LA-UR-04-2718

Approved for public release;
distribution is unlimited.

Title: Abstract: Evaluation of Power Costs in Applying TMR to FPGA Designs

Author(s):
Nathan Rollins1,
Michael J. Wirthlin1,
Paul S. Graham2
1Brigham Young University, Provo, UT 84602
2Los Alamos National Laboratory, Los Alamos, NM 87545

Submitted to:
2004 MAPLD International Conference
Ronald Reagan Building and International Trade Center
Washington, D.C.
September 8-10, 2004

Evaluation of Power Costs in Applying TMR to FPGA Designs

Nathan Rollins[1],Michael J. Wirthlin[1], and Paul Graham[2]
nhr2@ee.byu.edu, *wirthlin@ee.byu.edu*, and *grahamp@lanl.gov*
[1]Department of Electrical and Computer Engineering, Brigham Young University, Provo, UT. 84602
[2]Los Alamos National Laboratory, Los Alamos, NM

I Introduction

Triple modular redundancy (TMR) is a technique commonly used to make designs reliable in the presence of single event upsets (SEUs)[1]. This design hardening technique triplicates all of the resources used in a design and then uses a majority voter to vote on the outputs of the triplicated design.

While TMR protects a design against SEUs, this increased reliability comes at great cost[2]. Previous studies have shown that TMR can be used to make a design immune to SEUs[3]. However, this SEU immunity comes at great cost in terms of design area and speed. A completely SEU immune design comes at the cost of at least 3x in area, and a reduction in design speed.

In addition to costs in area and speed, applying TMR to a design will greatly increase the power dissipation. Power consumption is becoming a defining design criterion for semi-conductor devices[4]. FPGAs in particular, consume relatively more power than other semi-conductor devices such as ASICs. FPGAs are less power efficient than ASICs due to their flexibility. The re-programmability of SRAM-based FPGAs causes them to require a larger number of transistors than ASICs. A larger number of transistors leads to larger leakage current. Leakage, or static power, previously considered insignificant, can no longer be neglected. Power characteristics of an FPGA affect the density, performance, reliability, and cost of a device[5].

In space-based applications where device thermal behavior and cooling is an integral design consideration, power consumption is an important design consideration. In this design environment there exists a trade-off between improving reliability with techniques such as TMR and increasing the power dissipation. The goal of this study is to evaluate the costs of TMR in terms of the increase in power consumption. Triplicating an entire design suggests that the amount of power consumed may increase by a factor of 3x. Such an increase in power consumption may have a significant impact on the overall system design.

II Power Estimation and Measurement Tools and Designs

Reliable power measuring tools are an essential part of measuring the power cost of TMR. The two tools we use in our study are JPower and XPower. JPower [1] measures the amount of actual current flowing in a circuit using the SLAAC-1V FPGA computing board[6]. Xilinx's XPower tool, estimates the amount of power which a design would consume[7].

A set of testbench designs are developed for evaluating the cost of power due to TMR. A set of designs will also enable us to compare the output of the two different power evaluating tools. XPower and JPower are used to estimate and measure respectively, the power consumed by each design. TMR is then applied to each design and the power tools again measure the amount of power dissipated. By comparing the amount of power consumed in the TMR designs with the amount of power used in the non-TMR designs, we can see the cost of TMR in terms of power.

In previous TMR tests[3] two simple designs were used to study the area and speed costs of SEU-immune designs. The two designs used in these previous tests are an 8-bit incrementer and an 8-bit loadable counter. In this power study, we use these simple designs as part of our testbench designs to examine the power costs due to TMR. In the final paper, this study will provide results from a larger set of designs including an 8-bit CPU and a QPSK demodulator.

[1]JPower was named after Jason Zimmerman who made this tool functional while working at Los Alamos National Laboratory during summer 2003

	Incrementer						Counter (auto)		
	Place 1 (auto)			Place 3 (optimal)					
	No TMR (mW/MHz)	TMR (mW/MHz)	Ratio	No TMR (mW/MHz)	TMR (mW/MHz)	Ratio	No TMR (mW/MHz)	TMR (mW/MHz)	Ratio
XPower	3.84	13.07	**3.40**	4.76	11.90	**3.10**	23.14	97.57	**4.22**
JPower	3.85	18.43	**4.79**	4.78	11.94	**3.15**	27.70	118.83	**4.28**

Table 1: Cost of power due to TMR in terms of a factor of of power increase for an auto-placed and optimally placed incrementer design, as well as for an auto-placed counter design

III Power Estimation and Measurement Results

For each of the different testbench designs, the power evaluating tools are used to measure the power of each design at a range of frequencies on a Xilinx Virtex V1000 FPGA. Taking power measurements in a range of frequencies enables us to create a plot of frequency vs. power from which we can interpolate a slope with units of mW per MHz. The y-intercept of this power vs. frequency slope represents the static power consumed by the design. TMR is applied to each design and the power tools are again used to measure power at a range of different frequencies. Comparing the slope of a design with TMR implemented vs. the slope of a design without TMR determines the factor of increased dynamic power consumption due to TMR for a design (the cost of TMR in terms of power).

Power measurements and estimations are made on the replicated incrementer design and are shown in Figure 1. This graph shows two pairs of plots. The bottom pair of lines are the XPower power estimations and JPower power measurements without TMR. The upper pair of lines are the power measurements when TMR is applied. These plots indicate that XPower does an accurate job of estimating the power consumed by the design.

Table 1 lists the dynamic power consumption of the two designs for both the TMR and non-TMR version of the circuit. The dynamic power measurements are reported by the XPower estimation tool and the JPower measurement tool in units of mW/MHz. Note that the XPower estimation and JPower measurements are very close. This table also includes two placement alternatives for the incrementer design to demonstrate the impact of placement on power consumption. The effects of design placement will be more thoroughly discussed in the paper.

To measure the *cost* of TMR in terms of dynamic power, we calculate the ratio of TMR power consumption versus the non-TMR power consumption

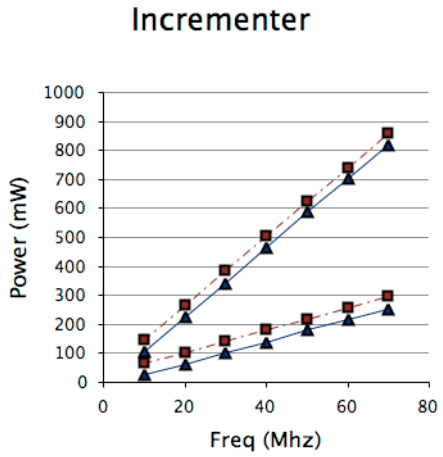

Figure 1: XPower (dashed lines) estimations and JPower (solid lines) measurements of replicated incrementer design with and without TMR (optimal placement)

(see Ratio column). In our measurements we found that the use of TMR increased the dynamic power consumption of these circuits by a factor of 3.1 to 4.8. These results suggest that the increased cost in power consumption due to TMR is design dependent and likely greater than a factor of 3. This increase in dynamic power consumption is due to the increase in area, signal route length, signal fanout, and other TMR design implications.

This paper will determine the cost in dynamic power consumption of TMR for a wide variety of circuits. In addition, it will discuss the impact of placement on power consumption and differences in power consumption for a variety of FPGA architectures. By completing this study, we will better understand the power requirements for FPGA circuits that employ TMR to improve the reliability of circuits operating in a radiation environment.

References

[1] F. Lima, C. Carmichael, J. Fabula, R. Padovani, and R. Reis. A fault injection analysis of Virtex FPGA TMR design methodology. In *Proceedings of the 6th European Conference on Radiation and its Effects on Components and Sysemts (RADECS 2001)*, 2001.

[2] Carl Carmichael. Triple module redundancy design techniques for Virtex FPGAs. Technical report, Xilinx Corporation, November 1, 2001. XAPP197 (v1.0).

[3] Nathan Rollins, Michael Wirthlin, Michael Caffrey, and Paul Graham. Evaluating tmr techniques in the presence of single event upsets. In *Proceedings of the 6th Annual International Conference on Military and Aerospace Programmable Logic Devices (MAPLD)*, September 2003. To Be Published.

[4] A. Allan D. Edenfeld W. Joyner Jr A. Khang M. Rogers Y. Zorian. 2001 technology roadmap for semiconductors. *Computer*, 35:42–53, January 2003.

[5] Xilinx. Fpgas power and packages. *XCell*, 1997.

[6] USC-ISI East. *SLAAC-1V User VHDL Guide*, October 1, 2000. Release 0.3.1.

[7] Xilinx, Inc. *XPower Manual*.