Automatic Insertion of Selective TMR for SEU Mitigation

O. Ruano, P. Reviriego, and J.A. Maestro.

***Abstract*— In this paper, a methodology is presented to perform automatic selective TMR insertion on digital circuits, having as a constraint the required reliability level. Such reliability is guaranteed while reducing the area compared with TMR.**

***Index Terms*— Single Event Upsets (SEUs), Triple Modular Redundancy (TMR), optimization, fault injection.**

I. Introduction

Fault tolerance on semiconductor devices has been a meaningful matter since upsets were first experienced in space applications several years ago. Integrate circuits operating in the space environment can be upset by charged particles that generate errors in the system. Therefore, the interest in studying fault-tolerant techniques in order to keep integrated circuits (ICs) operational has increased. In order to guarantee the circuit reliability against single event upsets (SEU), some mitigation techniques have been proposed in literature during the last few years. These techniques cover a wide range of methods which can be classified as fabrication process-based, design-based and recovery techniques. The fabrication process-based, such as Epitaxial CMOS processes [1] and Silicon-on-Insulator (SOI) [2], can reduce the radiation effects of Total Ionization Dose (TID) and single event latch-up (SEL). However, they do not eliminate SEUs. The design-based techniques, also called architectural techniques, are based on logic redundancy such as Triple Modular Redundancy (TMR) [3], Error detection and correction coding (EDAC) like Hamming [4] or hardened memory cells like HIT and DICE [5]. Finally, new techniques based on recovery have been proposed for SRAM-Based FPGAs. The main idea is to recover the original programmed information after an upset [6]. Among all these methods, the well-known TMR has become a common practice. In this case, the SEU sensitive logic (memory cell) is tripled and voters are placed at the outputs to identify the correct value (Figure 1). In [7], the TMR technique is shown to be a suitable solution to provide reliability against SEUs, for an 8051 microcontroller. The fault injection results showed that less than 1% of the bits upsets could provoke an error in the output of the TMR design. However, the TMR technique presents some overheads because of its full hardware redundancy, such as area and power dissipation.

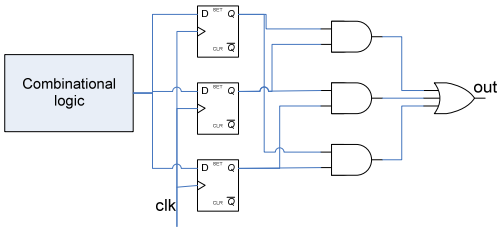

Fig. 1: TMR memory cell with single voter

As a result, both area and power consumption are increased three times, so many applications cannot accept this high cost. In order to make the TMR technique valid for many applications, a heuristic method called selective TMR is proposed in [8][9] to make a selective insertion of TMR in the nodes which present a bigger vulnerability to SEUs. The rest of works found in literature propose new fault tolerant techniques that try to minimize the area cost. This is the case of the DWC technique introduced in [10] that combines time and hardware redundancy reducing the area compared to TMR. Other examples can be found in [11][12]. Focusing on the idea of the selective TMR, and on systems that can allow a fixed error rate, an analysis of an innovative method for selective TMR insertion is presented in this paper, which provides an automatic selective insertion based on an iterative optimization method. The problem is addressed as an optimization problem where two variables must keep a balance; the reliability level demanded by the user must be assured in the final design, while, on the other hand, the area cost due to the modular redundancy must be minimized. This paper is organized as follows. Section II introduces the proposed methodology for TMR optimization and explains in detail the major factors that are involved. In section III, a tri-state pipeline stage is used as a case study where the methodology is validated by fault injection experiments in VHDL. Final remarks and future work are placed in Section IV, followed by the references.

II. Proposed Methodology

Currently, the design of radiation tolerant circuits is very expensive in terms of area and power consumption if the Triple Modular Redundancy (TMR) technique is applied. Therefore, this work proposes a selective TMR insertion adapted to the fault tolerance level required by the

This work was supported by the Spanish Ministry of Science and Education under Grant ESP-2006-04163, the Regional Government of Madrid and the European Union FEDER programme.

O. Ruano, P. Reviriego and J.A. Maestro are with Universidad Antonio de Nebrija,C/Pirineos, 55 E-28040 Madrid, Spain (phone: +34 914521100; fax: +34 914521110; email: {oruano, previrie, jmaestro}@nebrija.es).

application. Although TMR clearly provides a superior protection level, not all the applications demand the highest fault tolerance level. In some cases, an intermediate protection could be enough, what makes TMR excessive in terms of area cost. The proposed methodology is based on a combinatorial optimization problem which makes use of the reliability constraints that the target circuit has to meet. The issue is addressed like a graph partitioning problem which is included in the category of NP-complete problems. It consists in partitioning a given initial graph which includes all registers of the initial circuit, in two sub-graphs. One of them, stores the registers that the system has chosen to protect with TMR, and the second stores the remaining registers without TMR. The purpose of this methodology is based on finding the best partition that assures the minimum possible area in terms of the number of registers on which TMR is inserted, while meeting the reliability constrains specified for the circuit. Therefore, the optimization process looks for a solution among a finite set of alternative results which minimizes the area and assures the required reliability. The methodology workflow defined to drive the process is described in Figure 2:

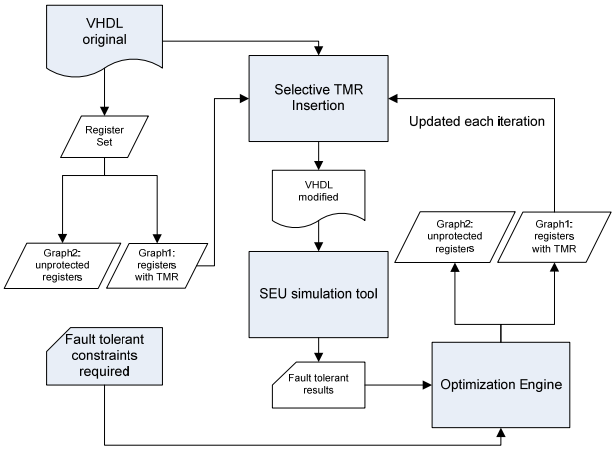

Fig. 2: Methodology workflow

The operational method is detailed next. First of all, data inputs required by the methodology must include an original VHDL description, a random initial partition of the registers grouped into two sub-graphs and the fault tolerance constraints required by the application. Once these data are available, the methodology is ready to run in order to obtain a valid partition which can satisfy the specified reliability constraints.

The first step consists in measuring the fault tolerance for the random initial partition in order to evaluate it and see if it is a valid solution or not. In order to achieve this objective, the process is detailed next. The original VHDL and the set of registers with TMR are given to the Selective TMR insertion module in order to insert TMR on the appropriate registers. Next, the modified VHDL is used to perform a fault injection campaign using the SEE simulation tool [13][14]. The fault tolerance level of the modified VHDL is reported by the fault injection results. These results are sent to the optimization engine module which evaluates them against the reliability constraints required by the application. Once the initial partition has been measured, the above process is repeated with a new partition of the registers selected by the optimization engine, which manages the remaining possible partitions through an optimization algorithm. For each new partition, selective TMR insertion to modify the original VHDL and fault injection is performed to evaluate the fault tolerance. The process is finished when the optimization engine cannot find a new partition with less area cost than the last valid solution found which can satisfy the reliability constraints. In the following, some issues which can have impact in the proposed methodology are discussed.

A. Optimization Algorithm

In order to guide the successive movements and the several partitions which are tested through the fault injection process, an iterative algorithm has been selected to perform the experimental results. Fiduccia-Mattheyses [15] is the meta-heuristic algorithm which has been chosen as it has a linear-time behavior O(N), so its worst case computation time grows linearly with the size of the problem (in this case with the number of registers). A final solution is achieved by moving one register at a time, for one graph of the partition to the other, updating in each step the partition. Another feature of this algorithm is that reduces the chance that the minimization process becomes trapped at local minima (hill-climbing) [15]. In an attempt to reduce the number of registers with TMR while satisfying the reliability constraints, it drives the intermediate partitions according to a cost function which will determine if each partition can satisfy the reliability constraints or not.

B. Cost function

The cost function must drive the algorithm in order to accept or deny each movement which is proposed by Fiduccia-Mattheyses inside the search space. This cost function is included in the minimization problem in the following way:

- Given: a function f: A $\rightarrow$R
- Sought: an element x_0 in A such that $f(x_0) \leq f(x)$ for all x in A.

The domain A of f is called the search space, while the elements of R are called candidate solutions or valid solutions. A valid solution that minimizes the objective function is called an optimal solution. With these premises, the selective TMR insertion issue can be characterized in the following way:

- A represents all possible partitions which can be generated from the whole set of registers that compose the design under test.
- X_0 represents a valid solution that is, a partition composed by two sub-graphs in which the first contains the selected registers to protect with TMR and the second the rest of registers without TMR, that meets the specified reliability constraints.

All valid solutions must satisfy a specified set of constraints. In our case, a set of reliability constraints on the circuit output are used. Consequently, in order to measure the quality of a

solution according to the parameters of the problem (number of registers, reliability constraints and real error rate obtained by the fault injection), the next cost function has been defined such that combines all these parameters:

$$F = \#registers + k_1 \cdot \max(\#error_1 - reliability\_constraint_1, 0) + \\ + k_2 \cdot \max(\#error_2 - reliability\_constraint_2, 0) + \ldots + \\ + k_n \cdot \max(\#error_n - reliability\_constraint_n, 0)$$

where #registers is the total number of register that composes the currently partition, $error_i$ the number of errors detected in the i output through the fault injection process, $reliability\_constraint_i$ is the error threshold that must be satisfied by each output and finally the k_i coefficients which represent the influence of each output.

C. Reliability Constraints

Reliability constraints are specified as the tolerance level which must be met by a valid solution in any case. The optimization engine and more specifically the cost function requires these parameters in order to accept or deny the movement of each register. These factors must be specified by the user as the tolerance thresholds which must be met by each output of the circuit and therefore, their choice is considered key for the optimization process. Depending on these constraints the circuit can tolerate more or less failures and as a result it will have a direct effect on the insertion of selective TMR.

D. Optimization Performance

It is known that the runtime required to obtain a valid solution, for the chosen iterative algorithm, increases with the size of the problem. In some cases, the complexity is so high that even heuristic methods are not able to obtain accurate solutions in reasonable runtimes (too many registers). With regard to this issue, it is detected that the initial partition ($P_{initial}$) has a major influence to reduce the execution time of the method as it will be shown in the experimental results of the next section. In this case, a previous knowledge or a static analysis which reports information of possible weak points of the circuit can improve the runtime performance reducing the number of iterations or movements needed, initiating the process from an initial partition near to a valid solution instead from a random partition. Also, from this information, the cost function can be customized through K coefficients in order to prioritize those weak points identified as sensitive to SEUs, converging earlier to a solution. Finally, other kinds of algorithms can be used in order to make a performance comparison.

For this work, all these points have been taken into account in the next section in order to present experimental results which optimize the insertion of TMR in a selective way ($P_{initial}$, reliability constraints and k coefficients).

III. Case Study

In order to verify the proposed methodology, a case study is analyzed in this section. The circuit under test computes the sum of three 8-bit data buses A, B and C, that is to say (A+B), (B+C) and (A+C) and supplies the results on a single 9-bit output bus in three consecutive clock cycles. Also, an extra output control signal (OutDataReady) should be at logic 1 for one clock cycle to signify when the three summed outputs are available. The proposed design (54 flip-flops) is shown in Figure 3 [16]:

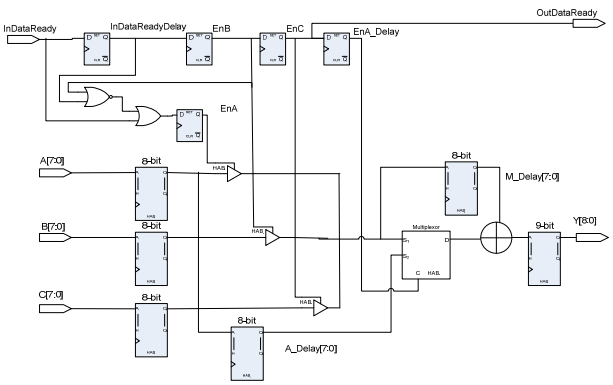

Fig. 3: Tri-State Pipeline Stage

The proposed methodology has been applied to this circuit in order to obtain results of the automatic insertion of selective TMR. Experimental results are collected in Table I, where three different scenarios regarding the requested reliability (0, 5, and 10 per cent respectively) have been proposed in order to show the area decrease compared with TMR. Also, for each reliability scenario three different tests have been proposed where the initial partitions are changed in order to study how the method performance can vary in terms of the explored solutions. The reliability constraint has been considered only for the Y output as all registers of the system have effect on it, included those which determine the value of the OutDataReady output. All tests have run simulations of 100,000 clock cycles injecting random campaigns of 10,000 SEUs with a minimum separation of 3 cycles between two consecutive ones. These results allow a comparison with the area cost of the standard TMR approach (Table I).

TABLE I: Automatic Insertion of Selective TMR for Area Efficiency

Reliability constraints	$P_{initial}$ (flip-flops)		SEUs detected	Partial solutions explored	P_{final} (flip-flops)		Selective TMR Cost reduction
	TMR	Without TMR			TMR	Without TMR	
Y= 0%	0	54	0	103	54	0	162 (0%)
	26	28	0	58	54	0	162 (0%)
	54	0	0	12	54	0	162 (0%)
Y:= 5%	0	54	450	73	30	24	114 (29.6%)
	19	35	490	28	30	24	114 (29.6%)
	54	0	493	39	30	24	114 (29.6%)
Y:= 10%	0	54	998	71	21	33	96 (40.7%)
	12	42	985	28	21	33	96 (40.7%)
	54	0	992	46	21	33	96 (40.7%)

With these results, it can be noticed that in the cases where the application does not demand the highest fault tolerance level (for example 5% or 10%), an intermediate protection could be enough to satisfy the reliability constraints, reducing the area cost in terms of flip-flops (29.6% and 40.7% respectively) compared with the standard TMR approach .

On the other hand, an interesting result offered by the methodology consists in the detection of critical nodes which show a bigger vulnerability to SEUs like in this case the control flip-flops (InDataReady_Delay, Ena, Enb, EnaC). In all tests, these registers have been tripled with a high priority in order to achieve the fault tolerance needed. In this way, the P_{final} illustrated in Table I for tolerances of 5 and 10% has included in all cases these control registers in the TMR sub-graph and in second place are included M_Delay and Y registers.

Finally, it also should be noticed that the initial partition is a significant factor in order to improve the performance of the process. Table I emphasizes that when beginning from a partition near to the final solution, a small number of searches is required. For example, in the case of 0% tolerance in order to achieve the same P_{final} (54 - 0), starting from a partition without any register tripled (0 - 54) 103 explorations are required while using a better partition like (26 - 28) or the best (54 - 0) 58 and 12 are needed respectively, improving the performance considerably.

These results can be contrasted with an analysis of the circuit. As it has been said, in all final partitions the control flip-flops have been included. It means that these flip-flops are considered more sensitive to SEUs by the process. An SEU in these flip-flops can alter significantly the good behavior of the circuit being able to propagate the error during three cycles in the worst case (SEU on InDataReady flip-flop). The right functionality is only recovered when the SEU goes out of the delay line of the control path (Ena_Delay). On the other hand, A_Hold, B_Hold, C_Hold and A_Delay data registers are considered less sensitive because they refresh the useful data each cycle and also are used 1/3 of the cycles, so an SEU can be observed at the output or it can not affect the output value. So, these registers usually have not been tripled by the system due to the low probability of observing an SEU in the Y output. Nevertheless, an SEU in the remaining registers M_Delay and Y_Delay, will be observed in all cases but only during one cycle. With this analysis, it has been proved why the methodology firstly triplicates control registers, secondly M_Delay and Y, and finally if it does not satisfy the reliability constraints yet, the rest of registers.

IV. Conclusions and Future Work

In this paper, a new methodology to insert selective TMR for SEU mitigation has been presented. The benefits of applying the automatic selective TMR methodology are clearly proved with the experimental results that have been obtained: the technique results in a lower circuit complexity than the traditional TMR approach, while meeting the specified reliability level. The next steps will focus on the study of other optimization algorithms (simulated annealing, genetic, mimetic, adaptive), alternative cost functions that can be adapted prioritizing the influence of each output, and the selection method to choose the initial partition in order to improve the performance of the model.

References

[1] Y. Moreau, S. Duzellier, J. Gasiot, “Evaluation of the upset risk in CMOS SRAM through full three dimensional simulation”, IEEE Transaction on Nuclear Science on Volume 42, Issue 6, Part 1, Dec. 1995, pp.: 1789 - 1796.

[2] F. Irom, F.F. Farmanesh, A.H. Johnston, G.M. Swift, D.G. Millward, ”Single-event upset in commercial silicon-on-insulator PowerPC microprocessors”, IEEE Transactions on Nuclear Science Volume 49, Issue 6, Part 1, Dec. 2002, pp.: 3148 - 3155.

[3] C. Carmichael, “Triple Modular Redundancy Design Techniques for Virtex Series FPGA”, Application Notes 197, San Jose, USA, Xilinx, 2000.

[4] R. Hentschke, F. Marques, F. Lima, L. Carro, A. Susin, R. Reis, “Analysing area an performance penalty of protecting different digital modules with hamming code and triple modular redundancy”, Symposium on Integrated Circuits and Systems Design, Proceedings, Sept. 2002.

[5] R. Velazco, D. Bessot, S. Duzellier, R. Ecoffet, R. Koga, “Two CMOS memory cells suitable for the design of SEU-tolerant VLSI circuits”, IEEE Transaction on Nuclear Science, Volume 41, Issue 6, Part 1, Dec. 1994, pp.: 2229 - 2234.

[6] L. Sterpone, M. Violante, “A New Partial Reconfiguration-Based Fault-Injection System to Evaluate SEU Effects in SRAM-Based FPGAs”, IEEE Transaction on Nuclear Science on Volume 54, Issue 4, Part 2, Aug. 2007, pp.: 965 - 970.

[7] F. Lima, S. Rezgui, E. Cota, L. Carro, M. Lubaszewski, R. Reis, R. Velazco, “Designing a radiation hardened 8051-like micro-controller”, Symposium on Integrated Circuits and Systems Design, Proceedings, Sept. 2000, pp.: 255 - 260.

[8] P. K. Samudrala, J. Ramos and S. Katkoori, "Selective triple modular redundancy (STMR) based single-event upset (SEU) tolerant synthesis for FPGAs", IEEE Transactions on Nuclear Science, vol. 51, Oct. 2004, pp.: 2957 - 2969.

[9] P. K. Samudrala, J. Ramos and S. Katkoori , “Method and apparatus for creating circuit redundancy in programmable logic devices”, U.S. Patent Pub. No. WO/2004/077260, Oct. 2004.

[10] F. Lima, G. Neuberger, R. Hentschke, L. Carro, R. Reis, “Designing fault-tolerant techniques for SRAM-based FPGAs”, IEEE Design & Test of Computers, Volume 21, Issue 6, Nov-Dec. 2004, pp.: 552 - 562.

[11] O. Ruano, P. Reviriego, J.A. Maestro, "A New EDAC Technique against Soft Errors based on Pulse Detectors", Proc. of IEEE International Symposium on Industrial Electronics (ISIE'08), Cambridge (U.K.), June 2008 (in press).

[12] P. Reyes, P. Reviriego, J.A. Maestro, O. Ruano, "New Protection Techniques against SEUs for Moving Average Filters in a Radiation Environment", IEEE Transactions on Nuclear Science , Vol. 54, No 4, Aug. 2007, pp.: 957 - 964.

[13] D. Gonzalez, “Single Event Upset Simulation Tool Functional Description”, ESA Report TEC-EDM/DCC-SST2, July 2004.

[14] O. Ruano, J.A. Maestro, P. Reyes and P. Reviriego, "A Simulation Platform for the Study of Soft Errors on Signal Processing Circuits through Software Fault Injection", Proc. of IEEE International Symposium on Industrial Electronics, 2007, pp.: 3316 - 3321.

[15] C. Fiduccia, R. Mattheyses, “A Linear Time Heuristic for Improving Network Partitions”, In Proceedings of the 19th IEEE Design Automation Conference, 1982, pp.: 175 - 181.

[16] D.J. Smith, “HDL Chip Design. A practical guide for designing, Sythesizing and Simulating ASICs and FPGAs using VHDL or Verilog”, Doone Publications, ISBN 0-9651934-3-8, 2001.