

RCC Technology Choice

Vince Cianciolo

07 April 2009

Introduction

The outer two layers of the PHENIX silicon vertex upgrade (VTX) are instrumented with novel “stripixel” detectors (described elsewhere [1]). Signals from each $3 \times 6 \text{ cm}^2$ sensor are digitized in twelve 128-channel ASICs developed by FNAL (SVX4, also described elsewhere [2]). The SVX4s for each sensor are mounted on a PCB known as the readout card (ROC). In order to reduce the required databus, local control of each ROC is implemented in a chip known as the ROC Control Chip (RCC), housed proximal to the ROC.

Initially, a decision was taken to implement the RCC as a rad-tolerant ASIC. R&D was carried out in which an initial RCC design was implemented in an FPGA, and then translated into an ASIC using rad-tolerant CMOS 0.25 μm cells. Since then, two of the constraints that drove the ASIC decision have eased; material budget and available area:

- In order to remove the large, geometrically-dependent pedestal variations, the ROC geometry was modified such that the area under the sensor is now a uniform ground plane. This forces the readout bus to run under the sensor, and provides more area in which to mount the RCC. Furthermore, ACTEL sells a version of its Axcelerator family chip (0.15 μm CMOS), with sufficient resources (pinouts and logic) in a chip-scale ($1.3 \times 1.3 \text{ cm}^2$) footprint that is $\sim 15\%$ the size of chips available at the time of the initial decision.
- The small size of the newly available chip, combined with the fact that the additional material required by the modifications needed to mitigate the pedestal corrections increased the material budget of the ladder, results in a relatively small difference between ASIC and FPGA solutions. (At 1.3cm on a side, and 1mm thick, the extra material from the RCC FPGA corresponds to an average of $\sim 0.1\% X_0$, out of a total expected material budget of 4-5% X_0 per layer.)

RCC implementation in FPGA technology would provide substantial budget savings and would accelerate project schedule. Thus, the goal of this report is to document the expected risk of radiation induced failure in an FPGA implementation, in order to provide guidance for a decision on RCC technology choice.

Radiation Damage Mechanisms

There are three types of radiation damage to consider: non-ionizing energy loss (NIEL), total ionizing dose (TID), and single event upsets (SEU).

Non-ionizing energy loss (NIEL) is the primary damage mechanism for the silicon sensors. It is frequently quoted in terms of the flux of 1 MeV neutrons, although hadrons with energies typical of those found in RHIC collision products have similar NIEL. For the silicon sensors, which we will operate at 0°C , we believe we can tolerate an integrated NIEL corresponding to $3 \times 10^{12} \text{ n/cm}^2$ with no additional downtime, and negligible increase in noise. We believe we can tolerate an additional factor of two if we double the preamp reset frequency (at the cost of $\sim 10\%$ downtime). For CMOS circuits, damage from NIEL sets in at much higher doses, $\text{NIEL} > 1 \times 10^{15} \text{ n/cm}^2$ [3].

As we will see below, the expected NIEL for 10 years of operation at RHIC II is estimated to be $2.5\text{-}5\text{e}12$ n/cm² for layer 3 ($r = 10$ cm); higher than one might like for sensor operation, but negligible for CMOS chips.

Total Ionizing Dose (TID) causes increased current draw, and affects functionality through degradation of signal propagation time in a chip. TID effects were measured by Ichimiya *et al.* [4], for the ACTEL Axcelerator family of FPGAs for doses of 150 kRad using a gamma irradiation facility. They found a modest increase in the chip's leakage current for doses above 100 kRad, and no loss of functionality. Although, note: the functionality of the FPGAs tested in [4] is limited to ring buffers that likely do not require a substantial fraction of the chip resources (whereas we use much of the chip resources). Also, they appear to be operating the chip at a maximum frequency of 10 MHz (whereas we use speeds up to 40 MHz). As a result, their functionality tests are not testing timing degradation conditions as stringent as the RCC will require.

As we will see below, the expected TID for ten years of operation at RHIC II is 60-125 kRad for layer 3 ($r = 10$ cm), and a little more than half that for layer 4 ($r = 14$ cm). This would lead one to believe that the expected TID is likely acceptable for these chips.

Single Event Upsets (SEU) come in two flavors, hard SEUs, in which irreversible damage may be done to the chip (at the least, a power recycle is required), and soft SEUs, in which a bit is flipped. SEUs are caused by very localized energy deposits, and may be initiated by heavy fragments produced in knock-off reactions. The cross-section for SEUs, in units of cm²/bit, depends on the chip technology and the incident particle energy, but has an asymptotic value for high-energy particles. In [4], SEU cross sections were measured by illuminating the chip with $1.4\text{e}12$ 70 MeV protons/cm²: $\sigma_{\text{SEU}}(\text{soft, R-cell}) = 1.6\text{e-}14$ cm²/bit; $\sigma_{\text{SEU}}(\text{soft, memory}) < 4.9\text{e-}14$ cm²/bit. No hard upsets were observed.

As we will see below, this cross section for soft SEUs would lead to an uncomfortably high rate of SEUs in the full stripixel electronics. However, we expect this to be mitigated through triplication of logic in the chip. We will also see that the rate of hard upsets is not well-constrained. However, hard SEUs seem to require a minimum dE/dx that is much higher than what would be expected from collision products.

Dose Estimates

PHENIX measured radiation doses in the IR during the Run-6 proton+proton run at $\sqrt{s_{\text{NN}}} = 200$ GeV [5]. NIEL (equivalent 1 MeV neutron fluence) was measured with a stripixel sensor and reference diodes. TID was measured with chipmunks and TLDs. The chipmunks and TLDs were also used to measure the radial and z-dependence of the dose and the fraction of the dose received after collimation of the beams.

Post-collimation Fraction: The chipmunks measured the dose in real time. Therefore the dose per unit luminosity could be extracted. This was found to be 1.5 rad/ pb⁻¹, almost completely independent of beam conditions, after the beams had been collimated. This implies that the post-collimation dose was 18 pb⁻¹, but the total measured dose was 29.6 pb⁻¹. This implies that the post-collimation dose from other detectors (all of which supply only an integrated dose) is only 61% of the corresponding integrated dose.

At RHIC II the instantaneous luminosity will be more than two orders of magnitude higher than when these dose measurements were taken. We expect the post-collimation dose to scale with luminosity. But, we expect the remaining dose to grow much more slowly. (The remaining (pre-collimation) dose is primarily due to beam losses at the start of a fill, which should scale with the fill frequency - which should decrease as beam lifetimes increase - and the beam intensity - which will increase much more slowly than the luminosity.) The pre-collimation dose is expected to be ~5% of the post-collimation dose at RHIC II luminosities.

Radial Dependence: Chipmunks and TLDs were used to determine the radial dependence of the dose, from $r = 10$ -63 cm. Both were consistent with $r^{-1.8}$ (from the chipmunks we could see that this scaling changed to r^{-2} with the magnet off).

Beam Species/Energy Scaling: We need to scale the measured post-collimation doses to different beam species and energies (we estimate the pre-collimation doses to be less than 5%, and so ignore it). To do this we assume that the post-collimation dose is proportional to the product of the integrated luminosity (L), the interaction cross section (σ_{pp}) and the rapidity density at $y = 0$ (dN/dy_{pp}) for proton+proton collisions at the selected energy, and the average number of participant nucleons in the collision (N_{part}). Conditions of interest are shown in Table I. RHIC II luminosity projections were based on CAD guidance [6]. To get the scale factor for ten years of operation (798), we assume one 500 GeV proton+proton run, and 3 years each of 200 GeV proton+proton, light-ion and heavy-ion running.

Table I: Dose scaling factors.

PHENIX Measurement Conditions						
Species	$\sqrt{s_{NN}}$ (GeV)	dN/dy_{pp} $y=0$	N_{part}	σ_{pp} (barns)	L (nb^{-1})	Multiplier
pp	200	3.6	2	4.2e-2	1.2e4	1
Projection for 20 week RHIC II Runs						
pp	500	4	2	4.7e-2	3.3e6	341.3
pp	200	3.6	2	4.2e-2	7.0e5	58.2
AuAu	200	3.6	100	7.0	50	35.1
CuCu	200	3.6	35	3.4	500	58.9
CDF NIEL Measurements						
pp	1800	4.3	2	5.2e-2	1e6	123

NIEL Measurements/Estimates: The stripixel sensor, placed at $z = 25$ cm and $r = 10$ cm was found to have received an NIEL dose of $9.5e9$ n/cm^2 , during a delivered luminosity of 12 pb^{-1} . The reference diodes (placed at $r = 10$ cm) received (in the range $|z| < 15$ cm) an average NIEL dose of $7.8e9$ n/cm^2 . Note: doses increased for $z > 20$ cm, so the sensor dose is perhaps an overestimate. To scale this to ten years of RHIC II running, we multiply by the post-collimation fraction (0.61) and multiply by the luminosity scale factor (798), and get $4.6e12$ n/cm^2 for the stripixel estimate and $3.8e12$ n/cm^2 for the reference diode estimate.

Silicon sensors in CDF (placed at $r=3\text{cm}$, $\sqrt{s} = 1.8\text{ TeV}$) received an NIEL of $3.5\text{e}12\text{ n/cm}^2/\text{fb}^{-1}$ [7]. First we need to scale this estimate to our radius of 10 cm $(3/10)^{1.8} = 0.11$. This dose was determined with much higher luminosity, so we assume the post-collimation dose is negligible. Therefore, we need to multiply by $798/123 = 6.5$, to get an estimate of $2.6\text{e}12\text{ n/cm}^2$.

Finally, we can estimate NIEL doses from expected particle fluxes. We can calculate the expected number of charged particles in one unit of rapidity under the measured conditions (pp @ $\sqrt{s} = 200\text{ GeV}$, $L = 12\text{ pb}^{-1}$): $N = 1.8\text{e}12$. At $r = 10\text{ cm}$, the cylinder corresponding to pseudo-rapidity $|\eta| < 0.5$ has an area of 654 cm^2 , so $\Phi = 2.8\text{e}9\text{ particles/cm}^2$. If we scale up to 10 years of RHIC II, we find $2.2\text{e}12\text{ particles/cm}^2$. At an energy of 200 MeV , both pions and protons cause roughly the same NIEL damage as 1 MeV neutrons, so this should be directly comparable (and agrees reasonably well) with the other estimates.

These estimates agree with those derived in [5]. These NIEL doses are on the edge for stripixel sensor operation with the SVX4s, due to leakage current into the DC-coupled preamplifiers. Even with the design operation temperature of 0°C , the leakage current at the upper range of these estimates will require resetting the amplifiers more frequently than the abort-gap frequency, thus incurring deadtime. It is expected that other effects will be minor (larger pedestal width due to additional leakage current) or will be mitigated with on-line pedestal correction (temperature-dependent pedestals).

These NIEL doses are three orders of magnitude below those which might damage CMOS circuits [3].

TID Measurements/Estimates: The chipmunk placed at $r = 33\text{ cm}$ was found to have a post-collimation dose of 18 rad . Scaling up to $r = 10\text{ cm}$ $(33/10)^{1.8} = 8.6$, and then to 10 years of RHIC II operation, this leads to an expected TID of 124 kRad .

The TLDs at $r = 10\text{cm}$ saw $\sim 130\text{ rads}$ integrated dose. Scaling the post-collimation fraction up to 10 years of RHIC II operation, this leads to an expected TID of 63 kRad . Note: In [5] it is stated that TLDs placed adjacent to the chipmunks were consistently found to have recorded a lower dose than the chipmunks. This presumably is due to a difference in calibration, but it is not clear which of these values is correct.

As with NIEL, we can estimate the TID from expected particle fluxes. As shown above, the expected charged particle flux for 10 years of RHIC II operations at $r = 10\text{cm}$ is $\Phi = 2.2\text{e}12\text{ particles/cm}^2$. If we assume all particles are minimum ionizing ($de/dx = 1.66\text{ MeV/gm-cm}^2$, a clear lower limit) we find TID = 58 kRad .

Note, expectations at $r = 14\text{cm}$, are reduced by $1.4^{1.8} = 1.8$, $32\text{-}69\text{ kRads}$.

These expected doses are below ($r = 10\text{ cm}$), or seriously below, TIDs successfully tested in [4].

SEU Estimates: As shown above, we expect a charged particle flux of $\sim 2.2\text{e}12\text{ particles/cm}^2$. In [4] the cross section for soft SEUs was measured to be $\sigma_{\text{SEU}}(\text{soft, R-cell}) = 1.6\text{e-}14\text{ cm}^2/\text{bit}$; $\sigma_{\text{SEU}}(\text{soft, memory}) < 4.9\text{e-}14\text{ cm}^2/\text{bit}$. The number of bits in the RCC design is $\sim 1\text{e}4$, so the expected number of SEUs in one RCC, over 10 years of RHIC II operation (without mitigation, and assuming the larger of the two quoted

cross sections) is $1.1e3$. We will have 224 RCCs, so the total number of (unmitigated) SEUs would be $2.4e5$.

For the most part, such errors would be relatively harmless. The bulk of the chip area is taken up by the four 16-bit beam clock counter registers. These serve as diagnostics for data integrity, so if we simply threw out events with mismatching beam counters we would lose a tiny fraction of collected data. Another possibility is that an ADC bit gets flipped. It would be impossible to recognize corrupted ADC bits, but at this level it doesn't seem like much of a problem.

There are some other error scenarios (much less likely because there are many fewer bits involved) which would cause more trouble: for instance, an RCC may forget its configuration (and therefore, which hybrids are enabled). In such a situation the detector would need to have its serial configuration partially re-downloaded. The digitization or readout processes could also be interrupted, probably requiring a run re-start to clear. Such situations can be recognized and reported to the DAQ operator by the FEM.

ALL such effects will be greatly mitigated by triple-redundant voting logic implemented in the RCC. Basically, each bit has three copies, and the final value is set to the majority value of the three copies. In this way, two bits need to be simultaneously upset to actually have data corruption. Every bit on the chip is refreshed at 40 MHz, and the SEU rate (assuming no correlated bit flips) would be $3(R_{bf})^2/R_r$. Here R_{bf} is the bit flip rate, R_r is the refresh rate, and the equation is true when $R_r > R_{bf}$. This last condition is easily seen to be true: the maximum bit flip rate will be encountered in 500 GeV pp running ($R_{bf} = 4.9e-14 \text{ cm}^2/\text{bit} * 1.9e5 \text{ particles}/\text{cm}^2/\text{sec} = 9.3e-9 \text{ Hz}/\text{bit}$). With $R_r = 1e7 \text{ Hz}$, this leads to a negligibly small soft SEU rate.

In [4] there were no hard SEUs observed, with a total fluence of $1.4e12$ protons ($E = 70 \text{ MeV}$). We expect a total fluence of nearly twice this number of particles, and we will have 224 RCC chips. So, this failure to observe any SEUs does not, in and of itself, provide a rate limit that is meaningful to us. On the other hand, hard SEUs seem to require a minimum dE/dX that is significantly higher than deposited by a MIP. The required energy loss density discussed in the literature seems to be generally greater than $40 \text{ MeV}/\text{mg}/\text{cm}^2$, more than 20,000 times greater than the energy loss density of a MIP. Such large energy deposit densities generally require a low-energy heavy ion, and it is not known what the flux of these may be at RHIC. Unfortunately the only way to truly test this would be to place some chips in the IR. However, to get a statistically meaningful answer we would need to irradiate a large number of chips for a long time.

Note: my impression is that the chip becomes damaged over some period of time as the short-circuit, triggered by the hard SEU, burns up the chip. Some manufacturers have implemented automatic power-down when excess current is detected, in order to mitigate hard SEU effects. The RCC board has such protection built in via polyfuses on all power supplies.

Cost and Schedule Implications

If we abandon the ASIC we save several hundred thousand dollars. This is due to the fact that if we keep the ASIC we would need to incorporate final RCC design changes into the ASIC design, fabricate a prototype round of chips with the final RCC design, design and fabricate an additional round of the RCC board, and purchase the production RCC ASICs.

In addition to budgetary savings by going with the FPGA, we would be substantially farther ahead on project schedule, for the same reasons cited in the budgetary analysis above.

Conclusions/Recommendations

- The expected NIEL dose is not a concern for an FPGA implementation of the RCC.
- Chip functionality has been successfully tested [4] at TIDs up to and exceeding expected TIDs, even for the inner stripixel layer. The outer layer TID will be roughly half that of the inner layer. A test with our actual chip design may be desirable, and is certainly feasible. However, since we can only test at most a few chips, tests that we run would not be much more stringent than those already performed.
- Soft SEUs should be mitigated to negligible levels with triplication of the RCC logic.
- Meaningful limits on hard SEUs have not been established, but it is expected that such events are only caused by particles with extremely high energy loss density, whose flux at RHIC is unknown, but expected to be quite small. It does not seem possible to establish useful limits.
- If some RCCs are lost due to excessive TID, or hard SEUs, they could be replaced at a tiny fraction of the cost of the ASIC solution, although actually replacing the affected ladders could be cumbersome.

Given all this, I would recommend changing the project design to incorporate an FPGA-based RCC.

References

[1] Z. Li *et al.*, Nucl. Instrum. and Meth. **A518** (2004) 300, **A518** (2004) 738, **A535** (2004) 404, **A541** (2005) 137.

[2] L. Chrisofek *et al.*, "SVX4 Manual" (2005) http://www-ppd.fnal.gov/EEDOffice-W/ASIC_Development/Assets/publications/fermilab-tm-2318-e.pdf

[3] http://en.wikipedia.org/wiki/Radiation_hardening

[4] R. Ichimiya *et al.*, "Irradiation Tests of ROHM 0.35 um ASIC and Actel Anti-fuse FPGA for the ATLAS Muon Endcap Level-1 Trigger System", 10th Workshop on Electronics for LHC and Future Experiments (2004) http://documents.cern.ch/cgi-bin/setlink?base=cernrep&categ=Yellow_Report&id=2004-010

[5] J. Asai *et al.*, "Radiation Damage Study for PHENIX Silicon Stripixel Sensor"

[6] A.D. Frawley *et al.*, RHIC II Heavy Flavor Workshop Report, <http://rhicii-science.bnl.gov/heavy/doc/white-paper/wp-draft-jan12.pdf>

[7] J.A.J Matthews, "Effective Silicon Damage Fluences Measured by SVX and SVX'", CDF Note 3937 (1996) http://www-cdf.fnal.gov/upgrades/svxii/CDFnotes/cdf3937_cdf_part_fluence.ps.gz.