

Configuration Tests of RHBD Library with DICE, TAG4, Dual Rail, TMR and New 10T Voting Latch using UMC 65 and XFAB 180 nm

Robert L. Shuler

Abstract-- We describe heavy ion tests of RHBD configurations (varying guard timing, spacing, number of strings), including DICE, TAG4, TMR, I/O pads and a new 10T voting latch, to identify strategies for a portable nanoscale library. The TAG4, now off patent, provides superior performance to DICE. Large critical node separation requirements in bulk 65 nm are inferred. The new 10t voting latch performs as well as conventional TMR. A voting by block placement strategy demonstrates rapid reconfiguration for increased critical node group separation, and near optimal layout efficiency.

I. INTRODUCTION

THIS paper describes and analyzes heavy ion test results for several Radiation Hardened by Design (RHBD) techniques, including some not in common use but having improved performance potential, as well as reference unprotected and Triple Modular Redundant (TMR) circuits. The primary RHBD cells used are the Dual Interlocked Cell (DICE) which is widely used, and the TAG4 (containing 4 Transition nAnd Gates, also known as Guard Gates). We also include TMR with a new ten-transistor Compact Voting Latch (CVL) used in association with a modular block technique for automatic coupling of the voting circuits. These circuits have been previously described [1], but for some of them this is the first report of test results.

The TAG4 is topologically a latch version of a “bias-coupled” memory cell patented by Dooley in 1994 [2], and therefore freely available for use since 2015. Adapted by the author and made compact using an efficient transconductance mux, it has been shown previously to have improved single event performance over DICE [3]. It produces a cell about the same size as DICE, same internal power use, and only slightly more switching load due to the presence of more gates. The current effort updates the results using newer device technologies, and adds tests of timing and spatial configuration options.

The CVL is a novel 10T voting latch based on the TAG or Guard Gate, providing smaller area, faster operating speed

and lower power than conventional voting latches which typically have 22 to 24 transistors [Ibid. 1]. These are the first reported test results of the CVL, which could be a significant advantage to manufacturers of TMR-based space processors. The CVL was developed by the author while working for NASA, and is available for use as an open access design.

While one effect of this work is to certify the RHBD library and evaluate techniques for varying levels of hardness with the two process technologies tested, the main goal is to collect data that will support evaluation of techniques for developing a rapidly portable nanoscale RHBD method. Currently, extensive exploration of device physics takes about 2 years, including the design of customized layouts of key cells to improve performance. Then as many as 4 years can be required to port a design to that library and certify it. If the reconfiguration of the library can be parameterized, and the logical structure of the library can be held sufficiently constant, at least half of that time could be saved, resulting in power and speed performance at least a generation ahead. Whether it now looks like that can be achieved, and what additional tests and developments may be necessary, will be discussed in the results section.

II. TEST METHODS

Each circuit for testing was implemented as a 24 bit register and associated logic representing approximately the logic of an arithmetic and logic unit (ALU). This included an input open-drain mux, and combinational logic capable of performing add, subtract or various logic functions. This circuitry was entirely duplicated and the results compared by XNOR with a failed compare designating an error. One bit slice of this logic is shown in Figure 1. All layout uses full P and N guard bars against latch-up.

The ALU is shown as a “logic cloud” and the particular circuit shows a symbol indicating two DICE latches arranged as a D-type flip flop. Whatever flip flop is used may have variations in different experiments such as different internal spacing between latches, even in one case within a latch to separate critical node pairs (the final inverter from the rest of the latch), and for the dual interlocked cells there may be timing delay circuits for one of the dual inputs. These parameters are given in the next section.

Manuscript received July 13, 2017. (Write the date on which you submitted your paper for review.) This work was supported in by the National Aeronautics and Space Administration.

R. L. Shuler is with the National Aeronautics and Space Administration, Houston, TX 77058 USA (telephone: 281-483-5258, e-mail: robert.l.shuler@nasa.gov or robert@mc1soft.com).




Fig. 1. Bit slice of dual test circuit and compare logic

For dual interlocked cells (DICE, TAG4) the error detection logic is not perfect. The two gates shown are subject to Single Event Transients (SET) which may or may not be captured by the error reporting circuitry. For TMR circuitry, these circuits are triplicated.

The error capture and reporting circuitry is TMR. It latches the error and holds an output pin high for 8 clock cycles, to allow a test control FPGA to run at a lower clock rate than the chip under test. For XFAB 180 there is no special spacing for the error reporting TMR. For UMC 65, the three modules are separated from each other by at least 25 μm to minimize the chance of a 2-node strike causing a false error report. Corresponding critical nodes in a voting group are separated by 50 μm .

Clocks are distributed through triple redundant voting and buffering networks to TMR circuits, which at the last stage are wire-OR'd (force voted) for non-TMR circuits. For the UMC 65 cases, clock buffers maintain the same spacing between TMR elements as the error reporting circuitry. Both chips have internal ring oscillators (RO) in TMR configurations, as well as provision for external clock. The internal RO was used for 300 MHz testing of UMC 65, but otherwise external clocks were used.

I/O pad designs were also included as part of the RHBD library and these tests, rather than using vendor pads of unknown design, possibly requiring license fees. These provided an unsophisticated level shifting, and used 5x redundant force-voted heavy buffers in both directions. The proximity of the buffers, however, does not preclude an SET. It was felt that for output, the pad and trace capacitance would mitigate most SET. For input, multiple pads could have been used (were used in some of the authors prior experiments) but were not used in this case. This means that the external clock pad could be a source of SET, which if handled differently by checking pairs could result in false errors.

Beam angles were used only for the purpose of increased linear energy transfer (LET). RHBD circuits are typically sensitive to angle, and as beam time for exhaustive angle testing was not feasible, the most neutral angle available was sought, though no angle is truly neutral. The configuration chosen was to mount the test rig itself at a 45 degree angle as

in Figure 2 so that when the beam was not normal to the plane of the chip, the angle would be neither cross-row nor down-row.




Fig. 2. Test rig mounting for neutral beam angle

III. LIBRARY AND TEST CASES

A. 10-Cell Portable Library

A very small cell count library (approximately ten) is used to achieve rapid portability to new processes. The general scheme of this library and the construction of common derivative cells (latch, XOR, tri-state inverter) from the base cells is shown in Figure 3.




Fig. 3. Base cells (upper left), RHBD (upper right) and common derivatives (lower) for 10-cell library

Layout blocks for some common derivatives are provided for efficiency, but they are constructed by instantiating the base cells and hand wiring them together in a way that doesn't interfere with normal over-cell routing. P and N guard bars fully separate the two CMOS regions, and are closed off on row ends only for efficient utilization of area. The base library is shown in Figure 4.




Fig. 4. Base cell library with row end caps

For the 65nm library (shown), the design rules require extra space on account of the guard bars, which is utilized for making a greater variety of layout blocks from the base cells to improve routing efficiency. This results in good timing

predictability for the 65 nm auto-routed layout, whereas the 180 nm layout relied more heavily on the router and in cases noted below produced inconsistent timing results among test cases that should have had identical timing. The extra layout blocks are therefore deemed necessary. The extra space between P and N regions at 65 nm is not only utilized by this routing of derivative cells, but provides greater separation for diagonal ion strikes that can produce longer SETs in combination logic by upsetting adjacent cells in which logic levels are inverted.

B. Dual Interlocked Cell (DICE/TAG4) Family

DICE and similar dual cross coupled (interlocked) latches have been a staple of RHBD for at least two decades. A good deal of success has been realized in extending its life through moving transistors around to increase separation between critical node pairs [4]. More generally we re-draw the schematic as in Figure 5 to show the two sub-latches as an upper half and a lower half. Whether this is a DICE or a TAG4 (4 Guard Gates or Transition nAnd Gates) is determined by which of the inverter stacks shown in the lower half of the figure replaces each 2-input interlocking inverter.

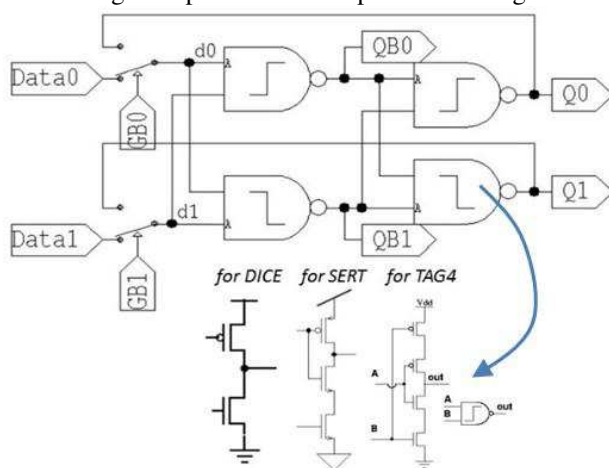


Fig. 5. Generic Dual Interlocked Architecture DICE/TAG4

Corresponding critical node pairs will always be split between the two halves. Each half is a schematic “symbol” that corresponds to a layout cell containing two gates of whichever type is desired, so that the halves are placed and routed as a unit. The default configuration was to include some spatial separation on both sides of these latch halves as shown in Figure 6.




Fig. 6. DICE (left) and TAG (center) half latch, and 2-inverter delay cell with slow gates (right)

Test cases with no space are designated with suffix $-ns$, and test cases with 2.5 times as much space are $-xs$. As the dual interlocked cells have two inputs, sometimes a guard delay is used on the 2nd input to stagger the arrival time of SETs and avoid their capture. Two dual slow inverter cells (one shown in Figure 6) provide 0.5 ns delay at 180 nm and 0.2ns delay at 65 nm. Test cases are designated with a “d” if the delay is present. It is also possible to put an entirely separate logic string on the second input, which we call “dual rail” [5] and list in the results with an “2” for two strings.

C. Triple Modular Redundancy (TMR) variations

The 180 nm chip contains a mixed route conventional TMR experiment, with all three strings thrown into a block together and no special separation. This is easier from a design flow perspective. A standard 24T voting latch was shown as in Figure 7. For the 65 nm chip the same voting latch was used for clock configuration, but placed in blocks with 50 μm separation.




Fig. 7. Conventional voting latch 22-24 transistors

The new 10T compact voting latch (CVL) was configured as the master in a D-type flip flop shown in Figure 8. For discussion of its operation see [Ibid. 1]. This was routed as three identical single-string blocks with voting signals on both left and right in a rotating fashion [Ibid. 1] so that voting was accomplished by block placement. Separation between critical nodes, which must be in corresponding positions, is determined by block width and was 25 μm for the CVL experiment. It can be changed by the simple expedient of changing block aspect ratio.

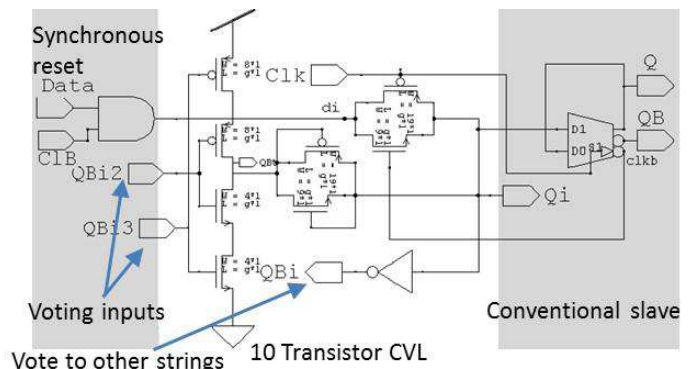


Fig. 8. Compact Voting Latch configured in D-Flip Flop

D. I/O Pads for Test purposes

It was desired not to use vendor input/output pads because they are an unknown in radiation testing and sometimes require license fees. The pads used were unsophisticated level shifting pads using transistor stacks in lieu of differential pairs, to give a high probability of working the first time. Their leakage was about 5 μ a per pad at 180 nm, which would only be important if operating below about 10 MHz.

Duplicate circuits and 4 large buffers were provided in each direction, which was thought to be sufficient to resist most SETs. This assumption likely was not true at 65 nm due to the large relative area inferred for strikes at greater LET than about 50 (due to circuit shrink and lower circuit energy). See results section for strike affected area analysis.

IV. RESULTS AND DISCUSSION

A listing of the test cases is given in Table 1. Testing was performed at the TAMU cyclotron.

Process	Circuit	Strings	Space bet FF	Space in FF	Delay
XFAB	Unprotected	1	0	0	0
XFAB	DICE	1	5 μ m	0	0
XFAB	TAG4	1	5 μ m	0	0
XFAB	dTAG4	1	5 μ m	0	0.5 ns
XFAB	d5TAG4	1	5 μ m	0	2.5 ns
XFAB	dTAG4xs	1	5 μ m	4 μ m	0.5 ns
XFAB	TAG4ns	1	0	0	0
XFAB	TMR	3	0	0	0
UMC	Unprotected	1	0	0	0
UMC	dDICE	1	0	0	0.5 ns
UMC	2TAG4	2	1.3 μ m	0	0
UMC	Clock cntrl	3	50 μ m	0	0
UMC	CVL TMR	3	24 μ m	0	0

Table 1. Experiment list

XFAB data were taken at 50 and 100 MHz, with similar data from each except that d5TAG4 and TAG4 were unable to operate at the higher speed. Apparently dTAG4 was able to operate because there is more than .5 ns variability in auto-route layout factors for the 180 nm library scheme as noted above. The same exact circuitry in dTAG4xs with less cell adjacency also operated at 100 MHz. It is also possible that adjacency of the latches reduces speed slightly, something the author has noticed in unpublished studies of interleaved inverter delay strings. The advantage in speed though, is probably what made it slightly more susceptible to SET capture. Error rates are plotted in Figure 9, with dTAG4xs

and dTAG4ns omitted for plot clarity. These are normalized by total fluence of 1×10^7 and also by 48 bit slices per experiment (for all experiments, both chips).

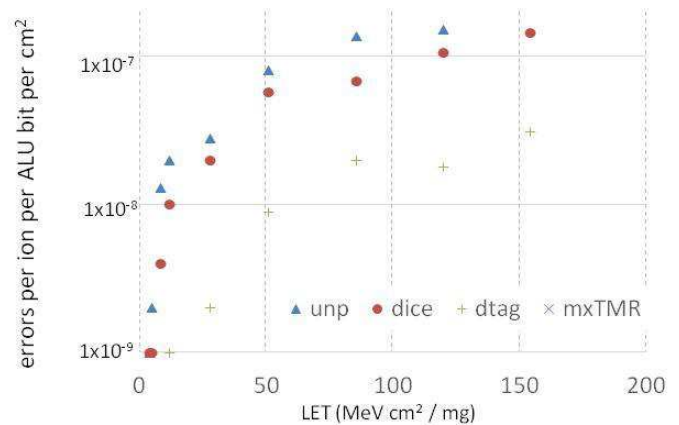


Fig. 9. XFAB 180 nm combined 50/100 MHz results

Unfortunately every experiment is not represented in this figure, but we can make three general conclusions:

1. The DICE (without delay guard) is hardly worth more than an unprotected (unp) latch in this process.
2. The TAG4 with 0.5 ns guard delay runs still at 100 MHz with significant combinational logic, and represents nearly an order of magnitude improvement in cross section and significant improvement in threshold.
3. The mixed-route TMR suffered no errors, which we would not expect to hold up generally, but is consistent with previous results at 350 nm that show it having nearly a further order of magnitude improvement over delay-based protection.

UMC 65 nm data were taken at 200 and 300 MHz, with combined data in Figure 10. The 200 MHz clock was external, through an I/O pad that though heavily buffered could be a source of error. The 300 MHz clock was from the internal ring oscillator, with frequency controlled by a 3-bit register/counter using TMR with spacing as previously described. An error in this circuitry implies most likely charge deposition over an extended distance, at least 25 μ m and probably 50 μ m (to affect corresponding nodes). It was not planned to take data on this circuit, but a scope was used to monitor clock frequency on a toggle bit, and when a change in frequency was noticed the data was captured from that point, and included in Figure 10.

Several general comparisons with the XFAB 180 nm data are evident:

1. The overall error rate is half an order of magnitude lower, presumably because of the smaller area of 65 nm circuits and consequent reduced number of strikes on each circuit. If errors were normalized by chip error, we'd find similar

or greater error counts as the 65 nm circuitry is over 3 times denser.

2. The DICE with delay has a lower cross section at LET below 50, but not a lower threshold. At higher LET, presumably because the affected area is great enough to upset a critical node pair, the DICE is barely better than unprotected latches.
3. The dual rail 2TAG4 performed very well as it has done in the past, only slightly worse than TMR. It would seem to be an efficient point for many applications. Unfortunately a quick procurement schedule for this chip required using a previously routed and verified design that had been pin out constrained to a small number of experiments, so no data on a dTAG4 circuit is available.
4. The data don't exactly fit a Weibull curve. In the case of the DICE we've stated this is because of a second mechanism, an area large enough to strike a pair of critical nodes. For the high reliability circuits it is probably an angle effect. Notice that the 120 and 154 LET points are at an angle, as shown in the beam data in Table 2.

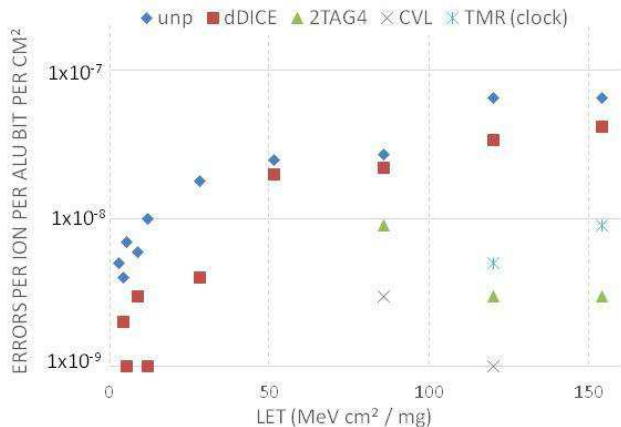


Fig. 10. UMC 65 nm combined 200/300 MHz results

LET	ion	angle	fluence
2.7	Ne	0°	1x10 ⁷
3.8	Ne	45°	1x10 ⁷
4.9	Ne	56°	1x10 ⁷
8.3	Ar	0°	1x10 ⁷
11.7	Ar	45°	1x10 ⁷
28	Kr	0°	1x10 ⁷
51	Kr	56°	1x10 ⁷
85.5	Au	0°	1x10 ⁷
120	Au	45°	1x10 ⁷
154	Au	56°	1x10 ⁷

Table 2. Ion beam, angle and fluence

Two conjectures are presented as an explanation of the CVL TMR upsets at an LET of 85.5.

All of the angled beams were at 45 degrees to the chip rows as noted above. The angles in Table 1 are with respect to the chip surface, with normal incidence being zero degrees. In the 2TAG4 experiment, pairs of critical nodes were mostly at either 0 or 90 degrees with respect to row position, so 45 degrees being in between might be less likely than a normal incidence of the same ion to affect both nodes. Thus the error rate consistently dropped for angled beams (the test was re-run several times, with data from the first run used for official results).

A similar explanation is postulated for the block separated CVL TMR experiment, shown in Figure 11.

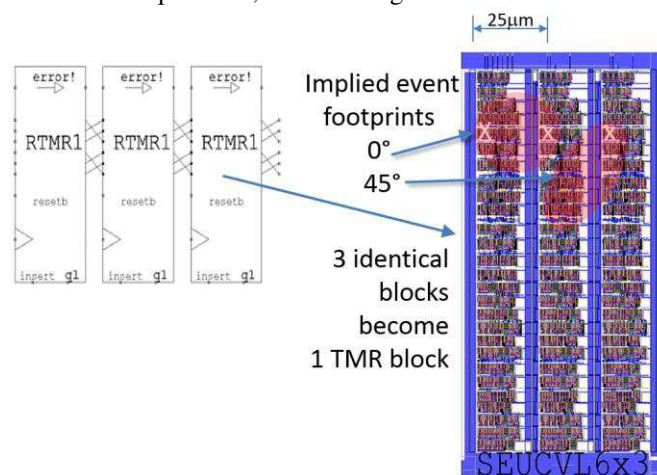


Fig. 11. CVL TMR block ion strike footprint inference

In the figure a white “X” marks locations of members of a critical node group. The red circle shows the possible pattern of a normal incidence gold ion with LET of 85.5 encompassing both nodes. Ion placement is critical, and in fact the first run was the only one with more than one strike, and many re-runs shown no upsets. But the first run data, also worst case, was used for official results. The elliptical red spot illustrates the estimated footprint of critical charge for the same ion at an angle to the rows. The width of the elliptical footprint is less, and from the data appears to have been less than 25 µm.

How then did the clock configuration with 50 mm separation become affected? This could have been due to:

1. SET on the single I/O pad driving the configuration register for the on-board ring oscillator. It had an asynchronous input which is very sensitive to SETs. Previous designs used 3 input pins for such clocks, but as previously mentioned, this layout block was pin out constrained.
2. The elliptical footprint could have reached across the greater separation of the mixed route TMR circuitry, and though it shouldn't have touched a corresponding node, there are numerous gates driving the voting nodes in the larger conventional voting circuit, any of which could be targets.

The second conjecture for the CVL TMR upsets is that they are due to upsets on the external clock I/O pad, because there were no upsets when using the internal 300 MHz ring oscillator (despite the fact that it was occasionally changing frequency, although as long as it did not exceed its top frequency, which was 300 MHz, that is the expected situation). Figure 12 shows the 300 MHz results for 65 nm. The dual rail 2TAG also performs better when restricted to the internal clock, despite the higher frequency, and at low LET the dDICE performs extremely well on cross section, presumably with the delay filtering out most SETs. (By definition, any error counts appearing in the combined result that don't appear here appeared in the 200 MHz results with the external clock.)

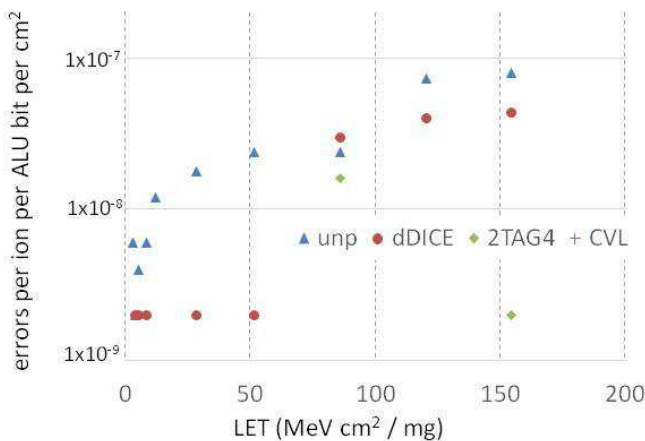


Fig. 12. UMC 65 nm results at 300 MHz internal clock only

V. CONCLUSIONS

The data show the CVL achieved its goal of being equivalent to TMR, and earlier analysis showed it to achieve layout efficiency goals [Ibid. 1]. It is therefore the best TMR method currently available.

The TAG4 consistently showed threshold and cross section improvements over DICE that are modest, but recommend its usage. It is known to be almost as good as the SERT, which is still under patent [Ibid. 3]. It has been reported that nodal separation in DICE at 90 nm offers an “order of magnitude decrease in upset cross section” [6], but our 65 nm results suggest that situation deteriorates as geometry shrinks. In addition to the direct shrinkage of circuits, there is an effect of decreasing critical charge for latch upset (Q-crit) in smaller geometry processes [7].

Future work is recommended on CVL/TMR spacing at newer process nodes and in FDSOI, whether a high performing single rail version of the TAG4 might be developed.

Currently it is not feasible to route dual rail designs in separate blocks per each rail, because 8 rather than 3 cross connections are required. Future work will also look into using the new 10T voting latch to devise a method of cross

coupling a dual rail design with only 3 connections, but there are many obstacles.

It has been our purpose to show that radiation tolerant designs can be developed primarily on a geometry rather than device physics basis, and ported rapidly with efficiency great enough for use in low volume deep space designs. The time from first access to a new process technology for small run research chips, to production of a space part, might be reduced from 6 years to 2. For the highest reliability designs, configuration based simply on block aspect ratio compares favorably in non-recurring costs to more customized node-separation strategies such as LEAP [8]. All the technology described in this paper is either developed as open access technology, or is off patent.

VI. REFERENCES

- [1] Shuler, R.L., “Porting & Scaling Strategies for Nanoscale CMOS RHBD,” *IEEE Trans. on Cir. and Sys.*, Vol. 62, No. 12, pp 2856-2863, Nov. 2015, DOI: 10.1109/TCSI.2015.2495779
- [2] Dooley, J.G., “SEU-Immune Latch for Gate Array, Standard Cell, and Other ASIC Applications,” U.S. Patent No. 5,311,070, May, 1994.
- [3] Shuler, R.L.; Balasubramanian, A.; Narasimham, B.; Bhuvan, B.L.; Neill, P.M.O.; Kouba, C., “The Effectiveness of TAG or Guard-Gates in SET Suppression Using Delay and Dual-Rail Configurations at 0.35 um,” *IEEE Trans. Nucl. Sci.*, vol.53, no.6, pp.3428,3431, Dec. 2006
- [4] Lee, H.K., Lilja, K., Bounasser, M., Relangi, P., Linscott, I.R., Inan, U.S., Mitra, S., “LEAP: Layout Design through Error-Aware Transistor Positioning for Soft-Error Resilient Sequential Cell Design,” 2010 *IEEE International Reliability Physics Symposium*, Anaheim, CA, 2010, pp. 203-212.
- [5] Shuler, R.L.; Bhuvan, B.L.; O’Neill, P.M.; Gambles, J.W.; Rezgui, S., “Comparison of Dual-Rail and TMR Logic Cost Effectiveness and Suitability for FPGAs With Reconfigurable SEU Tolerance,” *EEE Trans. Nucl. Sci.*, vol.56, no.1, pp.214,219, Feb. 2009
- [6] Amusan, O. A., Massengill, L. W., Baze, M. P., Sternberg, A. L., Witulski, A. F., Bhuvan, B. L., and Black, J. D., “Single event upsets in deep-submicrometer technologies due to charge sharing,” *IEEE Trans. Device Mater. Rel.*, vol. 8, no. 3, pp. 582-589, Sep. 2008
- [7] Chandra, V. and Aitken, R., “Impact of Technology and Voltage Scaling on the Soft Error Susceptibility in Nanoscale CMOS,” 2008 *IEEE International Symposium on Defect and Fault Tolerance of VLSI Systems*, Boston, MA, 2008, pp. 114-122. doi: 10.1109/DFT.2008.50
- [8] Lee, H.K., Lilja, K., Bounasser, M., Relangi, P., Linscott, I.R., Inan, U.S., Mitra, S., “LEAP: Layout Design through Error-Aware Transistor Positioning for Soft-Error Resilient Sequential Cell Design,” 2010 *IEEE International Reliability Physics Symposium*, Anaheim, CA, 2010, pp. 203-212.