The ATLAS Pixel Detector

M. Garcia-Sciveres
Lawrence Berkeley National Lab.
1 Cyclotron Rd.
Berkeley, CA 94720, USA

1 Introduction

ATLAS is a general purpose high energy particle detector presently under construction for the Large Hadron Collider (LHC) at CERN [1]. Charged particle tracking is performed with three concentric detectors inside a magnetic field sustained by a superconducting solenoid. The innermost detector, immediately outside the LHC beam pipe, is the pixel detector. No other detector technology presently available is capable of effective operation with the projected particle densities, radiation dose, and interaction rate that close to the LHC collisions. At the same time high rate pixel detectors for collider experiments are a new technology, and the ATLAS pixel detector will be one of a handful of first generation high rate pixel detectors ever constructed.

The basic unit of the pixel detector is the module. A module is a rectangular active device approximately 6cm by 2cm with 46,080 pixels, each 50μm in azimuth by 400μm along the beam. All modules are identical. Modules are arranged in 3 concentric cylinders with the axis along the beam (the barrel) plus 3 disks concentric with the beam at each end of the barrel (Figure 1). The barrel modules are shingled along the beam direction and held in a turbine arrangement in azimuth (Figure 2), while the disk modules are staggered on opposite faces of each disk. There are 1456 barrel modules and 288 disk modules covering pseudorapidity $-2.5 < \eta < 2.5$, in the radial range between 5cm and 12cm.

The beam collision rate at the LHC will be 40MHz with multiple interactions per crossing, and the detector must be able to resolve data from each crossing. In order to operate at such high rate every pixel must be read out by an independent electronics channel. Thus there are about 67 million channels in the barrel and 13 million in the disks. The readout electronics are implemented with custom integrated circuits (IC) fabricated in 0.25μm bulk CMOS. Transmission of data off the detector is zero-suppressed with a maximum rate of 160 Gb/sec, which at 1% occupancy corresponds to a trigger rate of 7.5 KHz with no deadtime. Data transmission uses the LVDS standard inside the active volume and custom optical links outside. Because of the

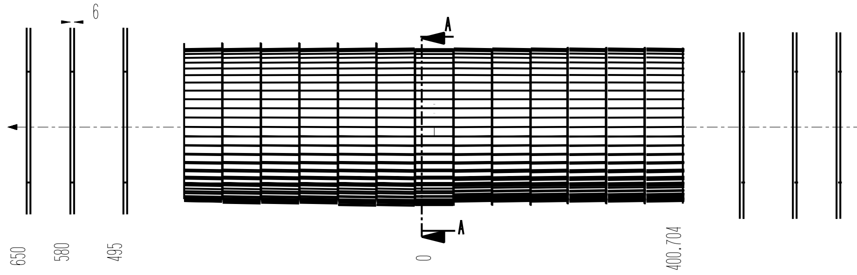

Figure 1: Geometry of the ATLAS pixel detector showing the barrel and 3+3 disks. Dimensions are in mm.

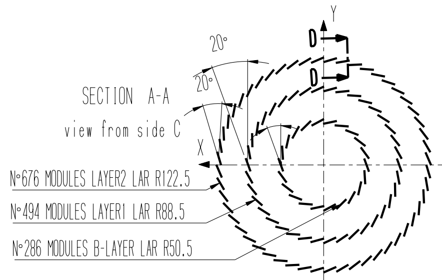

Figure 2: Turbine arrangement of barrel layers. Dimensions are in mm.

need to instrument every pixel the total area covered by electronics is larger than the area covered by active sensor elements, making for a very large power consumption to active area ratio. This places enormous requirements on the mechanical design and support services. The operating power in the active volume is 6.5 KW at 2 V, which must be supplied by a cable plant with minimal radiation length and in turn removed by the (also as massless as possible) evaporative C_3F_8 cooling system in order to operate the detector at less than 0^oC. In the end only so much performance per Kg is possible for the services and mechanical structure, and the complete pixel detector is about 10% of a radiation length(RL)normal to the beam (significantly more at shallow angles), with only 1.2% RL being active.

The projected radiation dose at 5 cm radial distance from the LHC beam crossings is 10^{15}cm^{-1} 1 MeV neutron-equivalent particles in the initial phase of operation of a few years. The dominant contribution is from minimum ionizing charged pions. The ionizing damage is approximately 50 Mrad silicon equivalent. Even at low operating temperature such dose will give rise large leakage currents in silicon sensors which can only be managed by connecting a small volume of silicon to each readout channel. Thus strip detectors are not a viable option. Similarly the large particle density per

event at LHC energies would result in too high occupancy for pattern recognition with a strip detector. This is not a coincidence as the radiation dose is mainly from physics particles. Thus at larger radius, outside the pixel detector, both radiation dose and track densities are adequate for a strip detector (SCT [2]).

2 Sensors

Aside from increased leakage current, radiation damage will first type invert the sensor bulk and then gradually increase the depletion voltage. The pixel sensors have n bulk and n+ implants on the readout side, with the p-n junction on the back side. This means that for unirradiated sensors the depletion begins from the back side and the pixels are not isolated from each other until full depletion. However, after type inversion the junction moves to the front side making the pixels isolated and operational even if the bulk cannot be fully depleted. Of course, maximum achievable depletion is still desirable to maximize the signal. Two key features make near-full depletion possible after the nominal full dose. A multiple guard ring structure on the back side of the sensors is capable of withstanding bias in excess of 600V (the design maximum operating voltage) without breakdown, and an oxygenation process step is applied during fabrication which curbs the increase in depletion voltage due to charged particle bulk damage. On the sensor front side each pixel is DC coupled to the readout electronics. There is also a bias grid structure using a punch-through connection to each pixel that permits biasing the sensor without any electronics present.

A detail of the multiple guard ring is shown in Figure 3. While this structure has the desired effect of allowing high voltage operation, it does necessitate double sided processing, which adds cost and lowers fabrication yield. Sensors are fabricated on 4 inch wafers (3 per wafer) with a nominal thickness of 250μm (Figure 4).

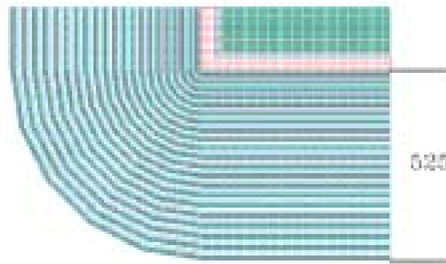

Figure 3: Detail of multiple guard ring structure on back side of sensor. Dimension in μm.

A simulation of the depletion voltage evolution with and without oxygenation for the inner barrel layer is shown in Figure 5. The simulation assumes the sensors are

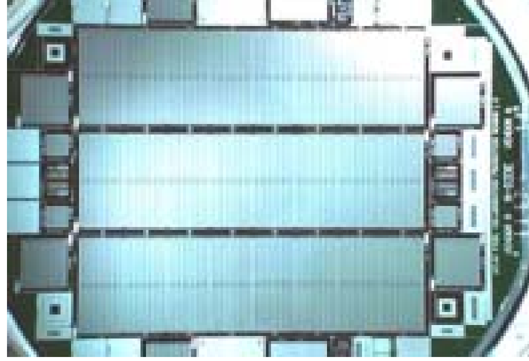

Figure 4: Photo of front side of pre-production wafer.

kept cold even during beam off, except for brief periods of warm up according to ATLAS projections. The charges particle fluence is taken to be 50% higher than nominal ATLAS projections. Note that the voltage for a 200μm depletion depth is shown, rather than full depletion.

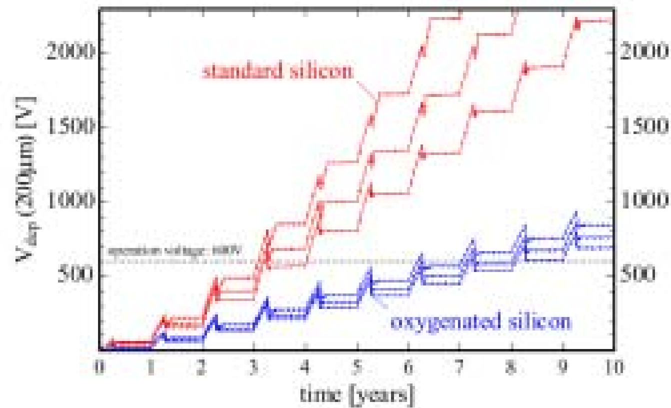

Figure 5: Simulated evolution of depletion voltage with and without oxygenation for inner barrel layer.

3 Electronics and Performance

The on-detector electronics are implemented in a suite of 4 IC's designed in a 0.25μm feature size bulk CMOS. These are the Front End chip (FE), the Module Control Chip (MCC), the DORIC and the VDC. The DORIC, connected to a PIN diode,

performs optical to LVDS control data and clock decoding and conversion, while the VDC, connected to a laser diode, performs LVDS to optical output data conversion.

The FE chip contains 2880 individual pixel channels, each with continuous reset charge-sensitive amplifier with leakage current subtraction, signal shaping, programmable threshold discriminator, and time over threshold (TOT) output. For every pixel that exceeds its programmed threshold (up to a maximum of 576 at any given time) a time-stamp and TOT are stored in local memory for later output or deletion based on time-stamp lookup. The hit efficiency is thus characterized by the fraction of in-time hits, with the time-stamp interval of 25 ns. Figure 6 shows the efficiency before and after full lifetime radiation dose as a function of trigger signal timing for detecting hits from minimum ionizing particles measured in prototype assemblies in a beam test. Unirradiated assemblies are over 99% efficient, whereas after irradiation the efficiency drops to about 97%.

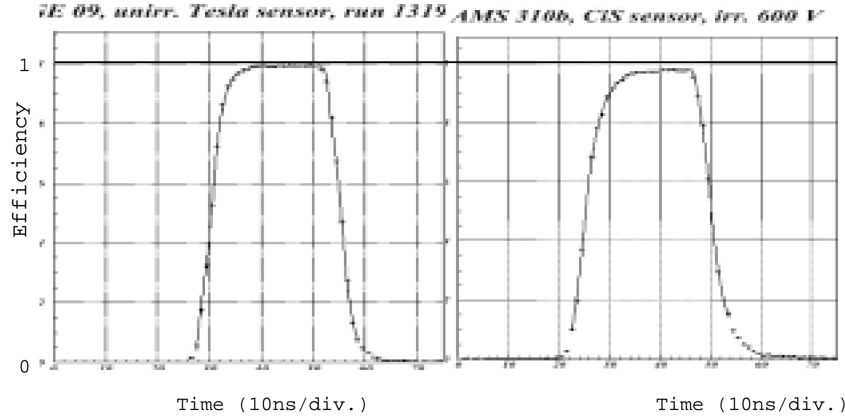

Figure 6: Hit efficiency vs. trigger timing from beam tests of un-irradiated (left) and irradiated (right) pixel assemblies.

The MCC chip manages the communication between the 16 FE chips of a single module and the upstream data acquisition system. The FE chips receive commands from the MCC on a parallel bus and are distinguished using a 4 bit address. Data output from each FE chip reaches the MCC on a dedicated serial LVDS connection. Data from all FE chips are staged in the MCC where module level event building and error handling take place. The MCC then sends data off the module over one or two serial LVDS links (depending on desired bandwidth), with each link being capable of 40 or 80 Mb/sec. The inputs to MCC (and hence the module) are a 40 MHz clock and a serial command line. The communication between the various chips is shown schematically in Figure 7

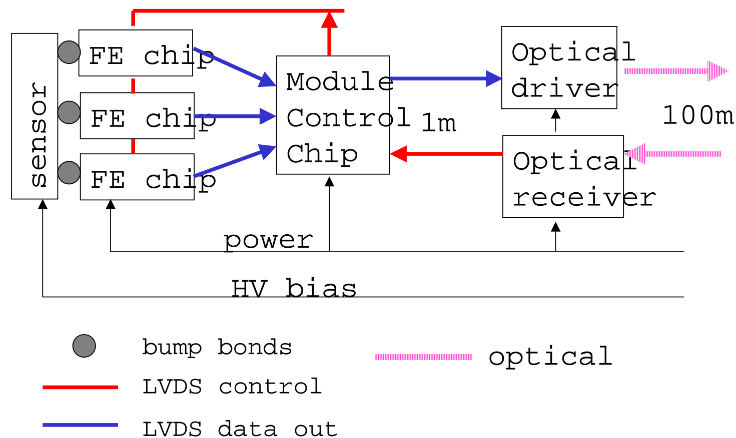

Figure 7: Diagram of connections between the various pixel integrated circuits.

All IC's have been designed in a commercial 0.25μm bulk CMOS process using a combination of full custom analog blocks and a range of design techniques (from auto-routed Very High Description Level, VHDL, to manually routed schematic level) starting from a custom digital library with enclosed geometry NMOS transistors for radiation tolerance. All chips have been tested to ionizing doses in excess of the 50 Mrad lifetime dose with little degradation. Another radiation effect that must be considered is the rate of Single Event Upsets (SEU), as there are of order 10^6 registers in every module. Studies indicate that some care must be taken to make circuits built in this technology robust to SEU at the level required for ATLAS operation, and there is ongoing work in this area.

4 The Pixel Module

The flip chip assembly of one sensor tile and 16 FE chips is referred to as the bare module. This is then combined with a copper on polyimide flex hybrid circuit holding one MCC chip to form the module proper. An flex cable ("pigtail") attached to the flex hybrid brings signals, power and sensor bias voltage in and out of the module. The module anatomy is illustrated in Figure 7. On top the flex hybrid covers almost everything, with the wire-bond pads of the FE chips protruding along the edges so that aluminum wire-bonds can connect the chips to the hybrid. The FE chip inputs are connected to the sensor pixels by means of bumps bonds, which are not visible. The flex hybrid is attached to the back of the sensor, while the back of the FE chips will be attached to and cooled by a carbon support structure (not part of the module). The sensor pixels are DC coupled to the FE inputs, which therefore provide the sensor

ground reference, and the back of the sensor is raised to the full bias potential. The isolation of this potential (which can be up to -600V) from the flex hybrid circuitry is provided by the sensor passivation plus a flexible cover layer on the back of the flex hybrid. In turn, the flex hybrid supplies the bias to the sensor via a wire-bond through an opening in both of these barriers.

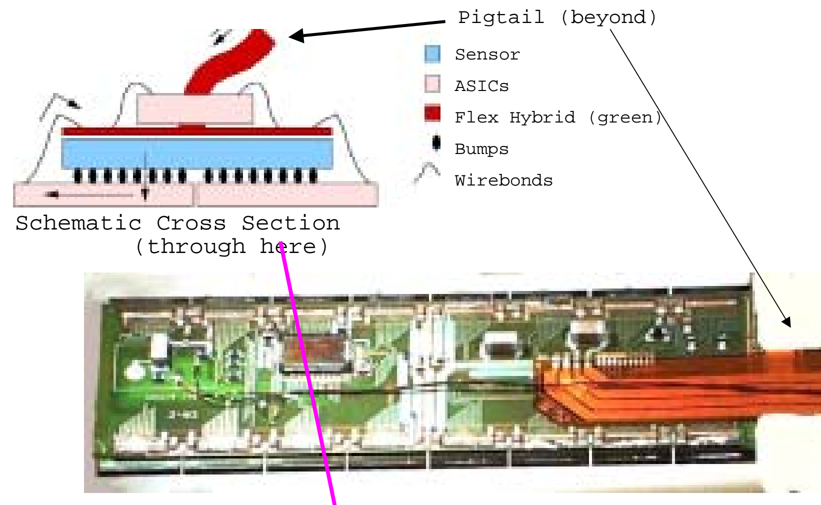

Figure 8: Photo and schematic cross section of a prototype pixel module.

Two technologies are used for the bump bonding of the FE chips to the sensor: indium and solder. Roughly half of the modules will be manufactured with each. This was done to provide technical redundancy and it was also is necessary to have two vendors to supply the required quantities. Figure 9 shows a detail of both types of bumps as well as a bonded assembly (bare module). The minimum bump spacing is 50μm. In order to reduce material and module thickness (which allows for a tight geometry in the barrel), the FE chip wafers are thinned to 200μm after the bumps are deposited, before the flip chip step. No underfill material is used between the bumps to minimize the capacitive coupling between pixels as well as the capacitive load on the FE inputs. Consequently the bumped assembly is mechanically held together only by the bumps which makes it vulnerable to stresses from thermal expansion mismatches. Glue interfaces between the flex and the bare module and between the module and the support structure must therefore be compliant. Low shear strength glues and special glue deposition patterns are used to achieve this.

In addition to meeting single module performance requirements, particular attention was paid to designing modules that can be mass produced with limited effort and consistently by different collaborating institutions. Except for bump technology and

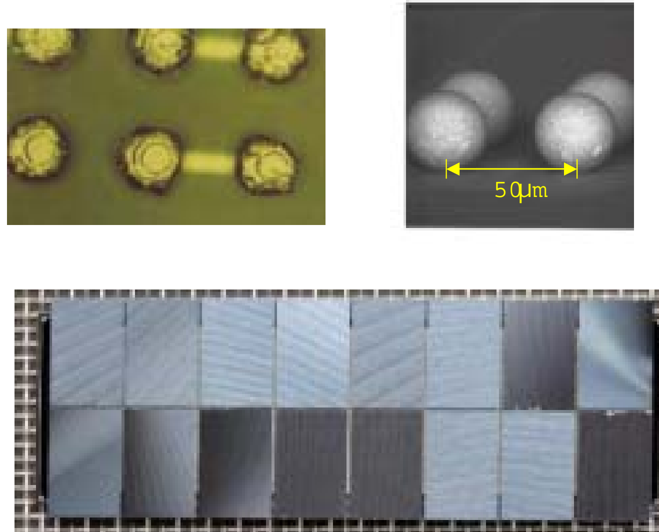

Figure 9: The bare module assembly seen from the back side of the FE chips (bottom) and the two bump technologies used, indium (left) and solder(right).

sensor manufacturer (neither of which affect assembly operations or physical characteristics) all modules in the detector are identical up to the pigtail, which varies between barrel and disks. A handling and assembly model was developed around a sacrificial printed circuit (PCB) frame, which is attached to the flex hybrid immediately after manufacture, fits various tools during assembly of the module components, and provides test connections at different stages up to the fully assembled module. A photo of the PCB frame is shown in Figure 10. The fully assembled module is only removed from its frame just prior to mounting on a support structure.

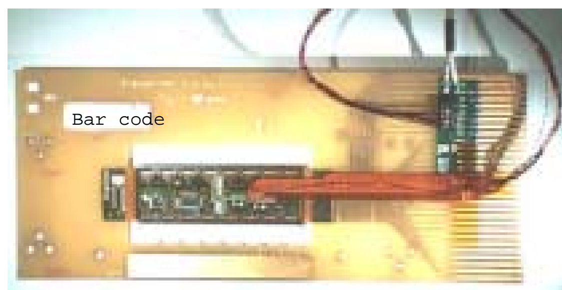

Figure 10: A fully assembled module with disk pigtail and service cable in its PCB frame

5 Mechanical Support and Services

The task of the support and services is to hold the pixel modules in the prescribed relative positions to an accuracy of order 10μm, even as they are cooled from room to sub-zero Celsius operating temperature, to supply electrically the large power required at controlled voltage, and at the same time remove it as heat. And all this must be done with as close to zero mass as possible using power supplies and cooling plant 100 m away. The detector must also be kept cold and dry even during idle times.

The mechanical structure is built out of carbon composites. The design relies heavily on computer modeling and simulation to achieve the desired global properties by controlling the composition of each section. The structure is hierarchical. The modules are first integrated on identical "local support" structures that are replicated on intermediate structures and finally held together by "global supports". The local support for the barrel (disk) section is the stave (sector), which holds 13 (6) modules. Staves are grouped into bi-staves, which share a cooling circuit. Photos of prototype local supports are shown in Figure 11. Aluminum cooling tubes are integrated into the local supports. The cooling system is based on C_3F_8 evaporation which is achieved by flowing the liquid through a capillary into the local support tubes.

Figure 11: Photo of a prototype bi-stave and sector.

A finite element simulation of the global support frame has been validated by building a life size section and measuring deflections and vibration modes. This allows the performance of the full frame to be known with high confidence before construction. Figure 12 shows a photo of a prototype section next to a solid computer model of the entire frame. The frame has high stiffness (first mode over 500Hz) with low mass and near zero thermal expansion coefficient.

Because the sensors must be kept cold at all times in order to achieve their projected useful life, the pixel detector must be environmentally isolated from the rest of ATLAS. This is accomplished by a support tube that houses the detector, the beam

Figure 12: Photo of a life size prototype frame disk section next to computer solid model of entire frame.

pipe, and the first 3m of services on either side, parallel to the beam. This support tube also allows installation of the pixel detector plus beam pipe and services (pre-integrated into a package) into the fully assembled ATLAS detector. The tube provides electrical shielding, active thermal isolation and mechanical support for the detector and services.

Electrical power connections inside the support tube use a combination of pure aluminum and copper clad aluminum wires to minimize material. Even so, the contribution to the total radiation length is significant because the operating current is large and the voltage drops must be small (the IC's are designed for 2V operation with only 1.2-2V overhead for maximum safe voltage they can tolerate). Voltage regulators are placed 10 m away from the detector. This is the closest that the radiation dose projections allow. Even with low cable losses, the power dissipated in the services beyond the detector is significant and requires dedicated cooling circuits. Each module has dedicated power lines inside the support tube.

Control and data are transmitted optically to a patch panel at the end of each disk section. At this point control signals are converted to LVDS and data from the detector is converted to optical. All conversion is done with fully custom components. Each module has a dedicated optical input and output.

6 Conclusion and Outlook

The ATLAS pixel detector will be among a few first generation high rate pixel detectors at particle colliders. It has been designed to operate in the high track density and radiation environment near an LHC collision point with a nominal useful life of 10^{15}cm^{-1} 1 MeV neutron-equivalent particle (mostly charged) fluence. Construction is already (October 2002) underway for the sensors and mechanical structures. The detector assembly is projected to be completed in 2005.

This paper was presented on behalf of the ATLAS Pixel Collaboration

References

[1] ATLAS Pixel Collaboration, ATLAS Pixel Detector Technical Design Report, CERN/LHCC/**98-13** (1998).

[2] ATLAS Collaboration, ATLAS TDR-5, CERN/LHCC/**97-17** (1997).