Design of a Novel Soft Error Mitigation Technique for Reconfigurable Architectures

S. Baloch[i]
i): School of Electronics & Engineering
University of Edinburgh, Mayfield Road, Kings Buildings
Edinburgh EH9 3JL, UK
+44 131 650 5592
Sajid.Baloch@sli-institute.ac.uk

T. Arslan[i,ii]
ii): Institute for System Level Integration
The Alba Centre, Alba Campus
Livingston, EH54 7EG, UK
+44 131 650 5592
Tughrul.Arslan@ee.ed.ac.uk

A. Stoica[iii]
iii): Jet Propulsion Laboratory, NASA
4800 Oak Grove Drive Pasadena
CA 91109, USA
+1 818 354 2190
Adrian.stoica@jpl.nasa.gov

Abstract—Commercially off the shelf (COTS) available reconfigurable architectures are becoming popular for applications where high dependability, performance and low costs are mandatory constraints such as space applications. We present a unique SEE (single Event Effect) mitigation technique based upon Temporal Data Sampling and Weighted Voting for synchronous circuits and configuration bit storage for reconfigurable Architectures. The design technique addresses both conventional static SEUs (Single Event Upsets) and SETs (Single Event Transients) induced errors which result in data loss for reconfigurable architectures in space. The design technique not only eliminates all the single event upsets and single event transients but eliminates all double event upset as well.

TABLE OF CONTENTS

1. INTRODUCTION 1
2. SEU MECHANISM 1
2. PREVIOUS WORK 3
3. PROPOSED MITIGATION TECHNIQUE 3
4. TEMPORAL SAMPLING 4
5. WEIGHTED VOTING CIRCUITRY 5
6. CASE EXAMPLE 6
7. RESULTS 6
8. CONCLUSION 7
REFERENCES 8
BIOGRAPHY 8

1. INTRODUCTION

Single event upset (SEU) is defined by NASA as "radiation-induced errors in microelectronic circuits caused when charged particles (usually from the radiation belts or from cosmic rays) lose energy by ionizing the medium through which they pass, leaving behind a wake of electron-hole pairs."[1]. SEUs are soft errors, and are non-destructive. An SEU may occur in analog, digital, optical components, or may have effects in surrounding interface circuitry.

Reconfigurable architectures are becoming increasingly popular with space related design engineers as they are inherently flexible to meet multiple requirements and offer significant performance and cost savings for critical applications. As the microelectronics industry has advanced, Integrated Circuit (IC) design in general and reconfigurable architectures (FPGAs, reconfigurable System on Chip (SoC) etc) in particular have experienced dramatic increase in density and speed due to decrease in feature sizes with which these devices are manufactured. These advancements have serious implications for the reconfigurable architectures when used in space environment where IC is subject to total ionization dose and single event effects as well. Out of these effects, single event effects are most difficult to avoid in space-borne reconfigurable architectures.

2. SEU MECHANISM

The effects of scaling on the single event response of microelectronics are a direct result of the physics of energy loss, charge collection, and upset due to a cosmic ray striking a junction in an IC device. The review here is brief and qualitative. Many good summaries exist [2], [3], [4] that review these concepts in more detail.

When an energetic ion passes through any material it loses energy through interactions with the bound electrons, causing an ionization of the material and the formation of a dense track of electron-hole pairs. The rate at which the ion loses energy is the stopping power (dE/dx). The incremental energy dE is usually measured in units of MeV while the material thickness is usually measured as a mass thickness in units of mg/cm2. The radiation effects community has adopted the term LET (Linear Energy Transfer) for the stopping power. An ion with an LET of 100 MeV-cm2/mg deposits approximately 1pC of electron-hole pairs along each micron of its track through silicon.

[1]0-7803-9546-8/06/$20.00© 2006 IEEE
[2]IEEEAC paper #1082, Version 4, Updated October 26, 2005

In the presence of electric fields, these electron-hole pairs quickly separate as they drift in opposite directions in the field and are quickly collected by whatever voltage sources are responsible for the field, thus producing a current transient. In bulk CMOS designs, such electric fields are present across every pn junction in the device. If an ion strikes a junction connected to a signal node, a current transient is subsequently observed on the signal node as the electric fields in the junction and funnel regions separate the electron and hole carriers[10]. The initial prompt current pulse is short lived, lasting on the order of only 100 to 200 picoseconds [11].

High energy protons and neutrons are also known to produce similar effects indirectly through nuclear reactions within the silicon. In these cases, a heavy ion recoil reaction by-product passes through a junction and produces a similar charge collection current pulse. In space, high energy protons primarily originate from the trapped proton radiation belts and from solar flares. For high-altitude aircraft, both high energy neutrons and protons are encountered as reaction by-products found in cosmic ray showers formed when an energetic heavy ion from space undergoes a nuclear reaction in the atmosphere. These induced currents are responsible for SEUs observed in space-borne circuits, typically static latches and SRAMs (Static Random Access Memories), over the last 10-15 years [9][10]. The effect that these currents have on a circuit depends on the response of the circuit to the charge collected on the signal node. Basically, the capacitance of the signal node (to first order) determines how large a voltage swing dV results from the collection of a charge dQ according to $dV=dQ/C$. For latches and SRAMs, positive gain feedback loops result into a data bit flip once the collected charge reaches a critical value sufficient to drive a node voltage past the switching voltage.

SEU in static latches and SRAMs became an important issue once feature sizes dropped below 10 microns and the critical charge for upsetting a circuit dropped below 1 pC (roughly corresponding to a particle LET of 50MeV cm2/mg and a collection depth of 2 microns). Static latch SEU vulnerability has been calculated [3] and measured [6] as a function of technology feature size to establish the relationship between the critical charge, needed to upset the circuit and the technology feature size. All results indicate that the critical charge needed to upset a latch decreases as the square of the feature size. If this relation holds as electronics feature sizes decrease from 0.8 micron (space-borne architectures) to 0.18 micron, the critical charge decreases by nearly a factor of 20 [10].

The FPGA's configuration bits (Bit-stream) are used to configure both the logic elements and the routing switches. Upset of a programming bit in a FPGA is much more serious than a conventional data bit upset. If a logic element control bit changes state, then the logic functionality of the FPGA is altered. If a control-bit of a routing switch experiences an upset, then the FPGA essentially becomes rewired. In either case, the programmed circuit function is no longer what was intended. For these reasons, the configuration storage of the FPGA must be totally immune to SEU.

Radiation which results into Soft Error (SEU and SET) can cause permanent faults in reconfigurable architectures, for example when it hits configuration memory of the architecture. The configuration memory contents control the overall functionality of architecture through routing of a signal as shown in Figure 1.

Figure 1: SEU in a signal routing path

This bit-flip due to SEU or SET remains effective till the time bit-stream (configuration data) is not reloaded (scrubbing technique) or not corrected by dedicated hardware. The results from bit-stream fault injection [5] and radiation ground testing [6] have confirmed the efficacy of Triple Modular Redundancy (TMR) structure combined with scrubbing, to recover upsets in reconfigurable architectures. However, the TMR with scrubbing technique has some limitations, such as an area overhead, does not cover SET faults and voting circuit faults. Figure 2 illustrates the general TMR scheme.

In this work, we propose a novel design technique to cope with both SEU and SET faults. The design technique is based upon unique temporal data sampling with weighted voting. This design technique not only gives 100% fault recovery from SEU but also gives 100% fault recovery from SET, dual event faults and 50% recovery from triple event faults. The technique has an auto correction mechanism and does not need scrubbing; hence, it improves the overall performance and speed of the system.

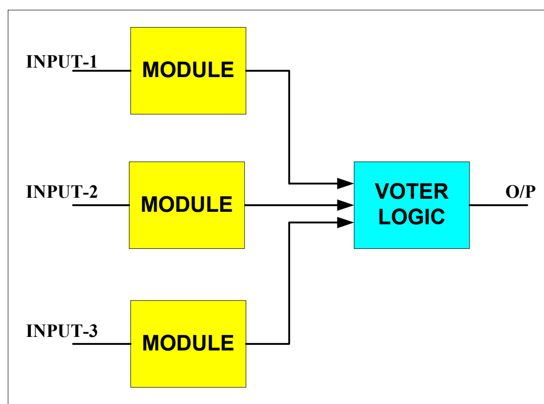

Figure 2: General TMR Scheme

2. Previous Work

Several SEU mitigation techniques have been proposed in the past few years to avoid faults in reconfigurable architectures. SEU hardened circuits may be accomplished through a variety of mitigation techniques based on redundancy. Redundancy is provided through hardware redundancy (extra components), time redundancy (extra execution time) or through a combination of both. Time and hardware redundancy techniques are largely used in ASIC [7]. Full hardware or time redundancy allows the voting circuits to vote for correct value in presence of single event upsets. The Time redundancy takes the advantage of the transient pulse characteristic to compare the signal's value at different intervals of time.

In the case of reconfigurable architectures, the problem of finding an efficient technique in terms of area, performance and power is very challenging due to high complexity of the architectures. An SEU is classified as a soft error but has permanent effects on a reconfigurable architecture, for example, the SEU can effect routing of a signal by effecting configuration memory and user's synchronous circuits by effecting the memory elements and combinational circuits. The consequences of SEU cannot be handled through standard ASIC fault tolerant schemes such as Error Correction and Detection Codes (EDAC) or standard TMR with single voter circuits because a fault in encoder/decoder or in voter circuit will invalidate the technique. TMR is an attractive scheme for reconfigurable architectures because it provides full hardware redundancy including user's combinational, sequential circuits, the routing and IO pads. However, TMR has limitations like area overhead, IO pad limitations, power dissipation and faults in voter circuitry itself.

In this paper, we present a new high level concurrent SEU and SET mitigation technique that is based on temporal data Sampling with Weighted Voting. This scheme copes with permanent effects of SEU and SET in the programmable matrix. In addition, this scheme copes with double-event faults and 50% triple-event upsets as well, which gives an edge to the proposed scheme over all other schemes.

3. Proposed Mitigation Technique

The transitions induced by cosmic rays have a pulse-width of 100 to 200 picoseconds [8]. As the feature size has dropped below 0.35 microns, SETs are no longer attenuated within the gates of the circuits and propagate as normal signals. This has serious implications for sequential and configuration bit storage circuits of reconfigurable architectures.

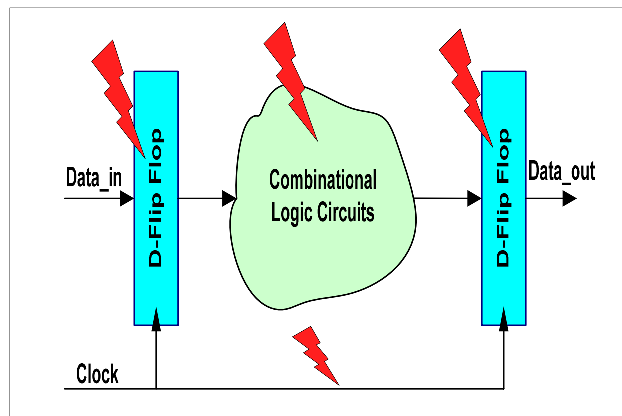

Figure 3: Typical Sequential Circuit Topology and SEU

While conventional SEU error rates are independent of the chip clock frequency, SET induced error rates increase in direct proportion to the operating frequency. This error rate relation actually compounds the SET problem as technology feature sizes shrink since smaller feature sizes result in smaller gate delays that permit circuits to be operated at higher clock frequencies. Not only does each combinatorial gate in the circuit contribute to the SET error rate (because transients are no longer attenuated), but the probability of storing any given transient error will also increase (because of the higher clock frequencies). Shrinking device sizes have had serious, if not grave, implications for FPGA devices used in cosmic-ray environments. We demonstrate that, for typical FPGA designs, SET induced error rates may actually exceed the SEU rates of unhardened latches as clock speeds approach 100 MHz for CMOS designs [7].

The typical circuit topology found in nearly all sequential and configuration bit storage circuits is shown in Figure 3.

If a heavy ion strike occurs within the combinatorial logic block of a sequential circuit, and the logic is fast enough to propagate the induced transient, then the SET will eventually appear at the input of a data latch where it may be interpreted as a valid signal. Similar invalid transient data might appear at the outputs of lookup tables (LUTs) and on routing lines due to SETs generated in the

programming elements, particularly for the case of EEPROM storage elements. Whether or not the resulting SET gets stored as real data depends on the temporal relationship between its arrival time and the falling edge of the clock. Similarly, SETs on clock, reset, and control lines can result in the storage of incorrect data within the data latches.

As shown in the Figure 3, the data is typically latched on the falling edge of the clock and then released to the combinational circuit. The output of the combinational circuit reaches second latch sometime before the next falling edge of the clock. At this edge, whatever data is present will be latched.

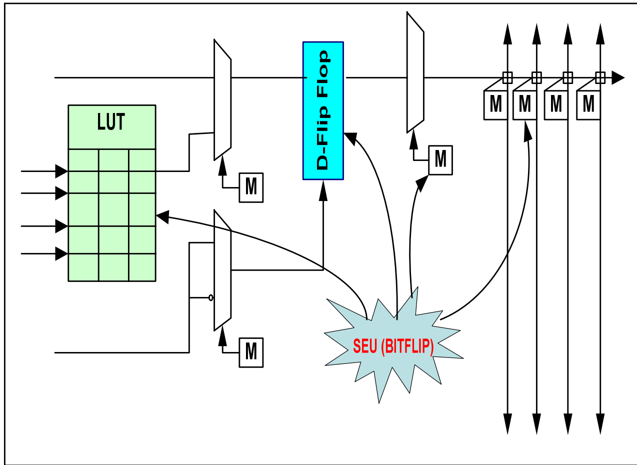

Figure 4: SEU Sensitive Configuration Bit Storage Circuits

Non-volatile, reprogrammable FPGAs typically use EEPROM transistors to store the configuration data. We describe how SET pulse broadening occurs in devices that use EEPROM transistors in an active fashion. The programmed threshold voltage of an EEPROM transistor is simply not large enough to provide charge dissipation drives comparable to normal CMOS transistors. It results in relatively wide SETs that are shown to propagate in FPGAs fabricated in technologies having feature sizes as large as 0.8 micron. The SET induced error rates in these cases can actually exceed SEU rates at frequencies of only a few tens of MHz.

SEU in the Configuration Bit storage circuits can result into routing of the signal to an undesired location and causing the whole reconfigurable system to crash as shown in Fig-1, 4. In order to overcome the problems, described above and limitations of the schemes introduced so far, we present a mitigation technique which is based on Temporal Data Sampling with weighted voting.

4. TEMPORAL SAMPLING

The first key step in our newly proposed technique is 'Temporal Data Sampling'. This sampling technique eliminates all SETs in the clock and data signal. A simple embodiment of the Temporal Data Sampling is shown in the Figure 3.

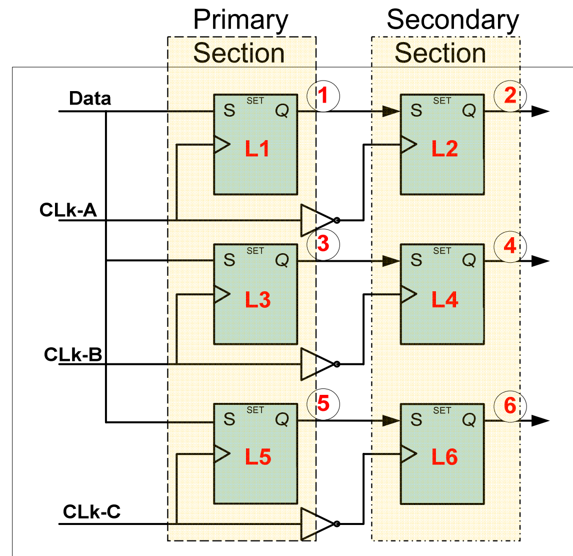

Figure 5: Temporal Data Sampling

The circuit consists of six level sensitive latches (3 at the input stage and 3 at the output stage) as illustrated in Figure 5. Each level sensitive latch is in sampling mode at the high level of its clock signal and is in blocking mode when its clock signal is low. In blocking mode the latch holds the data and the data changes at the input are blocked. In sampling mode the latch is transparent to the data. The set of three latches (L1, L2), (L3, L4), (L5, L6) operate in parallel and form the temporal sampling stage of the technique. Each set of level triggered latches constitute a negative edge triggered Flip Flop.

The temporal sampling stage stores data samples at different time intervals. These samples are used in weighted voting logic to eliminate single event upsets. Three different clocks (Clk-A, Clk-B, Clk-C) are used. These three clocks are derivative of the main clock and have a phase shift and 25% duty cycle to cope with the SETs, as shown in Figure 6. If an SEU is observed on any one of the clock lines, the phase shift in the remaining clock signals will help the respective set of latches to store the correct data at different time intervals, hence voiding the effect of a spurious glitch on the clock line due to radiation. Any transients due to radiation last for a small period of time and if happens at the negative edge of any one of the clock signals, then it will die out before the other temporal latches start their operation due to phase-shift in clock signals. So this clocking scheme will help cope with all the SETs either in the data line or in any one of the clock signals.

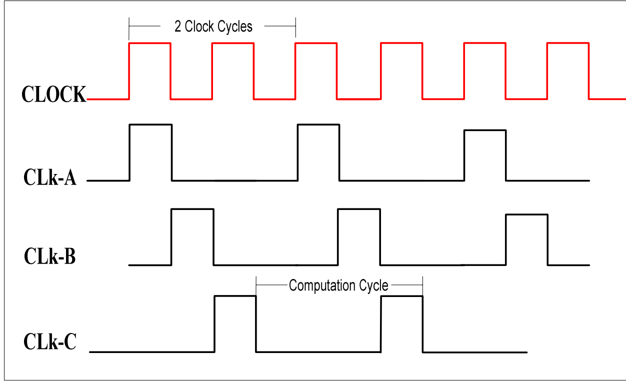

Figure 6: Clocking Scheme for the proposed Architecture

A conventional (SEU susceptible) sequential circuit (as shown in Figure 3) would satisfy timing constraints such that the maximum combinatorial logic transit time would be less that the period of the master clock minus setup time for the D-Flip-Flop.

In our proposed technique, data is released on the falling edge of Clk-C and must reach the next sampling stage before the falling edge of ClkA (minus the setup time). From Figure 6, this is just the period of the master clock minus the D-Flip-Flop setup time. Thus a conventional circuit that satisfies the timing constraints for the master clock will automatically satisfy the timing constraints of our new clocking scheme.

The insertion of the two extra clock phases (Clk-B and Clk-C) is needed for the additional temporal sampling. In Figure 6 we show these phases as occupying one additional master clock period. In this case the effective on-chip computational frequency is exactly one half the frequency of the master clock. Therefore a speed penalty of a factor two has been incurred to ensure upset immunity.

Normally, a conventional D-Flip-Flop can be constructed from two level sensitive latches. One might expect the temporal sampling latch to occupy more IC area than a conventional D-Flip-Flop based techniques (TMR, etc). For any given ASIC design, however, the total chip area will never grow by a factor of five. There are three reasons for this, which are explained below.

All of the combinatorial logic can remain unchanged and only the D-Flip-Flops in the design must grow. The amount that any given ASIC must grow will depend on the ratio of D-Flip-Flop area to total chip area. For example, fan anti-jam filter, has 7.5% of its area devoted to D-Flip-Flops [4] and an 8-bit ALU, contains addition, subtraction, bit shift, and bit logic blocks. The fraction of this chip devoted to D-Flip-Flops is 7.2% [8]. The total size penalty, expressed as an area increase factor, for using temporal sampling latches in these chips would therefore depend only on the number of latches. In comparison to the general TMR scheme for synchronous circuits as shown in Figure 2, the area penalty is negligible because D-Flip-Flop is conventionally constructed from two level triggered latches (same as our proposed technique).

5. Weighted Voting Circuitry

Majority voting is commonly used in TMR systems. The proposed mitigation technique is based upon unique 'weighted majority voting'. Each node has been assigned a voting weight based on a simple rule which is "voting weight of a node is inversely proportional to the probability of that node being disturbed by the radiation". The each output node (2, 4, 6) has two memory/storage elements in its data path, so the probability of being hit by radiations is double than that for the output node (1, 3, 5), which has only one memory/storage element in its data path. The output nodes (1, 3, 5) of the latches (L1, L3, L5) shown in Fig-5, has twice the voting weight than the output nodes (2, 4, 6) of the latches (L2, L4, L6). If P(n) is the probability of node n, being hit by SET/SEU then the voting weight of node n, VW(n) is defined as:

$$VW(n) = 1/P(n) \qquad (1)$$

Let's assume that the probability of nodes (2, 4, 6) is 1, which is double than the probability of nodes (1, 3, 5). The voting weight of nodes (1, 3, 5) is calculated as 2 through equation-1 and the voting weight of nodes (2, 4, 6) turns out to 1. Figure 7 shows the process of weighted majority voting. All outputs from the nodes (1, 2, 3, 4, 5, 6) are fed to three different 2/3 simple majority voters. The total possible voting weight for any logic value can be 9. The output of each 2/3 simple majority voter is fed into the correction unit. The correction unit checks the output and the input to correct voter circuit faults and to override in case of fault. These corrected outputs are fed into the last majority voter circuit to resolve the final output. Final output is decided on the basis of a 2/3 simple majority. The majority voter circuit gives one additional output as recovery. Recovery output is only asserted when more than one SEU/SET fault is detected and corrected. The recovery output may request the main microprocessor/DSP of the reconfigurable architecture for scrubbing. As the scrubbing request is only sent when more than one SEUs are detected and corrected, this improves the system performance as scrubbing frequency is halved than previously introduced SEU mitigation schemes [7][8][9].

The weighted voter circuitry has the capability to recover from internal faults. This unique feature enhances the reliability of the proposed scheme. Simulation results are shown in Figure 8. In case of an SEU, the down load process of bit-stream is no longer required due to auto correction feature of the proposed technique. The auto correction mechanism is activated through recovery command. This feature makes the design very flexible and

enhances the overall system performance. The complete process of SEU/SET mitigation is shown in Figure 7.

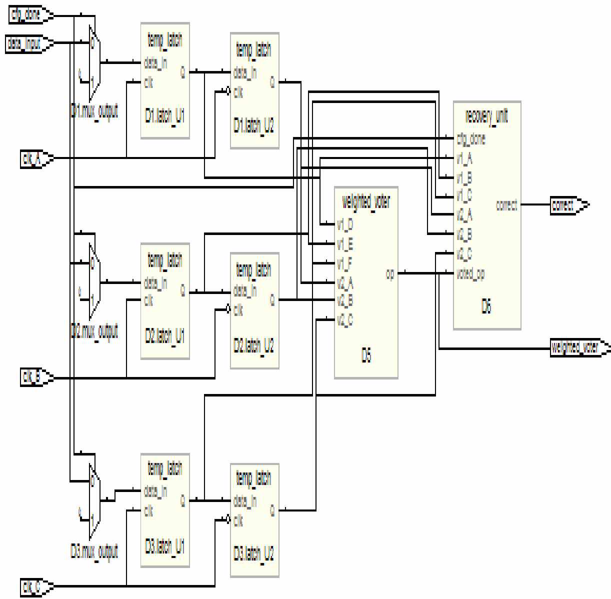

Figure 7: Proposed SEU/SET Mitigation Technique with Self Correction Mechanism

6. Case Example

Let's assume that the input data is 1 which is at the input of all the latches L1, L2, L3. At the negative edge of the relevant clock signal the same data will appear at node 2, 4, 6. In the following section, we will present different scenarios of fault recovery by keeping the initial conditions the same as described above:

a) Single Fault Recovery

A single fault can occur at any node. Let's assume that the node 2 (Figure-5) gets disturbed by the radiation and the data-value flips from logic '1' to '0'. Total votes are calculated on the basis of weights assigned to each node. Table-1 explains the mechanism. It can be seen that 8/9 votes for logic '1' eliminates the SEU fault.

b) Double Fault Recovery

Double fault can occur at any two nodes. Different combinations of double faults are radiation hitting two nodes, both having same weights (2 or 1) or one node with weight equal to 1 and the other with weight equal to 2. Let's assume that both nodes (node-3, node-5) with weight 2 are hit by radiation and theirvalues are changed. Our proposed scheme will vote for the correct value even if two nodes are disturbed, as illustrated by the table below.

Table 1: Single Fault Recovery

Node	Before SEU	After SEU	Voting Weights '1'	'0'	Total Votes For '1'	'0'
1	1	1	2	-	8	1
3	1	1	2	-		
5	1	1	2	-		
2	1	0	-	1		
4	1	1	1	-		
6	1	1	1	-		

Table 2: Double Faults Recovery

Node	Before SEU	After SEU	Voting Weights '1'	'0'	Total Votes For '1'	'0'
1	1	1	2	-	5	4
3	1	0	-	2		
5	1	0	-	2		
2	1	0	1	-		
4	1	1	1	-		
6	1	1	1	-		

The proposed technique is coded in C-program, which takes the net-list of the circuit under test as input. D-flip-flops are identified and structural modifications are made to the original circuit. The modified net-list is then fed into a software simulator to analyse the behavior of the circuit under SEUs/SETs. The net-list is can be mapped on any reconfigurable architecture (FPGA etc) through a suitable software tool (Xilinx Foundation Tool etc).

7. Results

The Experimental results are derived for a medium sized circuit ISCAS89 benchmark circuit S386 contains 12 Flip-Flops. We injected 100 random SEU, SET faults, double-event upsets (DEU) and triple-event upsets (TEU) to verify our proposed scheme. Faults were injected through software simulations. Simulation results are shown in Figure 10. Ref [8] has the limitation that scrubbing is required before the second fault occurs which limits the overall system performance. Ref [9] is based on a temporal latch scheme. It uses a simple majority voter circuit which does not cope with multiple faults and the whole scheme becomes void in case of faults in the majority voter circuitry itself.

Lima, F. et al.[8] proposed an architecture which is based on a "dice cell". The scheme uses four clock signals to take samples in contrast to our proposed scheme which takes only three. The reference scheme uses 9 memory elements and gives Single event Upset/Transient immunity while our

proposed scheme uses less memory elements (only 6 memory elements) and is immune to all SEU/SET and double event faults as well. Mavis, D. et al [9] proposed SEU technique based on 9 latches and 4 clock signals. The two techniques explained in ref [8] and ref [9] are very much similar apart from "dice cell"

Table 3: Fault Tolerance Comparisons

Mitigation	SET	SEU	DEU	TEU
Proposed Scheme	100%	100%	100%	50%
Ref [4] Scheme	63%	100%	-	-
Ref [5] Scheme	100	100%	32%	18%

There was no area or power analysis done for previous architecture to compare with the proposed scheme. So the authors implemented the proposed technique and Lima, F. et al's technique on 0.13 um CMOS technology. Figure-8 illustrates the area calculation. An area saving of approximately 40~45% is achieved for the circuit under test. Figure-9 shows approx. 42% Power saving through the proposed technique over the previously introduced techniques.

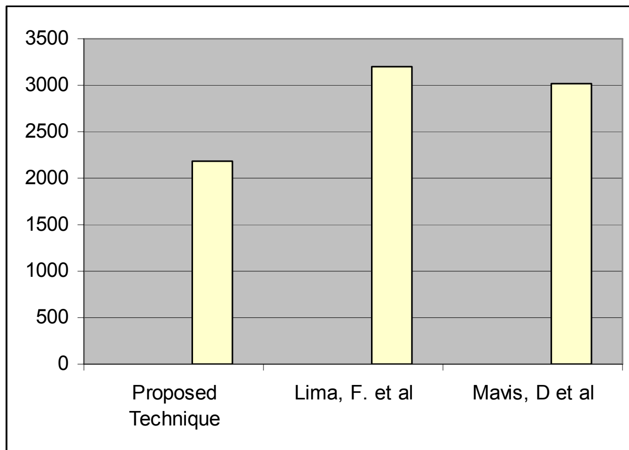

Figure 8: Simulation Results

The result clearly shows that our proposed scheme is much more efficient than other techniques. The scheme gives 100% fault recovery not only from all single event faults but also from double event faults.

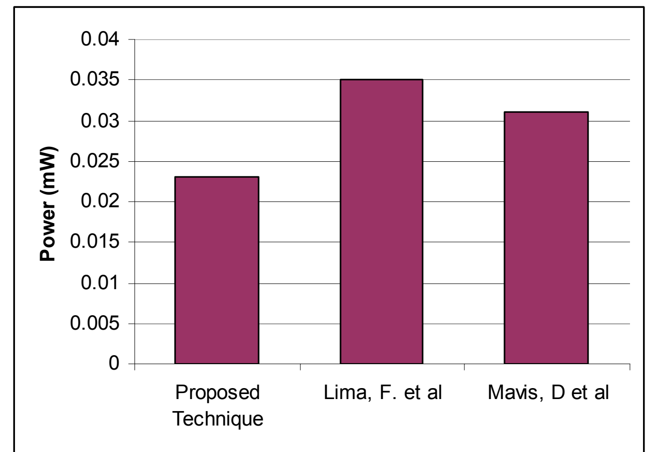

Figure 9: Simulation Results

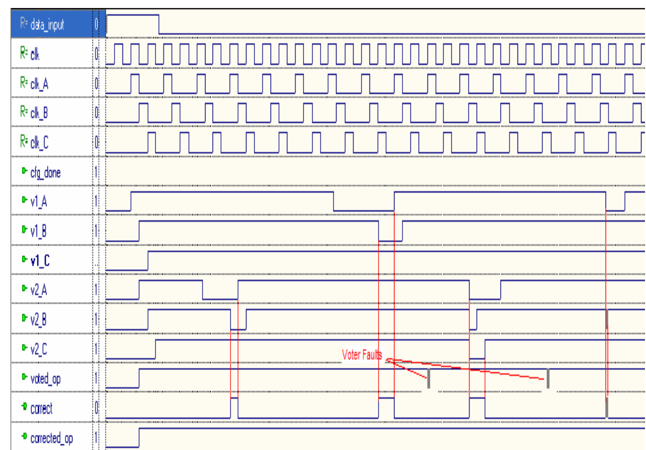

Figure 10: Simulation Results

8. Conclusion

This paper has described a new concurrent SEU/SET mitigation technique for reconfigurable architectures. The proposed scheme gives immunity against all single faults and double faults as well. Due to this added feature, scrubbing frequency can be reduced which doubles the system performance compared to other techniques. The consideration of using reconfigurable SoC in space is fairly recent and there still a lot of work to be done in this area. The technique proposed can be employed in commercial FPGA's as well with minimum speed and area over head.

REFERENCES

[1]http://www.sti.nasa.gov/thesfrm1.htm

[2] Peterson, E.L., "Single-Event Analysis and Prediction",IEEE Nuclear and Space Radiation Effects Conference Short Course Text, 1997.

[3] Massengill, L., "SEU Modeling and Prediction Techniques", IEEE Nuclear and Space Radiation Effects Conference Short Course Text, 1993.

[4] Sexton, F.W., "Measurement of Single-Event Phenomena in Devices and ICs", IEEE Nuclear and Space Radiation Effects Conference Short Course Text, 1992

[5] Lima, F., Carmichael, C., Fabula, J., Padovani, R., Reis, R., “A Fault Injection Analysis of Virtex® FPGA TMR Design Methodology”, Proc. of (RADECS), Sept. 2001.

[6] Carmichael, C., Fuller, E., Fabula, J., Lima, F., “Proton Testing of SEU Mitigation Methods for the Virtex FPGA”,Proc. of MAPLD, 2001.

[7] Dupont, D., Nicolaidis, M., Rohr, P., “Embedded Robustness IPs for Transient-Error-Free-ICs”, IEEE Design and Test of Computers, May-June, 2002.

[8] Lima, F., Carro, L., R., Reis, R., “Designing fault tolerant systems into SRAM-based FPGAs” DAC03,

[9] Mavis, D.G., Eaton, P., H, “SEU and SET mitigation techniques for FPGA circuit and configuration bit storage Design” MAPLD International conference

[10]J. Olsen, P. E. Becher, P. B. Fynbo, P. Raaby, and J. Schultz, “Neutron-Induced Single Event Upsets in Static RAMS Observed at 10 km Flight Altitude,” IEEE

[11] E. Normand, “Single Event Upset at Ground Level,” IEEE Transactions on Nuclear Science, vol. 43, pp. 2742-2750, 1996.

BIOGRAPHY

Adrian Stoica is a Principal Member of Technical Staff in the Biologically Inspired Technology and Systems (BITS) Group in the Avionics Equipment Section at the Jet Propulsion Laboratory, California Institute of Technology. He obtained his MSEE degree from the Technical University of Iasi, Romania, ranking the first among the graduates in Applied Electronics Specialty. He received his Ph.D. in EECS from Victoria University of Technology in Melbourne, Australia, with a thesis titled “Motion learning by robot apprentices – a fuzzy neural approach”, which was the earliest work on humanoid robot motor skill learning by imitation. Since early 1986 his engineering and research work revolved around adaptive and learning techniques for autonomous intelligent systems. He joined JPL in 1996 where he has been working along two main themes: adaptive hardware for space autonomous systems, including evolvable/reconfigurable hardware, adaptive/learning hardware and sensor fusion hardware, and next-generation robots, focusing on rover intelligence and robot sensory-motor control. His current projects in reconfigurable and evolvable hardware address analog/digital field-programmable reconfigurable computing, automated hardware/software co-design and novel search/optimization techniques. Dr. Stoica initiated, and has chaired/co-chaired yearly since 1999, the NASA/DOD Conference in Evolvable Hardware, the main conference in the field of evolvable hardware as judged by number of attendants and references to published papers. He was Conference keynote speaker at ISMVL, ANNIE, gave tutorials at GECCO 2001 and CEC2003, and taught the first short course on Evolvable Hardware in the summer of 2003 at UCLA Extension. He has published over 100 papers in the areas of evolvable hardware, reconfigurable computing, fuzzy logic, neural networks, robot learning and is serving in the editorial board of several journals. He has received the 1999 JPL Lew Allen Award for Excellence (highest JPL award for excellence in research) and the Tudor Tanasescu Prize of the Romanian Academy in 2001.

Tughrul Arslan BEng PhD CEng FIEE MIEEE received the B.Eng. and Ph.D. degrees from the University of Hull, Hull, U.K., in 1989 and 1993, respectively. Professor Tughrul Arslan holds the chair of Integrated Electronic systems in the School of Engineering and Electronics in the University of Edinburgh and is also a co-founder and the Chief Technical Officer of SpiralGateway Ltd. He is a member of the Integrated Micro and Nano Systems (IMNS) Institute and Leads the System Level Integration group (SLIg) in the University. His research interests include: Low Power Design, DSP Hardware Design, System-On-Chip (SoC) Architectures, Evolvable Hardware, Multi-objective optimisation and the use of Genetic algorithms in Hardware design issues. He has published more than 170 articles and is the inventor of 3 patents in these areas. Professor Arslan is an Associate Editor for the IEEE Transactions on Circuits and Systems I, a member of the IEEE CAS Committee on VLSI Systems and applications, and sits on the editorial board of IEE Proceedings on Computers and Digital Techniques and the technical committees of a number of international conferences. He is a principal investigator on a number of

projects funded by EPSRC, DTI, and Scottish Enterprise together with a number of industrial and academic partners.

He is currently a Professor in the Department of Electronics and Electrical Engineering at the University of Edinburgh, Scotland, U.K., where he leads the digital VLSI activity. He is also a Lecturer at the Institute for System-Level-Integration, Livingston, Scotland, U.K. His research interests include low-power VLSI and system-on-chip design, high-performance DSP architectures, and evolutionary algorithms and their application to circuit design

Sajid Baloch received his BSc Engineering degree from University of Engineering and Technology, Lahore, Pakistan. He did his MSc from Institute for System Level Integration, Livingston, Scotland, UK in 2003 and won runners-up prize for the best student of the institute. He is currently doing his PhD from University of Edinburgh, Scotland, UK. His research interests are Fault tolerant systems for space application and SEU hardened reconfigurable architectures.