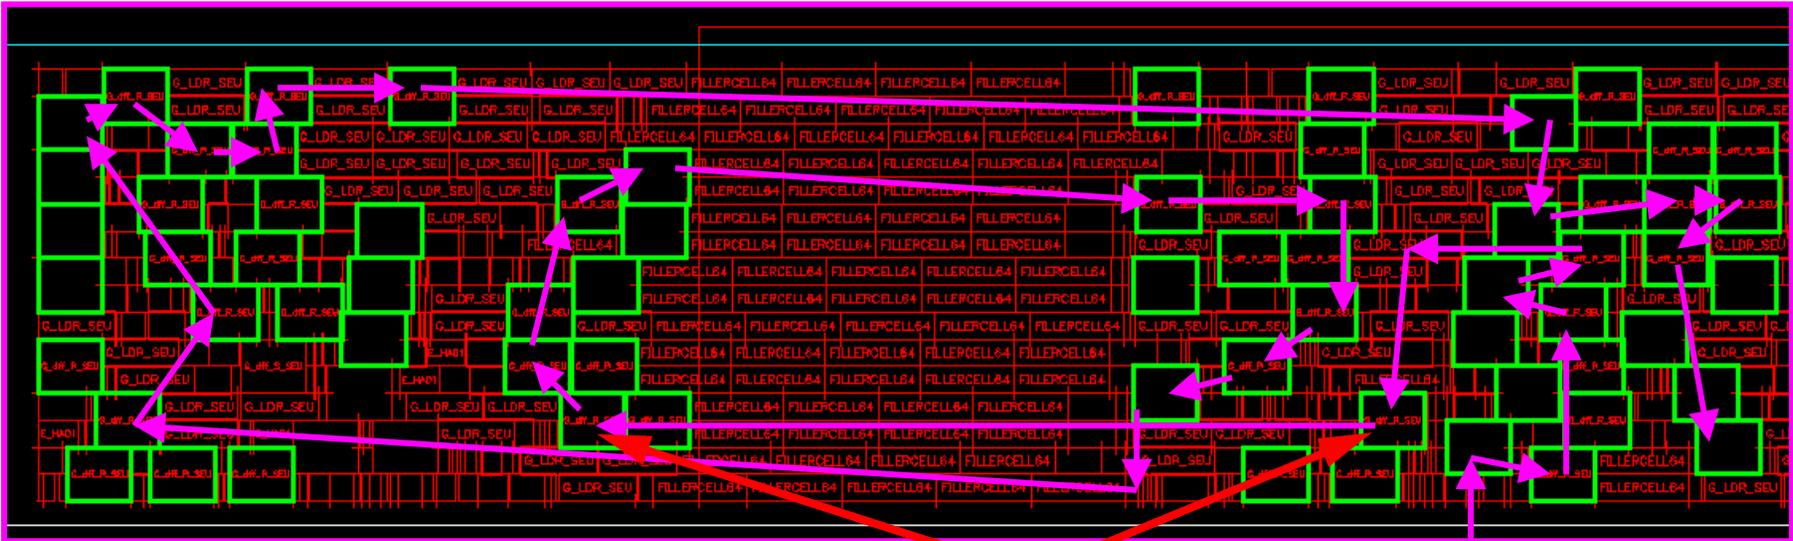
- In the MainCtrl block, there is a single-point defect in the two shift registers. For Register #1, the defect is between bits 6 and 7. For Register #2 it is between bits 5 and 6. Note that the SEU-tolerant implementation of MainCtrl requires bits to be present in both shift registers in order to properly latch a 1 into the Command Register itself, so a failure in one register alone is enough to cause the Command Register to fail.
- In the Global Register, there is a single-point defect in the shift register between the initial 30-bit ConfigReg block and the following 8-bit HitbusReadback block.

Correlations with Layout:

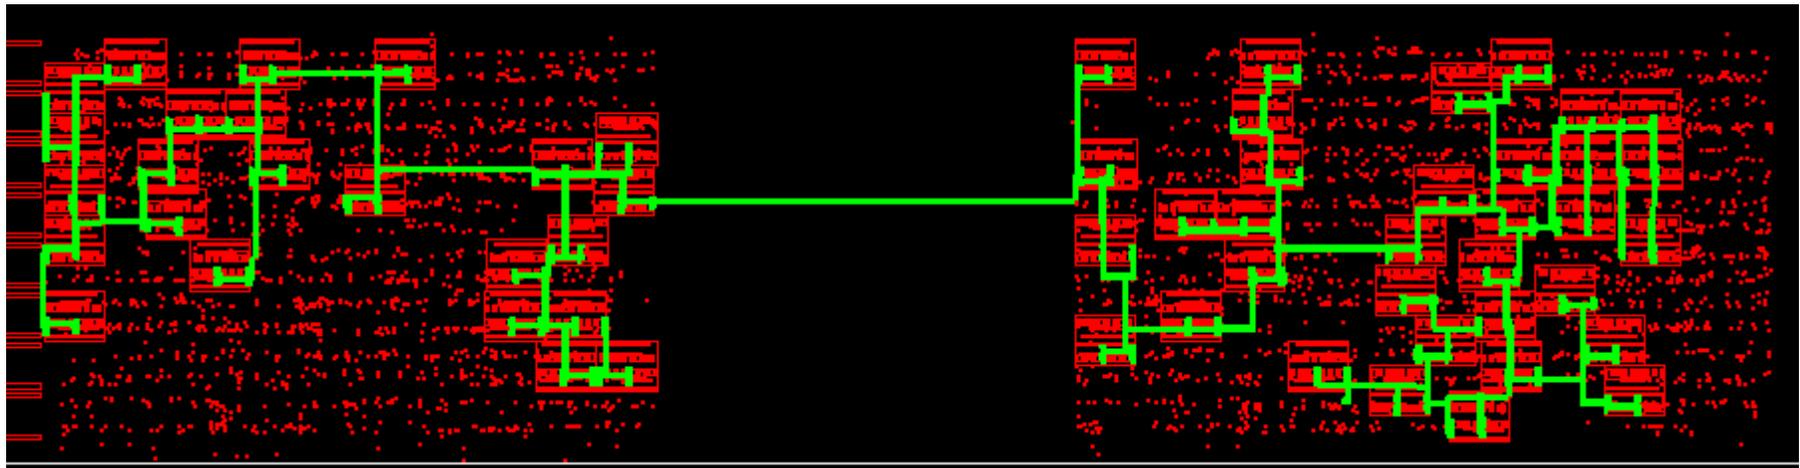
- The layout in the place and route blocks in FE-I2 is complex. A single block of circuitry is targeted into 6 different regions (5 place and route and one containing 50% place and route plus full-custom Trigger FIFO) with only bussing allowed between the regions. The pitch of the blocks is the column pair pitch of 800μ . For FE-I1, the circuitry fit into 3 regions (two place and route and one full custom).
- In FE-I1, there was a modest effort to re-route the clock tree for the shift registers by hand. This was not done for FE-I2.
- In FE-I1, each of the shift register fit into a single place and route region, and a register was no larger than about 300μ by 400μ . Only relatively small standard DFF and SEU-tolerant latches were used.
- In FE-I2, significant SEU tolerance was introduced, with triple-redundant SEU-tolerant latches and SEU-DFF for the shift registers, as well as redundant shift registers in the Command Decoder. Additional bits were also created in these registers (5 more Command Register bits, 14 more Global Register bits). The Global Register is spread over three place and route regions and is roughly 300μ by 2000μ , while the Command Register fits in two place and route regions and is roughly 300μ by 1500μ .
- The larger size creates significant additional problems for the timing margins of shift registers (and possibly other blocks) due to routing and wire-loading effects. The design flow used did not account for these complications, as no formal timing analysis tool (Pearl or PrimeTime) was used.

Command Register Layout:

- Data path for Register #1 (upper) and Register #2 (lower). Red arrows = problems:



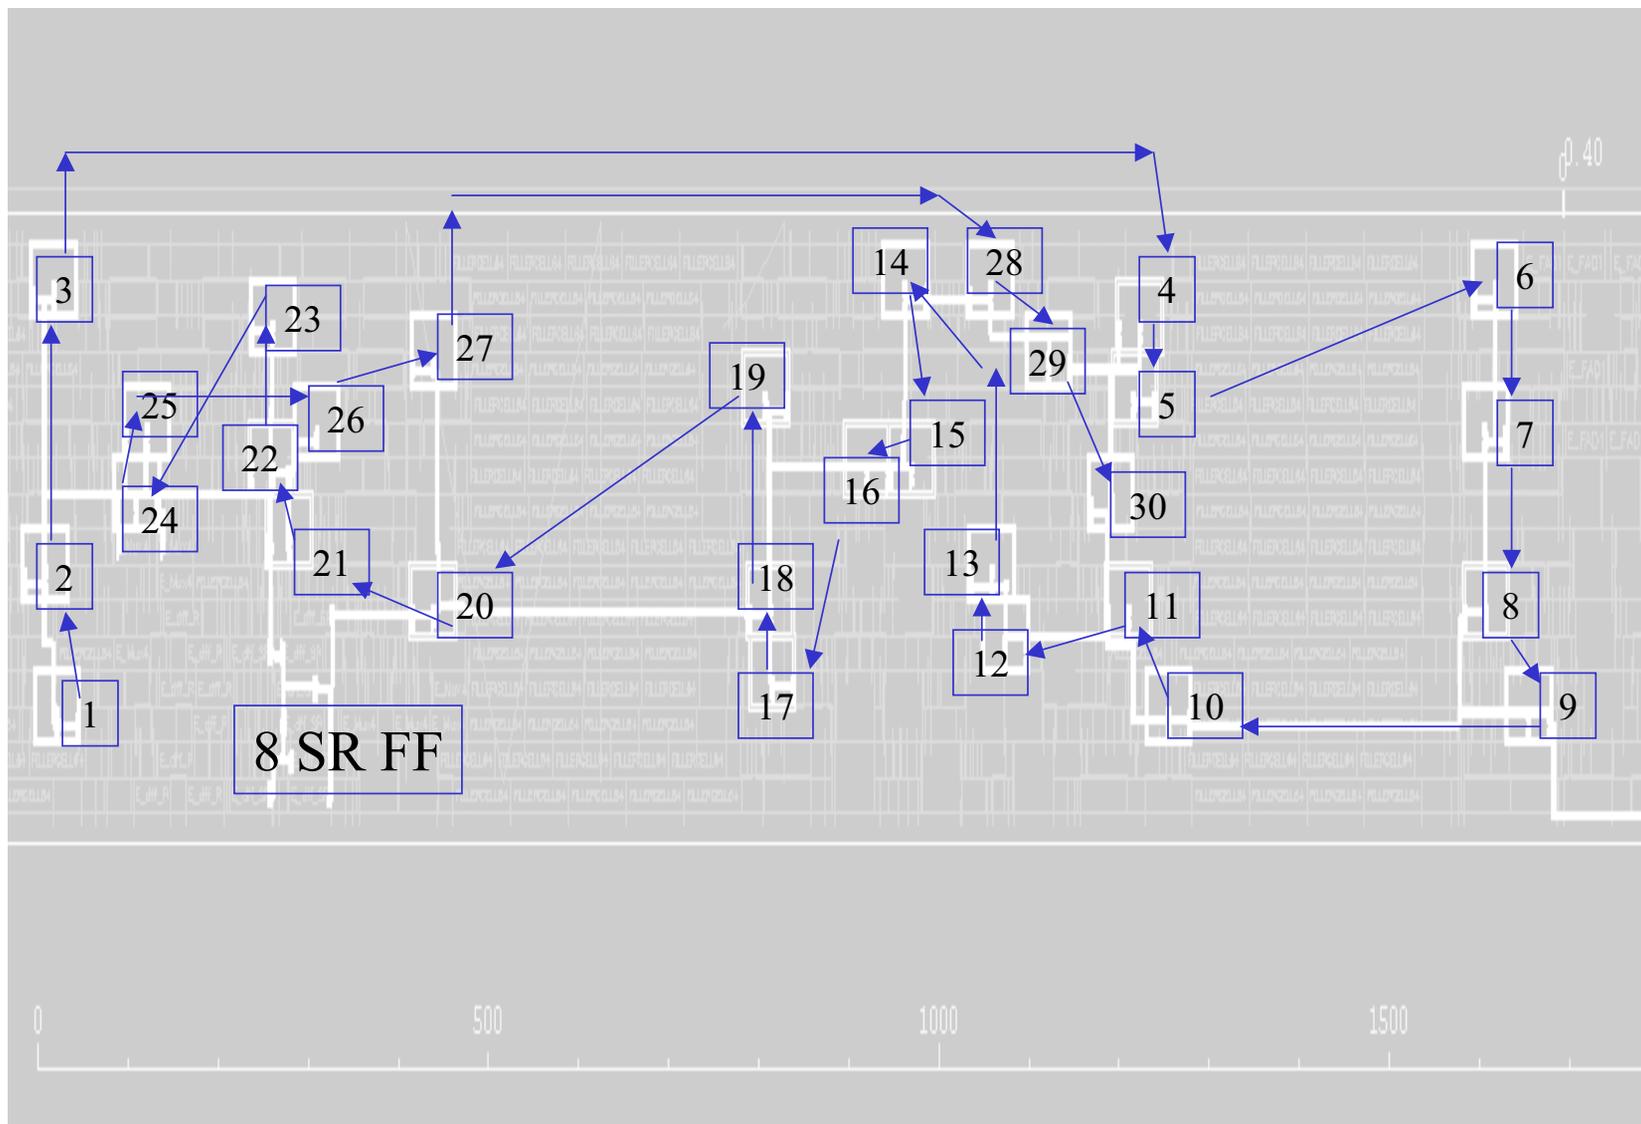
- Clock Tree for the registers (clock enters in lower right of tree):



- Points of concern are cases where there can be a large clock delay compared to the data. Normally, this will happen when the FF are far apart, and the data can follow a direct path, but the clock follows a less direct path. Also, the loading on the data lines is much lower and so their risetimes are much faster (sub-ns versus about 4-5ns), and longer lines become RC transmission lines with delays.
- Observe that there are long data paths in same direction as clock tree for Bit 6->7 for Register #1, and for Bit 5->6 for Register #2. These are consistent with the observed location of problems.
- However, there is a longer connection from Bit 15->16 in Register #1, and two long connections from Bit 14->15->16 in Register #2, with no observed problems.
- As VDD is lowered, clock timing is dominated by RC and less sensitive to FET g_m . Lightly loaded data lines slow down more, reducing skew between clock and data.

Global Register:

- Data path for Global Register (clock tree is white, clock enters at lower right):

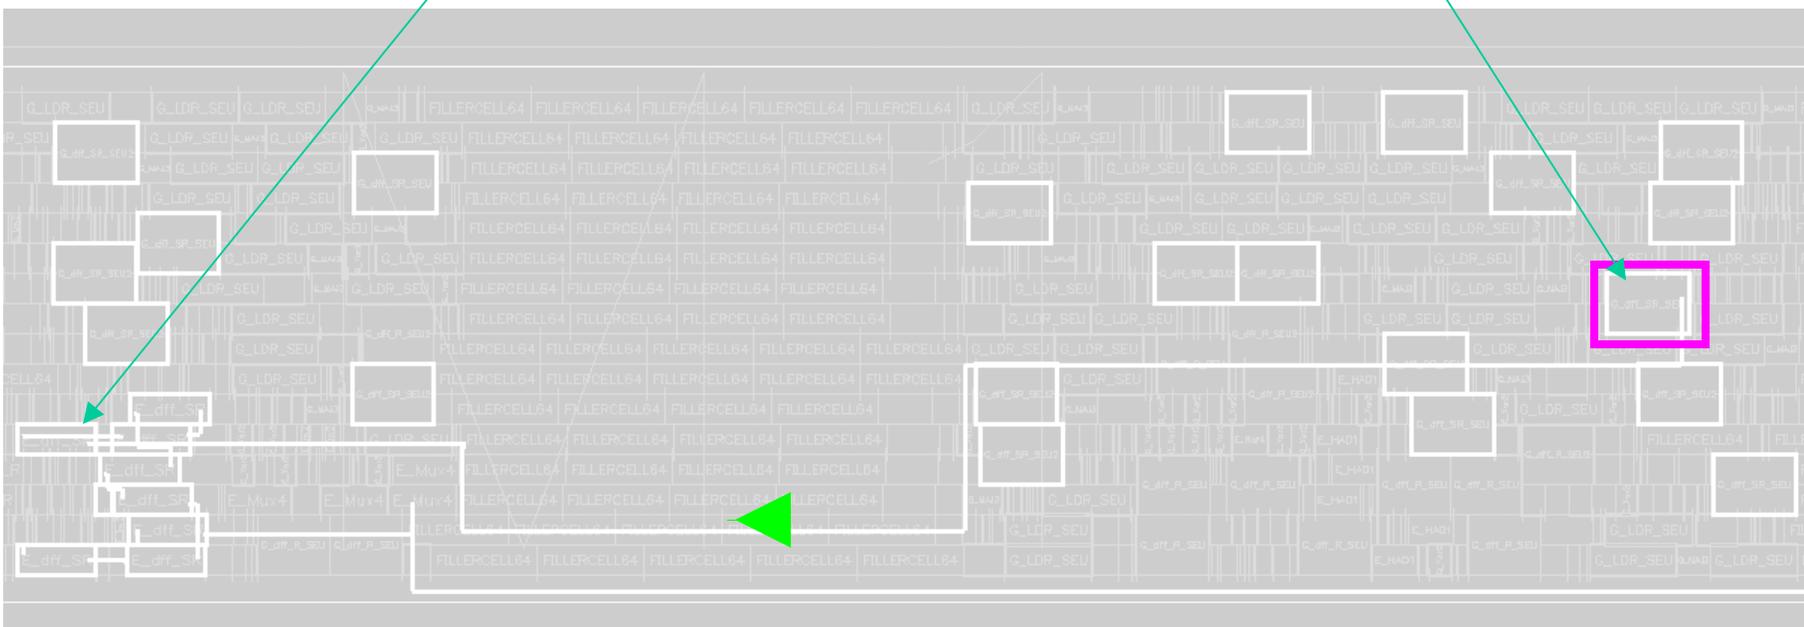


- Eight DFF in lower left are the standard cell DFF for HitBus Readback.

•Connections between ConfigReg (30-bits) and Hitbus Readback (8-bits):

First bit of hit bus reg

Last bit of config reg



•Once again, the data path is long, but the clock path is longer. This suggests that this location is the weakest point in the 38-bits of place and route Global Register.

Conclusions:

- Conclude that there is some layout evidence to support poor clock routing as the explanation for the VDD-dependent errors. For the moment, it is more qualitative and circumstantial than we would like.
- In particular, the three observed single point failures are among the weakest points in the shift register clock versus data routing.
- First, the blocks themselves are much larger in FE-I2 than in FE-I1. Also, because of the increase in bit count, and the larger loads presented by the SEU-DFF, the loading of the clock distribution is much worse, with an estimated risetime of about 5ns in both shift registers. This makes the possibility of a race condition even more probable.

Issues:

- There is as yet no quantitative demonstration of the major observations.
- First, the observed dependence on VDD has not been demonstrated in detail (but a plausible mechanism for failure to start as VDD is raised has been described).
- Second, the observation of lost 1's but never lost 0's over a rather large range of VDD has not been demonstrated in simulation, and seems somewhat inconsistent with a failure mechanism based on poor routing alone.
- Third, have not demonstrated that the actual trace lengths in the layout are sufficiently long to create a race condition between bits.

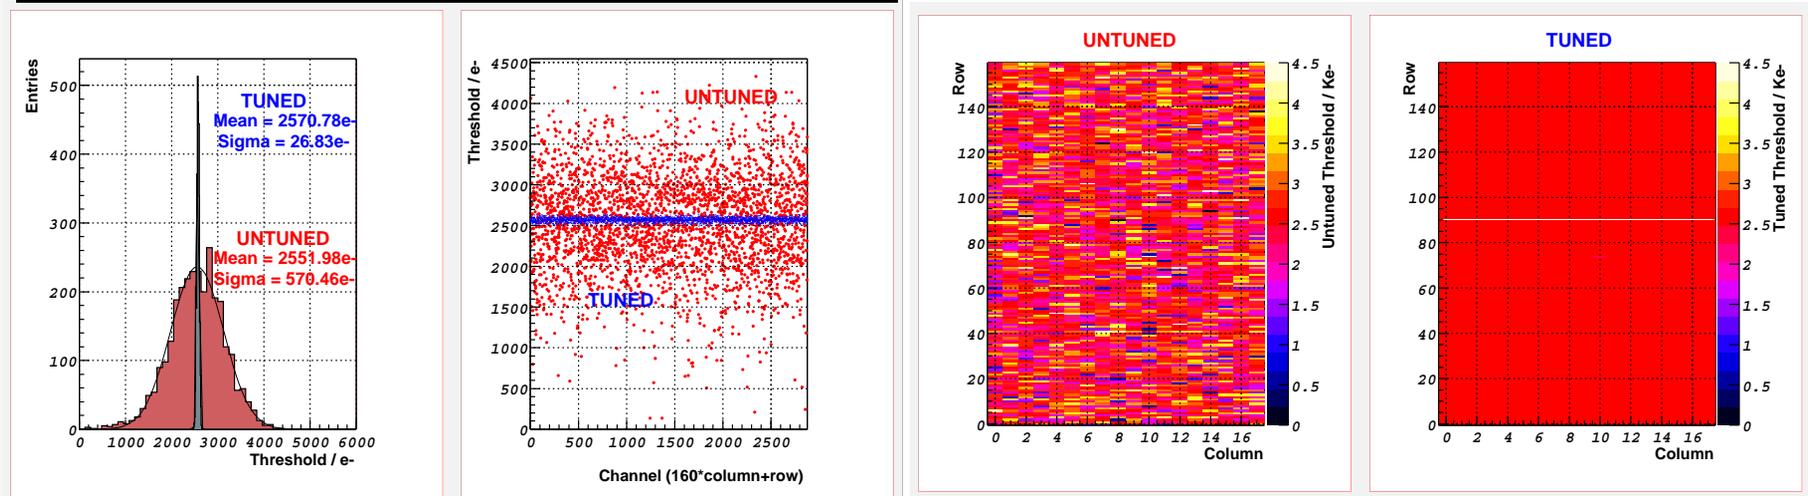
Other Measurements of FE-I2:

- Operating at VDD=1.6V, we have tested most of the functionality in the chip.
- We can say that everything that has been tested so far is working roughly correctly, and in most cases, just as we expected.
- The new threshold control, involving a 7-bit high-quality TDAC in each pixel works well, and gives a monotonic and quite linear threshold control over a wide range. The 5-bit GDAC in each pixel also works well, and gives a global threshold control which permits moving a tuned threshold distribution without significant change in dispersion.
- The raw threshold dispersion is about 600-700e, instead of the 900-1000e observed for FE-I1. A simple tuning algorithm produces a threshold dispersion of about 50e, a brute force algorithm produces a dispersion of about 25e.
- The bias compensation circuitry works well, and eliminates the large variations in IP and IL2 seen along a column in FE-I1. A monitoring circuit allows quantitative analysis.
- The auto-tuning circuitry works reasonably well, and usually produces a dispersion in the range of 50-100e. There are a number of residual systematics (odd/even TDAC effect, non-uniform threshold, etc) that remain to be understood, but this looks like it could provide significant help in module re-tuning.
- Current-mode DACs are quite linear (minor non-uniformity seen in 10-bit DAC).

- Known problems, such as the RCU bug and the marginal CEU=40 timing in FE-I1 have been fixed or improved.
- New digital feature of HitBus scaler for hot pixel finding is working correctly.
- New programmable latency self-trigger is working correctly.
- New SEU-management circuitry like the Hit Parity work as expected.
- SEU-hardened circuits, such as the Hamming-coded Trigger FIFO, and the triple-redundant latches, seem to work fine. Of course assessing the SEU-hardness of FE-I2 requires high intensity beam testing.
- New power management circuits like the analog and digital linear regulators work well also, with a dropout voltage of less than 100mV for operating a single chip.
- First tests in which configuration operations are performed at 1.6V and data taking operations are performed at 2.4V show that the chip works properly, with the exception of some problems with EOE status bits (not observed at lower VDD). This could be a further indication of poor timing margins in the Digital_Bottom block, and appears to be related to timing problems in writing the Trigger FIFO.
- One wafer has been probed in Bonn, and the yield for nine good column pairs was $272/288 = 94\%$. More complete analysis is underway, but even with more sophisticated cuts, the yield looks very high from the engineering run.
- First bumped (untested) wafers have been received, and single chip testing will go ahead quickly. Once we have single-chip assemblies on test boards, we can complete the analog evaluation of FE-I2.

Example of “brute force” threshold tune with 7-bit TDACs:

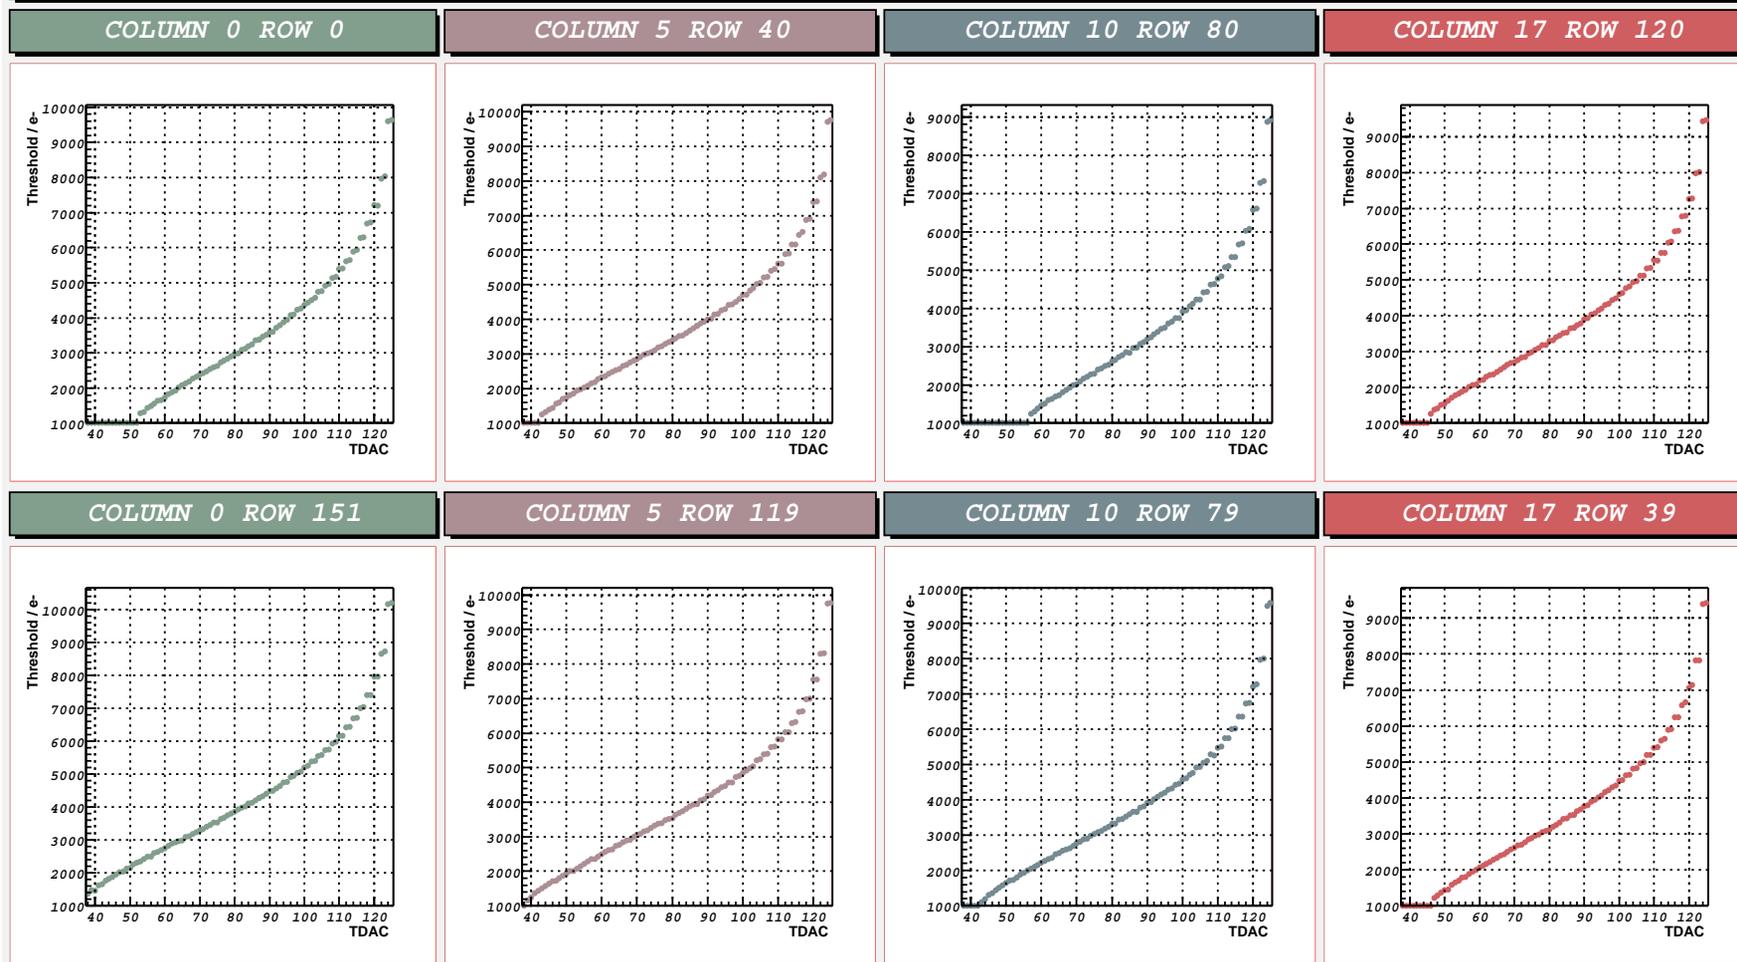
FE-I2 (probed on wafer). Draconian TDAC Tune (every point scanned)



- Initial dispersion is about 600e for TDAC=64 (compared to 900e for FE-I1B), which is roughly the expected improvement due to improved transistor matching.
- After tuning, it is possible to achieve 26e dispersion on threshold. Simpler scans with about 9 points and interpolation achieve 50e dispersions.
- Note that top pixels appear noisier without sensor due to their higher IP current which makes them faster for a given load.

Example of TDAC linearity:

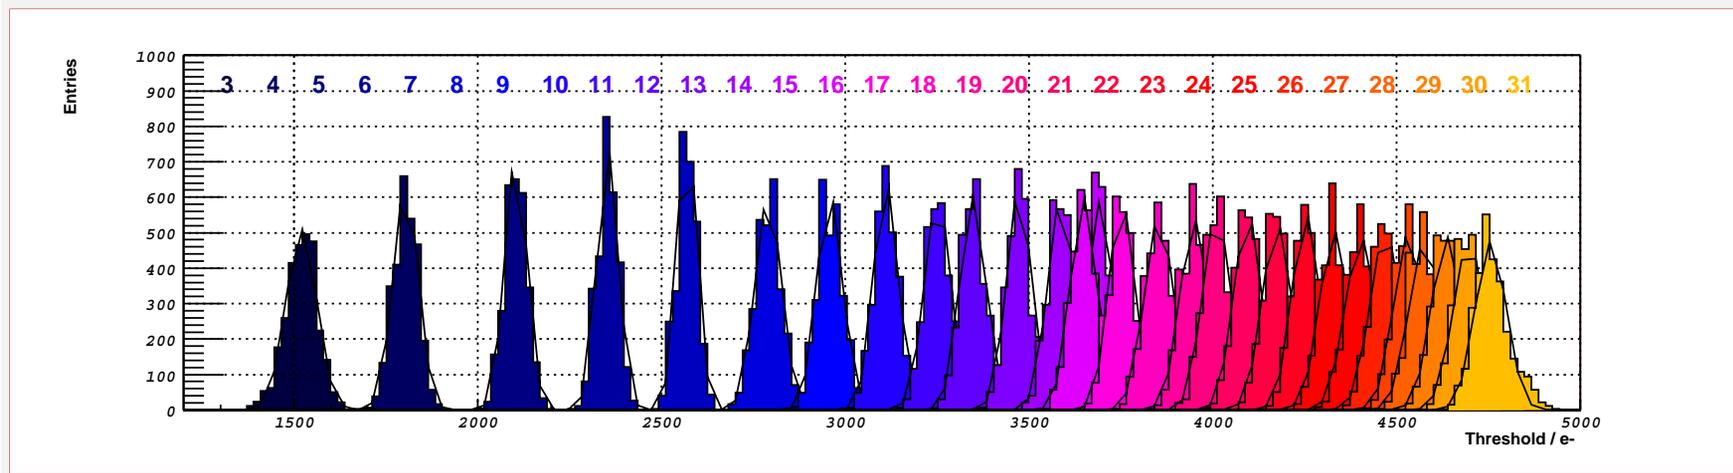
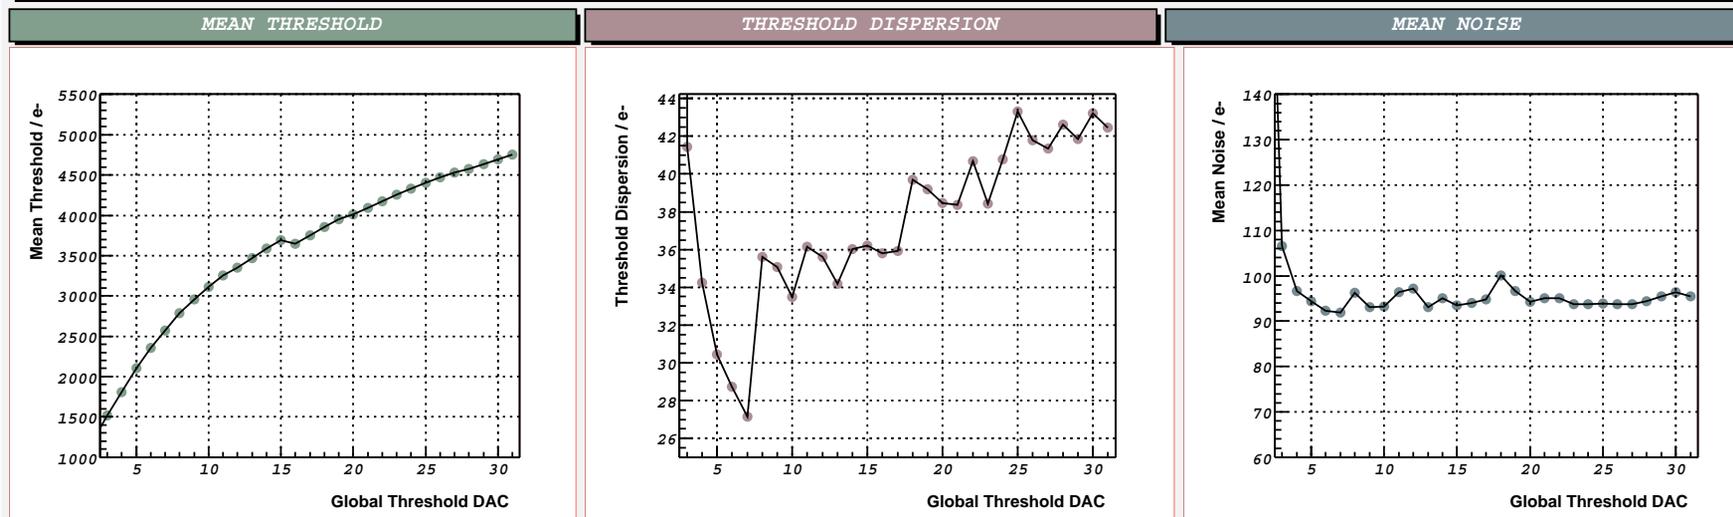
FE-I2 (probed on wafer). Examples of Threshold versus TDAC at GDAC=7



- It is clear that the new TDAC design is very linear over a large range close to the center of the scale.

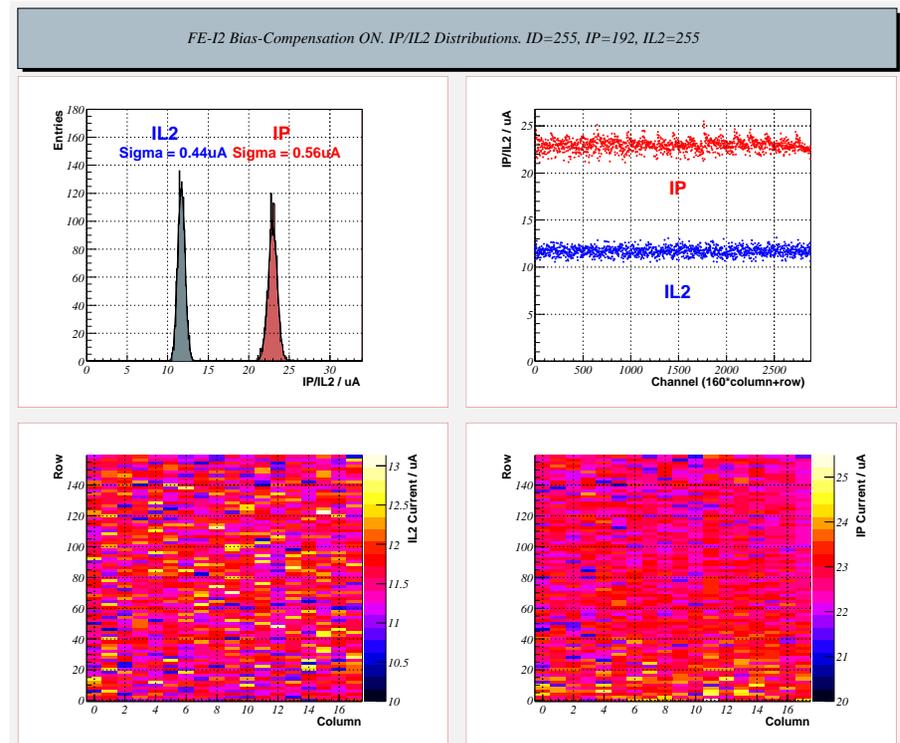
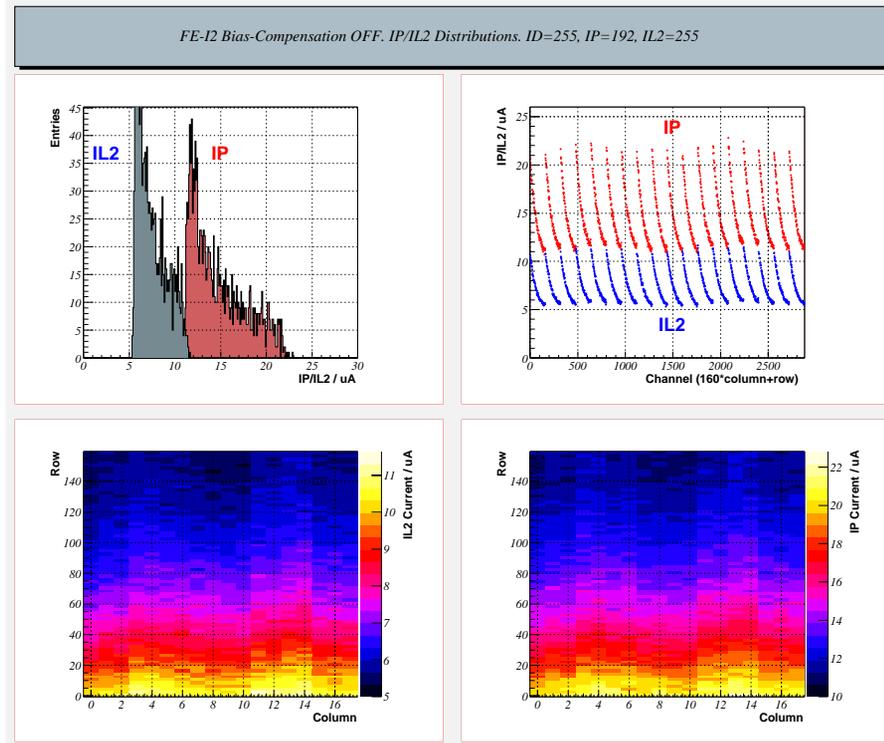
Behavior of tuned threshold distribution vs GDAC:

FE-I2 (probed on wafer). Effect of Global Threshold DAC. Tuned at GDAC=7



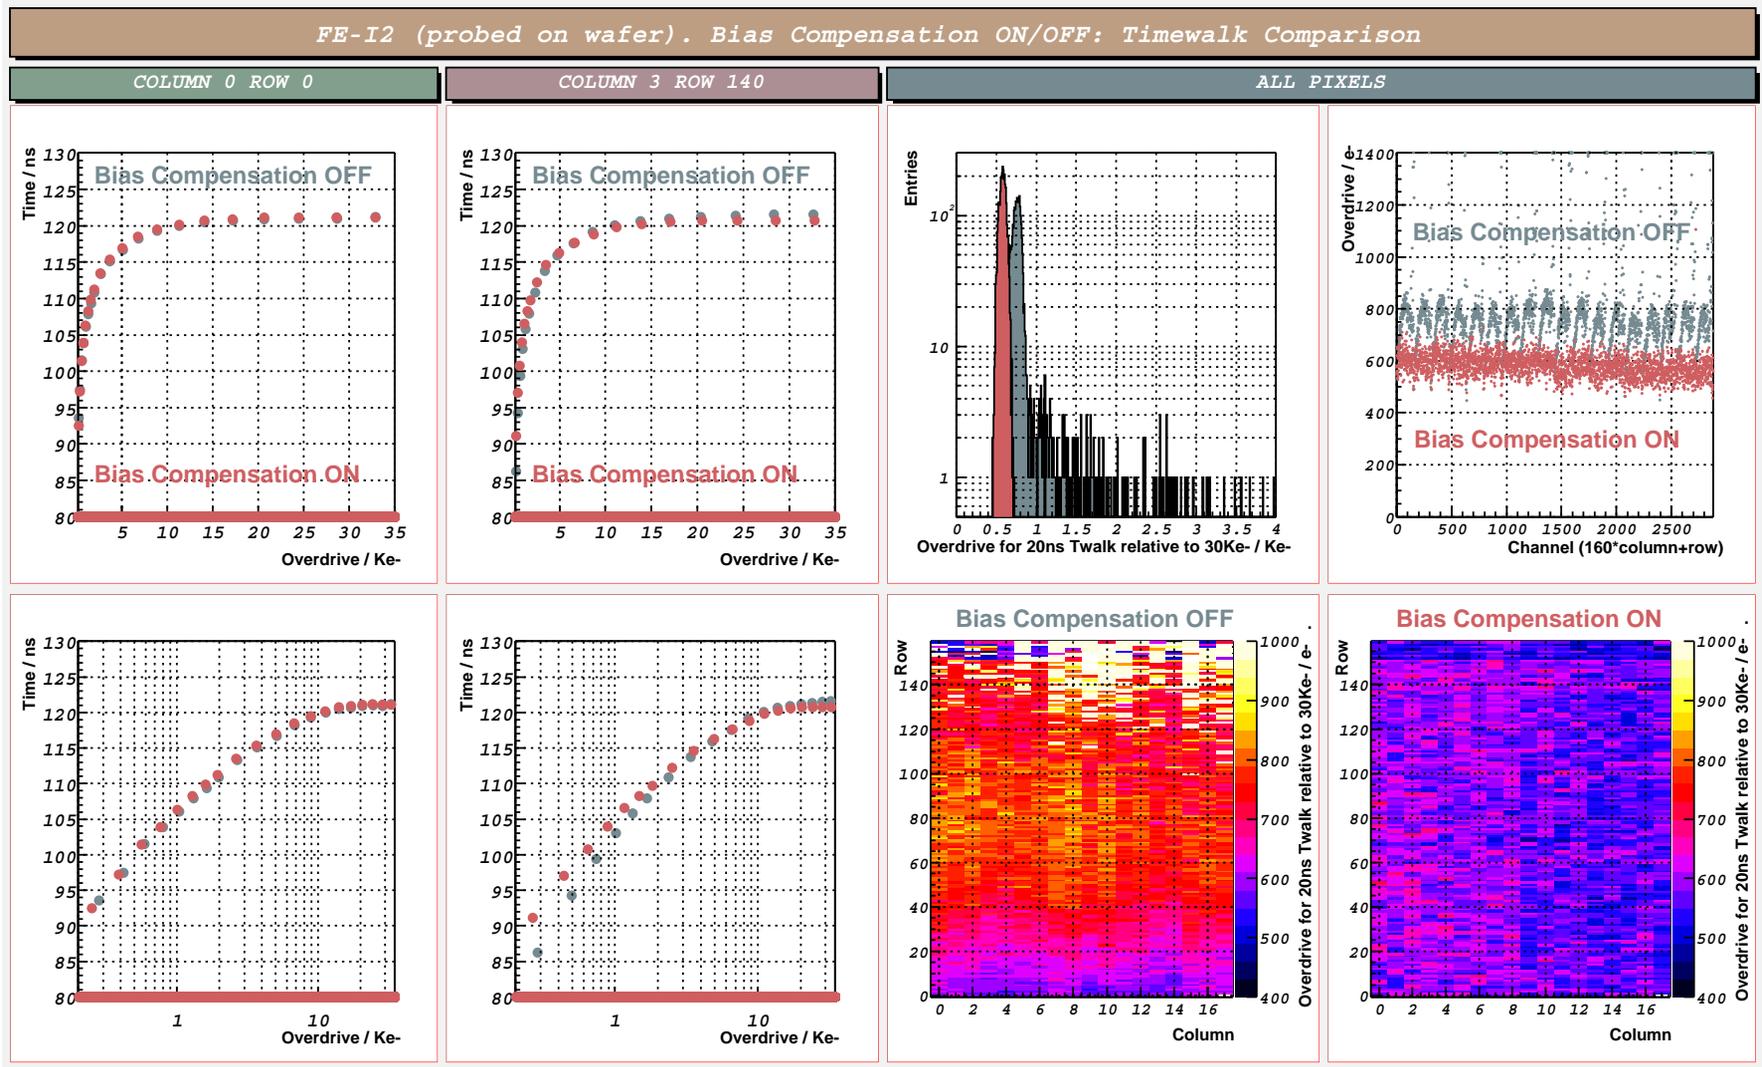
- Principle observation is that tuned distribution can now be moved without significant dispersion (here, tuned at GDAC=7).
- Problem at MSB transition which makes GDAC=16 lower than GDAC=15.

Explore performance of Bias Compensation using IPMon:



- In FE-I1, there were significant internal voltage drops on analog supply busses, particularly AGnd. This is turn caused row-dependence in the bias currents which were referenced to AGnd, principally IP and IL2.
- In FE_I2 there is one bias compensation circuit per pixel. Odd pixels compensate IP, even pixels compensate IL2. The local bias current is shared by two pixels.
- Left plots shows large systematic variations in IP and IL2 when circuit is turned off (note IP and IL2 currents set to very large values to enhance drops). Right plots show compensation eliminating systematic effects. IP = $24 \mu A$ and IL2 = $12 \mu A$.

Timewalk measurements with/without biascomp:

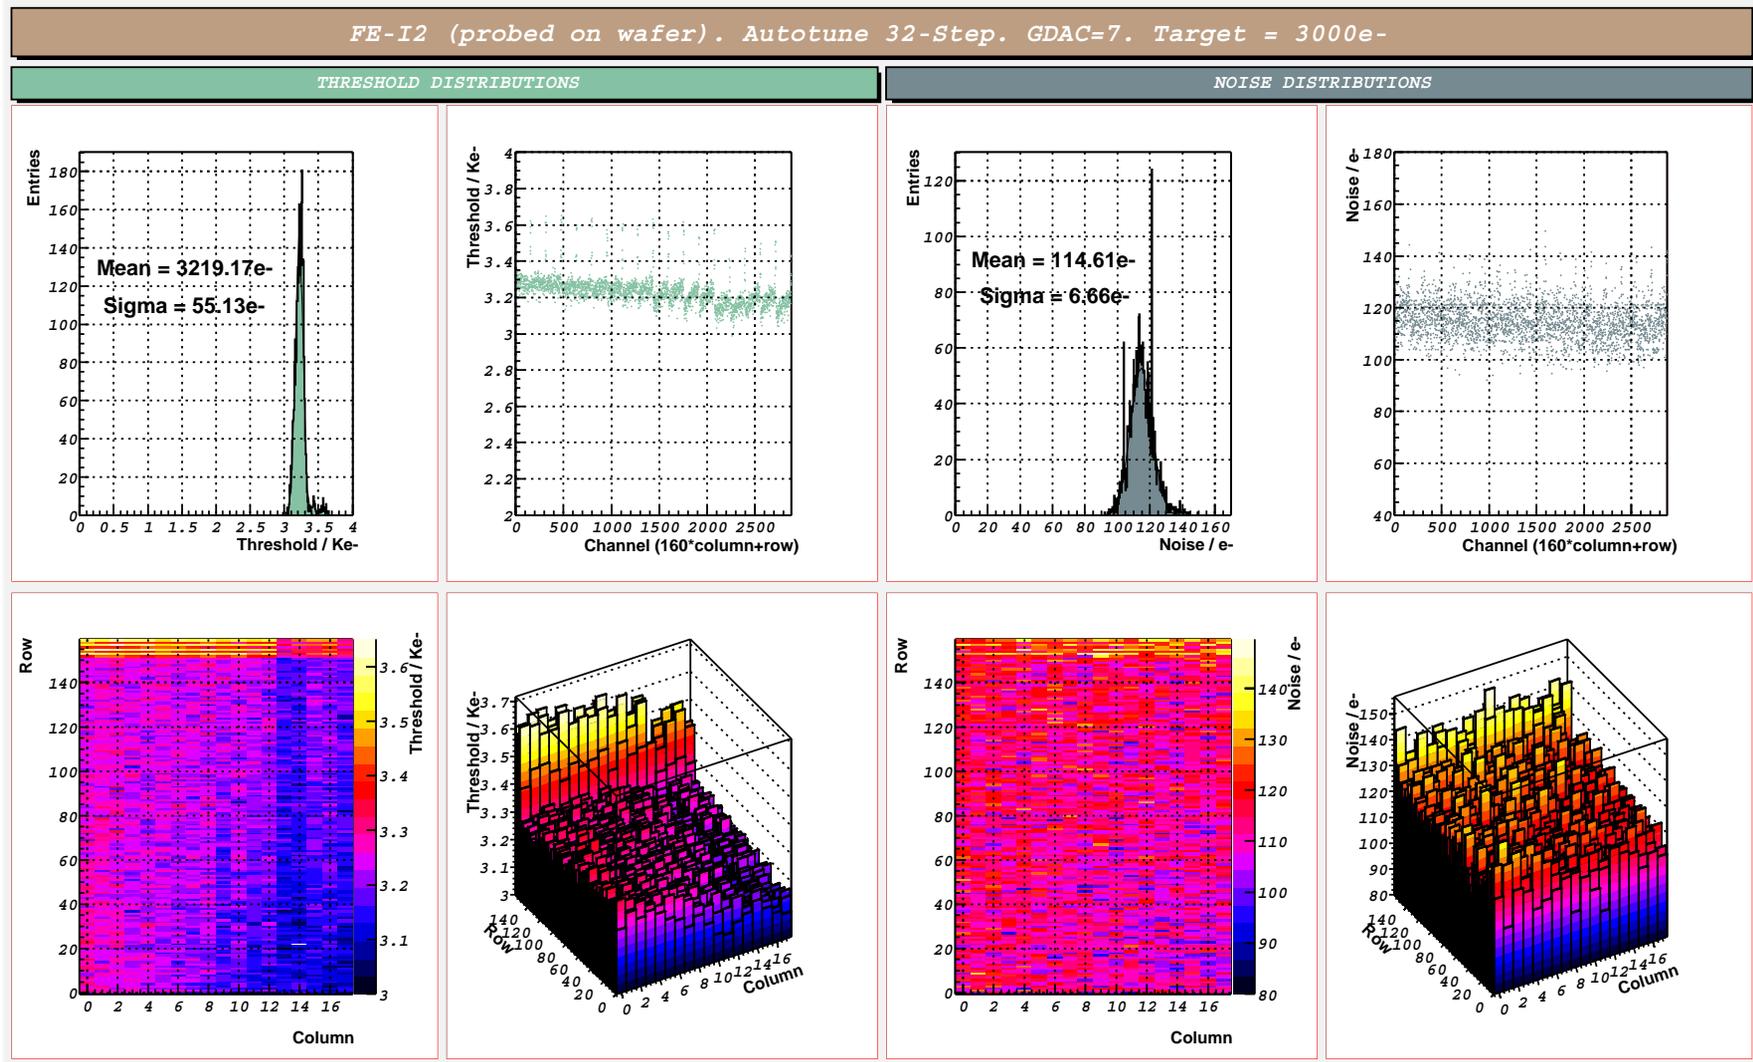


- Timewalk variations without sensor are small. However, can see the degradation in timewalk performance visible for high row numbers without compensation disappears when compensation is on, making all channels behave identically.

Auto-tuning Capability

- In FE-I2, Laurent and Ivan have added a simple auto-tune support circuit (a 5-bit up/down scaler and some control logic) into each pixel.
- The basic operation is simple. One does a TDAC scan while injecting a fixed charge, and chooses the first TDAC value for which each pixel is “above threshold” (meaning it fires more than 50% of the time).
- In practice, one does a staged scan (32_step works well). For the active part of the pixel matrix, one does a backwards scan of TDAC values (highest to lowest). For each TDAC value, the pixels are injected 100 times. The up/down counter in each pixel is preset to mid-scale (MSB is set). For each injection, if the discriminator fires, the counter increments, otherwise it decrements. If the MSB is still set at the end of the sequence, then the current TDAC value is latched in that pixel and cannot be overwritten until auto-tune is disabled. At the end of the full sequence, the “tuned” TDAC values will be stored in each pixel and can either be used directly, or read back and stored in a file.

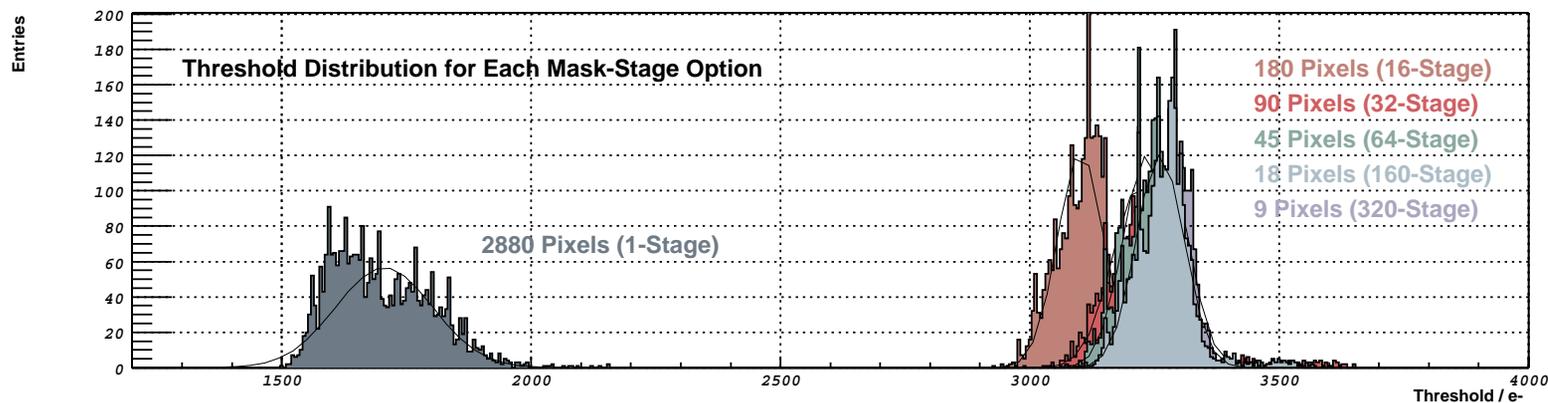
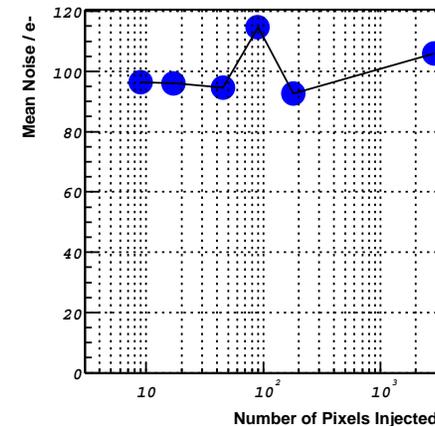
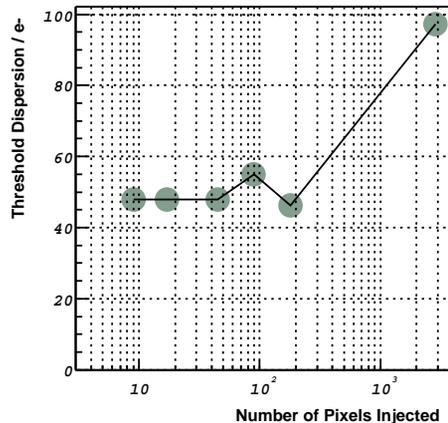
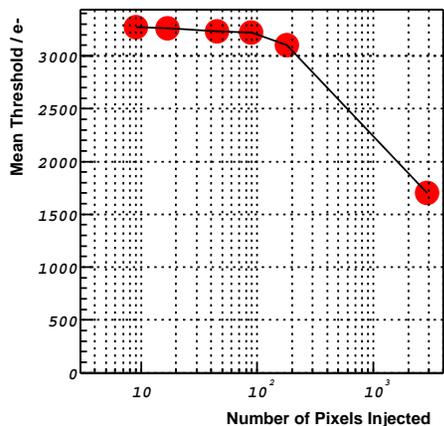
Example of Auto-tune scan:



- Plots show results of a verification threshold scan. The tuned dispersion is only 55e, but there are systematics. Ganged and inter-ganged pixels have higher threshold (due to higher noise ?). Some pattern dependence plus non-Gaussian TDACs also

Summary of Auto-tune scans for different mask patterns:

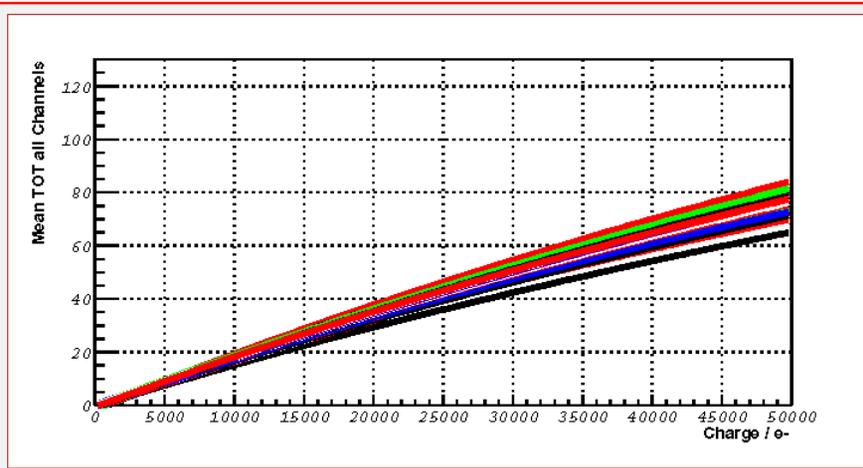
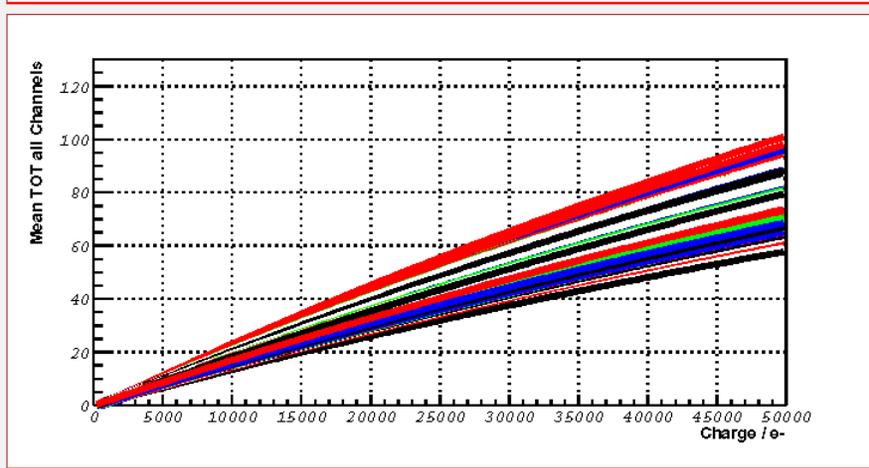
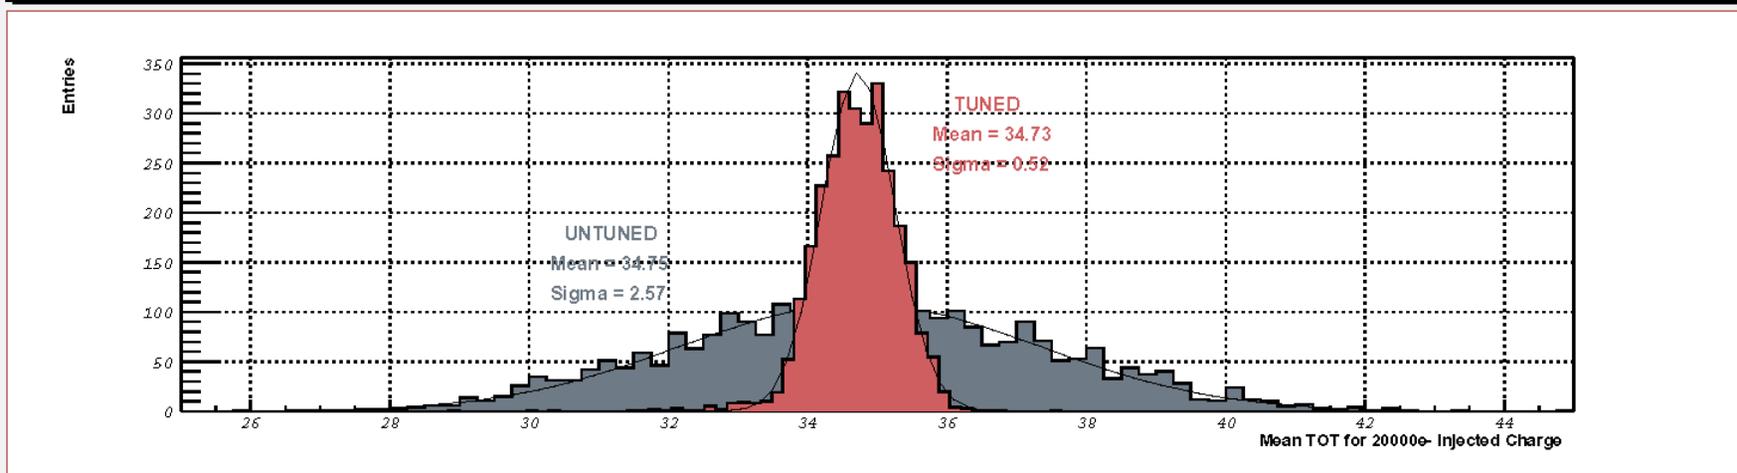
FE-I2 (probed on wafer). Autotune; Effect of Mask Stage Option. GDAC=7. Target = 3000e-



- Observe that 32_step scan is still OK, but 16_step scan shows systematic shift in target threshold (due to loading effects ?). Presently, 32_step scan takes about 30 seconds for a single chip. Suggests module will tune in about 1 minute.

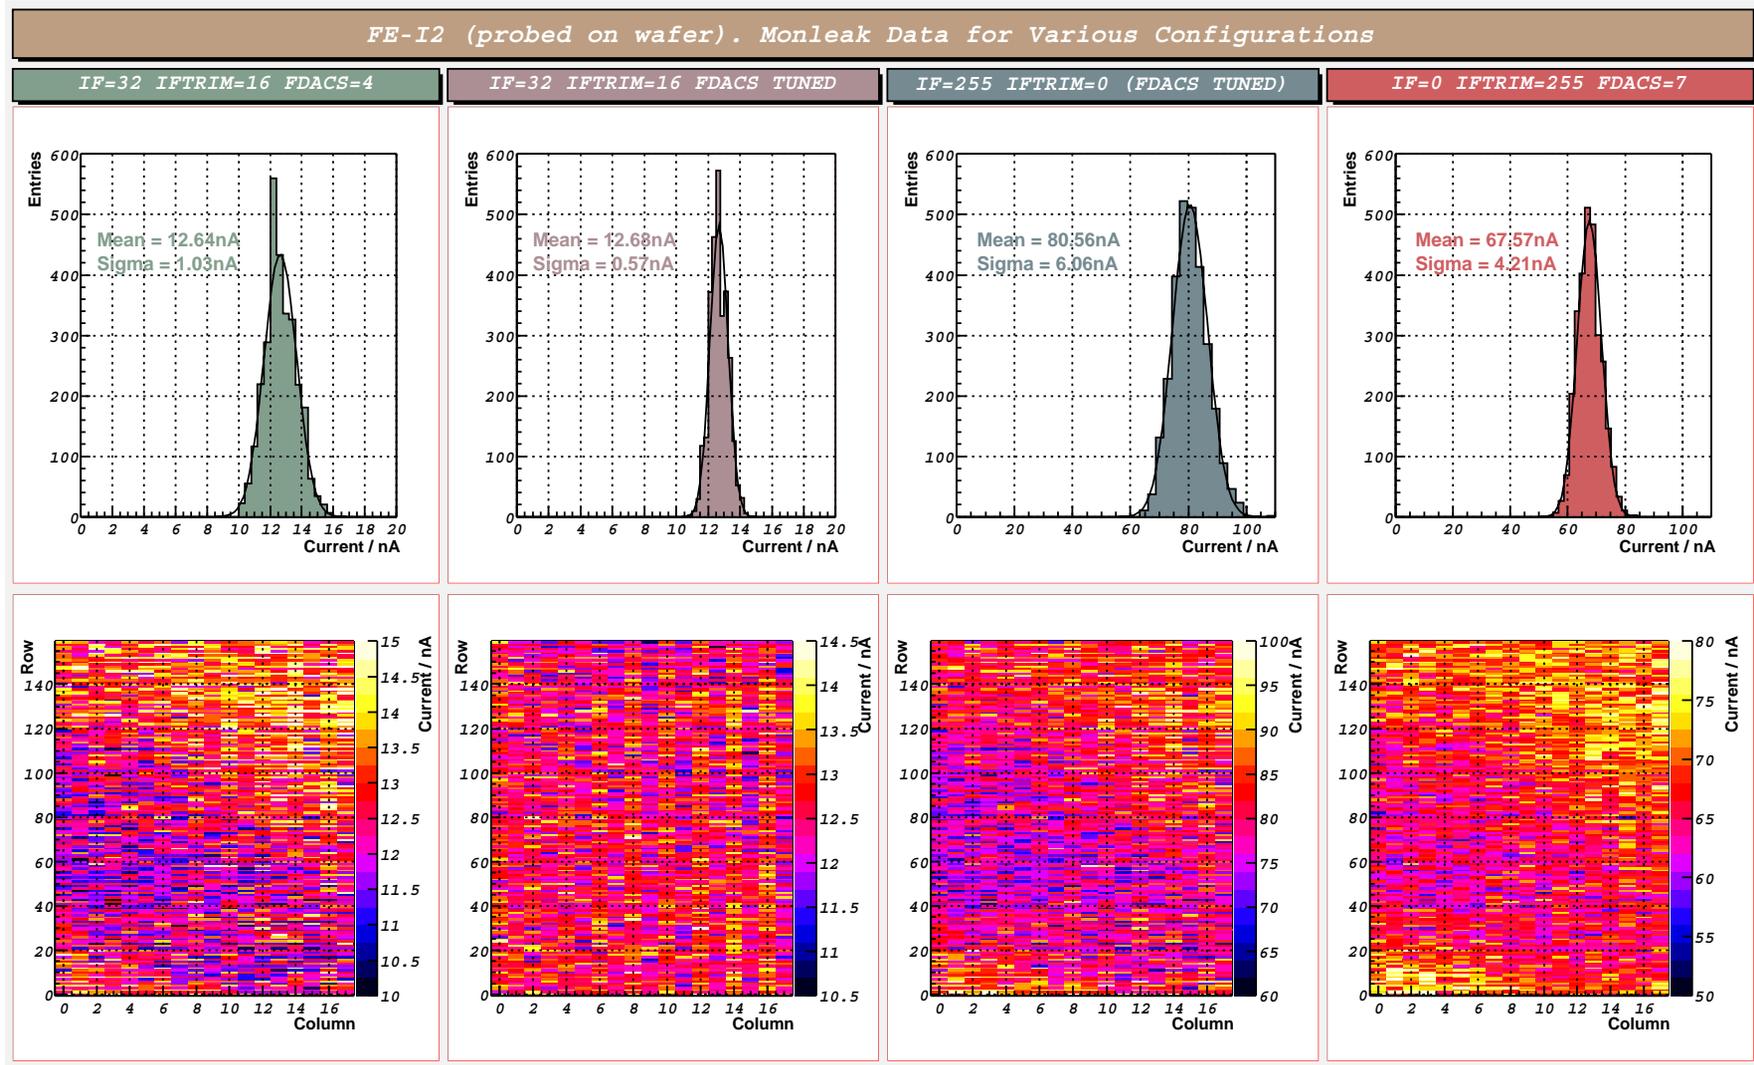
Feedback DAC tuning and TOT:

FE-I2 (probed on wafer). FDAC Tuning Performance



- Plots show that TOT response can be tuned to better than 2% using 3-bit FDAC. Although tuning is done for a particular charge (20Ke), tuning is valid over a larger range of TOT (compare left and right families of curves).

Monleak measurement and FDAC tuning:



- Leftmost plot shows raw Monleak with FDAC=4, next plot shows tuning FDACs also eliminates non-uniformity in Monleak. Finally, extreme cases with IF=255 or with ITrimIF = 255 show similar non-uniformity. Is three bits enough for the FDAC ?

Additional Problems:

- Have tried carrying out some data operations by using a modified TDAQ which does register config commands at $VDD = 1.6V$ and data taking commands at $VDD = 2.4V$. So far, testing has not been exhaustive, but one problem has appeared.
- For VDD of about 2V and above, there appears to be a problem in writing to the Trigger FIFO. This manifests itself as non-sequential BCID errors and Hamming Code errors. The former means that the BCID being written into the Trigger FIFO are not incrementing by 1 each trigger as would be expected. Instead, they jump around somewhat randomly. The latter means that the Hamming code redundancy bits in the Trigger FIFO are not consistent, and indicate that a data corruption problem has occurred.
- It appears that neither of these problems results in errors in the event structure. The hits always appear associated with the correct EOE word, and the BCID is consistent for all hits and the EOE word for a given 25ns crossing. This points at the problem occurring during the TFIFO write, an operation which must complete in 25ns. For example, if the Hamming bits themselves were late, the FIFO would record the data correctly, but the erroneous check bits would cause the output logic to reconstruct the wrong data value on readback.
- Although the problem at the single chip level is not that serious, for a module, it suggests that we will lose the ability to cross-check synchronization between FE chips, since the BCIDs for individual chips do not advance in the same way. This would be a significant loss in the original module architecture.

- In addition, the spurious Hamming code errors would need to be masked. However, this is less serious.
- Will continue testing all features of FE-I2 as a function of VDD in this way to look for additional problems.
- In addition, we will carry out elevated frequency tests, temperature cycling tests, and irradiations in July, to look for other problems.
- The challenge with timing problems of this sort is that it is difficult to know how much margin is left in the chip once it is operating. This is why we are eager to carry out timing analysis and re-do the place and route.

Fixing FE-I2, Next Steps:

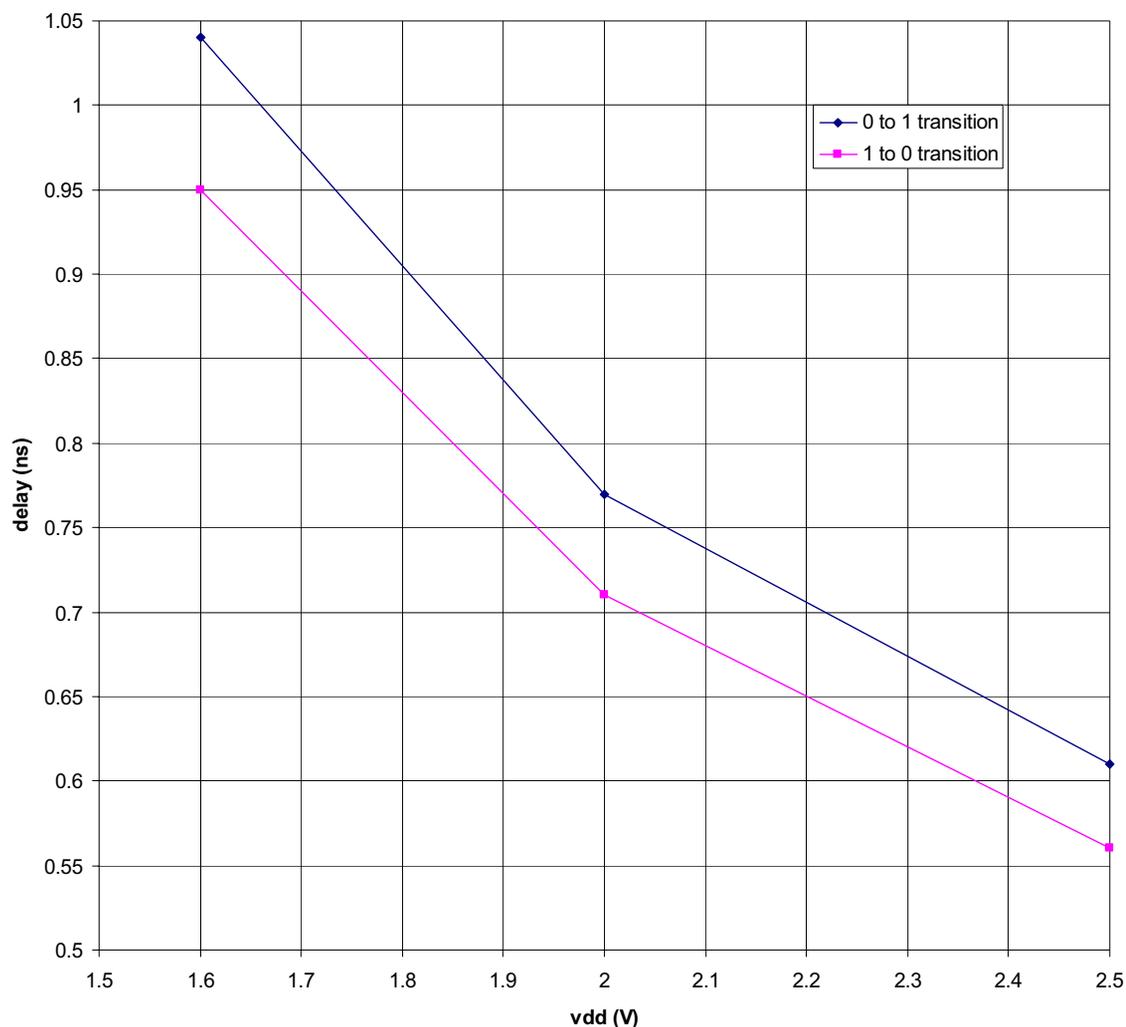
Ongoing Work:

- More complete quantitative analysis is underway, in two parts.
- One part (Laurent) is to study the performance of the basic circuit blocks and evaluate the timing margins, as well as the potential delays due to RC effects. This is a more “parametric” analysis to understand the real sizes of the various effects.
- A second part (Emanuele) is a more bottoms-up analysis of the actual layout with extracted R and C parasitics. This is very complex, because the number of parasitic R and C generated is huge.
- Neither analysis is complete enough to present at this time.

Open Issues:

- The timing margin for operation of the shift register is relatively large, and it seems a bit difficult to develop the large differential clock and data delays required to create a race condition. However, the problems do occur at places where the layout is particularly far from optimal. More quantitative analysis is required.
- The observation of lost 1's, but never lost 0's, over a rather wide voltage range (about 1.9V to 2.8V) seems somewhat inconsistent with the rather small 0/1 timing asymmetries observed in simulation up to this time. However, the simulations do show the correct sign for the asymmetry.

- The first analysis of a shift register made up of SEU-DFF shows very large timing margins, suggesting clock for bit n+1 needs to be later by almost 1ns:



- Do observe that 1->0 transition is faster, which will make the transfer of a one followed by a zero more marginal than the opposite case. This agrees with observations, but difference is not large, and can be compensated by $\Delta V=200\text{mV}$.

Next Steps:

- Have confirmed with IBM that there are six wafers from the engineering run waiting to begin the back-end processing (metallization). We can change a few mask layers (M2, V2, and M3) to re-route the clock by hand for the shift registers, and then resubmit the design. Still waiting for the official quote from IBM for this, but based on previous experience, changing three masks and getting 4 wafers back should cost about 35K\$, and have a fairly rapid turn-around time (perhaps we could get wafers back by early August). Propose that the present information is already convincing enough to suggest that this would be a useful intermediate step. It would allow reduction of risk for any follow-on submission, by demonstrating a fix. This step is clearly the critical path for beginning production, so it needs to go ahead as rapidly as possible.
- In parallel, would pursue a more robust solution, involving re-doing the place and route block (Digital_Bottom) using a timing analysis tool to eliminate marginal timing conditions. This FE-I3 would be prepared and ready to go when the modified FE-I2 wafers return. Depending on the performance of modified wafers, and on the results of additional simulations and measurements of the existing FE-I2, we would either submit the new version (means 300K\$ to do another engineering and another production run, as for FE-I2), or we would make more of the modified wafers (means about 130K\$ for another production run from existing masks). The schedule would be about the same for either path, but the cost is quite different, so this places a significant burden on the intermediate step to be a high-quality patch.

- Plan to start on the modified layout for the clock routing for the three registers this week. Estimate that layout modifications and verification could be completed in about one week. It is possible that the critical path for the intermediate step will be getting a new PO in place for IBM. For the moment, we are waiting for the quote.
- Would then proceed to work on new version of place and route block, with a slightly different approach for the shift registers. Would buffer the clocks better, and probably include a clock input and output as separate pins on the DFF blocks to allow us to force the desired clock routing from the schematic. Approach used in the Pixel Register (two-phase clocking) is less desirable, because it would require XCK to run a state machine, and we strongly prefer a design where the basic control of the chip can be done without relying on the LVDS I/O.
- Would also work on complete timing analysis of the Digital_Bottom block. This is more difficult because of the lack of good models for the new standard cells. This can be managed to some extent by forcing conservative margins. This is moving ahead now, using Hyperextract and PrimeTime.
- There are indications of a timing problem in writing to the Trigger FIFO at higher VDD (above 2.0V), which emphasize the importance of treating the place and route block with care.
- The creation of FE-I3 would take all of July and possibly longer. Note that a new submission at the end of July would not provide wafers before early October (engineering run) or early November (production run), so the net effect is at least a 4-month delay in the electronics schedule compared to our previous planning.