

Pixel Upgrade R&D Issues and Planning

K. Einsweiler, LBNL

Third Hit System:

- Assume this is a continuation of two-hit system, so there is no new electronics activity for this. This is covered by Gil.

B-layer Upgrade:

- Assume that this is the first significant upgrade opportunity.
- However, assume that it is based on rather similar module geometry to that of present detector, in order to preserve basic services and mechanical design.
- Concentrate on improved segmentation and reduced material.

Total Pixel Replacement for SLHC:

- The date and scope for this is largely unknown. Luminosity goal could be as high as 10^{35} , but time structure is unknown. Could be close to DC (almost unbunched beams) or could be bunch-based with higher crossing frequency (e.g. 80MHz).
- At this point, would re-examine all assumptions (chip size, module geometry, pixel size and layout). Would also consider L1 trigger functionality for VTX triggering.

LHC Accelerator Upgrades

Upgrades without major changes to the machine hardware:

- Collide only in ATLAS and CMS, increase bunch intensity and crossing angle slightly, and increase B to 9T (15 TeV collision energy). This could lead to a luminosity of about $2 \cdot 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$.

Upgrades with changes to insertion regions only:

- Increase both crossing angle and bunch intensity significantly, modify insertion quadrupoles to reduce β^* by a factor two, and improve RF to shorten bunch length by a factor two. This could lead to a luminosity of about $5 \cdot 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$.

SLHC upgrade (factor 10 in luminosity):

- Modify injectors to increase beyond “ultimate” beam intensity. Possibly increase SPS energy to 1 TeV with superconducting upgrade. Use super-bunches up to 300m long to increase luminosity still further for a given current in the machine.
- Note in the superbunch case, there would be essentially no short-term time structure in the beam, so the whole notion of “tracks associated with beam crossings” would most likely be lost. Would have to define arbitrary interval with acceptable pileup (order of 10ns ?)

Problems and Issues with Present ATLAS Pixel Detector

Material per measurement layer is very significant:

- Adding all contributions gives about 2.5% X_0 per layer at $\eta=0$.
- Major contributions come from electrical power services and cooling services.
- Electrical services are driven by relatively large currents for present designs (roughly 2A total current for 10 cm² active area).
- Cooling services are driven by relatively large power dissipation for present designs (almost 6W for 10 cm² active area, at end of lifetime, including leakage current and possible increases in supply voltages).

Services burden is enhanced by redundancy:

- Need for very high efficiency throughout the detector lifetime has caused us to be very conservative (every hit is important for every track).
- Each pixel module (1744 total in 3-hit system) has individual power service connections all the way back to the counting rooms.
- Each cooling unit (bi-stave and bi-sector, 80 in total) has a separate cooling circuit with inlet and exhaust lines all the way back to the cooling plant.
- First operational experience will provide feedback on how this can be optimized.

Bump-bonding production flow is critical:

- No mainstream commercial vendors offer bumping at our desired pitch.
- Prototyping has been underway with present vendors for more than five years. Each of them has a capability of approximately 1 m² per year of bumped assemblies.
- These vendors do not run a dedicated production line, but have only a single line on which all work is done. Quantities are sufficiently small that each time we provide more material, there are problems.

R&D Topics in Pixels

Mechanics and Cooling:

- Present structures based on carbon-carbon materials for very low CTE, and evaporative fluorocarbon cooling fluids. The evaporative system is challenging to operate, with complex feedback and control required to achieve stable operation over a wide range of conditions. However, it has very significant cooling capacity due to latent heat, and cooling liquid within sensitive volume is minimized to essentially a thin film on the tube wall, which evaporates and cools system.
- System requires large number of cooling circuits, with significant additional power expended to bring output temperature back up above the dew point.

Steps towards an improved design:

- Optimization of material and number of cooling circuits/connections was begun under Phase I SBIR.
- Concept involved use of cooling structures as both local and global supports (present design uses only local supports for cooling).
- It is based on heat pipes (largely carbon) with a condenser at the end of the volume to reduce the pipe count.
- Simple prototypes were tested under Phase I. Unfortunately, Phase II was not funded, so work related to pixels has stopped (some work continues for space applications).


Heat pipe cooling concept:



Groove wick, 31 total
0.3mm wide opening,
0.6mm deep



Heat pipes tied together with simple structure. Condenser at end fed by a few cooling circuits for redundancy.




- How can we support this type of R&D within our community ? Cooling and mechanics expertise within HEP community is somewhat limited, and the subject does not receive the respect or funding required to advance the state of the art...

Electrical interconnection (bump-bonding ?):

- For high density interconnects between sensor and FE electronics, presently rely on 50 μ pitch bumping from “non-standard” vendors (military aerospace, or R&D institutes). No change in this situation is foreseen in the next ten years.
- Commercial roadmap for such bump-bonding over area arrays is rather flat:

IC Technology Roadmap: Chip Interconnect Pitch




Source: ITRS

- Would benefit significantly from developing a US-based vendor, providing local access and closer collaboration.


- For longer term, there are more aggressive ideas for interconnection that can be explored. They take advantage of the fact that the individual channels may not require perfect electrical isolation.
- One example is to use a conductive thin-film (e.g. Indium-Tin Oxide (ITO) film) which can achieve several $K\Omega$ of inter-pixel isolation in thicknesses of a few hundred nm. This could be combined with a very fine-pitch anisotropic conductive film (z-axis film) to provide very inexpensive but fine pitch hybridization.
- Present state of the art from one of our current bump-bonders (IZM Berlin):

Future Requirements: Bump size and Pitch



Solder Bumping PbSn37/63
Diameter: 10 μm
Pitch: 20 μm

Staggered bumping pattern:



Allows connecting finer pitch pixels without matching sensor pitch to bump pitch.

Requires use of top metal layers in FE chip to form “patch panel”

Sensor Development:

- Present ATLAS sensor designs were carried out by German groups (Dortmund and MPI), in collaboration with the ROSE (RD-48) collaboration. Based on conventional double-sided processing, with n+ pixel implants in n- bulk material, and using a low-dose patterned p-spray to achieve inter-pixel isolation. High-temperature oxygenation step performed to improve radiation-hardness.
- These sensors have been extensively tested to NIEL doses of 10^{15} 1 MeV neutron equivalent (although NIEL hypothesis not strictly valid for oxygenated material). This corresponds to roughly 50MRad ionizing dose for hadrons, and is the present lifetime dose estimated for the outer pixel layers, and for 2-3 years for the B-layer.
- Observed charge trapping is relatively modest (less than about 20% charge loss) when operated at -600V. Main issue for operation of the present detector is managing leakage current (about 50nA per pixel at -7C). Operation at -20C would probably provide acceptable performance to higher doses (2×10^{15} NIEL ?).
- Further design studies for capacitance reduction (which reduces noise and power for FE electronics) would be useful.
- In world where micro-electronics has eight or more metal layers, can move beyond concept that sensor element and electronics channel must have identical geometry. Could imagine bricking schemes or non-rectangular sensor elements that could optimise results. However, there will always be a tradeoff between robustness (high performance operation with reduced charge after irradiation) and better resolution based on charge sharing in unusual sensor geometries (requires large signal).

- First R&D work would be irradiation of single-chip assemblies to higher doses, which can be achieved with dedicated PS runs. Unfortunately, these runs will no longer be available after 2004, and not clear we can muster the resources on that timescale. However, this would be the only chance to assess failure modes in the present design when pushed beyond 10^{15} NIEL, and as such would provide unique information.
- Possibility to work on more advanced pixel-specific sensor geometries (somewhat independent of basic design recipe for radiation hardness). Examples include bricking or other highly charge-shared geometries that could improve spatial resolution for adequate signal sizes.
- Longer term, expect RD-50 to provide some ingredients for a next generation design (their goal is sensors for 10^{16} NIEL). The US sensor community is limited by the lack of appropriate vendors to work with, so center of gravity of work is in Europe.

Power Management:

- Long service runs in ATLAS to reach areas where standard commercial power supplies can be operated (up to 140m from pixel system).
- This requires a hierarchical system, with supplies in counting rooms and rad-tolerant regulators at PP2 (inside muon system, 12m from the pixel detector).
- Low-mass services design attempts to minimize material inside tracking volume, but requires a large voltage drop budget to achieve this. This requires cooling of services, and use of regulators in remote sense mode to compensate for large drops (up to 4V provided at PP2 to achieve 2V on pixel module).
- Have also designed building blocks for power management into the FE chips. These include two integrated, programmable linear regulators (one digital, one analog) as well as a shunt regulator for current-mode operation. All have been tested, but none appear to be necessary for operation of the present detector.
- As maximum supply limits continue to decrease, need to protect fragile gate oxides will make this approach vital. Present linear regulators designed to operate continuously with input voltage of 4V, in spite of the fact that the 0.25μ process has a maximum voltage rating of 2.7V. Not difficult, but requires careful analysis.
- This approach would allow us to operate 0.13μ (or less) electronics with the present power distribution network, with essentially no additional risks. Of course linear regulators do not improve the power efficiency (they are not transformers !)
- This approach does not address the services issue.

- Reducing services requires significantly more “sequential” connections than we have in our present robust parallel power distribution schemes (serial powering), or high-efficiency non-magnetic DC-DC conversion that would allow bringing high-voltage low-current power into a module and converting it to low-voltage high-current needed for the electronics (charge pump regulators ?)
- One scheme discussed in ATLAS pixels involves sequential connection of multiple modules using constant current power supplies and internal regulators. Basically, recycle electrons by using the same ones for each module on a stave/sector.
- This scheme uses a single shunt regulator per FE chip (ideally per module) to provide an almost fixed voltage drop across a module (between VDD and DGND), in order to ensure that internal linear regulators can do their job. None of these building blocks need to operate at high voltage, so they can be built using rad-hard layout techniques in DSM CMOS processes, just like the rest of the FE electronics.
- Complications are that no common grounding scheme is possible (the FE substrates of each module sit at a different potential), the HV bias for the sensor must be floating, and referenced to the local module supply/ground, so individual HV supply channels are needed, or -600V must be treated as a common point as well (very peculiar grounding scheme). In addition, all “signal” communication to the modules must be AC-coupled to avoid the need for very wide common-mode range.
- Will use existing building blocks to carry out larger scale system test of serial powering scheme. This requires fabrication of a new Flex design to deal with different power distribution within module, but is otherwise straightforward.

Data Transmission:

- All data transmitted in and out of pixel modules is referenced to the LHC crossing clock for simplicity. It is transmitted using VCSEL and silicon PIN arrays on multi-mode fiber (only the first 12m is rad-hard).
- As the bandwidth requirements for B-layer operation were understood, we adopted data transmission on both clock edges (80Mbit/s), and dual link transmission (160Mbit/s). This is enough to provide a nominal factor two of margin at design luminosity for the present B-layer.
- Increased luminosity or increased trigger rate will require a major improvement in data transmission. Simply adding more links at 80Mbit/s is not cost effective (fiber is expensive). Also, components of optical data transmission (CMOS electrical driver/receiver, VCSEL/PIN optical driver/receiver), are capable of much higher bandwidth operation, without huge penalties in power.
- Next generation module architecture would probably include more sophisticated TTC link to provide high-frequency data transmission clock which might be independent of the crossing clock (or would most likely be a frequency-multiplied LHC crossing clock). Would need to develop corresponding “rad-hard gigabit” link technology (but low power), probably switching to single-mode fibers.
- More sophisticated approaches are being explored in which one uses multiple-frequency VCSELs to transmit multiple fast data streams on a single fiber.
- There should be a pixel-specific R&D effort, but not clear how much pixel-specific development would be required - perhaps rather little.

On-Detector Electronics

- Almost all aspects of pixel detector performance are driven by the performance of the on-detector electronics.
- Improved physics performance will be driven by improved point resolution and improved segmentation (and of course reduced multiple-scattering due to reduced material).
- Radiation-hardness, using modified layout techniques, for DSM CMOS seems intrinsic, and only improves with further thinning of gate oxide. Therefore, we can follow the roadmap of the micro-electronics industry. Since pixel electronics is driven by segmentation, this suggests the possibility of significant improvements due to evolution of commercial technology.
- Furthermore, modern processes have been optimized to produce excellent yield for large (several cm^2) designs, which is an advantage for pixel designs, as chip boundaries force us to invent “special pixels” (long or ganged) to cover the gaps.
- Yield example: 400K transistor 50mm^2 ABCD in DMILL 0.8μ yields less than 25% good die, 3.5M transistor 80mm^2 FE-I2 in IBM 0.25μ yields more than 90% good die. Present die size driven by early yield discussions for DMILL and Honeywell. Die size could be increased by a factor 3-4 for next generation, assuming yield is OK.

Advantages of newer “nanoscale” CMOS processes:

- Interconnect possibilities improve. For example, IBM CMOS6 0.25 μ offers 6-metal, with one thick metal in Aluminum (higher resistivity). IBM CMOS8 0.13 μ offers 8-metal, with three thick metal layers, all in Copper so resistivity is lower. This will be essential to implement larger pixel arrays without major voltage drop problems (present FE chip required clever compensation circuitry to cope with voltage drops on power busses).
- Speed and density of process improves dramatically. However, assuming we continue to use annular NMOS in order to achieve the most reliable radiation hardness, cannot take full advantage of this.

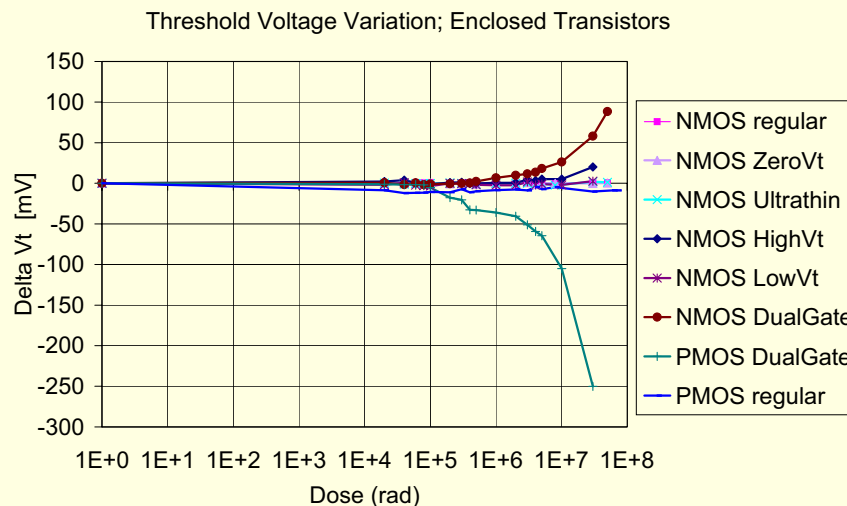
Disadvantages of smaller CMOS feature size include:

- Reduced supply voltage, with implications for power distribution, and for dynamic range in analog designs. For example, 0.13 μ is a 1.2V process.
- Reduced SEU thresholds and increased SEU cross-sections due to more compact geometry of devices (easier to induce transients on smaller capacitances) and lower device thresholds. Already observed in PSI irradiation of 0.13 μ SRAM.
- Possibility of gate rupture due to high-electric fields induced by heavily ionizing particles (not seen for 0.25 μ electronics in proton beams).
- We can probably live with these...

Radiation Hardness:

- First measurements for 0.13 μ process indicate significant improvement over 0.25 μ process. In particular, VT shifts for 0.13 μ thin oxide devices after 70MRad are almost negligible, as opposed to the 50-100mV which would be expected for 0.25 μ .

- W=10 μ m, L = min size:
 - Standard, HighVt, LowVt, Ultrathin = 0.12 μ m
 - DualGate = 0.26 μ m; ZeroVt = 0.42 μ m



First results from irradiation of test transistors in IBM 0.13 μ process.

Dual gate devices have 5nm oxide, similar to 0.25 μ process. Other devices are 2.2nm or less.

Thin oxide shows no shift, thick oxide shift is similar to present CMOS6.

- Present FE chips already suffer from VT shifts at large doses, which affect the threshold dispersion and the TOT measurement (feedback current). Significant tuning capability required to overcome these shifts. Design would be much simpler without them.

Scaling of designs:

- For matching purposes, the V_T mis-match in mV/μ^2 remains roughly constant, so the device size required to reach a certain level of matching remains about the same and does not scale. In pixels, essentially all analog transistors are in weak-inversion region, where current scales exponentially with V_T , so matching is quite challenging. Note also that a large device has higher capacitance due to the thinner gate oxide (roughly double between 0.25μ and 0.13μ). Therefore, analog designs may not scale much. Will increasingly rely on digital circuitry (“tuning”) to achieve performance.
- Have already observed in FE-I1 that squeezing the layout of SEU-tolerant circuitry is dangerous. Appears that for PS energies, there is a characteristic size for an upset, which is roughly 5μ . Redundant circuitry which is squeezed into such small regions will often see correlated upsets, which severely limit the achievable SEU-tolerance. Finally, SEU thresholds for small devices are lower, and cross-sections higher. Therefore, SEU-tolerant storage locations may also not scale significantly with the process feature size.
- Additional issue is power consumption. For digital readout, will have to work hard to avoid consumption scaling linearly with the channel count. For analog readout (dominant current consumer in present chips), not clear how capacitance will scale with area reduction. Expect that capacitive load will not scale down with area, but have significant constant terms. This will cause the analog consumption to increase in order to achieve constant timing performance with greater segmentation.

Charge Measurement:

- Present design uses timing infrastructure required to tag hits at 40MHz to make a TOT measurement in each pixel (effectively a simple Wilkenson type ADC in each pixel). This is a very effective approach, but requires tradeoffs between charge resolution and deadtime, which will not be acceptable at higher occupancies.
- In addition, the present preamplifier design has a TOT which is quite linear in charge. Because of the $3.2\mu\text{s}$ latency limit of ATLAS, by which time all hits must be stored in on-chip buffers, this behavior imposes a maximum charge limit as well as a minimum charge threshold. The range for the present generation of FE is tunable, but is roughly 3Ke to 100Ke.
- A peak-sensing approach would provide reduced deadtime, and a natural saturation mechanism, but digitization of the charge with decent resolution would be much more challenging.

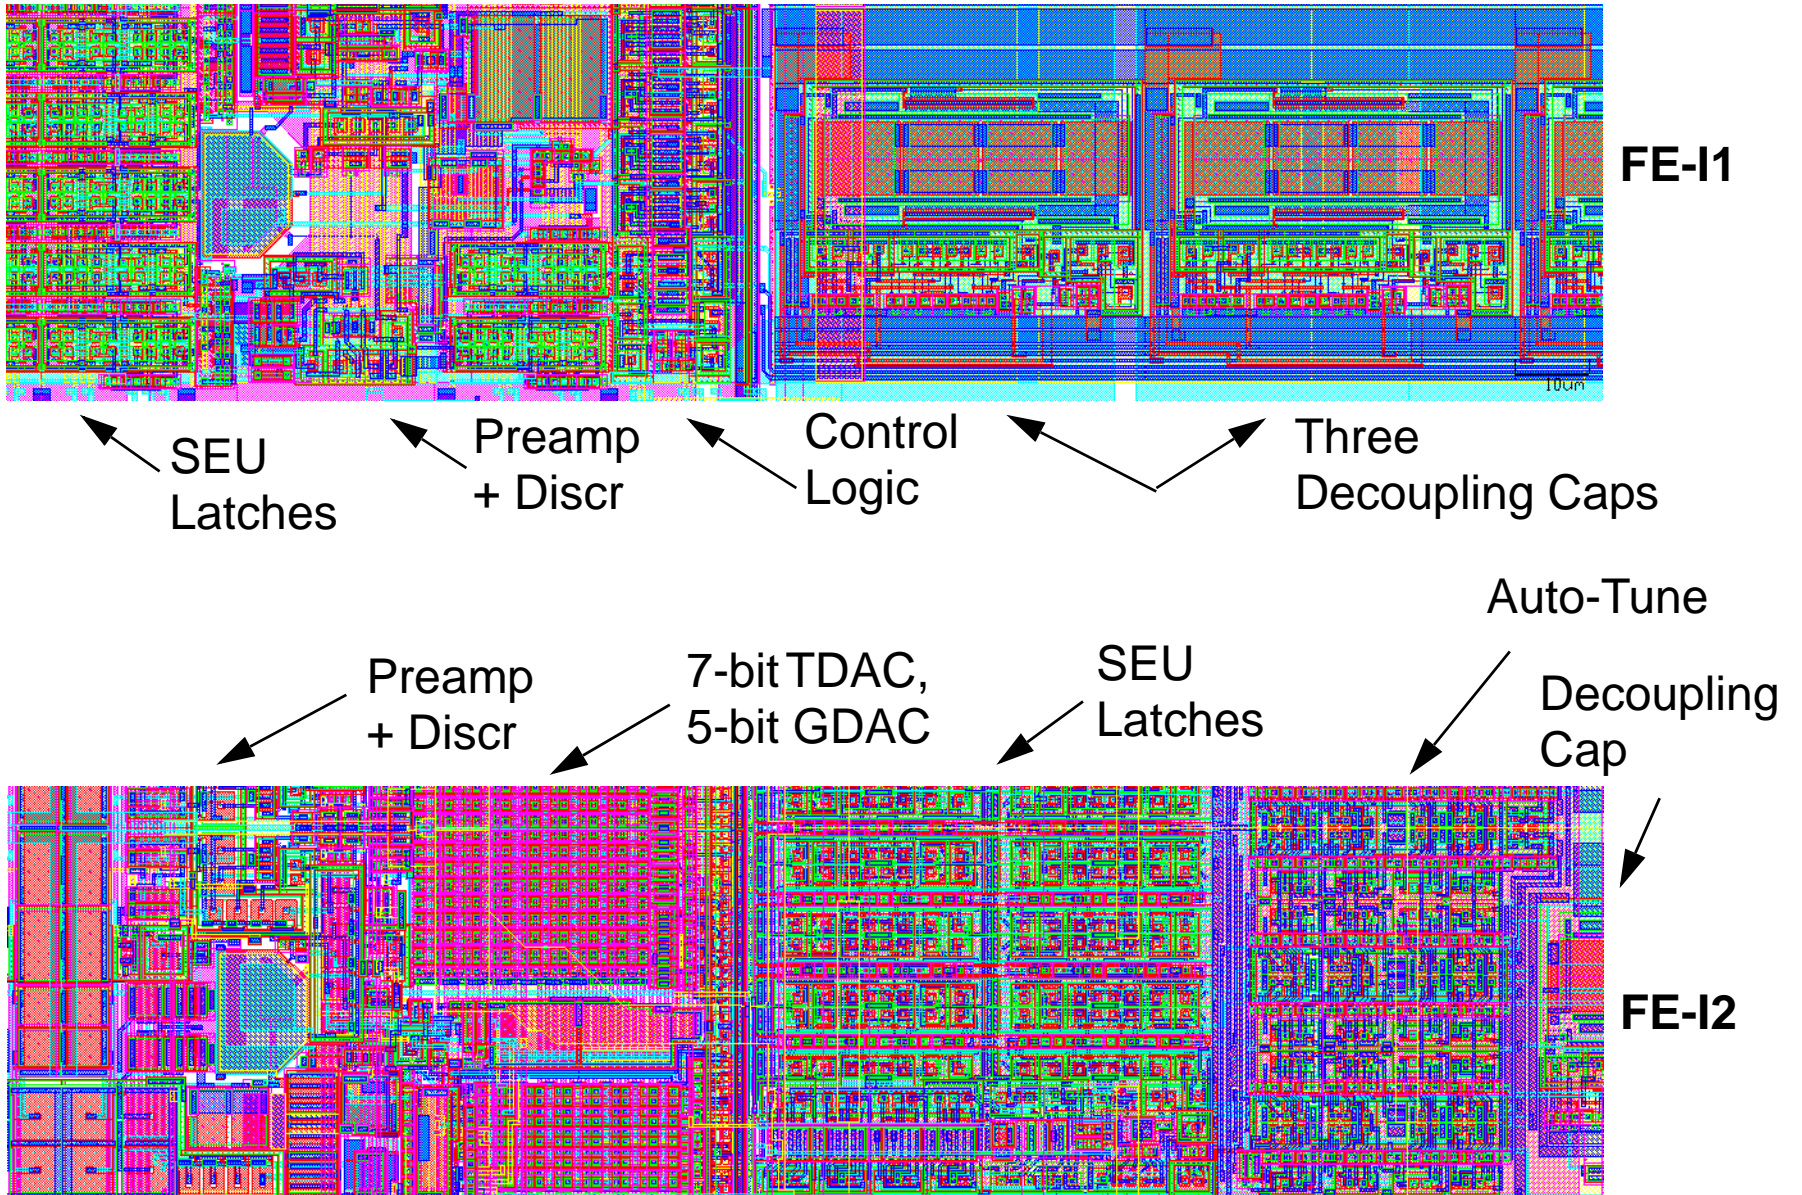
Timing performance and power consumption:

- Basic analog timing performance is strongly related to analog current budget. In particular, it is very challenging to achieve small pulse height slewing between large and small charges given the requirement for single-crossing readout. This is the so-called timewalk problem. Present chips achieve 5Ke in-time threshold.
- Have implemented a simple digital correction circuit which uses a single TOT threshold to copy a hit into its “tagged” crossing as well as into the previous crossing. The idea is that hits with a TOT below a certain threshold must be out of time by one crossing due to timewalk. More such tricks could be used in the future.
- Some indications that a given timing spec can be met with lower power using a zero-crossing type design, where peaking time is reduced, but peaking time is carefully tuned for each channel to preserve good timing uniformity. Only a complete design will confirm this hypothesis.

Digital design issues:

- Need to address SEU-tolerance in the standard cells and the architecture from the very beginning in order to achieve reliable operation at the highest luminosities. Already, standard cell SEU-tolerant latch and DFF for 0.25μ work very well.
- Design for fault-tolerance to the extent possible in order to improve yield. Present design uses independent column pairs with minimal glue logic. Column pair enable controls multiple features to ensure defective column pairs do not kill complete chip. Further improvements might be possible in the future.

- Pixel front-end + control in FE-I1 was about 100μ long. In FE-I2, it is 200μ long.



- New layout is extremely dense. Hit logic and digital readout fills remaining 200μ .
- It will be very challenging to squeeze the area of such a cell by a factor 2-3.

B-Layer Upgrade

Assume same basic module geometry and specifications:

- B-layer would still be built with similar geometry staves, similar services and similar current budget for module operation. Module dead area might be improved through use of larger FE chips. Could investigate use of etched dicing of thinned chips to reduce dead regions at chip boundaries.

Basic Goals:

- Improve segmentation. Evaluate nominal target of $30\mu \times 200\mu$ for pixel size. Further simulation needed to validate whether this is optimal. This needs design studies for both electronics and sensors. Evaluate whether this was best achieved using staggered bumps of the present pitch (larger capacitance, but known technology) or trying to scale bumps with vendors (requires new rounds of prototyping, with significant numbers of good parts consumed)
- Reduce material budget by using larger FE chip size to create module with reduced dead area, allowing reduced overlap.
- Explore more advanced power distribution concepts such as serial powering to reduce electrical services budget. This would first be done with “system tests” based on modifications of existing modules. Depending on the results, further prototypes might be needed.
- Explore improved cooling concepts based on heat pipes to minimize cooling services burden.

- Try to improve overall radiation hardness of electronics and sensor by at least a factor of two. Also develop a more SEU-tolerant design through improved standard cells and improved architectures. The first step involves a more detailed understanding of failures of the present generation of sensor and electronics.
- Try to implement higher frequency data transmission system using higher performance MCC incorporating the necessary link support circuitry.
- In general, try to improve physics performance by better segmentation and less material. Try to implement first versions of improvements in all areas to continue pushing the envelope, and developing necessary ideas and experience for SLHC upgrade (which will be truly challenging).

R&D Program:

- If it is possible to find adequate support for engineering, it would be reasonable to begin R&D in some of these areas already in 2004.
- IC effort for pixels now being wrapped up, so by next calendar year, there will be manpower available.
- Two key sets of measurements would be interesting to work on using present production electronics. First would be the irradiation to double dose at the PS. Second would be construction of a stave or sector prototype using current-mode or serial powering.
- Much more difficult to find people and funding for cooling R&D.

Assume we will migrate to the next generation of CMOS:

- Present CERN plan is to skip 0.18μ and jump to 0.13μ (IBM CMOS8). This seems like a reasonable step for pixels. This generation is already a commercial process for micro-processor giants (Intel, TI, Xilinx are in high-volume production), and is stabilizing from IBM. CERN has not yet started to move forward on corresponding frame contract, but expects to do so shortly.
- Major issue is high cost: mask costs go from about 125K\$ for 0.25μ to about \$600K\$ for 0.13μ , so barrier to prototyping becomes very high.
- Access is available through MOSIS, and we have an NDA in place already. Minimum die size is 10mm^2 , and cost is 50K\$ for 40 die.

Issues for the FE chip design:

- What is the right front-end concept: Standard preamp/discriminator or zero-crossing for timing ? TOT or peak-sensing for charge ? This depends critically on scaling of current budget with pixel size/count. Would plan to do some initial design studies, but will want to generate some silicon fairly soon to understand analog design issues as early as possible. This would probably involve building a test chip with modest changes to existing designs in order to gain experience before investing too much design effort.
- What is the right pixel size ? Do we give priority to improving segmentation in the narrow or the long direction ?

- What is the right chip size ? Present total module electronics footprint of 22.2mm x 60.6mm gives an active area of 16.4mm x 60.8mm, or only 74% live fraction.
- Maximum die size is 21.4mm in height, so could imagine a single chip with column pairs containing 512 rows of 32μ pixels with perhaps 3mm of overhead at the bottom of the chip (present chip has 2.8mm). The pixels could be 200μ wide, with 32 column pairs. The FE chip would have 32K channels in an area of 19.4mm x 12.8mm (3 times the current FE chip). Such a chip would probably have 50-70M transistors. A module based on such a chip would use only 5 FE chips to provide 160K pixels and about 85% live fraction. All of the ganged pixels would be eliminated. This concept could serve as a starting point for more detailed design.
- Difficult to predict the power consumption of such a chip, but expect it will be difficult to keep the total module current the same. Expect to use internal regulators in each FE chip. Can one imagine a current-mode powering scheme instead of the traditional voltage mode for perhaps pairs of modules ? This would already allow a significant reduction in services (or alternatively, better segmentation with the same services).
- Would also be worthwhile to survey existing high-speed link development and study different approaches to low-power higher speed data transmission. This could involve collaboration with others, as it is a more generic problem.

Next steps for electronics:

- Work with CERN/RAL/FNAL on next generation design environment.
- Further study on radiation-hardness issues, and allowable transistor geometry. Initial results suggest the process (dominant effect is V_T shifts in PMOS) is much more radiation hard than 0.25μ . Appears to be no need for guard-rings around NMOS, and for wider devices, no need for annular NMOS at all.
- Begin studying analog design issues. Matching coefficients are similar, and so expect that many analog transistors cannot shrink (but their capacitance will grow compared to 0.25μ). Limited supply range imposes new approaches to dynamic range management (even bandgap references sit beyond the supply rail).
- Begin study of “nanometer” design flows for digital design. Wiring issues and signal integrity become even more critical than for 0.25μ , since not only placement but actual routing critically affects performance (dominant capacitances are to neighboring wires). Look at developing SEU-tolerant and very TID-tolerant standard cells.
- Put together a test chip submission in 0.13μ , involving some analog blocks (DACs, current references, simple front-end prototype) and some digital blocks (SEU tolerant cell designs). This would be a first necessary step before embarking on a full design program.

Total Replacement Scenario

- This scenario would arise once the accelerator performance had increased to the point where major tracker improvements were required. It would be part of a larger picture of extending pixel-like detectors to larger R, replacing straws with strips, etc.
- Would naively imagine a “large area” pixel optimization for larger R, and a “high performance” pixel optimization for small R. Both would evolve from present design in different directions. The “large area” path would need to minimize power/current per unit area to reduce material, and optimize for very large die size in a mature (low cost, high yield) process. The “high performance” path would need to use a cutting-edge process to get maximum segmentation and performance, and would follow from the B-layer upgrade development.

Inner pixels:

- This would follow on the development of the B-layer upgrade, and would continue to use all the improvements developed in this area.
- On the timescale of SLHC, this is likely to be 90nm process for electronics, or even the next technology node beyond 90nm, depending on how rapidly the industry development proceeds.
- It would need to benefit from RD-50 sensor development, as no appropriate sensors exist to go to 10^{16} NIEL doses.

Outer pixels:

- Large area region likely to cover about 10m^2 . Needs to provide segmentation much better than present strip systems, but not equal that of present pixel systems.
- Something like $50\mu \times 5000\mu$ could be interesting (about 40 times lower area per channel than present SCT). A chip could be a single column pair wide, as would like to avoid covering full sensitive area with electronics when it is not needed.
- This might lead to a chip of 2×512 channels that might be roughly $4\text{mm} \times 20\text{mm}$, thereby covering only 20% of the area. One could use commercial bump-bonding (e.g. IBM C4 is 100μ bump on 200μ pitch) which can be ordered as a standard option at the time of wafer production. Sensor side just has UBM, and gets readout chips flip-chipped onto appropriate locations. Sensor substrate could be about $3.2\text{cm} \times 10\text{cm}$, with two rows of readout chips, requiring 3000 modules for a full system.
- Lots of optimization will be required to reduce cost and material per unit area (emphasize reducing power per unit area).
- Could almost use present generation of pixel sensors, but operating temperature would probably need to be lowered. Total dose would be about 2×10^{15} NIEL for 20cm inner layer and 10^{35} luminosity ? For -7C would see up to $1\mu\text{A}$ of leakage per pixel, or about 4A at -600V for 10m^2 of area.

Triggering in Inner Detector:

- A major issue for this upgrade would be the triggerability of the pixel system.
- From the FE electronics point of view, a triggerable system would be a simple extension, just providing direct outputs from the EOC buffers instead of storing data until L1 appear.
- However, either the core trigger functions would be off-detector, in which case massive improvements in data transmission technology would be needed, or core functions would be on-detector, in which case massive R&D would be needed to create very sophisticated digital processors in the pixel volume.
- Personally, believe the trigger problem for LHC is so challenging that one would need to transmit the data over multi-Gbit optical links. This might require SiGe development for high-radiation environments, plus a conversion to single-mode fiber infrastructure, and would clearly be very challenging within a limited volume and power budget.

Proposal for FY04 Start

- Believe that this is not too early to start for B-layer upgrade R&D. It is still very early for SLHC R&D, because so little is known about the accelerator configuration.
- Propose that initial focus be on-detector electronics, as experience shows this is always the critical path for silicon detectors. This is also the one area where work can proceed without major conflicts with the ongoing project.
- Initial work would be primarily in the analog area, including design studies in 0.13μ . The major advantage of starting design so early is that we could re-consider some of the aspects of the current design which have been largely frozen for more than 5 years. This would be about 0.5 FTE, and might be partially leveraged by LBL Division funds (however FY04 looks grim for us).
- Parallel work in the digital area would proceed, emphasizing SEU-tolerance and looking for lower power readout solutions. This would be about 0.3 FTE
- This work would culminate in a small test chip submission, either to MOSIS or via CERN if some early MPW possibilities become available there in 2004. This would range from perhaps 10K\$ (CERN) to 50K\$ (MOSIS).
- It would also give us a start on some key technical decisions that must be made early, including pixel and die size, power distribution and data transmission schemes.
- The total cost would be in the range of 100K\$ to 150K\$.
- If additional funding available, consider cooling R&D and bump-bonding prototypes.