

Article

Optimal Physical Implementation of Radiation Tolerant High-Speed Digital Integrated Circuits in Deep-Submicron Technologies

Jeffrey Prinzie ^{1,*}, Karel Appels ¹ and Szymon Kulis ²

¹ Department of Electrical Engineering (ESAT), KU Leuven, 3000 Leuven, Belgium; karel.appels@student.kuleuven.be

² European Center for Nuclear Research (CERN), 1211 Meyrin, Switzerland; szymon.kulis@cern.ch

* Correspondence: jeffrey.prinzie@kuleuven.be; Tel.: +32-14-72-13-38

Received: 14 March 2019; Accepted: 12 April 2019; Published: 14 April 2019



Abstract: This paper presents a novel scalable physical implementation method for high-speed Triple Modular Redundant (TMR) digital integrated circuits in radiation-hard designs. The implementation uses a distributed placement strategy compared to a commonly used bulk 3-bank constraining method. TMR netlist information is used to optimally constrain the placement of both sequential cells and combinational cells. This approach significantly reduces routing complexity, net lengths and dynamic power consumption with more than 60% and 20% respectively. The technique was simulated in a 65 nm Complementary Metal-Oxide Semiconductor (CMOS) technology.

Keywords: single-event upsets (SEUs); digital integrated circuits; triple modular redundancy (TMR); radiation hardening by design

1. Introduction

Digital integrated circuits are important in many of today's complex integrated circuits and systems. A wide range of digital circuits are used as stand-alone digital systems such as microprocessors and digital signal processors which are purely digital systems. A second portion of digital blocks can be found in mixed-signal integrated circuits where finite-state machines or counters assist or interface with their analog counterparts.

The operation frequency of a digital module mainly depends on the system requirements such as data throughput. In mixed-signal circuits such as analog-to-digital converters (ADC), phase locked loops (PLL) and clock and data recovery (CDR), the digital logic is clocked at or a derivative of the mixed-signal sampling time which can be as high as several GHz in recent technologies. The latter digital blocks are intrinsically timing critical and only have little timing overhead for circuit level redundancy in harsh environments.

It has been widely known that ionizing radiation can cause Single Event Effects (SEEs) in CMOS integrated circuits [1], especially in scaled technologies [2]. Single Event Upsets (SEUs) in flip-flops result in corrupted data or logic states [3]. An SEU on a flip-flop can occur at any time and may be unrecoverable if no redundancy is applied. Single-Event Transients (SETs) occur when a particle upsets combinational logic which has no memory [4]. The transient is only temporary. Thus, the digital system is only sensitive to SETs when the SET propagates to the data-inputs of a flip-flop during the setup- and hold-times of that flip-flop [5]. Many aspects can be considered regarding the propagation of the SETs such as logic masking [6,7] and pulse shrinking. However, a common conclusion can be made that digital logic which is clocked at high clock frequencies is more sensitive to SETs since the probability to capture an SET in the clock period dramatically increases in the GHz range [8].

Fortunately, Triple Modular Redundancy (TMR) can be added to digital logic to overcome these effects. TMR triplicates the logic and uses majority voters to correct logic signals [9]. TMR relies on the fact that only one logic signal can be upset at once and would therefore fail if two out of three triplicated nets are upset. This was less of a concern in old CMOS technologies where single particles only affected single digital cells. However, as technologies have scaled down, Multi-Bit Upsets (MBU) have become a serious concern such that one particle can affect multiple gates simultaneously [10–12]. With improper placement, the fault tolerance can dramatically reduce, especially in fast designs.

Many forms of TMR have been presented in recent decades, some of them compromising triplication effectiveness for power consumption or area. The most complete form of TMR, and the one which is addressed in this paper, is full TMR [13] where both the flip-flops, clock-tree and combinational logic are triplicated ([14,15]) as is shown in Figure 1. This method is the most reliable but also uses the highest number of resources and power. A competing method is temporal time-redundancy [16]. This smart approach does not triplicate the combinational logic but only triplicates flip-flops which are clocked with 3 delayed clocks [17]. The delay between the clocks is set to be larger than any possible SET, such that only one flip-flop could possibly capture an SET which is later on corrected. This method has proven its usage in space applications but its major drawback is its limited clock frequency. The intentional clock skew places serious timing constraints on the design typically resulting in sub-GHz designs. As such, many high-speed mixed-signal digital module implementations prefer the original TMR approach. Several other methods were reported in [18–22].

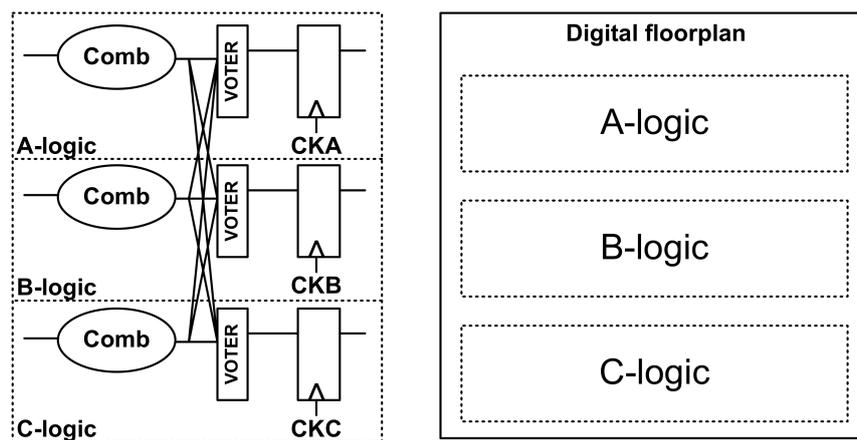


Figure 1. Full TMR design with correcting voters and its associated physical floorplan.

This paper is organized as follows: In Section 2, a novel placement method for high-speed digital TMR designs is presented. In Section 3, a performance analysis is made and compared with conventional methods. In Section 4, a conclusion is drawn.

2. Physical Implementation for TMR Circuits

2.1. Conventional 3-Block Approach

TMR-protected circuits only operate well if only one signal of a TMR signal is upset. Therefore, MBUs should be avoided. Historically, a floorplan was made as is shown in Figure 1. Only flip-flops were constrained to be placed in 3 physical groups respectively A-C and a spacing of 10 μm is ensured between these groups. The idea behind this approach is that the combinational logic and clock trees will follow the constrained flip-flops. This has proven its usage in many high-speed radiation hard designs in the past [23].

The drawback of this approach is the complex routing to interconnect the voters. Each voter has a connection to both A, B and C blocks such that each logic net has 6 cross domain connections. This becomes problematic when the design size increases. Firstly, connections from A to C would require to cross the entire B-logic which results in long nets. Automatic place-and-route tools will always optimize the design to meet the required timing constraints. Therefore, large buffers will be inserted to meet the timing constraints which increases the power consumption of the design. Secondly, these long vertical connections result in a dense metal routing. This may result in a sub-optimal design and may result in worse signal integrity. As the design size increases, so does the problematic physical implementation with this method.

2.2. Novel Interleaved Approach

Ideally, each cell is guided with a spacing constraint to place the standard cells distributed. Such features have become available in the newest place-and-route releases. This can be used to ensure the spacing between flip-flops but not between combinational logic in the data-path. For moderate to complex data-paths such as a fast 8-bit counter, the spacing between logic in a common TMR data-path cannot always be ensured to prevent MBUs in the combinational logic.

The approach proposed in this paper allows a semi-distributed placement to allow maximal freedom to place-and-route tools to optimize the design. It is based on the conventional 3-block approach but uses an interleaved placement constraining method. This is shown in Figure 2a. Multiple repeating small regions allow cells of A-C branches to be placed at different vertical locations in the design. These regions have fixed heights. As the design size increases, only the number of vertical regions increases, not the height of it. As a result, the vertical connections between voters only have to cross a narrow placement region and not 1/3 of the design since the place-and-route tool has more freedom to place the standard cells vertically in the design. The spacing between each region is sufficient to prevent MBUs (10 μm in our case). Grouping the flip-flops to the correct regions (A-C) is straightforward since the naming of these registers includes the TMR group through automatic TMR insertion before synthesis [23]. However, after synthesis, the data-path cells cannot be grouped by name anymore. Therefore, an algorithm was designed to trace-back the combinational logic that drives a flip-flop. As such, if each flip-flop can be segmented by name, the combinational logic can be as well by tracing back the input logic tree. This algorithm finds the fan-in of the flip-flop by searching for the drivers of the input nets of the cells as is shown in Figure 2b. This is done iteratively to find the full fan-in logic tree and stops at a register. A special exception is required when searching the fan-in of a voter block. Only the respective input is used since otherwise, the entire TMR data-path (A-C) is grouped in the same region due to the cross coupled connections before the voters. By means of this method, all cells, including combinational logic can be efficiently constrained to regions. If only the flip-flops were allocated to their respective groups, unconstrained data-path cells might be placed in incorrect regions resulting in MBUs in the combinational logic. This algorithm is executed before placement of the cells and is re-executed after clock tree synthesis and design optimization to.

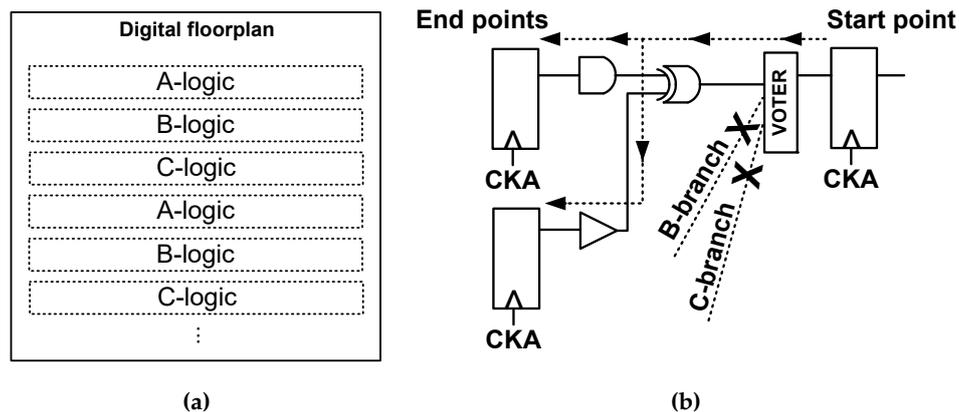


Figure 2. Logic placement of an interleaved design. (a) Interleaved physical floorplan; (b) Logic fan-in cone traceback.

3. Simulated Performance Analysis

To assess the performance of this new placement and floorplanning method, different comparing tests have been done between the proposed interleaved approach and the conventional 3-block approach. Three digital designs, each with 8 identical independent high-speed counters, were used. The counter sizes of the three designs were varied to implement more complex standardized data-paths. A summary of the designs and associated timing constraints is shown in Table 1. The designs were implemented and benchmarked using Innovus CAD (Computer Aided Design) tools with optimization efforts high. The timing constraints were chosen to be close to the technology limits to ensure a timing critical design. For each design, the power consumption, net length, net capacitance and routing density was analyzed and compared. The designs were implemented in a 9-track standard V_T 65 nm CMOS library. The interleaved method has a slice height and spacing of $7.2 \mu\text{m}$ while the 3-block method has the same block spacing. The results discussed below are extracted from the timing, power and area reports from place-and-route tools.

Table 1. Benchmarked designs.

	Counter Size	Clock Constraint	Inst. Count	Design Width
Design 1	8×8 bit	1.2 ns	2808	$175 \mu\text{m}$
Design 2	8×16 bit	1.7 ns	4560	$250 \mu\text{m}$
Design 3	8×32 bit	1.9 ns	9120	$350 \mu\text{m}$

The routed designs of “Design 2” are shown in Figure 3a,b, where the 3-block and proposed interleaved placement method are used respectively. It is clear that the proposed approach results in significantly reduced complexity compared to the conventional 3-block approach. More specifically, the vertical routing difficulty is highly reduced since the place-and-route tool has more freedom to optimally place the cells in the design. To quantitatively analyze this, a histogram of the total net length is shown in Figure 4 for both implementations. The histogram shows that a significant portion of the nets has lengths between $1/3$ and $1/2$ of the design size for the 3-block implementation due to the voter interconnects. This peak is not present anymore in the proposed interleaved implementation method. Most interconnected nets now have a net length of approximately $25 \mu\text{m}$. As the design size increases, the difference becomes more significant. However, some long nets are still present in the interleaved placed design. By analyzing pre-Clock Tree Synthesis (CTS) and post-CTS histograms, it becomes clear that these longer nets originate from the clock tree. Clock trees of clocks A-C are now distributed across the entire design while in the 3-block approach, the clock tree was only placed locally in each of the 3 regions. Similar results were obtained from both Design 1 and Design 3.

The standard-cell spacing between cells of the same TMR branch is shown in Figure 5. This plot shows a histogram of the distances between cells of A-branch to respectively B- and C-branch cells. The distances were measured only between cells which implement the same connected logic tree. Again, these results show a significant reduction of the spacing for the proposed method which correlate with the net lengths in Figure 4. With the interleaved method, most cells are spaced within a 45 μm distance which corresponds to roughly 3 small, vertical interleaved banks. From this result, it is clear that the placement engine is given more freedom to place cells closer without compromising radiation hardness. For to the 3-block method, the average cell distance between A-branch and B- or C-cells shows two peaks which correspond to the large A-B-C block distances in the floorplan. For the proposed interleaved method, the reduced cell spacing leads to a reduction of the power consumption and an improved routability of the design.

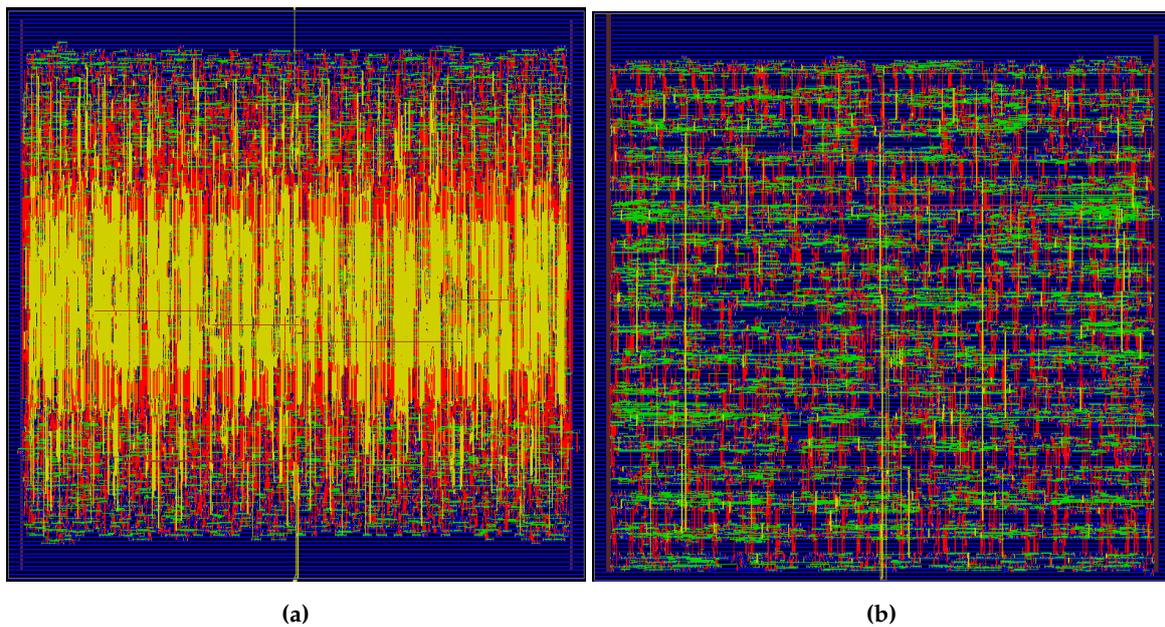


Figure 3. Routed 8×16 bit designs with a floorplan size of $250 \mu\text{m} \times 250 \mu\text{m}$. (a) 3-Block approach; (b) Proposed interleaved approach.

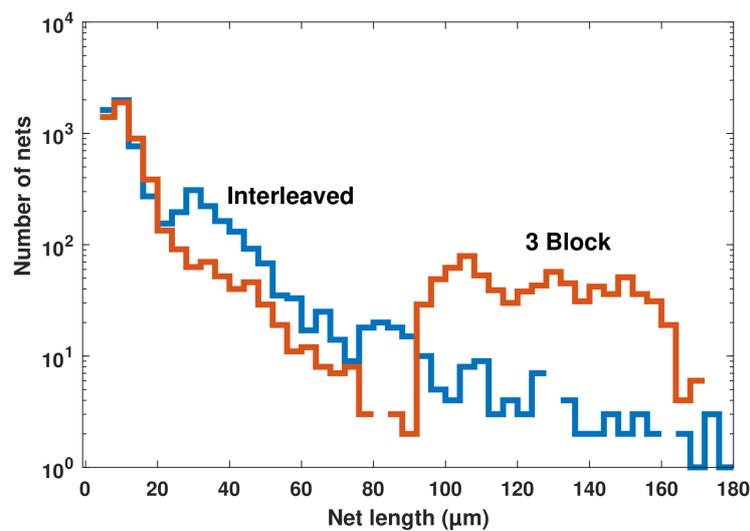


Figure 4. Net length histogram for the 8×16 bit design.

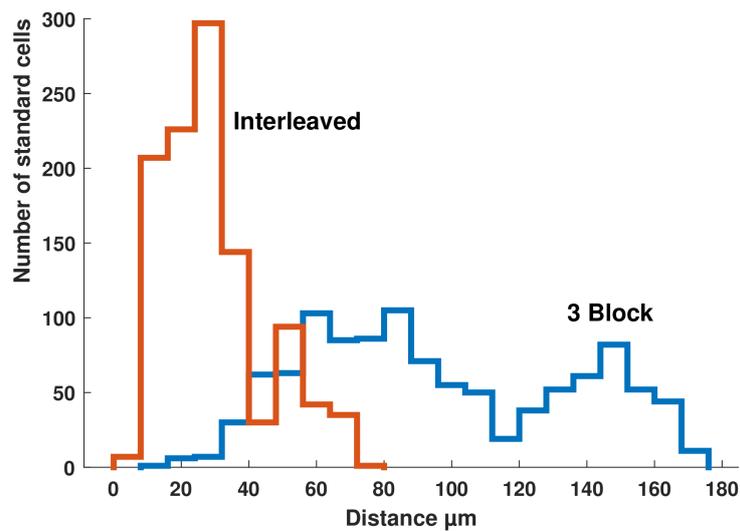


Figure 5. TMR branch logic spacing for the 8×16 bit design.

The routing complexity was analyzed by measuring the local metal density for each metal layer on a 100×100 square grid across the design. The design examples used a metal stack with 5 routable layers, of which M2 and M4 were vertical routing layers. Figure 6 shows the average metal density across the design floorplan for all different layers. All horizontal layers and M1 (standard cell routing and power rail) show no significant differences between both methods. A significant difference can be found in M2 and M4. Vertical inter-TMR branch connections are routed in these layers such that the density is relatively high for the 3-block implemented designs. However, the interleaved design shows a 50% reduction in M2 and nearly no routing in M4 which is a significant improvement. Near the vertical middle of the design, the peak density of M2 drops from 22% to 8% while that of M4 reduces from 30% to 0.5%. A vertical cross section in the middle of the design is shown in Figure 7. Here, the average metal density as a function of the vertical location in the design is shown. These results clearly indicate that the routing complexity is dramatically reduced which also lead to a reduction of the overall power consumption.

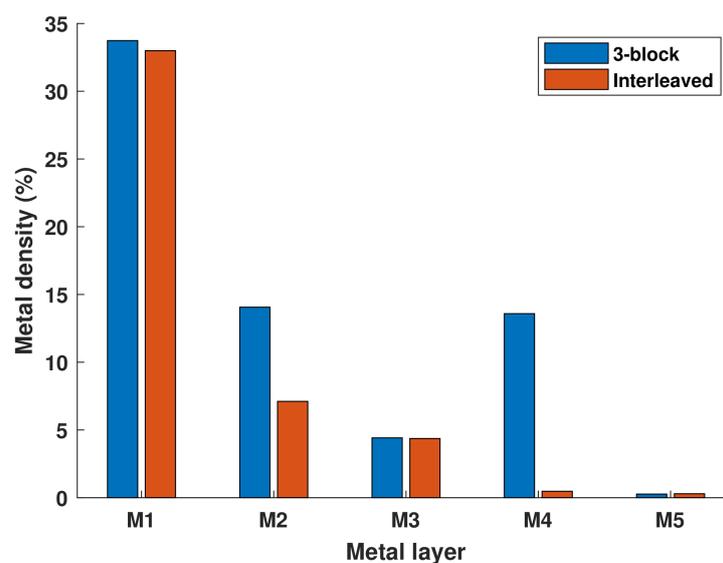


Figure 6. Average design metal density for different metal layers for 3 Block (left) and interleaved (right) methods.

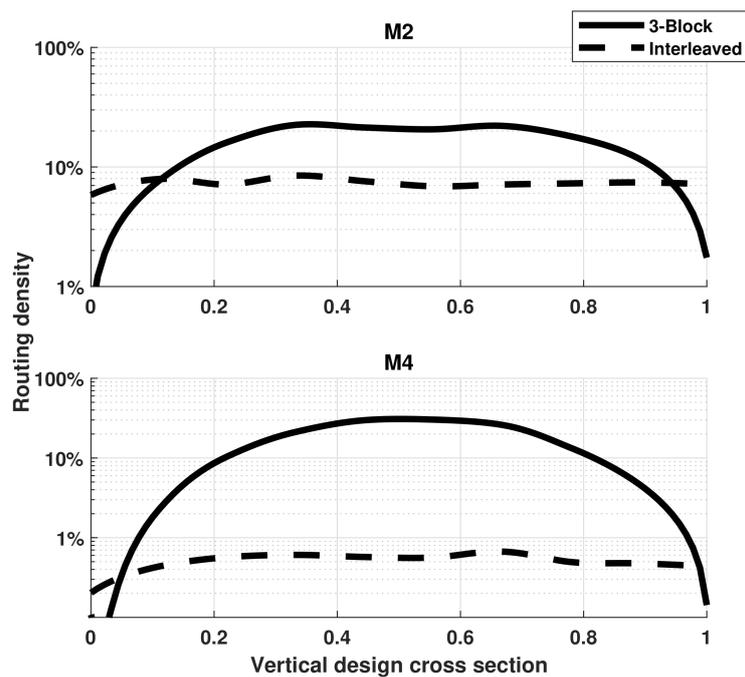


Figure 7. Vertical cross section of metal routing density for M2 and M4 vertical routing layers.

A summary of the power consumption and net lengths is shown in Table 2 for the 3 designs. These figures are extracted after routing and Clock Tree Synthesis (CTS). The internal power is the power consumption of the unloaded standard cells, switching power is the dynamic power consumption due to the switching of the capacitive loads. Total capacitance is the sum of all net and input capacitances of the cells. From these figures, it is clear that the internal power does not significantly change. A small reduction comes from smaller buffers required in the design. However, the switching power changes significantly and scales proportionally with the total capacitance of the nets. A comparison of the results between the 3-block and the proposed interleaved method clearly indicates that a significant reduction of 14% up to 47% can be achieved with our method as the design size increases. These results illustrate the advantages of the proposed placement method. Additionally, it can be seen that the total and average net length of the design reduces by 36% to 65% which is a significant improvement.

Table 2. Performance comparison.

(a) 8 × 8 bit			
	3-Block	Interleaved	Difference
Internal power	20.59 mW	19.46 mW	−5.4%
Switching power	10.55 mW	8.712 mW	−17.42%
Total Power	31.14 mW	28.18 mW	−9.51%
Total capacitance	32 pF	26 pF	−14.8%
Total net length	74,425 μm	47,436 μm	−36.26%
Average net length	17 μm	11 μm	−35.46%
(b) 8 × 16 bit			
	3-Block	Interleaved	Difference
Internal power	21.95 mW	22.58 mW	+2.8%
Switching power	11.67 mW	9.224 mW	−20.95%
Total Power	33.62 mW	31.8 mW	−5.41%
Total capacitance	45 pF	38 pF	−15.3%
Total net length	142,468 μm	89,096 μm	−37.46%
Average net length	23.9 μm	14.2 μm	−40.47%
(c) 8 × 32 bit			
	3-Block	Interleaved	Difference
Internal power	55.17 mW	51.96 mW	−5.82%
Switching power	34.82 mW	18.26 mW	−47.55%
Total Power	89.99 mW	70.22 mW	−21.96%
Total capacitance	108 pF	60 pF	−44.28%
Total net length	368,032 μm	128,320 μm	−65.13%
Average net length	26.8 μm	9.1 μm	−65.89%

4. Conclusions

This paper has presented a novel method for physical implementation of Triple Modular Redundant high-speed digital circuits. The method uses a distributed constraining approach for TMR branches to avoid long interconnects between voters. A TMR logic fan-in search algorithm is used to segment combinational logic in TMR A, B and C groups. The method was tested with increasingly complex digital modules and shows results which improve as the design size increases. For the tested circuits, the total net length reduced up to 65% while the switching power consumption reduced by 44%. Furthermore, the routing complexity was significantly simplified compared to a bulk 3-block physical floorplan.

Author Contributions: Conceptualization, J.P., K.A. and S.K.; methodology, J.P. and K.A. and S.K.; validation, J.P. and K.A.; investigation, J.P. and K.A.; writing—original draft preparation, J.P., K.A. and S.K.

Funding: This research was funded by FWO—Research Foundation Flanders.

Conflicts of Interest: The authors declare no conflict of interest. The founding sponsors had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, and in the decision to publish the results.

References

1. Harrington, R.C.; Kauppila, J.S.; Warren, K.M.; Chen, Y.P.; Maharrey, J.A.; Haeffner, T.D.; Loveless, T.D.; Bhuva, B.L.; Bounasser, M.; Lilja, K.; et al. Estimating Single-Event Logic Cross Sections in Advanced Technologie. *IEEE Trans. Nucl. Sci.* **2017**, *64*, 2115–2121.

2. Nsengiyumva, P.; Ball, D.R.; Kauppila, J.S.; Tam, N.; McCurdy, M.; Holman, W.T.; Alles, M.L.; Bhuva, B.L.; Massengill, L.W. A Comparison of the SEU Response of Planar and FinFET D Flip-Flops at Advanced Technology Nodes. *IEEE Trans. Nucl. Sci.* **2016**, *63*, 266–272. [[CrossRef](#)]
3. Chen, C.-H.; Knag, P.; Zhang, Z. Characterization of Heavy-Ion-Induced Single-Event Effects in 65 Nm Bulk CMOS ASIC Test Chips. *IEEE Trans. Nucl. Sci.* **2014**, *61*, 2694–2701. [[CrossRef](#)]
4. Benedetto, J.; Eaton, P.; Avery, K.; Mavis, D.; Gadlage, M.; Turflinger, T.; Dodd, P.E.; Vizkelethy, G. Heavy ion-induced digital single-event transients in deep submicron Processes. *IEEE Trans. Nucl. Sci.* **2004**, *51*, 3480–3485. [[CrossRef](#)]
5. Benedetto, J.M.; Eaton, P.H.; Mavis, D.G.; Gadlage, M.; Turflinger, T. Digital Single Event Transient Trends With Technology Node Scaling. *IEEE Trans. Nucl. Sci.* **2006**, *53*, 3462–3465. [[CrossRef](#)]
6. Dodd, P.E.; Shaneyfelt, M.R.; Felix, J.A.; Schwank, J.R. Production and propagation of single-event transients in high-speed digital logic ICs. *IEEE Trans. Nucl. Sci.* **2004**, *51*, 3278–3284. [[CrossRef](#)]
7. Chen, R.M.; Diggins, Z.J.; Mahatme, N.N.; Wang, L.; Zhang, E.X.; Chen, Y.P.; Liu, Y.N.; Narasimham, B.; Witulski, A.F.; Bhuva, B.L. Analysis of temporal masking effect on single-event upset rates for sequential circuits. In Proceedings of the 2016 16th European Conference on Radiation and Its Effects on Components and Systems (RADECS), Bremen, Germany, 19–23 September 2016. [[CrossRef](#)]
8. Reed, R.A.; Carts, M.A.; Marshall, P.W.; Marshall, C.J.; Buchner, S.; La Macchia, M.; Mathes, B.; McMorrow, D. Single Event Upset cross sections at various data rates. *IEEE Trans. Nucl. Sci.* **1996**, *43*, 2862–2867. [[CrossRef](#)]
9. Ruano, O.; Reviriego, P.; Maestro, J.A. Automatic insertion of selective TMR for SEU mitigation. In Proceedings of the RADECS, Gothenburg, Sweden, 16–21 September 2018; pp. 284–287.
10. Evans, A.; Glorieux, M.; Alexandrescu, D.; Polo, C.B.; Ferlet-Cavrois, V. Single event multiple transient (SEMT) measurements in 65 nm bulk technology. In Proceedings of the 2016 16th European Conference on Radiation and Its Effects on Components and Systems (RADECS), Bremen, Germany, 19–23 September 2016. [[CrossRef](#)]
11. Zhang, H.; Jiang, H.; Assis, T.R.; Ball, D.R.; Narasimham, B.; Anvar, A.; Massengill, L.W.; Bhuva, B.L. Angular Effects of Heavy-Ion Strikes on Single-Event Upset Response of Flip-Flop Designs in 16-nm Bulk FinFET Technology. *IEEE Trans. Nucl. Sci.* **2017**, *64*, 491–496. [[CrossRef](#)]
12. Bhuva, B.L.; Tam, N.; Massengill, L.W.; Ball, D.; Chatterjee, I.; McCurdy, M.; Alles, M.L. Multi-Cell Soft Errors at Advanced Technology Nodes. *IEEE Trans. Nucl. Sci.* **2015**, *62*, 2585–2591. [[CrossRef](#)]
13. Ulloa, G.; Lucena, V.; Meinhardt, C. Comparing 32nm full adder TMR and DTMR architectures. In Proceedings of the 2017 24th IEEE International Conference on Electronics, Circuits and Systems (ICECS), Batumi, Georgia, 5–8 December 2017; pp. 294–297.
14. Balasubramanian, P.; Prasad, K. A Fault Tolerance Improved Majority Voter for TMR System Architecture. *WSEAS Trans. Circuits Syst.* **2016**, *15*, 108–122.
15. Shinohara, K.; Watanabe, M. A double or triple module redundancy model exploiting dynamic reconfigurations. In Proceedings of the 2008 NASA/ESA Conference on Adaptive Hardware and Systems, Noordwijk, The Netherlands, 22–25 June 2008; pp. 114–112.
16. Hasanbegovic, A.; Aunet, S. Heavy Ion Characterization of Temporal-, Dual- and Triple Redundant Flip-Flops Across a Wide Supply Voltage Range in a 65 nm Bulk CMOS Process. *IEEE Trans. Nucl. Sci.* **2016**, *63*, 2962–2970. [[CrossRef](#)]
17. Schmidt, R.; García-Ortiz, A.; Fey, G. Temporal redundancy latch-based architecture for soft error mitigation. In Proceedings of the 2017 IEEE 23rd International Symposium on On-Line Testing and Robust System Design (IOLTS), Thessaloniki, Greece, 3–5 July 2017; pp. 240–243.
18. She, X.; McElvain, K.S. Time multiplexed triple modular redundancy for single event upset mitigation. *IEEE Trans. Nucl. Sci.* **2009**, *56*, 2443–2449. [[CrossRef](#)]
19. Tooba, A.; Abdus Sami, H.; Hossein, M.; Jeong A.L. Input vulnerability-aware approximate triple modular redundancy: Higher fault coverage, improved search space, and reduced area overhead. *Electron. Lett.* **2018**, *54*, 934–936.
20. She, X.; Li, N. Reducing Critical Configuration Bits via Partial TMR for SEU Mitigation in FPGAs. *IEEE Trans. Nucl. Sci.* **2017**, *64*, 2626–2632. [[CrossRef](#)]
21. Matush, B.I.; Mozdzen, T.J.; Clark, L.T.; Knudsen, J.E. Area-Efficient Temporally Hardened by Design Flip-Flop Circuits. *IEEE Trans. Nucl. Sci.* **2010**, *57*, 3588–3595. [[CrossRef](#)]

22. Nagpal, C.; Garg, R.; Khatri, S.P. A delay-efficient radiation-hard digital design approach using CSWP elements. In Proceedings of the 2008 Design, Automation and Test in Europe, Munich, Germany, 10–14 March 2008; pp. 345–359.
23. Kulis, S. Single Event Effects Mitigation with TMRG Tool. *J. Instrum.* **2017**, *12*. [[CrossRef](#)]



© 2019 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).