

Performance Measurements of the ATLAS Pixel Front-End

John Richardson
Lawrence Berkeley National Laboratory
1, Cyclotron Road
Berkeley, CA 94596
USA
On behalf of the ATLAS Pixel Collaboration.

1 Introduction

The ATLAS [1] Pixel [2] Collaboration recently developed a full-scale, front-end read-out ASIC (FE-I1 [3]) in a $0.25\mu\text{m}$ technology. This technology, whilst inherently non-radiation hard, may be rendered radiation tolerant upon application of a special set of design rules developed at CERN [4]. FE-I1 chips from the initial batch of wafers delivered by the foundry early this year were used to develop many prototype pixel MCMs (multi-chip-modules), incorporating 16 FE chips along with a single Module Controller Chip (MCC[6]) and single-chip assemblies. I report here on the laboratory-based evaluation of these assemblies, with a specific emphasis on the analogue performance. FE-I1 was fabricated in two flavours, FE-I1A having a nominal 10fF feedback capacitance and FE-I1B with 5fF. The next generation front-end (FE-I2) will have a feedback capacitance close to, or the same as FE-I1B (which was measured to be $\approx 6\text{fF}$). Therefore, the measurements described here are confined to FE-I1B devices.

2 The TurboPLL Test System

The system which was utilised in testing FE-I1 and FE-I1 assemblies was developed at the Lawrence Berkeley National Laboratory during 2001 and early 2002. This ‘*TurboPLL*’ system will be used throughout the ATLAS Pixel Collaboration for the evaluation of the single-FE chip to single-module scale performance. It will also be used for quality checking procedures during the module production phase. Such procedures include bare wafer probing, probing of diced-and-bumped FE chips (prior to flip-chip), bare module probing (i.e. testing before hybrid attachment), detailed analogical evaluation of assembled modules and testing of devices at test-beam and

irradiation-beam sites. A similar test system was developed in early 1998 for a previous generation of Pixel ASICs. Experience with this earlier system proved that the ‘fully integrated’ approach in which identical hardware and software is used at all evaluation stages is highly beneficial.

Figure 2 depicts the architecture of the current test system, the basis of which is a 6U VME-based board known as the TurboPLL. At the next level down there is optionally a TurboPCC (Pixel Control Card) or PICT (Pixel I.C. Tester). This piece of the system interfaces directly to either bare pixel FE chips or fully instrumented MCMs. The purpose of the TurboPLL is to generate all of the necessary waveforms




Figure 1: *The Test System Architecture (left) and a Photograph of the TurboPLL.*

for the configuration of downstream entities such as the MCC, Pixel FE chip or the TurboPCC, each of which utilises its own protocol. The TurboPLL is also responsible for issuing calibration hit-strobes and level-1 triggers along with upstream data reception. Trigger generation may be fully autonomous (following strobe issuance with a precise, programmable latency), or prompted by an external source. Event data transmitted by Pixel FEs or the MCC (which differ in format), is decoded by the TurboPLL and optionally sent as a 32-bit parallel stream to the host PC (via a 512K-deep ‘data FIFO’), otherwise being histogrammed using an available 16MBytes of SRAM. The intelligence to perform these tasks is programmed into a large FPGA which responds to 32-bit command words from the host PC fed via a 64K-deep ‘control FIFO’. At the front-end of the TurboPLL reside two 512K-deep FIFOs for signal transmission and reception to and from pixel devices under test. These FIFOs serve to divorce the device operation clock from the FPGA clock which always operates at 40MHz. In this way the operation frequency of pixel FEs or modules may be varied from $\approx 15\text{MHz}$ to 115MHz. This multi-frequency range capability is useful in ascertaining the marginality of devices in order to gain a handle on how robust their digital performance would likely be in the experiment.

The TurboPCC piece of the system incorporates a charge injection scheme in which a chopper circuit is used to create a precisely timed voltage step in conjunction with two 16-bit DACs. This voltage step is transmitted to one or both of a pair of calibration capacitors that are connected to each pixel preamplifier in the FE array. The FE-I1 design includes its own internal charge injection machinery, (which must exclusively be used in the experiment for threshold tuning and calibration). The more ideal external charge injection facility of the TurboPCC provides the means to evaluate and calibrate the internal circuitry. The timing (relative to the Level-1 trigger) of either this injected pulse or of the digital strobe, which is sent to the FE chip to effect internal voltage chopping, may be adjusted on the TurboPCC with a granularity of 0.66ns over a range of ≈ 170 ns. This enables the time response of pixel front-ends to be accurately determined. There are four pixel module ports on the TurboPCC which allow for the possibility to provide power, configuration data, operation clock and hit strobes to three devices whilst reading out and thus testing a fourth. This has the benefit that a single system may be used to facilitate the irradiation of up to four devices and conduct ‘mini-system tests’ on support structures instrumented with a small number of modules. The ‘PICT’ is a more complex version of the TurboPCC which was developed to perform complete parametric testing of digital FE-chip integrity during wafer probing. A host of DACs and delay chips enable the amplitude and relative delays of all of the control signals to be varied. Meanwhile, window comparitors are employed to investigate the margins of upstream data returning from FE chips under test.

3 Digital Performance of FE-I1

The FE-I1 readout logic includes the provision for artificial hit creation in each pixel cell upon application of a simple digital strobe to the back-end of the discriminator. In this way the entire readout circuitry may be tested without ever having to enable the analogue front-ends. Having verified full functionality at the basic control-register level, the first test to be performed usually involves ‘injecting’ digital hits into each pixel and verifying that the return serial data stream is composed of all of the expected data. In each pair of columns, the maximum capacity for hit registration is 64, corresponding to the number of buffer locations at the end of the column pair. In order to examine each of these buffers, the digital test involves simultaneously creating a hit in every 5th pixel according to the ordering of the pixel register. This pattern is stepped through the array 5 times with many hit strobes issued for each (typically 100). In order to transmit this many hits in to the buffers in the available time, (which corresponds to the difference between the trigger latency and the time-over-

threshold (TOT)¹), the column readout clock (i.e. the Φ clock) must be operated at the maximum 40MHz which effects a hit transfer rate of 20MHz.

For MCMs, the amount of data which may be processed from a single FE chip is limited by the available receiver FIFO space in the MCC. When FE chips within MCMs are tested in this way therefore, a mask pattern is typically used in which only 10 pixels per column pair are strobed. For this pattern, the Φ clock may be operated at 20MHz. Digital tests of FE-I1 in which the operation frequency of the chip (XCK) is varied, typically reveal perfect operation up to 81MHz with Φ running at 20MHz. The first errors are evident in the TOT field, the correct hit pattern meanwhile is produced up to 90MHz. If Φ is operated at 40MHz the performance margin is somewhat less comfortable with TOT corruption occurring at 43MHz and a small number of hit address' becoming corrupt at 48MHz. The digital supply voltage margin extends down to 1.3V at XCK=40MHz and 1.8V with XCK running at 80MHz. The correct loading of the control registers of FE-I1 has been verified for XCK frequencies in excess of 100MHz.

Digital hit creation has also been employed in proving the correct operation of the special TOT modes available in FE-I1. Two tunable *TOT thresholds* may be applied, one of which is used to reject low-TOT hits (to eliminate noise tails). The other threshold is intended to act as a digital timewalk correction by reading out hits with low TOT twice in contiguous beam crossings. Another mode exists in which the TOT field in the serial hit data transmitted from FE-I1 is replaced by (optionally) the 8-bit leading or trailing edge timestamp of the hit for diagnostic purposes.

4 Analogue Performance of FE-I1

The first stage in understanding the analogue behaviour of FE-I1 is to calibrate the charge injection circuitry. This involves measuring the magnitude of the calibration capacitors and the voltage scale of the VCAL DAC, which for internal charge injection is used to define the DC level which the chopper steps to from the analogue supply voltage (AVDD). Each pixel FE has a small (C_{inj-lo}) and a large (C_{inj-hi}) injection capacitor connected to its preamplifier and one has the option of applying the calibration voltage step to either the small capacitor or both. The small capacitance is used when detailed measurements are made at small charge scales of $<^2$ 1MIP, e.g. noise and threshold. The large capacitance is useful in performing measurements which require the injection of charge up to very high values (\approx 200,000e⁻). Such measurements include timewalk, crosstalk and TOT calibration. Incorporated into the bottom of FE-I1 is a special charge-pump circuit along with arrays of capacitors

¹For digital hit creation the TOT is derived from the strobe width in Beam Cross-Over (BCO) units

²Expectation charge arising from the interaction of a Minimum Ionising Particle

which match in design the three critical front-end capacitors C_{inj-lo} , C_{inj-hi} and the feedback capacitor (C_f). The DC current of the charge pump upon application of a voltage gives the capacitance value according to $C = \frac{dQ}{dt} \cdot \frac{dt}{dV} = \frac{\Delta I / (fV)}{\Delta n}$ where n is the number of capacitors in the array, selectable from 0,1,2 or 4, and f is the applied frequency which is derived from XCK with magnitude XCK/4, XCK/8, XCK/16 or XCK/32.

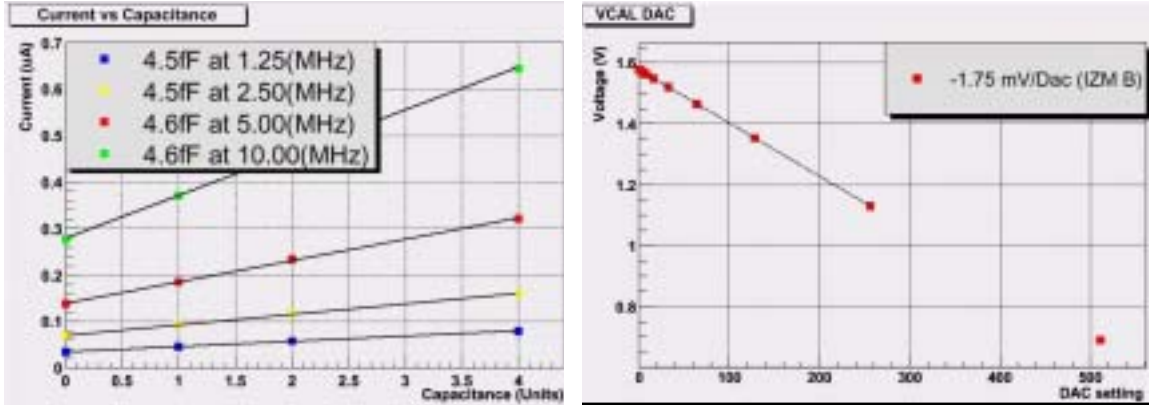


Figure 2: Measurement of the Small Injection-Capacitance Magnitude (left) and Characterisation of the VCAL DAC.

Figure 2 shows the result of a charge pump measurement on the C_{inj-lo} replica capacitors for the four available frequencies and for all possible numbers of capacitor units. A value of around 4.5fF is consistently measured. An example characterisation of the 9-bit VCAL DAC is shown on the right from which a gradient of 1.76mV per DAC count is derived. The combination of these two numbers leads to an internal injection calibration of 44.7e- per DAC count.

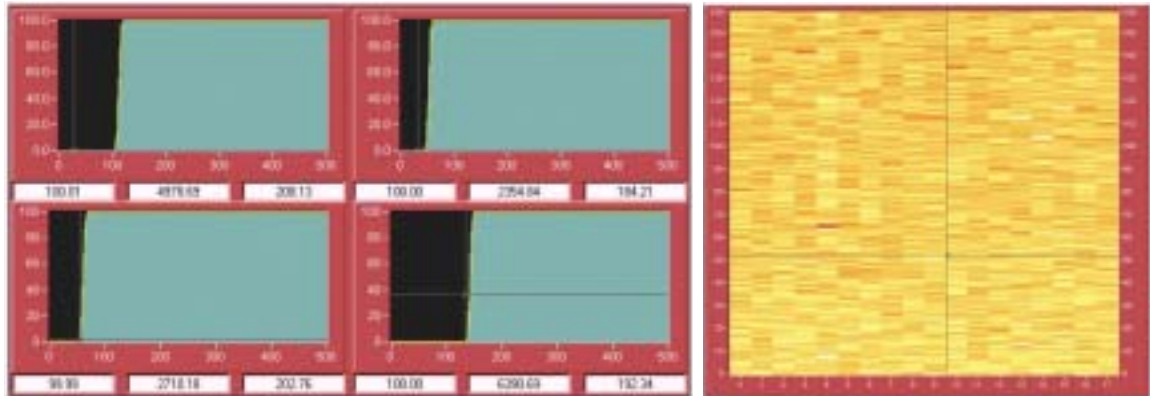


Figure 3: 'S-curves' of Occupancy vs. Charge and Threshold Scan Pixel Hit-Map.

In order to determine the threshold and noise of each individual pixel, the amount of injected charge is scanned by either varying the VCAL DAC setting (for internal injection) or the value of one of the VCAL DACs on the TurboPCC. For each value of charge within the scan, 100 strobes are issued. As the charge magnitude passes the discriminator threshold the pixels start to generate hits, eventually reaching a plateau corresponding to the number of given strobes. A histogram of occupancy versus charge is built up for each pixel which has an ‘s-curve’ profile. The s-curve arises from the fact that the noise present in the pixel channel causes the threshold measurement to statistically vary in a Gaussian manner. This Gaussian distribution is integrated in this type of measurement and since there is no functional form for the integral of a Gaussian, an approximate error function is used in order to fit to the data and derive the threshold (given by the median) and the equivalent noise charge (ENC) which is given by σ .

Figure 3 illustrates four example s-curve histograms from a threshold scan performed on an FE-I1 chip. On the right of the figure the integral of all hits in these histograms is plotted as a geographical colour-scale map for a whole FE-chip, column number on the horizontal axis and row number in the vertical. The colour variation indicates the degree of threshold dispersion over the chip for a case in which the individual threshold trim DAC settings have not been optimised in order to minimise the width of the threshold distribution. Note that every pixel in the chip is responsive, this is generally the case for FE chips which pass the most fundamental selection criteria at the wafer probe stage.

In order to meet the required performance demands in terms of fake occupancy and efficiency in ATLAS, the thresholds in all pixels need to be matched at the level of $\approx 100e^-$. The FE-I1 pixel cell design incorporates a 5-bit threshold tuning DAC (TDAC) which provides a mechanism for making small relative threshold adjustments at the single channel level. Figure 4 shows some example threshold and ENC distributions for a single chip assembly in which a special single-FE-scale production-style sensor is bump-bonded to an FE-I1 chip. A Gaussian fit to the initial untuned dispersion has a σ of $868e^-$. This is reduced to $83e^-$ after tuning. The post-tune distribution has a slight upper tail but no channels at thresholds too low which would cause them to be inoperably noisy. The RMS of this distribution is $\approx 100e^-$. In the lower half of this figure are the ENC distributions corresponding to the untuned and tuned cases. Before tuning the noise distribution peaks at $248e^-$, this comes down to $231e^-$ after tuning since there are no longer any pixels at extremely low thresholds (in an oscillation condition e.g.) to influence the overall noise level of the chip.

In previous generations of ATLAS Pixel FE chip [5] the strategy for finding the best TDAC settings for each pixel involved performing a single threshold scan at a middling TDAC setting. The resultant threshold distribution was then carved up into equal slices in order to decide the TDAC assignment for each pixel by virtue of the slice in which it was to be found. Such an algorithm relies on the assumption that the

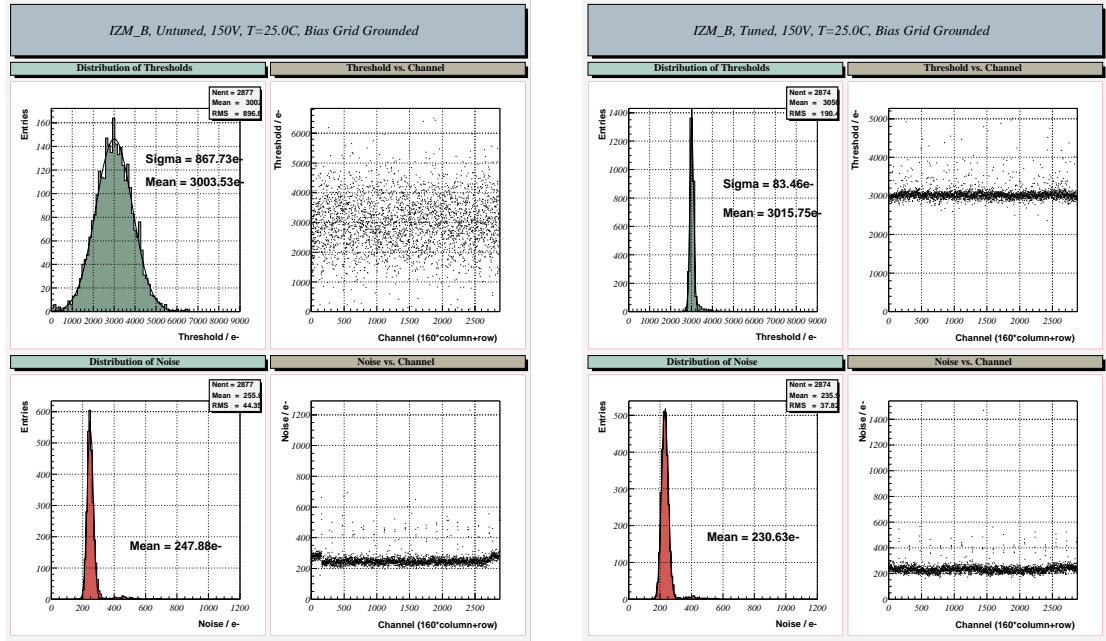


Figure 4: *Threshold and ENC Distributions before and after Threshold Trim-DAC Optimisation.*

threshold vs. TDAC function for each pixel is linear with a gradient which is uniform throughout the chip. In the case of FE-I1 it was found that these functions tend to become exponential in nature at step-size scales which are commensurate with the scale of the initial dispersion. Furthermore, the trim DACs themselves turned out to be non-monotonic and there is a systematic variation of the ‘trim-current’ which is used to feed these DACs, along the pixel columns. The result is that one is forced to make threshold scans for every available TDAC setting in order to obtain the best TDAC tune and then, for each pixel, choose the value which is closest to the target threshold. This results in a factor 32 increase in the amount of time required to prepare modules for testing with a source or for test-beam, etc. This would have serious implications both for the module production schedule and for the calibration of modules *in situ* in ATLAS. Each of these issues is being addressed in the design of the next generation front-end chip, FE-I2.

Figure 5 shows threshold and ENC distributions for an entire 16-chip MCM. With careful TDAC tuning a threshold dispersion not dissimilar to the single-chip-assembly case (113e-) is achieved. The noise evaluation reveals an ENC of 263e- in the tuned state which also compares very favourably with assemblies constructed from single FEs.

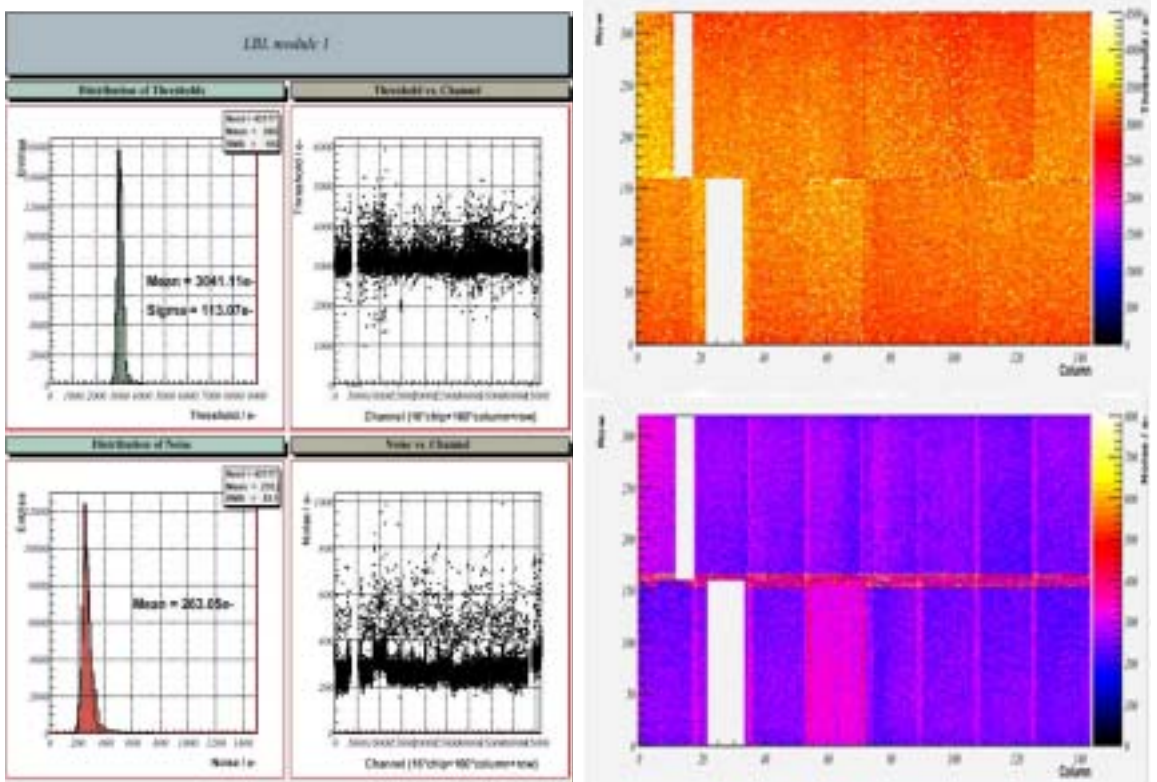


Figure 5: *Full MCM-Scale Tuned Threshold and ENC Distributions.*

The ability to properly associate hits with their originating beam interactions is critical to the efficient operation of the ATLAS Pixel Tracker. The BCO at the LHC is 25ns, therefore it's important to verify that the range of times at which the discriminator is seen to fire, (with respect to the time at which the particle traverses the silicon sensor), is acceptably within this time. For smaller charges the analogue chain in the pixel cell takes longer to respond. This tends to be dominated by the preamplifier stage when a finite load capacitance, (i.e. from the sensor), is present. In the absence of any capacitive load, the discriminator speed tends to be the limiting factor. The convention adopted here is to express this *timewalk* as the amount of *overdrive* (or charge-above-threshold) for which the discriminator fires 20ns more slowly than for an overdrive of 50,000e-. The choice of 20ns is made in order to allow some contingency for other sources of timing uncertainties, (e.g. jitter in the trigger distribution).

The strategy for evaluating timewalk in FE-I1 is to determine the relative response time for a large range of input charges. For each charge the level-1 trigger delay is set in order to be slightly too late (e.g. by 1 BCO) and the precise delay of the calibration hit-strobe is scanned with respect to the trigger. As this delay is increased (in 0.66ns

steps) the hit is gradually pushed forward in time until it eventually matches the trigger and is read out of the chip. The resultant histograms of occupancy versus strobe-delay essentially have the appearance of a step function, except the step has a finite width due to the projection of noise in the channel onto the time axis, (as per the derivative of the timewalk function). Fitting an error function to such a histogram yields an accurate relative time measurement which is given by the error-function median. The timewalk function is derived by plotting these median points versus the difference of the input charge and the known threshold for each pixel. A process of interpolation is then used to find the 50,000e⁻ overdrive point and the 20ns timewalk point in order to extract the minimum in-time overdrive limit.

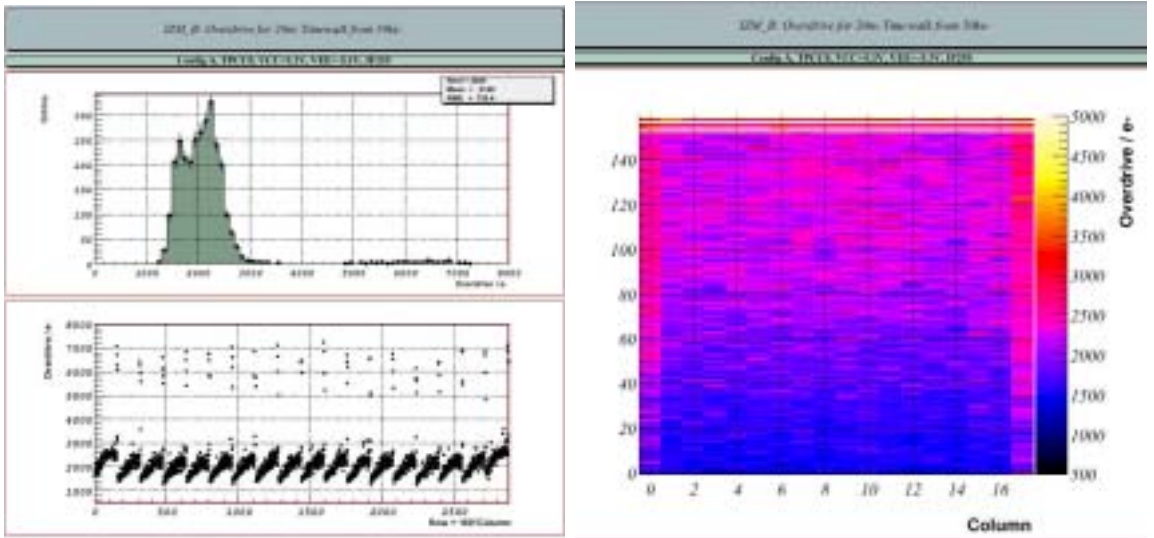


Figure 6: *In-Time Overdrive Distribution for an FE-I1 Single-Chip Assembly.*

Figure 6 show the distribution of in-time overdrive for an FE-I1 single-chip assembly as a histogram, a scatter plot (overdrive versus channel-ID) and a colour-scale map representation. Clearly there is a systematic worsening of the timewalk with increasing row number which is due to a deficiency in the distribution of the main preamplifier bias current (IP). This issue is being addressed in FE-I2. The overdrive values for the very lowest row numbers are the most relevant since those channels are in receipt of the correct biases. For these channels an overdrive of $\approx 1300e^-$ is recorded indicating an ‘in-time threshold’ of 4300e⁻ for the nominal global threshold of 3000e⁻. The ATLAS Pixel module concept involves a subset of special pixels which are implemented into the sensor layout in order to cover the narrow regions in between FE chips. In the $r\phi$ (row-wise) direction there are four extra rows of pixels (per FE chip) which are ganged together with four other rows using direct metal bridge connections on the sensor. Therefore for four of the row IDs on the electronics chip (153,155,157 and 159,) one expects to encounter twice the capacitive load (and twice

the leakage current). In reality there is appreciably more than twice the load due to the extra parasitic capacitance presented by the metal bridges themselves. For the z -direction the outer two columns are lengthened by 50% (to $600\mu\text{m}$). These pixels also have an enhanced load capacitance therefore, thus higher noise, crosstalk and timewalk. Eight of the channels on each FE chip fall into both categories.

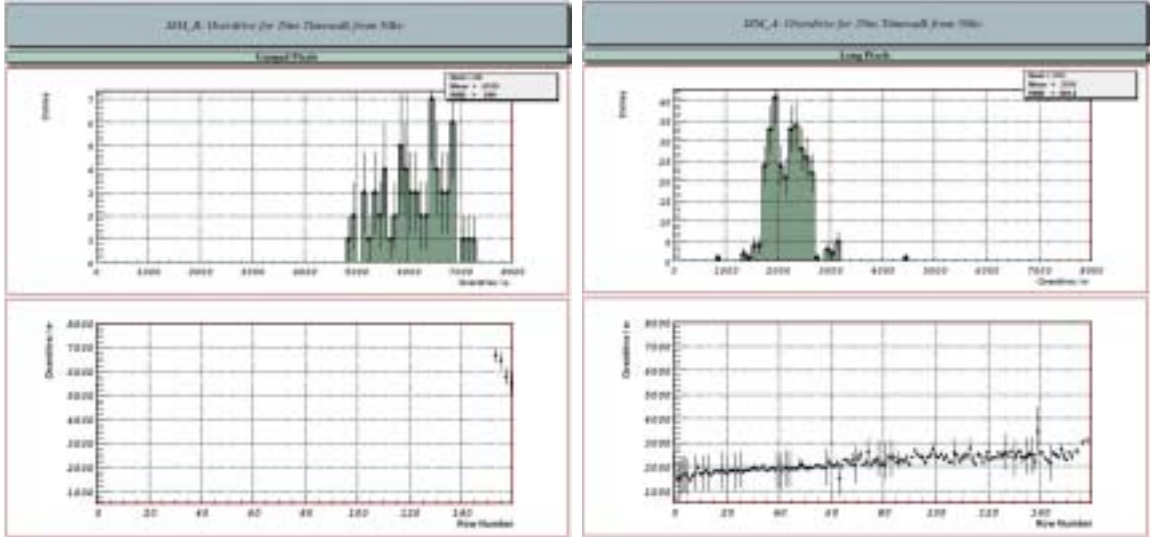


Figure 7: *In-Time Overdrive Distributions for ‘Ganged’ Pixels (left) and Elongated Pixels.*

In Figure 7 these special pixels are individually examined. Since the ganged pixels are at the top of the chip the fact that the load capacitance is much higher is exasperated maximally by the poor preamplifier bias distribution. An in-time overdrive range of 5ke- to 7ke- is measured. For the long pixels the figures are between 1.7ke- to 2.7ke-. In FE-I2 provision is being made to supply the ganged pixels with a much enhanced preamplifier bias current to compensate for their high load. Since they only account for a small percentage of channels, the impact on the power budget in taking this measure is minimal.

The method for determining the degree of charge loss to neighbouring pixels (i.e. the *analogue crosstalk*) involves enabling a pixel for which the threshold is known and injecting a range of charge into its neighbours up to very high ($\approx 200\text{ke-}$) values. As the magnitude of the injected charge is increased, eventually the degree of charge which couples into the readout-enabled pixel is sufficient to cause its discriminator to fire. The percentage of crosstalk is then simply evaluated as the quotient of the threshold and the median charge for this to occur (by fitting an error-function in the usual manner). Figure 8 shows the distribution of crosstalk for an FE-I1 assembly. For the regular $50\mu\text{m} \times 400\mu\text{m}$ pixels the crosstalk is determined to be 2.4%

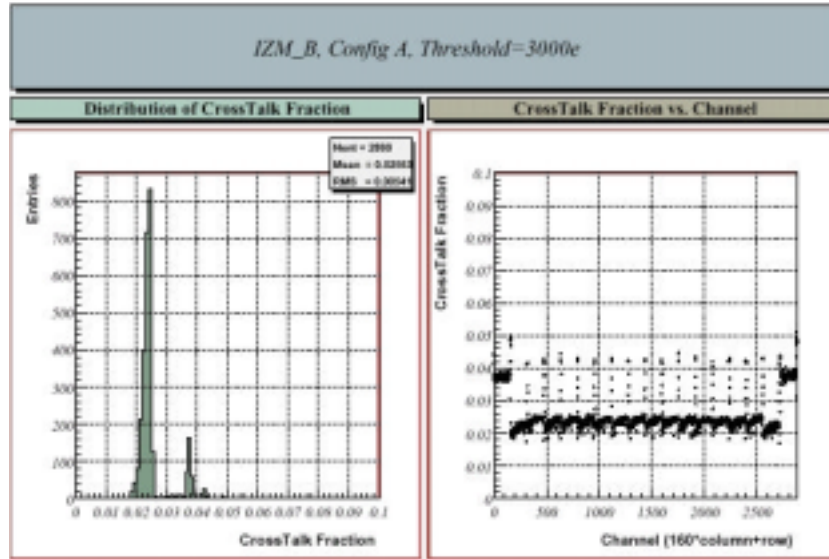


Figure 8: Crosstalk for an FE-I1 Single-Chip Assembly.

for the special channels (ganged pixels and long pixels) the figure is 3.9%, comparing very favourably with the 10% requirement.

In the left of Figure 9 the relationship between the 8-bit Time-Over-Threshold charge measurement and input calibration charge is expressed as a set of functional fits for an entire FE-I1 assembly and as a distribution of Mean TOT for an input charge of 50ke-. For these data no attempt has been made to match the TOT calibrations channel to channel by adjusting the 5-bit feedback current trim DACs. Without any tuning the matching is already better than 10% with a dispersion of 9.4BCOs recorded for a distribution which peaks at 112.5BCOs.

One of the special features of the FE-I1 design is the implementation of a special leakage current ADC (MONLEAK) which may be utilised to measure the leakage current to an accuracy of 9 bits for any chosen combination of pixels. After irradiation, this feature may be used to measure the individual leakage current in each pixel (and to correlate it to the noise). Before irradiation the leakage is too small to measure and the circuit is actually measuring the feedback current $\times 3/2$. For individual pixels this current is also too small to measure, however several pixels may take part in the measurement in order to derive an average. Advantage has been taken of this feature in order to verify that the channel to channel matching of the actual feedback current itself is excellent. On the right of Figure 9 an example leakage current map is shown for an assembly which has been irradiated to 300kRad at the T7 irradiation beam at CERN. The device was at room temperature for this measurement giving a typical current of 30nA per pixel. The actual beam profile is clearly visible and

measurements of this kind were used in order to examine the uniformity of exposure during the irradiation[7] of several FE-I1 assemblies in May 2002.

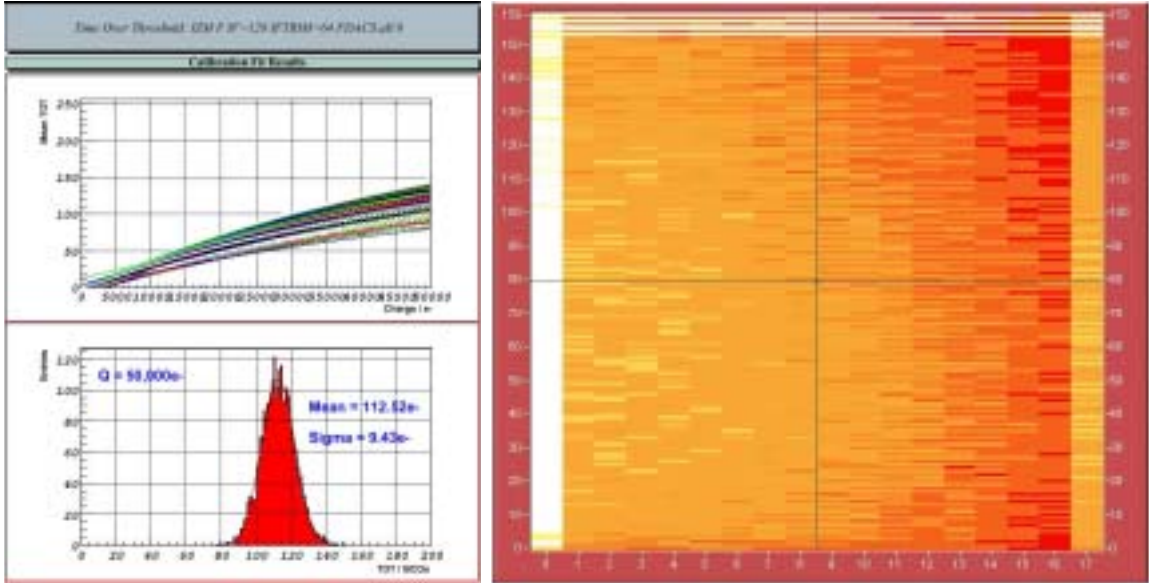


Figure 9: *TOT Calibration Curves and Distribution at 50ke- (left) and a Leakage Current Map for an Irradiated FE-I1 Assembly Produced Using the MONLEAK ADC.*

5 Conclusions

FE-I1, the first ATLAS Pixel Deep-Submicron front-end readout chip, performed very close to expectation in most respects, meeting the specified requirements for noise, threshold dispersion and crosstalk at the single-chip-assembly and full multi-chip-module scales. The timewalk is acceptable for standard and elongated pixels which are in receipt of the correct preamplifier bias current (IP). A bias current distribution issue has the effect of worsening the timewalk with increasing row-number. This problem is understood and will be fixed in the next generation front-end chip (FE-I2). Also the special ganged pixels will be brought into specification, with regard to timewalk, by providing them with an extra degree of preamplifier bias current. The digital readout logic is very robust up to frequencies well in advance of the 40MHz standard operation frequency provided the column readout clock is operated at 20MHz. At 40MHz the frequency range is more marginal and this performance aspect is being addressed in the FE-I2 design.

References

- [1] ‘ATLAS Technical Proposal’, *The ATLAS Collaboration* **CERN/LHCC/94-43**
- [2] ‘Inner Detector Technical Design Report’, *The ATLAS Collaboration* **CERN/LHCC/97-17**
- [3] ‘The Front-end ASIC for the ATLAS Pixel Detector’, *K. Einsweiler for the ATLAS Pixel Collaboration*. This conference.
- [4] ‘An Introduction to Submicron Electronics’, *M. Campbell*. Vertex 2000, 9th International Workshop on Vertex Detectors, Sleeping Bear Dunes MI, September 2000.
- [5] ‘The ATLAS Pixel On-Detector Readout Electronics’. *J. Richardson for the ATLAS Pixel Collaboration*. 5th Workshop of LHC Electronics, Snowmass CO, September 1999.
- [6] ‘MCC: the Module Controller Chip for the ATLAS Pixel Detector’, *R. Beccherle et al.* **ATL-INDET-2002-002**. Submitted to Nuc. Inst. Meth. A
- [7] ‘Irradiation Tests of ATLAS Pixel Electronics’, *A. Saavedra for the ATLAS Pixel Collaboration*. This conference.