WBS 1.1.1.3 Pixel System

On-detector Electronics

Major Topics:

- On-Detector Electronics and Test System (WBS 1.1.1.3): Status and Issues

Very brief technical status and next steps:

- Motivate near-term schedule and goals

Describe cost and schedule status, particularly for FY03

- Discuss FY03 manpower estimates, costs, and milestones

Summary and Conclusions

Electronics Challenges and Requirements

Main challenges are in FE chips:

- Operate properly after total dose of 60 MRad (nominal ATLAS 10 year dose). Cope with reduced signal of 10Ke and sensor leakage current of up to 50nA per pixel at end of lifetime. For the B-layer, this corresponds to a lifetime of about 2-3 years at design luminosity.
- Operate with low noise occupancy (below 10^{-6} hits/pixel/crossing) at thresholds of about 3Ke with good enough timewalk to have an “in-time” threshold of about 4-5Ke (hit appears at output of discriminator within 20ns of expected time). This requires a small threshold dispersion (about 300e) and low noise (about 300e).
- Associate all hits uniquely with a given 25ns beam crossing. Contributions to this timing come from timewalk in the preamp/discriminator, digital timing on FE chip, clock distribution on module, and relative timing of different modules.
- Meet specifications with nominal FE analog current budget of 60mA/chip and digital current budget of 25 mA/chip, and MCC digital current budget of 100mA. This leads to the total module budget of 970mA analog and 600mA digital. The worst case values used for the services design are 80mA/chip analog and 40 mA/chip digital, and an MCC budget of 160mA, leading to the total worst case module budget of 1290mA analog and 800mA digital.

Deep Sub-micron Approach:

- One of dominant effects of irradiation of CMOS devices is creation of trapped charge in the critical gate oxide layers. Below about 10nm oxide thickness, the charge trapping largely vanishes due to quantum tunneling effects. Modern 0.25μ processes are the first to operate fully in this regime (they have 5-6nm oxides).
- The RD-49 collaboration has studied details, confirming that if one controls leakage paths using layout, then a commercial 0.25μ process can be very rad-hard (circuits tested to 30MRad). Many technical concerns addressed, but basically little experience with full-scale devices, so some concerns still remain.
- All experience so far with analog and digital designs suggests that the silicon behaves almost exactly like the SPICE BSIM3 simulations. Nevertheless, given our lack of experience with these processes, we are planning several prototypes.
- CERN has negotiated a frame contract for LHC with IBM for their CMOS6 0.25μ process which extends through Mar 2004. This fixes prices and terms for engineering and production runs, and would provide the basis for our production procurement. Work is underway for an extension of this contract. We can also access the TSMC 0.25μ process in production quantities via the MOSIS consortium as a back-up should problems arise.
- This path places us into commercial mainstream, where we can be assured of low prices and availability in the future. IBM and CERN frame contract moving towards 0.13μ process, providing a technology path for upgrades to the B-layer.

What is included in the FE-I Engineering run ?

- Complete module chipset, including two versions of FE-I1 and new MCC-I1.
- Second generation of opto-chips, including VDC-I2, VDC4-I2, DORIC-I2.
- LVDS Buffer chip for use in support cards
- Test chips, including Analog Test Chip, modified CERN STS (standard test structure) and CapTest chip including capacitance measurement circuit for about 60 different M1/M2/M3 structures.
- Alignment marks for bump-bonding vendors.

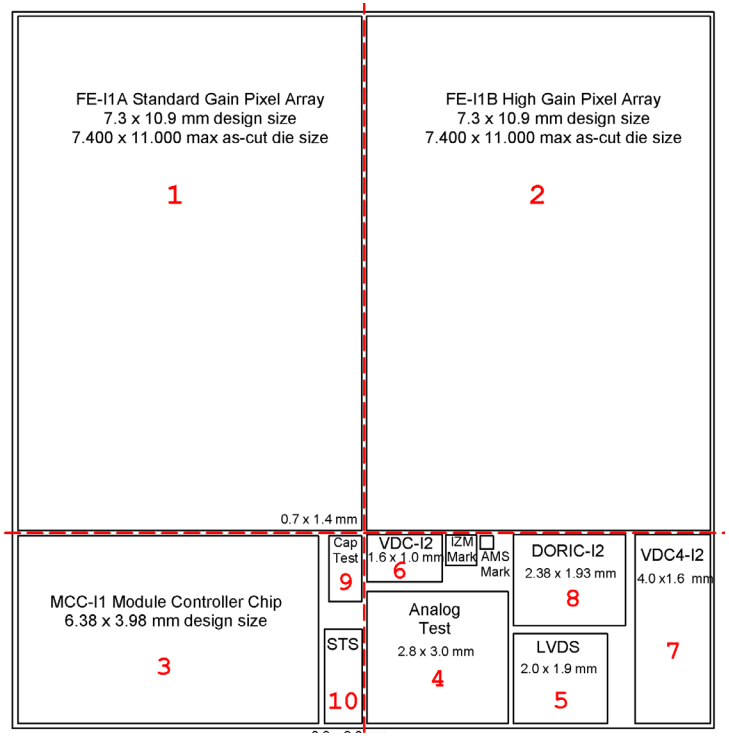

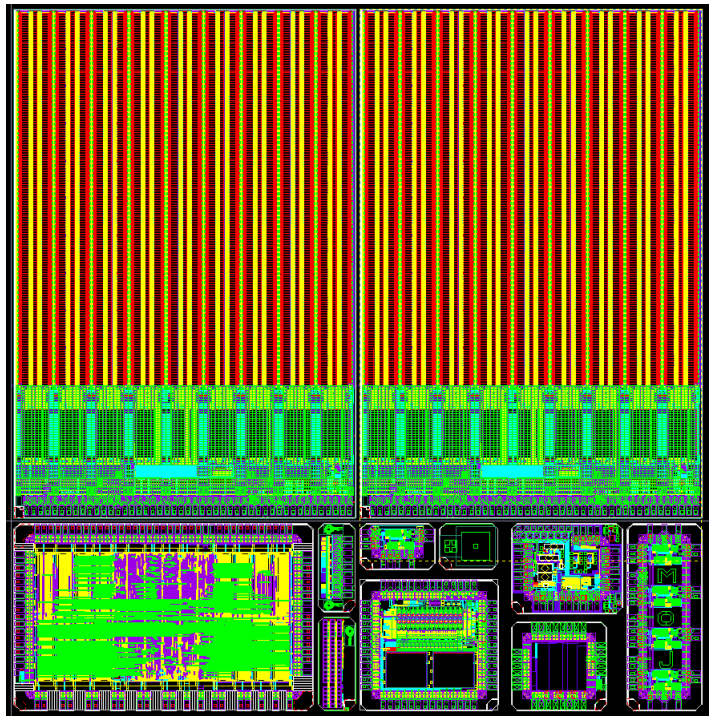

- LBL has taken responsibility for global integration and verification of reticle.

Comments on overall LBL role

- We are the critical player in this area.
- For the FE-I1 chip, the only collaborating institute was Bonn. They had a senior physicist/designer and one EE student working on the chip. They have made significant contributions in the analog blocks, and assembled the analog test chip. However, they have required significant help from our analog designer to complete some of this work properly, and we do not anticipate a large role for them in future submissions. In addition, their senior physicist (P. Fischer) is leaving for a C4 position in Mannheim this Fall. Their student will spend 3 months in LBL working on FE-I2.
- We have developed the extended 0.25μ design kit used for all of the chips in this run. In addition, a major effort has been invested in designing a high quality pad library for both all digital and mixed-mode chips.
- We have taken responsibility for all of the final verification, reticle assembly, and technical contact role for the submission. This is significant extra work and responsibility.

Yield results from FE-I1 run:

- Initial order was for 12 wafers, nominally considered two back-to-back engineering runs (we do not know what IBM actually did).
- Wafer probing gave average yield of about 15% for FE-I chips, with most problems occuring in a small number of circuits, particularly in the Pixel Register (a 2880-bit shift register, segmented into 9 independent pieces). Yield was very position dependent, with characteristic “rings of death”.
- Received 4 additional wafers from this run. They were initially held after front-end processing, then metallization was completed. These wafers had different, and more significant pathologies, and we found a yield of only 3% for the FE chips.
- Finally, IBM agreed to process two additional lots of wafers (25/lot). We have received the wafers from the first of these lots in early June, on an evaluation basis. Both Bonn and LBL have confirmed a high yield of 80% for basic digital tests. A more complete menu of selection criteria, including digital/analog current at reset and during operation, 100% efficiency for digital inject on all pixels, current reference and DAC slope cuts, etc. leads to a yield of 70% averaged over 5 wafers probed so far in LBL.
- CMS APV25 project has suffered a similar history, with more runs. Initial engineering run gave yield of 70-80%, followed by two bad lots (yield as low as 5% on some wafers, with “rings of death”), one good lot, and now two medium lots (30-40% yield).

- Foundry has performed significant analysis and does not understand cause. Claims no other customers (besides CERN) see such widely varying yields.
- For production, assume “average” yield of about 50%. Do not anticipate that we have seen the end of the bad lots. In the frame contract, there is a limited contractual basis for returning wafers based on user-defined target yield (requires a minimum of 3 lots to establish). For example, if we can establish a target yield of 50%, then CERN can return wafers with less than 10% yield, and if we can establish a 70% target yield, CERN can return wafers with less than 28% yield. Not clear how this will be handled by CMS and ATLAS, but informal discussions are starting. CMS APV needs about 500 wafers, ATLAS pixels about 300 wafers, and they are foreseen to be the major initial IBM customers. ATLAS TRT and LArg, plus CMS pixels, and many other smaller projects are also customers at a smaller scale. Initial TRT engineering run for DTMROC had very good yield (more than 80%).

Production assumptions:

- Use assembly yield projections from Maurice, assumption of 50% wafer probe yield, plus projection that production wafers containing only FE chips would provide 300 die/wafer. Maurice calculates that we need a total of 41 wafers for disk production and 171 wafers for barrel production for two-hit system. Numbers would be 62 and 255 for full 3-hit system.

Next Steps towards a production FE-I

- After several months of study in the lab, with irradiated assemblies, and in the testbeam, we start to have a fairly clear picture of the performance of FE-I1.

FE-I1 Testing:

- Using bare single chips, and single chip assemblies, have carried out extensive measurements of electrical performance.
- Have irradiated 8 bump-bonded single chip assemblies to nominal dose (10^{15} n equivalent and 65 MRad) at the PS using 20GeV protons. These assemblies have been evaluated in detail, and the electronics performance is essentially unchanged.
- All basic performance parameters have been measured, including threshold dispersion, noise, threshold tuning performance, TOT measurements, inter-pixel cross-talk, timewalk, and timing uniformity.
- Still some areas requiring further work, including absolute TOT calibration, and more complete measurements of irradiated assemblies.
- Most recent summaries of measurements can be found in EDMS document for CPPM Pixel Week in June. Many new measurements made in July at CERN.
- **Conclusion:** chips are close to meeting ATLAS requirements.

Comments on Irradiated Sensor Performance

Some issues with sensor performance after irradiation:

- Observe some changes in both CIS and Tesla sensors. These include modest increase in noise (at -30C where leakage is less than 2nA/pixel), increase in inter-pixel cross-talk, and degradation of timewalk performance. It is difficult to isolate what is an electronics effect and what is a sensor effect. However, data suggests that inter-pixel capacitance increases somewhat after irradiation to 10^{15}. This could be related to overcompensation of the low dose p-spray which surrounds the n+ pixel implants by the trapped charge in the sensor oxide layer.
- Observe significant increase in noise in most Tesla assemblies.Use of internal leakage monitoring circuitry indicates that there is an excess current (beyond that supplied by the HV bias supply), which is most significant at low bias voltage. This suggests a problem on the n-side of the sensor. Further study indicates that the problem is the bias grid. Normally, this grid is connected to analog ground in order to turn off the punch-through biassing mechanism (the pixel implants are held at roughly +1.0V by the preamplifier biassing). It appears that the punch-through voltage is significantly reduced by irradiation in the Tesla assemblies, causing the bias grid connection to inject current into the preamplifiers, simulating excess leakage current. A better strategy would seem to be to let the bias grid float, and to decouple it to analog ground to prevent injection of high-frequency noise. This needs to be studied in more detail in the testbeam to make sure there are no major problems, such as loss in charge collection efficiency.

Remaining areas of concern for FE-I1:

- Before irradiation, the major area of concern is the threshold dispersion and threshold control. The threshold tuning procedure for FE-I1 is painful due to the non-ideal behavior of the threshold control. First, the threshold in an individual pixel is a non-linear function of the adjust current, second the TDACs are non-linear (and even non-monotonic) due to layout systematics, and third, the performance of different pixels is very strongly row-number dependent due to problems with the bias current distribution.
- Once the threshold dispersion has been tuned for a given chip, it is very sensitive to many things. Significant changes in temperature degrade the tuned dispersion, small increments in radiation dose (even after a large total dose) degrade the tuned dispersion, and finally small changes in the global threshold degrade the tuned dispersion.
- The performance of the “ganged” pixels is sufficiently different, due to their much larger capacitive load, that we will provide an optimized preamplifier design just for them in the future. They require several times larger preamp bias to achieve the nominal performance specifications.
- After irradiation, the threshold dispersion increases significantly, and the threshold control becomes even more difficult. The timewalk appears to be further degraded, although this seems more likely to be a sensor problem than an electronics problem. In addition, the SEU-tolerant latches are observed to have a much higher upset rate than expected.

Major issues for FE-I2:

Threshold dispersion:

- Will begin from FE-I1B front-end design, and possibly further increase the gain by decreasing Cfb from present observed value of about 6fF to about 5fF.
- Analyzing contributions of individual devices to dispersion, and consider increasing size of some transistors to improve matching in those cases where this does not have a significant speed penalty.
- Reworking threshold control in pixel, implementing high-quality “bipolar” trim DAC, and new bias distribution scheme. This should improve predictability and flexibility of threshold adjustment, and reduce row-dependent systematics.

Optimized preamplifier for ganged pixels:

- Large load presented by these pixels results in double the noise, and about 3 times the timewalk for nominal bias conditions. Long pixels however have only slightly worse performance. Therefore, modify preamplifier optimization for 4 pixels per column that are connected to ganged pixels.

Improve SEU management:

- Present design uses SEU-tolerant latch for Global and Pixel configuration information. Measured upset rate at PS was larger than expected.
- Restore Global Register checksum in real-time to detect critical SEU in the chip.

Improve pixel control design:

- Present design has power sequencing problems with VDD/VDDA, and several other minor flaws.

Improve sense amplifier margin:

- At highest column hit transfer rate, needed for B-layer operation, present design only works to about 47MHz. Improve sense amplifier design to increase this margin, to be sure the readout will work after high radiation doses in the B-layer.

Improve chip-wide signal distribution:

- Indications that some yield issues arise from very large, minimally buffered nets used in pixel chip. Improve these by buffering and reducing antenna size.

Miscellaneous:

- Fix timing problems in RCU and PixClk generator
- Fix mirror sizing errors
- Improve self-trigger
- Improve leakage monitoring control
- Improve overvoltage protection circuits
- Improve integration of linear regulator to allow better evaluation of combining VDD/VDDA at FE chip level.

Milestones for FE-I (tentatively revised)

Milestones exist for FE-I1 evaluation:

- Major milestones all completed now, but still analyzing irradiated assembly data, module data, and testbeam data.
- Given the delays in FE-I2 compared to the original schedule (July 02 submission), and the high yield of the recent new lot, we have decided to use the present generation of chips to prepare a large number of “pre-production” modules, approximately 50 with each bumping vendor.

Milestones for FE-I2 (almost exactly FE-I1 + 1 year):

- FE-I2 FDR: 10/8/02, Joint MCC-I2 and FE-I2 FDR to be held at CERN after October pixel week
- Submit FE-I2: 11/25/02
- FE-I2 wafers return: 1/20/03, assuming 8 week turnaround time
- FE-I2 wafer testing complete: 2/17/03, assuming 4 week testing period
- FE-I2 bumped assemblies available: 3/24/03, assuming 4-6 week turnaround
- FE-I2 module irradiation at PS completed: 5/03
- FE-I2 module evaluation in lab and at SPS completed: 7/03
- Complete relevant problem/change list: 7/03

Milestones for FE-I Production (FE-I3):

- PRR requirements are to demonstrate electronics performance in system. Basic characterization of single chips and modules using FE-I2 are clearly required. Consider that system test with at least 1-2 sectors and 1 bi-stave should be required. Consider that irradiation and operation of 8 complete modules in PS and characterization in lab and SPS is also required.
- FE-I PRR: 8/03
- FE-I production submission: 9/03
- This aggressive schedule assumes that essentially no changes are required from FE-I2, so that preparation of the FE-I3 production run would proceed in parallel with most characterization work (unlike with FE-I2, where characterization and next submission are largely sequential efforts).
- This is roughly three months later than the schedule we worked out in the Nov 01 US ATLAS Pixel review. It continues to increase the pressure to have the module assembly proceed with great efficiency and no unpleasant surprises. This is the strongest justification for the fabrication of 50+50 modules with FE-I1 - it will advance the schedule of all module assembly work (as well as providing first real system tests with many modules).

Test System for FE Chips and Modules

- New generation designed. Incorporated experience with present system, and optimized to cover complete range of production needs with one modular set of hardware and software, keeping same basic interfaces to provide flexibility.
- Includes upgrades for greater range of test capability (vary amplitudes and timing), plus optimized buffering and variable frequency testing:

Architecture is directly based upon the original PLL approach, which had proven so ideally tailored to our needs and which represents the model upon which the ATLAS ROD design was developed.

System Architecture

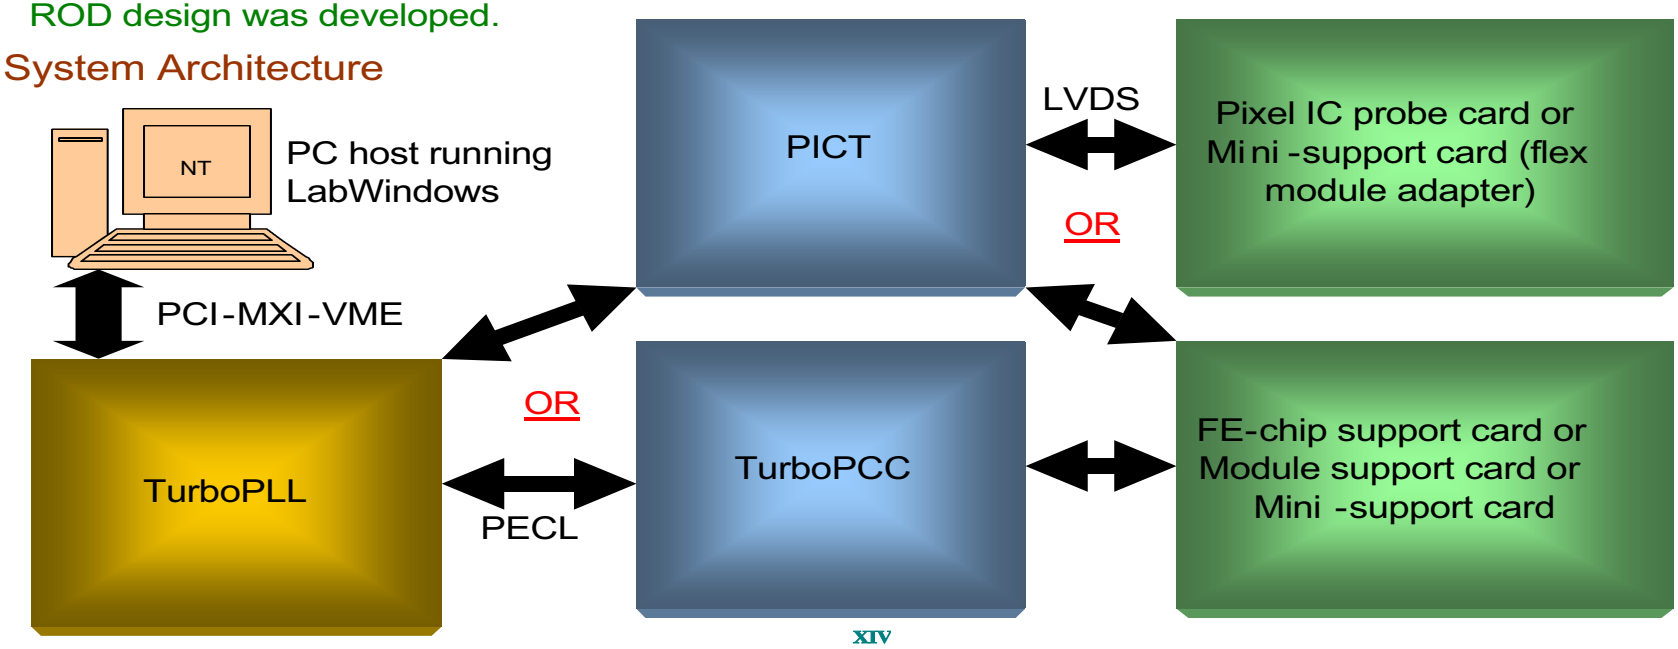

Design Goals:

- New system uses current generation FPGA for much more programming flexibility (old FPGA completely full, making code revisions very difficult).
- New system allows partial (TPCC) or complete (PICT) evaluation of operating margin available in each chip. Optimized cuts can then be used to select die and modules that should continue to work properly after full lifetime radiation doses.
- Cover wider range of needs, including parametric testing at all stages from initial wafer probing, to module testing, and module burn-in.
- System designed to allow operation over wide range of supply voltages, from a minimum of 1.6V up to 4V, to cover testing of 0.8μ and 0.25μ (and below) chips.

Test System Status and Schedule:

TurboPLL:

- TurboPLL design is complete. Two versions fabricated, and extensively tested. Have now produced full 22 final modules. First seven have been tested and used extensively in the collaboration for FE-I testing in the lab, PS, and SPS. Remaining boards loaded, and will be tested at LBL over the next month or so.
- TurboPLL VHDL is almost complete, with a relatively short missing feature and bug list. We hope to clear this list by the end of July to have a first, complete version of VHDL. Significant flexibility remains in FPGA for future improvements.

PICT:

- PICT design completed in late 2001. Eight boards fabricated, three loaded, and one largely debugged and tested. This board has been operated extensively in the lab, and allowed us to get all the problems out of the core circuits used in the TPCC. It has not yet been fully commissioned for wafer probing.
- Expect PICT to be used only by probing groups (Bonn and LBL), with both groups needing two functional setups plus a spare, or a total of 6 boards in the field. Need to complete this work once time pressure for FE-I1 probing is reduced. For now, Bonn and LBL have developed simpler probing systems that allows testing most features, but not evaluating the margins. Expect to commission production probing on remaining FE-I1 wafers, and be ready to go for FE-I2.

TurboPCC:

- Design was completed in March, and first prototypes were fabricated for use during the PS irradiation. They worked well, but an additional problem was found in multiplexed mode.
- The board was revised, and production parts orders were completed in June. We have loaded three of the production TPCC boards, and they have been extensively used in Geneva during the recent testbeam period.
- The parts are waiting at the loaders for the final go ahead to load the remaining 19 boards. We have issued that approvall yesterday, after resolving the last two issues with the new board. loaded boards will return in two weeks, and then debugging starts. Expect to distribute cards during August.

Evolution of testing systems

- Present testing has emphasized single chips or single modules. This is what the PLL/PCC design was optimized for. Simple tests can now be done with up to 4 modules using TPLL and TPCC, but the return data is only available from one module at a time. This will take us through a “3-module” test of a half-sector or partial stave.
- Next major step is basic system test with a sector or a half-stave, using ROD. Given present state of pixel opto-links, propose to proceed with copper-based BOC. This is called the SimpleBOC, and uses LVDS links over twisted pair cable for two groups of 7 modules. It would operate comfortably over 10 meters. The module end would be plug-compatible with the opto-daughter card interface, and would plug into PP0. This copper link could be transparently replaced by real opto-links as the necessary hardware/software becomes available.
- This scale of test system will evolve into the object used during macro-assembly (module mounting onto sectors, and assembly of sectors into disks). Expect that ROD-based testing could be quite important during assembly. Alternative approach would be to use PLL-based system to exercise individual modules.
- Intend to build a 10% scale system for commissioning purposes beginning in about 2004 in the SR building at CERN. Not clear when the center of gravity of system testing would move from macro-assembly institutes to CERN. The CERN system might not be set up significantly before delivery of barrels and disks to CERN.

Production Testing Plans

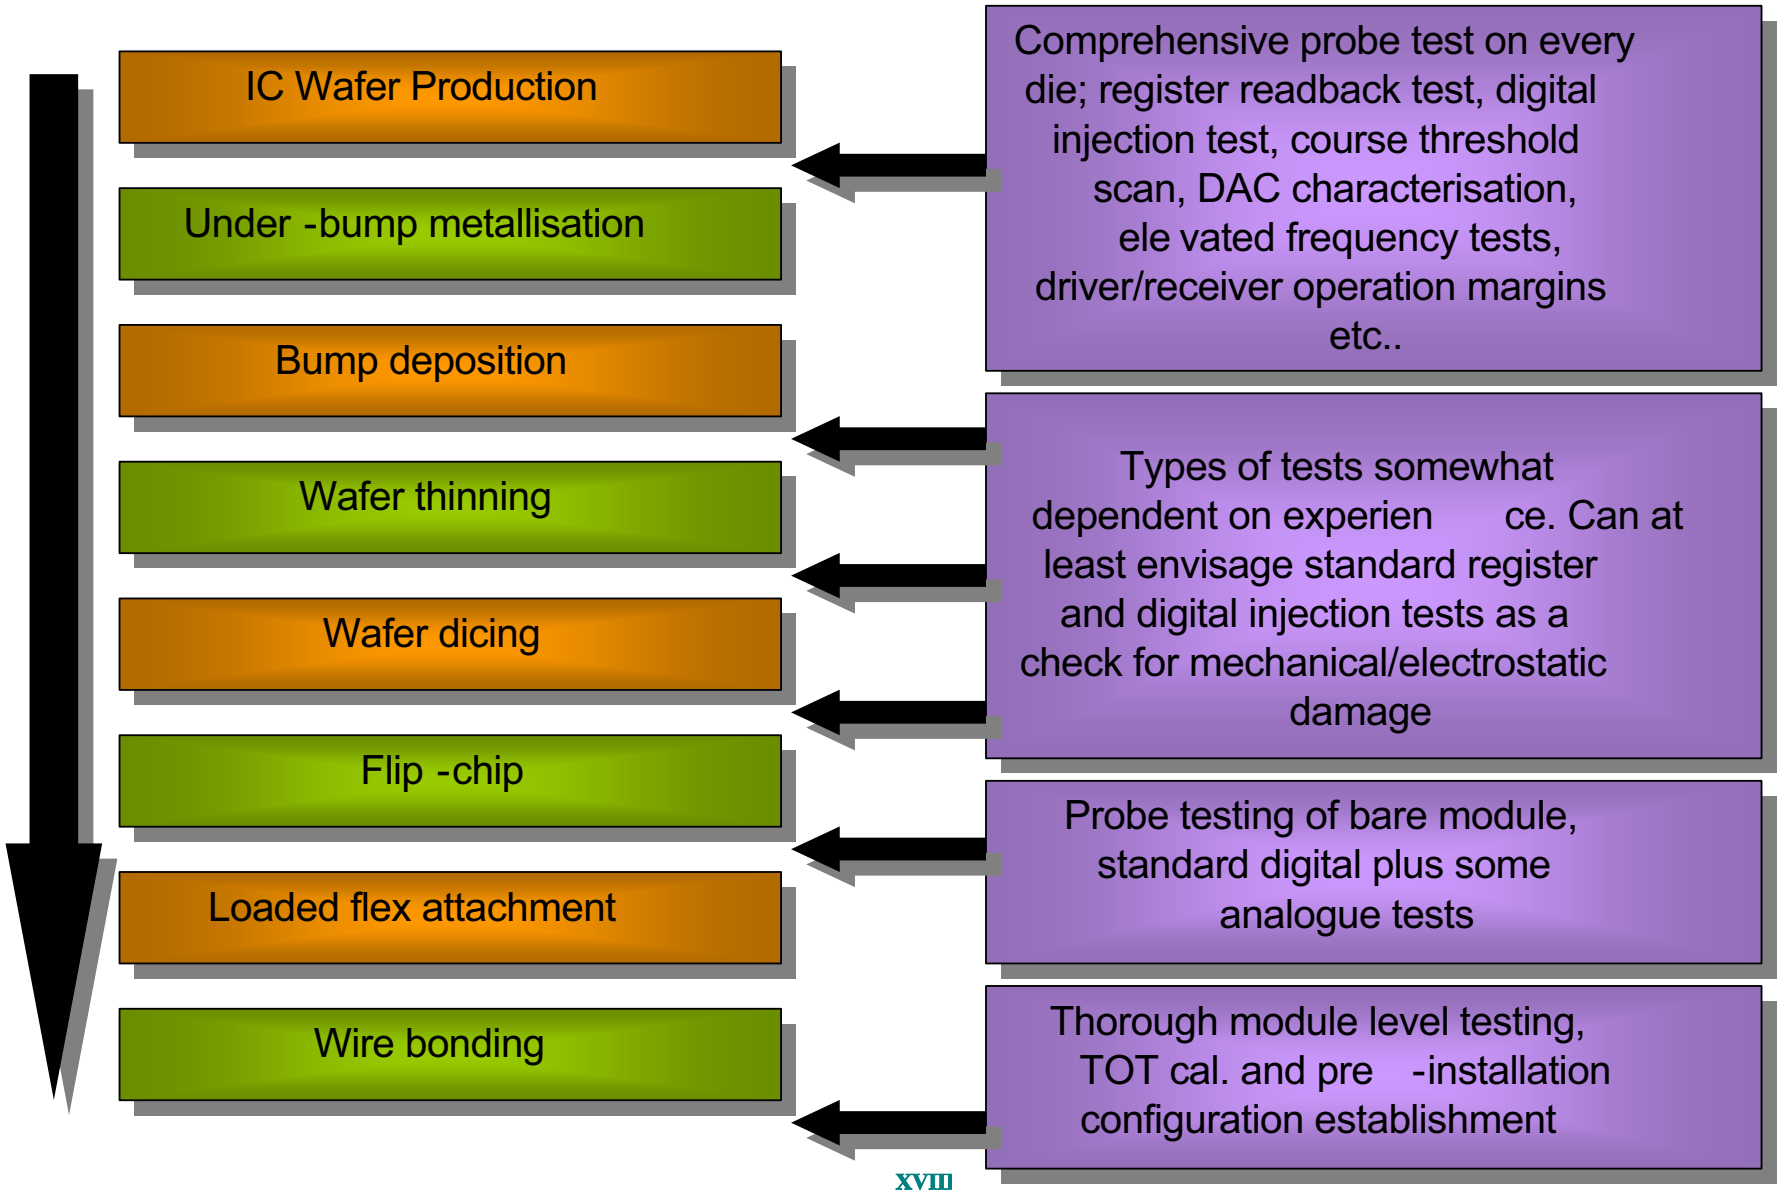

- New system addresses range of needs from wafer probing to module burn-in.

On-detector Electronics Deliverables:

US Responsibilities include:

- FE chip design, testing and production (LBL): Contribute roughly 20% towards the common procurement of the series production. Test roughly 50% of FE ICs at wafer level.
- Opto-link chip design, testing and production (OSU): Contribute approximately 50% towards the common procurement of the series production.
- Design and provide hardware/software for lab/testbeam single chip and module testing, production FE wafer probing, production module testing/burn-in (LBL).

LBL engineering personnel estimate:

Person	Months in FY02	Months in FY03	Months in FY04
Laurent B	3	12	2
Jean-Marie B	1.5	3	2
Peter D	0.5	1	0
Emanuele M	1	3	0
Gerrit M	1	4	0
Ivan P	2	1	0
John R	3	6	6
Bryan H	1	1	0
John J	2	1	1
Chinh V	1	2	1
Engineer Total	16	34	12
George Z	0.5	2	0
Helen C	0.5	1	0
Technician Total	1	3	0

- IC tasks include FE-I2, and FE-I3 (production chip). Continuing manpower is concentrated on front-end issues, as other parts are essentially final.
- Board tasks include completion of TPLL/PICT/TPCC plus SimpleBOC in FY02, and completion of production burn-in system in FY03.
- Assume that all chip testing and most board testing is performed by physicists, with use of engineering manpower only has needed for technical problems.

Electronics Cost Estimate:

Electronics Design manpower (1.1.1.3.1):

- Engineering personnel estimate above leads to about 290K cost in FY03, compared to the ETC02 estimate of 205K. With our present (perhaps optimistic) estimate for FY02 rollover of 60K, we should increase the FY03 budget by 25K.
- For FY04, the current estimate would be 116K, compared to the ETC02 estimate of 11K. This is an increase of 105K, assuming no rollover from FY03.
- Major cause for the overall increase relative to ETC02 is continued reduction of base program support for ATLAS pixels at LBL, forcing us to support additional engineering personnel on the project budget. Relative size of FY03 and FY04 increases are related to schedule shifts compared to ETC02.

Electronics Prototypes (1.1.1.3.2):

- Plan for FE-I2 run is a standard minimal production run of 48 wafers. US share would be 21.5% of frame contract price of roughly 245K, or 53K. Already allocated 64K for this purpose in FY02, and transferred to OSU.
- Additional costs for test equipment (1.1.1.3.2.4) cover wafer probing, test cards, and miscellaneous hardware and software. Most of this work is completed in FY02. ETC02 had no funds for FY03, but we estimate 5K is required here for FE-I2 prototype support. Module-related test cards are contained in different WBS.

Electronics Production (1.1.1.3.3):

- Present estimate is based on 212 wafers calculated to be required for production, with the US paying MOU fraction of 21.5% of the total (very close to the disk fraction of the electronics). There is a major price break point in the frame contract at 250 wafers (10% reduction), and we will need some spares, so base the estimate on a 250 wafer order. This is a total of 650K, so the US share would be 140K.
- Additional costs for wafer storage, probe cards, probe station maintenance, and miscellaneous hardware and software are estimated in ETC02 to be 11K in FY03 and 6K in FY04. See no reason to adjust this at the present time.