Electronics Planning for FY04

K. Einsweiler, LBNL

FE-I2 Status and Next Steps:

- Problems in Command and Global Registers
- Re-spin submission (FE-I2.1) and new engineering run (FE-I3)

MCC-I2 Status and Next Steps:

- Problems in MCC-I2 config mode
- Re-spin submission (MCC-I2.1)

Irradiation and Test plans:

- First results from May and program for the rest of the year

Longer Range Schedule:

- Production plan and issues

FE-I2 Status

Concurrent Engineering Run and Production Run:

- Submitted GDS file to CERN on April 4, and forwarded to IBM on April 8.
- Wafer has two identical FE chips per reticle, and 144 good reticles per wafer.
- Received first 6 engineering run wafers at CERN on May 16. Two wafers immediately sent out, one to each bumping vendor.
- Hand-carried 4 remaining wafers to Bonn for initial testing on May 21.
- Production order of 48 wafers delivered to CERN on June 20. They were divided, 24 wafers were carried to Bonn, and 24 wafers were carried to LBL.

Total Costs:

- Basic elements of total cost of each run are included in the frame contract. However, there are additional charges that change from run to run, depending on other submissions through penalty charges for multiple starts in one quarter.
- Cost for FE-I2 run was $303,582.96, with US share of 21.5%. Breakdown was $184,531.44 for the engineering run and all NRE and extra charges, and $119,051.52 for the 48 production wafers.
- However, all transactions passing through CERN are converted to CHF, so in an era of somewhat volatile exchange rates, exact numbers can change.

Summary of FE-I2 Testing

Problems in Control Section:

- Initial testing showed that the chip could not be addressed if the digital supply voltage (VDD) was above about 1.8V
- After extensive testing, this has been traced to timing “defects” in the shift registers used in the Command Register (dual, redundant 29-bit shift registers) and the Global Register (231-bit shift register). The defective parts of the registers are all inside a large synthesized place and route block.
- It turns out that over a wide range of VDD (about 1.8V up to maximum test voltage of 2.8V), the problem is a single defect in each of the three shift registers. This results in the loss of isolated 1’s, and the loss of a 1 followed by a zero.
- 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.
- We interpret this problem as a race condition between clock and data created by poor control of the clock routing in the place and route blocks.

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 FE chips in a module would work, and operated the MCC at a higher VDD. However, the minimum digital regulator voltage is about 1.85V, and FE-I2 does not work there. The major feature of FE-I2 that does not work at the allowed VDD of 1.6V is the highest frequency column clock readout, which is required for high luminosity operation in 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. Due to problems with MCC-I2, this would most likely be done using Flex V5 plus MCC-I1.

Conclusion of study of correlation with layout:

- Conclude that there is significant layout evidence to support poor clock routing as the explanation for the VDD-dependent errors. For the moment, this conclusion is more qualitative and circumstantial than we would like.
- In particular, the three observed single point failures in the shift registers are among the weakest points in the shift register layout for 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 from initial analysis:

- There is as yet no quantitative simulation 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, although there is an observed asymmetry between 1's ands 0's in the SEU DFF in simulation, the observation of lost 1's but never lost 0's over a rather large range of VDD seems very striking, and 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.

Timing Analysis

- Have managed to carry out a first static timing analysis, using PrimeTime, to assess the state of the Digital_Bottom block. In such an analysis, one estimates all of the parasitic R and C (using Hyperextract) in the clock and signal paths, then looks at the clock for each clocked cell and estimates risetimes, and setup and hold times.
- First analysis used crude Hyperextract parasitics file, and used simple timing models from CERN standard cell library, with SEU-DFF represented by standard DFF model.
- Already, this first analysis showed problems with timing margins in the shift register. There were a large number of hold violations (86) observed. For the CCK net (the Command Registers), the top two hold violations were the two single-point defects seen in the lab. The next two hold violations were for locations with no known problems. For the RegClk net (the Global Register), the top violation was for the observed single-point defect, and the next was for a location with no known problems. Thus, there is quite a strong correlation with the observed problems.
- It also showed significant problems with 40MHz clock distribution (no real clock tree was used). This net has a total capacitance of 5pF, but much worse, a total resistance of 3.9KΩ. A histogram of the risetimes at each circuit node shows a cluster at worse than 10ns risetime, for circuits close to the end of the long XCK net.
- Overall, the problems can largely be traced to lack of proper clock tree generation, which is essential once one has resistive traces longer than a few mm. In addition, CERN cell models assume good clock risetimes, and are not otherwise valid.

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. It has a modest non-uniformity between the values 15 and 16.
- 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 with configuration operations performed at 1.6V and data taking operations performed at 2.4V show the chip works properly, with exception of some problems with BCID labeling and EOE status bits (not observed at lower VDD). This is most likely 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 has been performed, and leads to yields in excess of 90% for full testing of analog blocks, as well as basic digital tests.
- First wafers from AMS and IZM have been bumped and diced. First diced chips have been tested and flipped at both vendors, and assemblies from both bumpers are on the way to LBL. Expect to have them on single-chip boards within a week.

Additional Problems:

- Carryied out some data operations using modified TurboDAQ 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 (4 LSB of timestamp) 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 occured in either writing or reading the FIFO.
- 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 occuring 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 implies that we would 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.

Fixing FE-I2, Next Steps:

Re-spin submission (FE-I2.1):

- Confirmed with IBM that there are six additional wafers from engineering run waiting to start back-end processing (metallization).
- We have modified the clock routing for the 96 bits of place and route shift register to make sure the clock is always conservatively routed against the data (clock enters at the end of the register and propagates towards the start). Each of the three registers has a separate clock trace to optimize the routing. In addition, two inverters were stolen from another location to buffer the clock for Global Register.
- We submitted a new GDS file with these modifications on June 25. CERN performed final DRC checks and forwarded it to IBM on June 26. We expect to get a forecast for a wafers out date in the near future. Hope that the date will still be in July, as this is treated as an engineering run and therefore uses a turbo slot in the foundry.
- Initial quote received was very high. After some negotiation, an intermediate value was received of $57,321.44 for processing six wafers with 3 mask changes.
- This submission should eliminate the major problem with the three shift registers, allowing us to operate the FE-I2.1 chips over a wider range of power supply voltages, and making it almost certain that we can build I2.1 modules. This appeared to be the lowest risk path towards the final production FE, making sure that we can continue searching for problems.

FE-I3 Submission:

- In parallel with the fabrication of the new FE-I2.1 wafers, we have begun work on a corrected version of the place and route block. This involves additional steps in the design flow which were not done for FE-I2 because of the relatively simple circuitry contained in the place and route block (we were wrong...)
- Our chosen design flow for FE-I3 involves using HyperExtract (part of the Silicon Ensemble place and route toolset) to evaluate routing parasitics, and then PrimeTime (part of Synopsys synthesis toolset) to analyze timing margins and guarantee conservative setup and hold times throughout the design. This is the only reliable way to be sure this type of block will work across the full range of operating conditions of the chip once the block size has grown significant.
- The weakest link in this design flow is the relatively low quality models available for the timing of the standard cells, and the fact that we have introduced a number of new cells that will have to be characterized.
- Meanwhile, we have understood the CERN models better, and if one generates proper clock trees, their accuracy appears to be acceptable (since signals normally have rather light loads). However, if the clocks in the design have poor risetimes, then the models are no longer correct.
- In addition, we have found a tool that is part of the Nanosim (TimeMill) toolset which allows simple setup of test vectors for standard cells, and convenient generation of most required output formats. This has allowed us to quickly check the CERN models, and will allow us to produce simple models for our new standard cells.

- Based on our initial timing analysis, we intend to proceed with a new place and route for the Digital_bottom block. This will involve using CTGen to generate clock trees for CCK-related clocks and the XCK crossing clock. We will also use a higher-level cell for the shift register slices, with a clock-in and a clock-out pin, plus inverter delays for the clock inside each slice. This schematic will force the clock to be routed against the data and guarantee conservative timing. Finally, after routing the block, a static timing analysis will again be performed, and any (hopefully minor) setup/hold violations resolved by hand.
- A first look at this task suggests that FE-I3 should be ready in 4-6 weeks. Tentatively expect submission by the end of August. This would provide first engineering run wafers by end of October, and production wafers by end of November.
- Of course, we will carefully evaluate the FE-I2.1 wafers. If it appears that these wafers can produce useful modules in less critical parts of the detector (disks), then we could consider beginning advanced production with them. However, the timing analysis which highlights poor clock slopes due to resistive traces without clock trees, and the observation of the Trigger FIFO problem seem severe enough that this already looks like an uncomfortable path. Nevertheless, it would probably gain us about 3 months...

Status of MCC-I2 Run

- Significant SEU-hardness improvements made in MCC-I2, along with a number of minor bug fixes. Die area increased, but still just barely fits in Flex V5.
- GDS file was sent to CERN on April 21, and forwarded to IBM on April 23.
- Reticle contains 4 identical MCC-I2, two DORIC-I5e, and two VDC-I5e. There are a total of 134 good reticles, or 536 good MCC-I2.
- Five wafers arrived at CERN on June 3. One was shipped to Bonn for immediate testing with their probing setup. The other four were shipped to Genova for testing with Delta.
- The wafer tested in Bonn was later sent to GDSI for thinning and dicing. Maurice has now distributed diced parts to Genova. Several tests have been done on first packaged parts in Genova, showing good results.

Total Costs:

- Cost for MCC-I2 run was $148,701.04, with US share of 21.5%

Summary of MCC-I2 Testing

- Initial testing, using Bonn system with special probe card and FE-I chip on card indicated only one significant problem. Unfortunately, this problem prevents all readback of configuration from FE chips in a module.
- This is a fatal problem for a system as complex as ours (almost 1 Mbit of configuration data per module). In addition, it would prevent proper use of the new auto-tuning circuitry in FE-I2.
- It can be effectively fixed by rather simple modifications to M1 mask in one place. The fix will not provide the feature of blocking any inputs from the FE chips during testing of the MCC-I2 in playback mode. However, this was really a safety measure, and is not expected to cause any operational problems.
- It also appears possible to implement this fix in several chips using FIB. Proposed program is to attach MCC-I2, tested using TurboDAQ, to Flex V5. The entire flex frame would then be used as a support for FIB modification of the MCC-I2. The fixed Flex would be attached to FE-I1 bare modules for testing of the fix, and checking for any other outstanding problems in the MCC-I2 design.

- The fix requires cutting a trace in M2 and connecting one side of the cut part to VDD by a jumper down to the thick M1 VDD trace.

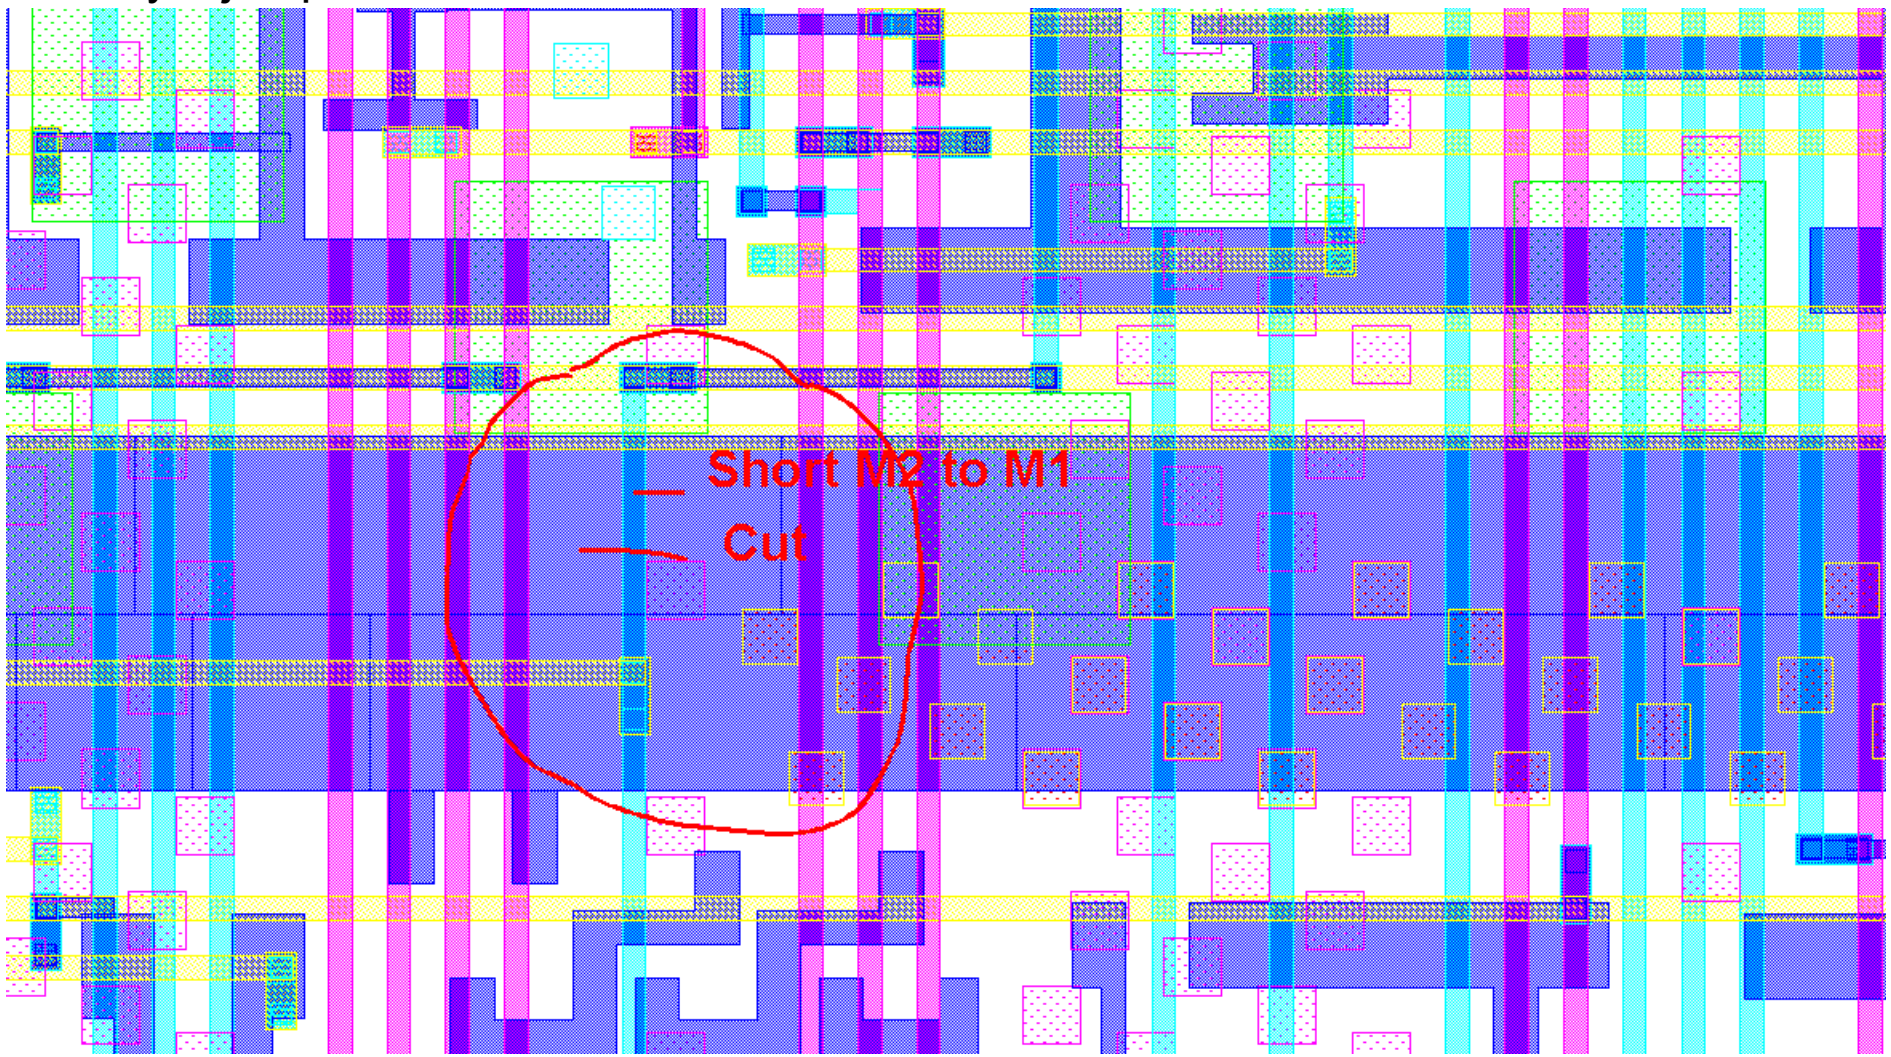

Re-spin submission (MCC-I2.1):

- Have verified with IBM that there are also six additional wafers from this run waiting to start back-end processing.
- Have received quote for processing of six wafers with single mask change (M1), with cost of $46,121.04
- Present plan for re-submission is trying to maximize the testing that occurs before submitting the fixed GDS. Since there are only six wafers left, once they are used, the next step would require a new engineering run. Genova is emphasizing complete testing of packaged parts using their test-vector based system. LBL will attempt to build FIB-fixed modules using Flex V5 and FE-I1 bare modules.
- Our present understanding suggests that no further problems will need to be fixed, and so these new wafers could be considered MCC-I2 production.
- Discussion ongoing with Genova about how much MCC-I2 information to accumulate before going ahead to commit remaining six wafers to metalization. Balance off desire to make MCC-I2.1 a production submission and avoid further costs of another engineering run, against schedule requirement for set of FE-I2.1 modules for high intensity PS irradiation in Oct. The latter requires submitting the MCC-I2.1 no later than about the end of July.

Irradiation and Test Plans This Year

- Carried out irradiation of seven FE-I1 modules in May at PS. Preliminary dosimetry indicates we only reached about half of lifetime dose (1.1 10^{15} p/cm^2) on average, despite initial signs that we had full dose. However, dose was very non-uniform. Most aspects of system worked well. Need improved scanning hardware for next modules irradiation (2D scanning). In addition, efficiency for module irradiation seems to be only about 10% (need 10 times more dose than the target average module dose), which makes it almost impossible to reach the full lifetime dose.
- Modules transferred to pixel lab, and undergoing re-characterization. All 7 modules survived, although there are miscellaneous problems which are disturbing (intermittent data connections, HV bias holdoff problems, thermal problems). Initial experience with tuning inhomogenously irradiated modules indicates significant difficulties. In particular, threshold adjustment system and feedback current (TOT) adjustment system cannot cope with demands of large dose variation.
- Expect that these modules will be tested in mid-July at the SPS.
- Next irradiation is in late July, and should concentrate on FE-I2 single chip assemblies and MCC-I2 packaged parts. We are planning on something like 6+4 parts of each type, which is limited statistics for the complex FE chip. Would be based on similar test setups to those used in I1 irradiations last year. This will give us a detailed understanding of the individual chip performance. In particular, we will focus on SEU testing of registers and readout, and performance of analog blocks, with frequent scans to look for changes.

- August irradiation will concentrate on opto-electronics and opto-boards. This will be very similar to the irradiation of 2002, but should have 4 complete opto-boards with final opto-packages mounted. This is critical for the final production qualification of the opto-boards, opto-packages, and opto-chips.
- August H8 testbeam will include about 7 days of high-intensity running. It is not yet clear whether this can be carried out using FE-I2 modules or not. FE-I2.1 and MCC-I2.1 will almost certainly not be available in module form at that time. Modules built with MCC-I2 will not allow SEU studies of FE-I2. Modules built with MCC-I1 will not allow study of additional SEU-hardening in MCC-I2. It may not be possible to operate FE-I2 modules reliably at VDD=1.6V. Need further experience to understand which parts will be evaluated.
- October high-intensity PS run is intended for FE-I2 module irradiation. Should be possible to have modules made with FE-I2.1 and MCC-I2.1 on this timescale, so these would be irradiated. FE-I3 will almost certainly be unavailable in module form on this timescale, so this most likely cannot be a final production module qualification irradiation, and will have to be repeated in 04. However, expect that differences should be so minor that we can begin production without any risk.

Expect to assemble all information together in PRR in Dec 03.

- This would cover the FE and MCC chips, and all associated issues. The opto-electronics is covered in a separate PRR.
- Should be appropriate moment to cover all performance and radiation hardness issues for full on-detector system.

Production Issues

Wafer procurement:

- Would propose to initially purchase about 50% of the required wafers. MOQ is 48 wafers, and it is possible to write a single PO with multiple items (wafer lots) with individual delivery dates, requesting IBM to stagger the delivery to make sure there is never a large back-log of untested wafers.
- Once we have the final maskset, it is just $120K per 48 wafers (13824 potentially good die).

Yield:

- Significant new information from work on low yield designs with IBM-Burlington. Major driving factor in two recent cases (APV25 and Medipix) was excess ILD thickness and corresponding loss of via yield. There is a recipe in Altis for improving this, and first new wafers for APV25 have highest yield yet seen (90%). Burlington recipe is different, and appears not to suffer from the same problems.
- This suggests that our 90% yield (for 3.5M transistor chip !) is not an anomaly. We will check this by probing two production lots of wafers.
- Propose to still plan for yield as low as 50% in financial aspects, but expect that such yields should not be seen.

Wafer probing:

- System and corresponding tests are still evolving.
- Work is done in Bonn and LBL. Both use a similar setup, and LBL has worked hard to produce a common framework to run the tests (TurboDAQ with probing enhancements) and to analyze the results (Pixel Wafer Enlightenment), with much work done by our postdoc Aldo Saavedra.
- Both teams are presently upgrading to final FE probing system, which will allow a very complete set of analog and digital measurements to be performed. Estimated probing time is likely to be only slightly better than one wafer per day, in part because the high yield causes all tests to be run on almost all die. However, at present, each wafer gives enough die for 15 modules, so the wafer probe is not a bottleneck.

Equipment:

- Situation at LBL is not very solid. Expect 2/3 of modules will be AMS, which requires very significant single-chip probing effort at LBL.
- Presently use very ancient Alessi prober for single-chips, and very unreliable Micro-manipulator for wafers. Recently suffered multi-week downtime on MM due to power supply problems (third problem since it was purchased). Suggest we should consider adding another prober to avoid potentially lengthy downtimes during production.

- Also, probing system relies on GPIB instruments. Duplication would cost about 15K$. This should be done, and would allow us to do wafer and single-chip probing in parallel, or to perform single chip probing on two systems in parallel, if the schedule requires it at some stage.

Test Suite and Data archiving:

- Have had good experience with FE-I1 probing, with essentially no surprises with chips containing defects that were not tested for (except for those chips where the selection was skipped).
- Test suite for FE-I2 is not yet finalized, but most tests have been coded. Will still require some cost/benefit optimization to reach final list of tests for wafers.
- Present system makes very limited use of rigid PDB. However, probably need to make slightly more use of it. Plan to produce revised list of test results soon.
- Access to raw data, and generation of ROOT file archives, as well as all analysis of the data, is performed by very flexible environment called Pixel Wafer Enlightenment. This tool, combined with Module Analysis Framework, should be adequate for analysis of all relevant data generated during production. However, significant additional software effort is required to automatize many aspects.