Status of Functional and Radiation Testing of the Deep Submicron SCA Controller

S. Böttcher, S. Negroni, J.A. Parsons, S. Simion
Nevis Labs, Columbia University

October 22, 2001

Abstract

We describe the status and results of the program to functionally test and radiation qualify the deep submicron (DSM) SCA Controller ASIC designed for the ATLAS LAr Front End Board.

1 Introduction

The Front End Board (FEB) of the ATLAS LAr readout uses switched capacitor array (SCA) analog memories for analog storage of the calorimeter signals during the Level 1 trigger latency. The addressing of the SCAs, and other control functions of the FEBs, are managed by two SCA Controller (SCAC) chips on each FEB.

The SCAC has been implemented using a radiation-hardened standard cell library targeted on a 0.25 μm CMOS process. Details of the DSM SCAC design can be found in reference [1]. The SCAC chip, measuring 4×4 mm^2, was submitted for prototyping in the Feb. 28, 2001 MPW run. Bare die were received in July.

2 DSM SCAC Fabrication and Packaging

The DSM multi-user run was organized and ordered through CERN. The processed wafers were delivered to CERN, who arranged the dicing of the wafers and delivered bare die to the individual customers.

It was reported to CERN by the vendor that the run had been processed using "striping", a technique in which the vendor purposely varies the process parameters in a systematic way over the various reticles on the wafer. This technique allows a direct evaluation of the robustness of the design for the range of parameter variations to be expected during a large production run. Unfortunately, CERN did not recognize this information until after the wafers had been diced, and therefore information concerning the position of particular die on the wafers has been irretrievably lost, removing the possibility of correlating individual chip performance with the associated values of the parameters. On a more global level, it would still be possible to conclude that a high overall yield would imply a design which is robust against the expected parameter variations.

Upon reception of the bare die at Nevis in July, we discovered via visual inspection a problem with the processing of some die, which were missing some critical layers. We reported this observation to CERN, who then investigated the problem with the other customers and the vendor. It is now understood that certain of the processing steps are not implemented for reticles which are not complete due to their position near the wafer boundaries. However, these partial reticles were not correctly identified to the wafer dicing company, with the result that some not fully processed chips were also diced and included in the delivered die. Typically 10 - 20% of the delivered die suffer from this problem, and are as a result completely non-functional. These die can be identified via visual inspection, and we have removed them from the sample of chips to be packaged and tested.

The tests described here have been performed so far on a total of 5 DSM SCAC chips which were manually bonded into the 100-pin PQFP package. Industrial packaging of the remaining (> 100) die has been ordered, with packaged parts expected in October.

3 Functional Testing

The DSM SCAC chips have been functionally tested in two different environments, a standalone test jig designed specifically for SCAC testing and evaluation, and integrated on the 1/4 Digital FEB which incorporates many of the other custom digital ASICs.

3.1 Standalone testing

The DSM SCAC chips were functionally tested in the lab using a test jig designed to be operated with the radiation test setup described later in section 4.3.

During the tests, the DSM SCAC was operated at 40 MHz, with pseudo-random triggers at a mean rate of 78 kHz. Proper operation of the SCAC was verified by checking all I/O signals every clock cycle against the corresponding signals of the SCAC emulator embedded in the FPGA running the test (see section 4.3 for more details).

The SCAC is designed for several different operating modes. s of several different types have been successfully performed:

- Autogain mode - the "standard" five-sample autogain operation of the DSM SCAC was verified to function properly.
- Fixed gain mode - the chips operate properly in fixed gain mode, provided the delay counter is appropriately configured (see section 7 for more details).
- Variable number of samples - it was verified that one can select a configurable number of samples to be read out, between two (three for fixed gain) and 32.
- Address sequencing - the DSM SCAC works properly in both the "standard mode", where the addresses are kept in sequence, and in the mode where the sequencing is switched off and the address sequence becomes gradually randomized.
- L1A encoding - is was verified that the DSM SCAC can be configured to either interpret single L1A pulses as triggers or to require two consecutive pulses.
- L1 latency - the L1 latency can be programmed to any value between 13 and 127 clock cycles

- BCID and EVTN counters - proper operation of the BCID and EVTN number counters, and use of the corresponding BCR and ECR signals, was verified. (However, it was noticed that the BCR and ECR pins of the DSM SCAC were inadverdently swapped on the pinout of reference [1]; a corrected pinout diagram will be prepared.)

DSM SCAC chips have been operated continuously for more than 24 hours without any errors.

To test the speed of the DSM SCAC, the clock frequency was increased above 40 MHz. Up until 55 MHz, the FPGA was able to verify that the DSM SCAC functioned properly. At higher frequencies, the FPGA no longer functioned. However, it was possible via oscilloscope measurements to verify proper operation of the DSM SCAC for frequencies up to 66 MHz.

3.2 Testing on the 1/4 Digital FEB

The so-called "1/4 Digital FEB" is a test jig designed to allow integration tests of most of the custom digital ASICs designed for the FEB. The various ASICs are mounted on separate daughter boards which plug into the 1/4 Digital FEB; this modularity allowed tests using FPGA emulators to proceed before the first versions of the ASICs were available.

The DSM SCAC was tested on the 1/4 Digital FEB in operation with the DMILL TTCRx, DMILL SPAC Slave, DMILL Configuration Controller, and DMILL Gain Selector. The testing verified the proper functionality of the DSM SCAC interfaces with those chips, including:

- receipt of 40 MHz CLK, L1A, BCR, ECR from the TTCRx,
- downloading, and reading back, of configuration parameters to the DSM SCAC via the chain SPAC Slave → Configuration Controller → SCAC,
- control by the SCAC of the Gain Selector in response to L1 trigger

3.3 Design imperfections

During functional testing, some imperfections in the design of the DSM SCAC were observed:

- In fixed gain mode only, it is necessary for the SCAC to introduce a delay of 5 RCLK cycles (200 ns each) between events in order to respect the FEB output event format, which requires some words of all zeroes to be present between events. The current DSM SCAC design inserts a delay of only 3 RCLKs. However, the additional 2 cycles can be achieved by properly configuring the programmable delay counter of the SCAC (for details, see reference [1]).
- The SCAC seeks to maintain the SCA write addresses in monotonic sequence, in order to minimize address-related noise in the readout. However, the address re-ordering will be lost in the current DSM SCAC if at any time there are so many triggers waiting for processing that the FREE FIFO contains fewer than 3 available addresses. Assuming proper implementation of the "leaky bucket" in the L1 trigger, this condition should never occur. However, system problems such as SEU-induced fake triggers could lead to this condition. Recovering the correct address sequence requires the SCAC be issued a Reset.

The detailed cause and behavior of each of these imperfections has been studied and verified in the Verilog simulation of the DSM SCAC, and solutions have been identified. While neither of these "bugs" is critical, it would be preferable to correct them in a future iteration of the design.

4 Radiation Testing Procedure

The FEBs will be located on-detector, in crates mounted on the feedthroughs in the gap between the barrel and endcap cryostats (for the EM Barrel) and at the back of the endcap cryostats (for the EM endcap and Hadronic endcap). All FEB components, including the SCAC, must therefore undergo a radiation qualification procedure to ensure their suitability for use in the ATLAS environment.

4.1 Radiation tolerance requirements

The simulated radiation levels in the EM Barrel crates is listed in Table 1 for 10 years of LHC operation for the cases of Total Ionizing Dose (TID), neutron fluence (NIEL) and hadrons with sufficient energy (> 20 MeV) to potentially cause single-event effects (SEE). The corresponding radiation levels for the endcap crates are approximately ten times smaller than in the barrel.

Radiation Type	Simulated Level for 10 yrs	Safety Factors			Total RTC
		SIM	LDR	LOT	
TID	5 kRad	3.5	5	2	175 kRad
NIEL	1.6×10^{12} n/cm^2	5	1	2	1.6×10^{13} n/cm^2
SEE	7.7×10^{11} h/cm^2	5	1	2	7.7×10^{12} h/cm^2

Table 1: Summary of the expected radiation levels, and corresponding safety factors, for the EM barrel crates for 10 years of LHC operation.

The ATLAS Policy for Radiation Tolerant Electronics [2] assigns safety factors (SF) which are to be applied in the radiation qualification process. The safety factors account for the uncertainties in the simulated radiation levels (SF_{SIM}), in possible low dose rate effects (SF_{LDR}), and for possible lot-to-lot variations (SF_{LOT}). Since the DSM process is not a radiation-guaranteed process, we consider the DSM SCAC as a commercial-off-the-shelf (COTS) component in assigning numerical values of the safety factors. The final DSM SCAC production would involve delivery of a homogeneous batch of 48 wafers; we therefore use values for SF_{LOT} appropriate for pre-selection of COTS components to be produced in series from a *known* batch. The SF values we use are summarized in Table 1.

The final Radiation Tolerance Criteria (RTC) are then calculated as the simulated levels multiplied by the product of the three safety factors. As shown in the last column of Table 1, the corresponding values are

- RTC_{TID} = 175 kRad,
- $RTC_{NIEL} = 1.6\times10^{13}$ n/cm^2 (1 MeV n equiv.), and
- $RTC_{SEE} = 7.7\times10^{12}$ h/cm^2 (hadrons > 20 MeV).

4.2 Radiation facility

The irradiation was performed at the *Harvard Cyclotron Laboratory (HCL)* [3] in Cambridge, Massachusetts. The HCL uses a proton synchrocyclotron to generate proton beams, primarily for radiation therapy. The facility can be used during the weekends for other purposes, including radiation qualification of electronics.

The HCL proton beam has a maximum energy 158 MeV on target; irradiation at lower proton energies can be achieved by insertion in the beamline of lucite absorbers of varying thicknesses. The beam energy resolution is typically 2 - 5 MeV for energies above 100 MeV, and about 10 MeV for lower energies.

The HCL high-intensity beamline delivers a maximum flux of about 3.5×10^{10} protons/cm^2/s. The uniformity of the flux is within $\approx \pm 20\%$ over a beamspot of $\approx$ 1 cm diameter. The proton fluence is measured online using a calibrated ion chamber positioned in the proton beam, the signal from which is integrated and used to generated "Monitoring Pulses" (MON) after given increments of integrated proton fluence. The MON are used by the cyclotron control room to operate the machine, and were also read by our DAQ PC in order to provide online measurements of the integrated exposure. Typical running conditions corresponded to rates of several MON per second.

Using a 20 - 160 MeV proton beam allows us to simultaneously accumulate total ionizing dose (TID), 1 MeV-equivalent neutron fluence, and hadrons with sufficient energy to potentially cause SEE. The conversion factors are energy-dependent. For protons in this energy range, the 1 MeV-equivalent neutron fluence is approximately twice the proton fluence. Table 2 summarizes the conversion for calculating the equivalent TID from the proton fluence.

Proton Energy (MeV)	Proton Fluence per MRad TID (10^{13} p/cm^2)
158	1.47
100	1.06
50	0.67
20	0.31

Table 2: Fluence of protons required to accumulated 1 MRad total ionizing dose, as a function of proton energy.

4.3 Radiation test setup

The SCAC radiation test setup utilizes the portable DAQ we developed for a variety of electronics testing and radiation qualification purposes. The DAQ is driven by a portable PC equipped with two National Instruments DAQ boards, one for digital I/O and the other for analog I/O and power connections. Each DAQ board is connected via its own 50-pin 75-foot twisted-pair cable to a "Sample Driver" board.

The Sample Driver contains voltage regulators for providing power to the Device-Under-Test (DUT) fixture; the power supply currents and voltages are measured through the analog I/O. In addition, four analog inputs and two DAC outputs are available to the DUT fixture. The digital I/O implements two 8-bit TTL ports, with strong TTL drivers and terminations on both ends of the cable.

The SCAC logic operates at 40 MHz. Due to the long cables required to remove the PC from the radiation environment, it is not possible for the PC to directly verify the correct 40 MHz operation of the SCAC. Instead, the real-time control of the SCAC DUT is driven by an FPGA mounted on the "FPGA testjig". The Sample Driver is connected via a 50-pin twisted-pair ribbon cable to the FPGA testjig. There is a 4-bit parallel full duplex byte stream protocol implemented on the radiation cabling setup. On top of that we have a message protocol. The computer is the master initiating each message, and the FPGA receives, or returns data bytes. The message request contains an address and a byte count. Thus we can address resources inside the FPGA. (The FPGA testjig can also be connected directly to the digital I/O board in the PC, without all the other cabling, in the case of bench testing without radiation, in order to provide faster I/O. In this case, the 4-bit protocal is replaced by an 8-bit wide protocol, basically i8255-mode-1, again to increase bandwidth.)

The FPGA testjig was connected to the DUT fixture by two short ($\approx$ 20 cm) 40-pin cables, long enough to allow the FPGA to be sufficiently shielded from the proton beam. All I/O pins of the SCAC were connected in this way to the FPGA, which could therefore monitor all SCAC activity in real time. The FPGA code included a complete emulation of the SCAC logic.

During operation of the test, the SCAC was operated at 40 MHz and triggered by a pseudo-random trigger generator at 78 kHz (40 MHz/512). The FPGA compared all SCAC outputs cycle-by-cycle with its own emulated outputs in order to search for discrepancies signalling SEU or other error conditions. The test sequence would run until the FPGA detected a mismatch between any DUT output and the corresponding output from the SCAC emulator inside the FPGA. In case of a mismatch, the FPGA immediately stopped the test so that the internal registers of the FPGA could be read in order to diagnose the cause of the mismatch.

The DAQ PC polled the FPGA approximately once per second to check if an error had occurred. Once an error was detected, the captured FPGA signals were read by the PC, after which the test setup was initialized and then the test started over with the same trigger sequence.

5 Radiation Test Results

As described above, the real-time monitoring of the SCAC DUT included both measurement of its current and complete checking of its operation for occurrences of SEU.

During the weekend of July 28/29, 2001, three SCAC chips (denoted samples C, D, and E) were irradiated. A fourth SCAC chip (sample A) was irradiated during the weekend of August 18/19, 2001. Some different proton beam energies were used in order to investigate the energy dependence of the SEU cross-section. The total proton fluence, and the corresponding total ionizing dose (TID), to which each sample was exposed is summarized in Table 3.

5.1 Current monitoring

The current drawn by the DSM SCAC DUT was monitored during irradiation. As the integrated dose increased, the current exhibited a slight decrease, totally a few per cent after the equivalent of several MRad; a plot of the current versus proton fluence is shown in Figure 5.1 for Sample C as an example. Similar behavior has been seen in other chips produced

Chip ID	Proton Energy (MeV)	Proton Fluence (10^{13} p/cm^2)	Total Ionizing Dose (MRad)
C	158	6.60	4.5
D	50	2.78	4.1
E	22	1.40	4.6
A	158	14.5	9.9

Table 3: Proton energies and fluences, and corresponding total ionizing dose (TID), to which each of the DSM SCAC chips were exposed.

with the same DSM process (for example, see reference [4]), where it has been attributed to an increase (in absolute value) of the threshold voltages of both the n-channel and p-channel transistors. This very small effect is of no concern for operation in ATLAS.

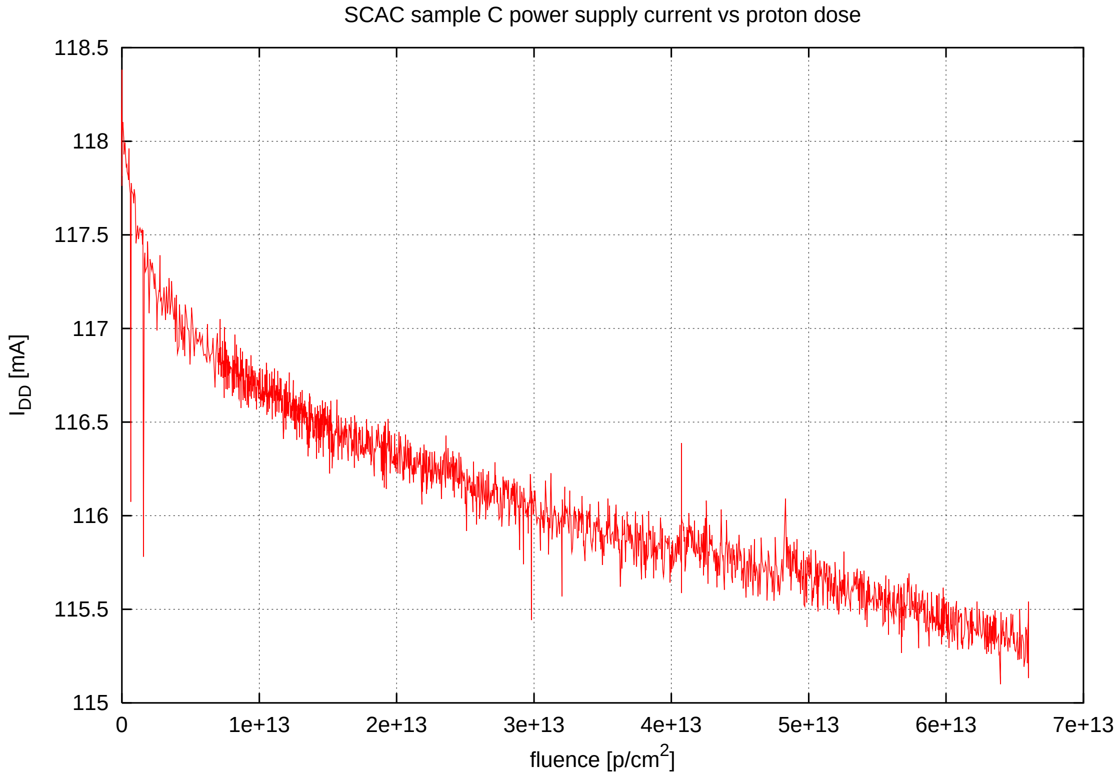

Figure 1: Current drawn by DSM SCAC Sample C as a function of the proton fluence.

No occurrences of latch-up were observed, and it was not necessary at any time to cycle the power to the SCAC.

5.2 SEU cross-section measurements

As described in section 4.3, whenever the FPGA detected a mismatch between the outputs of the DSM SCAC DUT and the outputs of the SCAC emulator logic within the FPGA, the test was interrupted so that the PC could read internal registers of the FPGA in order to determine the type of mismatch. This method allows measurement not only of the total rate of SEU-induced errors, but also diagnosis of the type of error that occurred.

The FPGA counted the number of clocks between the INIT and the occurrence of an error, so that the live time of the test could be accurately computed. The typical time between errors was several seconds, and a mean time of about one second was required to read out the FPGA and re-initialize, leading to typical deadtimes of 20 - 30%. In determining SEU rates, it is necessary to normalize the number of errors observed against the proton fluence accumulated during the live time in which the test was in fact operational.

Two broad classes of errors can be distinguished: those with single bit errors, and those with multiple bit errors. The single bit errors can be further broken down by the location of the error, be they in the SRAM (SRAM), in the Latency FIFO (LFIFO), in the storage of the bunch crossing id (BCID), or in the storage of the event number (EVTN). Table 4 summarizes the total number of observed errors for each chip tested, and also the contributions of multiple bit (MBIT) errors and of the individual sources of single bit errors. The fluences listed in Table 4 correspond to those accumulated during the live time of the test (and are therefore smaller than those listed in Table 3 due to the deadtime incurred reading out the FPGA trace information after an error occurred).

Chip ID	Proton Energy (MeV)	Proton Fluence (10^{13}p/cm^2)	Occurrences of Various Types of Errors					
			MBIT	SRAM	LFIFO	BCID	EVTN	Total
C	158	4.24	14	336	104	5	5	464
D	50	2.19	0	167	58	4	2	231
E	22	1.18	0	21	10	1	0	32
A	158	4.93	7	626	215	11	3	862

Table 4: Proton energies and fluences to which the sample DSM SCAC chips were exposed, and the number of various types of errors detected, during the live time of the tests. See the text for more details.

It is possible that at least some of the MBIT errors are not in fact due to SCAC errors, but result from stray radiation upsetting the operation of the FPGA used to monitor the SCAC performance. However, in the analysis presented here we (conservatively) assign all MBIT errors to the SCAC.

The data of Table 3 can be used to calculate a total error cross-section. This cross-section is plotted in Figure 5.2 as a function of proton energy. Note that the error cross-sections calculated for Samples C and A, each obtained with 158 MeV protons, differ by approximately 60%. This discrepancy is not yet fully understood, and should be taken as representative of the size of the current systematic uncertainties in our understanding of the error rates.

To compare a subset of our results with those published about another DSM ASIC using basically the same SRAM, we normalize the SRAM single bit error rates reported in Table 4 to a rate per SRAM bit, taking into account that the rates in the table correspond to errors within an effective SRAM block of $144 \times 13 = 1872$ bits. The results in Table 4 correspond to SEU rates of $\approx 1 - 5 \times 10^{-15}$ cm^2 per SRAM bit. In reference [5] concerning SEU measurements on a chip designed for use in CMS, an effective upset cross-section of $\approx 2 \times 10^{-11}$ cm^2 was reported for an 18 kbit SRAM block, corresponding to a cross-section per bit of $\approx 10^{-15}$ cm^2. This number was based on measurements with heavy ions and protons, and then using a model to extrapolate the results to the LHC environment. Given the different methods, the agreement between our results and those of reference [5] is very

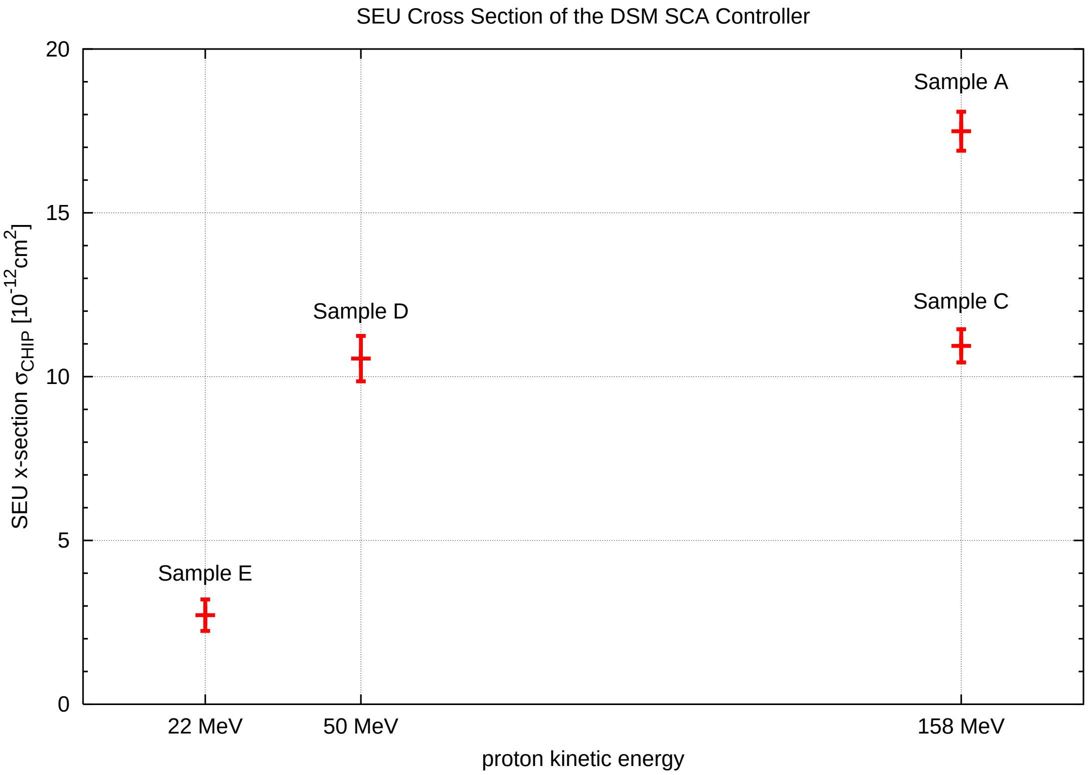

Figure 2: Total error cross-section as a function of proton energy.

good.

5.3 Impact of SEU-induced errors on ATLAS operation

Recall from Table 1 that, including the appropriate safety factors, $\mathrm{RTC}_{\mathrm{SEE}} = 7.7 \times 10^{12}$ h/cm^2/10 years (hadrons > 20 MeV). To conservatively estimate the number of SEU-induced errors expected for a single DSM SCAC chip during 10 years of ATLAS operation, we take the highest error cross-section measured (namely, that for Sample A using 158 MeV protons) of $\sigma = 1.75 \times 10^{-11}$ cm^{-2}, and assume it applies to ALL hadrons with energies above 20 MeV (even though we know from figure 5.2 that, as expected, the cross-section drops below approximately 50 MeV). With this conservative assumption, we anticipate 135 SEU-induced errors per DSM SCAC chip in 10 years (10^8 s) of ATLAS operation, or an average of approximately 8.5 days of operation between errors.

This error rate can be scaled to the entire EM Barrel system by taking into account the fact that a total of 1792 SCAC chips are required to instrument the EM Barrel (2 SCAC chips per FEB, with a total of 896 FEB). Scaling by 1792, we expect 242k SCAC SEU-induced errors in the entire EM Barrel readout in 10 years of ATLAS operation, or an average of approximately 413 seconds (6.9 minutes) of operation between errors.

Consideration of the impact of these SEU-induced errors on SCAC operation in ATLAS must take into account the result of the occurrence of each of the different types of errors. The SCAC design incorporates redundant memory bits with Hamming code, and many other features to harden the design against SEU (for details, see reference [1]). Many types of single bit errors are, in this way, correctable with no impact on running. Detection of other types of errors relies on the fact that each 128-channel FEB will be equipped with two SCAC chips, each of which will operate 64 channels. Data from the two sets of 64 channels, including the addresses supplied by each SCAC as well as various SCAC status and error bits, are combined and transmitted over a single optical link to the ROD. Monitoring and detection of SEU-induced errors in SCAC operation will be done in the ROD by comparing the addresses from the two SCAC chips, which should be identical. If a mismatch is detected, the FEB in question must be re-initialized via SPAC before its data can be considered valid.

We now discuss the impact of each of the observed types of errors:

- Multiple bit errors - these cannot be automatically corrected in the SCAC logic, and require that the FEB with the affected SCAC be issued a Reset in order to initialize and resynchronize the two SCAC chips.
- SRAM single bit errors - these are fully corrected by the EDC logic on chip, and have no impact on FEB operation. We only know that these errors occur because, as they are being automatically corrected within the SCAC chip, a flag is set to indicate that an error has occurred.
- LFIFO single bit errors - these are errors in redundant high order cell number bits, and have no impact on FEB operation
- BCID and EVTN single bit errors - these errors corrupt the BCID and EVTN numbers inserted in the FEB output dataflow to allow identification and alignment of complete events. Corruption of these counters would corrupt the data for a single event only, assuming appropriate re-alignment schemes are implemented in the ROD (these are necessary anyway due to the known fact that the optical links will occassionally temporarily malfunction due the SEU-induced effects in the links). Therefore, no FEB Reset would be required.

The combination of SRAM and LFIFO single bit errors correspond to typically 97% of the observed errors (see Table 4); as described above, these types of errors have no impact on the operation of the FEB. Occurrences of BCID or EVTN errors corrupts the FEB data for a single event, and requires the ROD re-align subsequent events to recover; these amount to about (1 - 2)% of the total error rate, and are therefore expected at a rate of approximately one such error in the EM Barrel per 5.7 hours, many orders of magnitude smaller than the SEU error rate for the optical links. The most serious type of error is the multiple bit (MBIT) error, since it requires the affected FEB to be issued a Reset via SPAC. However, the MBIT errors represent only about (1 - 3)% of the total, and are therefore expected to occur only once every approximately 3.5 hours.

6 Summary of Test Results

We have verified proper operation of the DSM SCAC, both in standalone tests and in operation together with the DMILL TTCRx, DMILL SPAC Slave, DMILL Config Controller, and DMILL Gain Selector. Four out of five SCAC chips tested are fully functional, with the other suffering from one defective output, corresponding to a preliminary yield (albeit it based on low statistics) of 80%.

The four fully functional chips have been irradiated with high-intensity proton beams of energy 22 - 158 MeV. No evidence was seen for degradations in performance due to accumulated ionizing dose or hadron fluence. The corresponding TIDs were in the range of several MRad, at least a factor of 20 above the $\mathrm{RTC}_{\mathrm{TID}}$. The integrated proton fluences were also a factor of several higher than the $\mathrm{RTC}_{\mathrm{NIEL}}$. SEU measurements lead to a conservative estimate of one FEB requiring a reset on average once every several hours (less than once per day before applying the Safety Factor of 10), integrated over the entire LAr readout.

7 Next Steps

Assuming a decision is made to utilize the DSM implementation of the SCAC in ATLAS, the design would need to be updated to include the new ESD-protected I/O cells. Given the lack of maturity of those new cells, a final multi-user run would seem appropriate. The very minor modifications required to correct the small design imperfections (see section would also be introduced at that time.

The next multi-user run is scheduled for submission in mid-November. This run would provide sufficient chips for the 30 FEB pre-production required for the "System Crate Test" milestone of December 2002. A DSM engineering run, including also the DSM Gain Selector and CLKFO chips, could proceed by mid-2002, on a similar timescale as the engineering runs foreseen for the analog and digital DMILL chips required by LAr.

8 Acknowledgements

We thank Ethan Cascio and the cyclotron personnel for their excellent operation of the Harvard Cyclotron Facility used for the irradiation measurements. We thank our colleagues at BNL for manually bonding the chips.

References

[1] S. Böttcher, J. Parsons, W. Sippach, and D. Gingrich. "Deep Submicron Implementation of the ATLAS LAr SCA Controller". Available at `http://www.nevis.columbia.edu/~atlas/electronics/asics/scac/scac-dsm.ps`

[2] For more details, see `http://atlas.web.cern.ch/Atlas/GROUPS/FRONTEND/radhard.htm`

[3] For more details, see `http://neurosurgery.mgh.harvard.edu/hcl/flyer.htm`

[4] Ph.D. thesis of G. Anelli, Grenoble (2000).

[5] F. Faccio, K. Kloukinas, G. Magazzu, and A. Marchioro. "SEU effects in registers and in a dual-ported static RAM designed in a 0.25 μm technology for applications in the LHC". Available as CERN-LHCC-99-033 and in the Proceedings of the 5th Conference on Electronics for LHC Experiments, Snowmass, 1999.