Response to the
Report of the US ATLAS Pixel Baseline Review

January 19, 2001

The response to the final recommendations of the Review Committee are given below. A preliminary response, prepared before the completion of the final Report of the Review Committee, may be found near the bottom of the page at
http://www-physics.lbl.gov/~gilg/USATLASReviewNov00.html
and contains some additional information.

Rec 5-1 The committee supports the development of the 0.25um CMOS Front End chips as a viable option. The recent results of the RD-49 CERN research project has indicated that this process is radiation tolerant in the LHC radiation environment. To minimize schedule risk, the committee also recommends working with the two identified vendors that manufacture ASICs with this process. The design of the Front-End ASIC chip is very challenging and the ATLAS collaboration should support the effort with its top experts in this area.

We are developing the front-end chip design in an environment that supports both TSMC and IBM, the two potential vendors. However, since a contract exist between CERN and IBM, we are targeting production to IBM. Should there be fatal problems with this approach, we believe it possible for production to occur at TSMC(through MOSIS).

We agree the design is challenging and would welcome any useful support from the collaboration.

Rec 5-2 To keep all the viable options open, work should still continue on the 0.8um-DMILL process, which was also shown to be radiation tolerant in the LHC environment. Other chip types that have been submitted for preproduction runs and prototyping have acceptable manufacturing yields. Every effort should be made to help the vendor resolve the manufacturing process problems through the use of outside experts in the chip manufacturing and processing area.

The situation with DMILL has evolved since the review. The FE-D2 run included all chip types and preliminary tests have been completed on all of these. Based on this information, we intend to continue the characterization, including after irradiation, of all four of the required chips. However, design priority is being given to the 0.25 micron prototypes, particularly to the front-end IC, until this characterization is complete. A decision will be then be taken to continue or not continue development in DMILL for each of the four relevant chips, one by one.

We believe we have made every reasonable effort to resolve manufacturing process problems with Atmel. This has involved considerable contact with experts outside the pixel community, as partly presented at the review. We are continuing this process and have transmitted to Atmel the preliminary results of our most recent tests of the FE-D2 run. A similar exercise is also underway for the SCT ICs(ABCD3) and we are in close contact with this effort.

Rec 5-3 SEU performance needs to be predicted using cross-section estimates and then measured for each design. There are still open areas that still have to be addressed in the post-radiation performance of the pixel design.

We agree that SEU effects are a significant concern for microelectronics located very close to the IP. We are pursuing a significant design, prototype, and measurement program to address this concern. In our first 0.25m test chip, we will already include several different FF designs, with different levels of SEU tolerance. One will be a standard cell design, one will involve three-fold majority logic to resolve bit flips, and one will use a special SEU-tolerant FF design developed by CERN and RD-49. We expect to compare the measured upset rates for structures in our test chip with the estimated SEU cross-sections and thresholds predicted by RD-49. This work will initially concentrate on the 55MeV protons available at LBL, and then be extended to the CERN PS (24GeV protons) and ultimately to PSI where the 300MeV pions are an excellent model for an LHC underlying event.

In our 0.25m design, we will use a very robust technique such as redundant storage (perhaps three-fold majority logic) for our critical configuration information. The exact choice on what approach to use will depend on our measurements of relative sensitivities, but we believe that a redundant scheme is the only way to achieve the desired reliability (integrity) of the stored information. For less critical information, whose disruption would affect only a single channel in a 50K channel module, we presently plan to use the SEU-tolerant FF design, which offers reduced sensitivity but no redundancy.

In addition, we are specifically examining the effects of upset events (small amount of charge deposited in a specific FET) on the performance of our critical digital logic blocks using simulations, and depositing spurious hits directly in the Verilog simulations. We are attempting to ensure that all state machines and logic blocks will suffer only minor transient effects after an upset, rather than longer term disruption of the chip operation.

Rec 5-4 The committee believes that more than 2 submissions of test chips are likely to be required to obtain a satisfactory ASIC. A more detailed plan on how many submissions is being requested.

We have worked out a more detailed plan for all activities required to reach production quality 0.25m pixel electronics. We have prepared a revised summary schedule and milestone list based on this detailed work plan. Our planning now includes two TSMC test chips(via MOSIS) prior to the submission of the first IBM engineering run in July 2001.

The first test chip, referred to as a digital test chip even though it also contains some analog blocks, will contain many individual circuit blocks that will be used in the pixel chips. It would be submitted in the Jan 8 2001 MOSIS run, to provide early feedback on our evolving design. We intend to include the current reference and current-mode DACs used for bias control, the revised LVDS I/O blocks, several shift registers based on different FF designs for SEU studies, the present SRAM storage cell used in the pixel, and several synthesized digital blocks from the bottom of the pixel chip (timestamp generator for example). The intention of this first test chip is to check our understanding of the design rules and basic SPICE parameters, and then to evaluate in detail the analog performance of some actual design blocks. Finally, it will be irradiated and serve to check our understanding of the process and some basic design blocks under irradiation.

The second test chip is similar to ones we have built for previous submissions to DMILL and HP. It would be an analog test chip, including a small array of the preamp/discriminator design to be used in the final pixel array, as well as the complete analog support for these circuit blocks (current reference, current DACs, bias generation, calibration circuitry). Our schedule would have first versions of all of these blocks available in time to submit a test chip to the Mar 5 2001 TSMC MOSIS run. With the expected 10-week delivery time from MOSIS, there would still be adequate time to characterize the performance of this chip before submitting the engineering run.

Finally, we have delayed the submission date for the FE-I1 engineering run in order to provide adequate verification time for the final design. We have also shifted the planned second engineering run date to allow for final test beam characterization of complete FE-I1 modules in the 2002 testbeam. The original schedule had very little float before the 2001 testbeam window would be missed, and would very likely have forced us to compromise on our verification of the FE-I1 submission in order to get devices in time for 2001 testbeam. We prefer a schedule where the testbeam verification comes at the end of the FE-I1 evaluation period, in order to be able to test fully irradiated modules and to benefit from the accumulated understanding of the chips which we would have by this time.

Rec 5-5 The committee recommends that additional testing be performed on the 0.25um process to insure that the chip will work in the high radiation “B” location especially in regards to single-electron effects (SEE), upset, and latch-up.

We would like to first emphasize that in terms of total dose, there is not a large difference between the B-layer and the outer layers. The outer layers should withstand a total dose of about 50Mrad (including several safety factors), over their design lifetime of ten years. The B-layer was never anticipated to survive ten years, but was expected to be replaced regularly as technology improved, and to have the same total dose lifetime as the outer layers. In addition, the performance of the B-layer, in terms of spatial precision and efficiency, has a large impact on ATLAS physics capabilities, and hence it is expected that there will be very strong interest in

replacing this layer before its lifetime has been reached. A second issue is whether our initial electronics design can cope fully with the high occupancies and rates at the small radius of the B-layer. As we noted at the review, the fallback position is to use exactly the same electronics for the B-layer as for the outer layers in case (a) the radiation resistance is found inadequate or (b) insufficient design time is available to modify the outer layer design to meet the more demanding requirements for the B-layer. The initial (lower luminosity) operation of the B-layer should be adequate with electronics designed for the outer layers. The ability to replace the B-Layer during a short access scenario is built into our design. Thus an upgrade would be the most natural path in case there are difficulties with the B-layer-specific design.

Rec 5-6 System test with detectors, hybrids, and support chips are essential to investigate system behavior and digital/analog cross talk. The committee recommends that a thorough test of system functionality including beam tests should be planned with milestones for doing so and incorporated into the schedule. The beam test and large- scale system test are not necessarily identical.

We have performed typically three beam tests at the CERN SPS per year on pixel prototypes since 1997, and plan to continue doing so into the future. In particular, we will be testing complete irradiated modules in the 1.5T magnetic field of the SPS H8 beam at the latest by June 2002. This work is carried out using the same VME-based readout electronics that we use in our labs, and allows us to carefully cross-correlate testbeam and lab results. The SPS provides special 25ns bunched-beam running periods as well. We have already used the first such run to verify that the performance we have measured in non-bunched beam tests is really identical to that measured with the real beam time structure. Finally, the next major step in the testbeam, originally scheduled to occur in 2001, but now delayed by the problems with DMILL chips, involves readout of multiple modules using the real off-detector electronics (what in ATLAS are called RODs), and the opto-link prototypes. This complete system test would occur in 2002 in our current planning.

Rec 5-7 The committee recommends more planning of the LV power distribution system. Major technical concerns about cabling, noise pickup, over current and transient protection remain unresolved. Their impact on the cost and schedule should be thoroughly analyzed.

We agree completely with this recommendation and note that the fabrication of multiple prototype, full-length, low-voltage cables is underway with the intention of testing modules with prototype power supplies in the next few months. This will provide a natural test bed for transient protection studies. For the original DMILL electronics, we had selected a ceramic surface mount varistor part for transient protection. This appeared to work well in the lab, and sample parts were irradiated in the SPS in May 2000 and were shown to work well after irradiation. However, this technology does not have a steep enough I/V curve to protect devices running at 2.0V (our planned 0.25m operating voltage) as opposed to the 3.5V-4.0V operating voltage planned for our DMILL chips. We have located a new type of transient protection, using an optimized punch-through diode, which looks ideal for protecting 0.25m chips. The particular design used for this diode looks promising in terms of radiation hardness, and so we have

ordered a large number of these devices for evaluation and irradiation. Once they pass basic lab characterization tests, we plan to use the prototype power cables and prototype power supplies to study the types of transient waveforms that can be induced in our power distribution system, and how well they are suppressed by these devices. We expect that this series of measurements will have largely validated our power distribution design within the next 6 months.

Rec 5-8 Radiation qualification of the custom ICs is a major task. The committee recommends that the manpower and milestones must be more carefully planned. It is the opinion of the committee that the costs to accomplish these tasks may have been underestimated.

We completely agree on the scale of this effort. The present activity in this area is largely carried by our European collaborators. During this year, the following irradiations have been performed: (1) total dose and SEU studies of opto-link components at PSI, in collaboration with ATLAS SCT; (2) irradiation of single DMILL transistors, and complete Analog Test chips in April at the PS, (3) irradiation of opto-link electronics in May at the PS, (4) irradiation of MCC chips in Oct at the PS, including real-time evaluation of functionality and SEU rates. These tests were all performed on devices contained on our FE-D1 wafers. The FE-D chips themselves were not irradiated in these periods because there were too many yield problems with the existing devices to warrant such effort. We intend to perform similar irradiations on most of the die from the FE-D2 run. We also have inserted appropriate milestones into our overall schedule for this work for the FE-I1 wafers. We are establishing an irradiation effort based at LBL as well, using 55MeV protons, in order to allow irradiation studies to continue during the 5-6 months per year that the CERN PS is not operating for irradiation work. The detailed studies of the MCC under irradiation have already been performed with their complete lab test system, with additional boards added to support irradiation studies in situ. Similarly, for the FE chips, we have already designed a rad-hard support card to interface single chips or modules to our existing test system, and allow full in situ characterization to be performed during irradiation, in the same way that we do this work in our labs. We intend to perform these irradiations using 55MeV protons (LBL), 24 GeV protons (CERN SPS), and ^{60}Co gammas. In addition, we will try to perform realistic SEU studies at PSI using their 300MeV pion beams.

Rec 5-9 In general, the schedule needs to include more detailed task descriptions to be able to develop a complete schedule and work plan so that the costs can be validated.

We have developed a preliminary detailed schedule for all of the electronics activity over the next 18 months, including the FE-I1 submission, its characterization, and the submission of the revised FE-I2. We have inserted additional milestones into the overall pixel project schedule for the test chips associated with the FE-I1 design effort, and for the evaluation of the FE-I1 wafers. The detailed schedule for the FE-I1 design includes schedules for the individual circuit blocks, both at the schematic and layout level, as well as top-level simulations at the schematic and post-layout levels. The revised review and submission dates in the summary schedule reflect this more detailed planning. An ATLAS Design Review has also been included prior to submission of FE-I1. The submission date for FE-I has been delayed by approximately 6 weeks to provide additional verification time in the design flow.

Rec 5-10 Management should increase the effort in the area of project management. Milestones should be established to allow for thorough IC verification. More manpower will be necessary to accommodate this transition to a production-oriented methodology.

As we noted in the review, we have included funding for systems/project engineering. Peter Denes, who has considerable experience in electronics management and who is now a senior member of the LBNL Engineering Division (as of November 1), has just started working with the ATLAS team in the area of IC and overall system design. In addition, the pixel Project Leader is exploring the possibilities of enhancing the senior engineering presence in Europe.

Our overall approach to design verification includes the following:

- *specifications for each chip*
- *the design is judged to be successful when simulations indicate that the specifications are met*
- *a complete top-level(full-chip) simulation will be performed at the "schematic" level and a design review held*
- *a complete top-level post-layout simulation will be performed and similarly reviewed*

Each time we submit a new version of the pixel chips (we have now made five versions in three different processes), we improve our verification methodology. The present plan for FE-I1 includes the following steps: (1) simulate and verify all individual blocks using SPICE and Verilog with the vendor corner models; (2) simulate and verify at higher levels using Verilog, including annotation for interconnection parasitics; (3) simulation of large SPICE netlists, including all layout parasitics, for large critical blocks like the column-pair; (4) many Verilog simulations at the top level, using random data patterns to look for potential design problems. In addition, we have worked with a newer verification tool (TimeMill) which allows simulation of very large circuit blocks with SPICE-like accuracy. If time and manpower permit, we will try to use this tool during the final verification steps to insure there are no timing errors at the highest level of the chip integration.

In our first DMILL submission (FE-D1), the global verification using Verilog did not include proper back-annotation of parasitics. In addition, the post-layout SPICE simulations of critical blocks were not completed prior to chip submission. The result was a chip which worked correctly, but had several internal buffering problems occuring at the top level in the chip integration. These errors were not fatal, but were not caught prior to submission because of the incomplete verification methodology used. We believe that the revised methodology proposed here does provide an adequate safety net to catch any significant errors in the design.

Rec 5-11 It would be prudent for the collaboration to develop fallback plans in case additional IC iterations are needed, or if planned levels of base manpower do not materialize.

The only realistic fallback plan is the one that we presented at the review: installation of the pixel system in 2006. We have increased the contingency on the design effort in the US to respond to this recommendation. In addition, every effort is being made to utilize design resources outside the US, and some modest additional manpower in Europe has been identified since the review and is being integrated into the effort.

Rec 5-12 There is still no project electronics engineer allocated for the electronics. This position should be filled as soon as possible.

Done

Rec 5-13 In view of the critical nature of the pixel electronics to all ATLAS, the committee found that overall coordination of the electronics effort needs to be strengthened. An effort should be made to improve communication and make collaboration-wide resource allocation decisions, otherwise disconnects could have an adverse schedule impact.

We believe this can be accomplished to some extent by adding systems/project engineering at LBNL, as we proposed at the review and have described above. However, we note that the electronics coordination within the global collaboration has been a rotating position, and that the US does not, and cannot, have final authority in all areas of the pixel electronics system.

Rec 6-1 As a result of the recent change to implement a fully insertable design, the committee recommends that this design be reviewed again next year when it will be more mature. It is the opinion of the committee that the mechanics can not be baselined until the redesign effort is completed in March 2001.

We respectfully disagree. The conceptual design of the pixel support tube(called Thermal Barrier below) has been completed during a few month visit of LBL engineers to CERN and RAL. A baseline concept, cost and schedule has been established.

Rec 6-2 The U.S. Atlas group should take the leading role in the design of the Thermal Barrier, which will be an integral part of the pixel detector installation process.

Done.

Rec 6-3 The amount of heat generated in the support services cables and optical hybrids needs to be understood. Also, CDF experience has shown that the Elco connectors need keeper plates to keep the connectors together for the cable bundles.

We agree entirely with the comments about thermal management and studies of this are part of our prototype program. We take note of the specific comment about connectors.

Rec 6-4 Support structure costs may be reduced by reevaluating the learning achieved from the old design. The costs did not reflect the large amount of learning from the old design.

We have scrubbed the WBS elements 1.1.1.1.1.2.2(Disk Support Rings and Mounts) and 1.1.1.1.1.2.3(Support Frame). The cost of 1.1.1.1.1.2.2 has not been reduced. The reduction in radial envelope(beyond that presented at the review) may require a partial redesign of the support ring to meet space requirements, and the only practical means to assess the viability of such a design is by analysis by Hytec. This may require more analysis than originally planned.

The cost of 1.1.1.1.1.2.3 ascribed to Hytec has been reduced from about 369K to 325K. The test program of the support frame prototype attached to a prototype support cone and endplate has been shifted to LBL. These tests are assumed to be carried out by LBL mechanical technicians(supported by base program). We assume the use of tooling already made by Hytec and the TV holography system now being commissioned at LBL.

In reviewing the estimate for this item, we identified tasks that had not been included. In particular, the modifications(precision holes) that must be put into the support frame. We have determined that this work should be done at LBL and have made an estimate of the tooling required based on preliminary design sketches. We have included these costs, assuming one prototype round for the existing support frame and one for production.

We understand the reviewers concerns that these design costs may appear high after some years of development. But we note that

- *some months have been added to the schedule at the request of the pixel Project Leader.*
- *a substantial design revision has occurred(going to an insertable-layout) with very tight clearances*
- *the design of disk ring mounting has not yet been tested*
- *the design of the frame has not been completed and significant uncertainties remain in some areas(services supports, interface to barrel shells being designed in Europe, overall mounting scheme....)*
- *all design effort on the pixel mounting scheme within the SCT remains to be done, and this is critical for the stability of the system*
- *the frame will be a "build to print" bid overseen by Hytec throughout the bid and construction. All fabrication drawings must be completely redone as a result of the recent layout change and the final dimensions of the frame are still uncertain. All tooling drawings created for the prototype must be redone.*

- *the disk ring fabrication will be done in a similar manner(although the fabrication may be sole-sourced)*

In short, it is the strong opinion of the subsystem leader that considerable engineering effort must be devoted to these critical items, and that LBL engineering cannot be substituted, given the very significant other mechanical commitments in the project. We will, of course, make every effort to minimize the overall expenditure for design, consistent with meeting requirements and schedule.

Rec 6-5 In addition, the support structure tolerance requirements should be reevaluated to determine if they could be relaxed.

We agree but this will naturally follow during the design and early fabrication process. We note that a requirements document exists so far only for the local supports(disk sectors and barrel staves) but will be created for reviews of the support structures. This will provide an opportunity to systematically review all tolerances and reduce fabrication costs if appropriate.

Rec 6-6 It should be noted that the Thermal Barrier's inclusion in the management contingency does not meet the criteria for inclusion in management contingency (a deliverable where the decision can be deferred). There cannot be a pixel system, even at the reduced deliverables, without the thermal barrier. It is the committee's recommendation that the thermal barrier design and production effort be included in the baseline and something else be put into management contingency in its place, or the mechanics be baselined at a later date. The PIXEL Group estimate to complete the design is roughly $0.5M plus design and engineering labor.

Since the pixel support tube(then Thermal Barrier) has become part of the US baseline, this is no longer relevant.

Rec 7-1 The group should carefully consider the limiting factors in production rates and try to maximize their capacity. The possible delays in the Electronics area will probably have to be compensated for in the module/test area to keep the overall program on schedule.

We agree and have(as described at the review) launched an extensive dummy module assembly program to debug all of the mechanical assembly and QA steps possible without actual electronics.

Rec 7-2 Wirebonding is a critical step, and the committee recommends that it either be concentrated at LBL or other sites. If performed at other sites LBL should extensively qualify and monitor these sites. LBL should explore the purchase of an additional wirebonder, either used or new.

We concur with these concerns but believe it premature to include now in the baseline cost estimate another wire bonder at LBNL. We intend to move forward with qualifying university sites(and LBL). We note that UC Santa Cruz is planned to share wire bonding for the SCT and some capacity there may also be available for pixel construction.

Rec 7-3 Multiple probing of die can lead to failures in the wirebonding operation and rework. Care must be taken to insure that the wirebonding site is not over the pad area where the die was probed. Wirebonding to these sites will cause a high impedance connection over time and effect circuit performance.

Understood. We are recommending larger bond pads for the front-end IC.

Rec 8-1 The committee recommends that the hybrid group address cleanliness in its process. An unclean process usually causes many unexplained manufacturing rejects during production.

We agree. Plasma cleaning is being done on the hybrids before wire bonding.

Rec 9-1 The performance and repeatability needs to be documented in more detail.

We are rather earlier in the prototype phase for the optical hybrids. The first hybrid, with both optical components and ICs, will be done in the first half of 2001. This recommendation will be addressed as part of the prototype program that will continue for some time.

Rec 9-2 Once the Optical Hybrid is installed, the electronics is difficult to repair. The committee recommends performing burn-in prior to module assembly to weed out any infant mortality hybrids.

We agree. The design we presented has the feature that the optical hybrid is a separate board that can be tested, burned-in, stress tested(thermally), etc.

Rec 9-3 The committee recommends additional package tests (accelerated aging tests) to determine the effects of –20C operation and radiation on the performance and the long-term failure rate of the hybrid.

The reliability of the opto-packages has been studied using accelerated aging tests AFTER irradiation. The PIN diode has passed the irradiation and accelerated aging test up to to the pixel dosage. The VCSEL has been tested to the SCT dosage and tests up to the pixel dosage are planned. The cap of the opto-pack is fabricated uses ULTEM, which according to CERN's Compilation of Radiation Damage Test Data (CERN 98-01), is rad-hard to 10 Grad. We irradiated the cap with Co^{60} up to 50 Mrad and no degradation was observed. We plan to do

accelerated aging test after irradiation of the completed packages. We also plan to do temperature cycling between -20 and +20 C on the packages and the hybrid.

Rec 9-4 Since a single point failure of the optical hybrid disrupts signals from 6 modules, the committee recommends performing rigorous reliability qualification testing

We agree, of course, with this concern and understand that the design of the optical hybrids must be as fail-safe as possible. This includes both the optical and electrical components, and power distribution to these components. We expect through the extensive prototype program(three rounds included) to be able to arrive at a very robust system. We also note that the design concept for the optical hybrids allows them to be replaced(unplugged and new ones inserted), albeit with some considerable time for access.