Microprocessors and Microsystems 37 (2013) 313–318

Contents lists available at SciVerse ScienceDirect

Microprocessors and Microsystems

journal homepage: www.elsevier.com/locate/micpro

A new methodology for single event transient suppression in flash FPGAs ☆

Farouk Smith *

Nelson Mandela Metropolitan University, Summerstrand, Port Elizabeth 6031, South Africa

ARTICLE INFO

Article history:
Available online 4 January 2013

Keywords:
Field programmable gate array
Single event transient
Single event upset
Tripple Modular Redundancy
Application Specific Integrated Circuits

ABSTRACT

This paper describes a Single Event Transient (SET) suppression design technique for hardening combinational circuits against SETs in non-volatile Field Programmable Gate Arrays (FPGAs). The proposed method adds a SET suppressor circuit that is insensitive to SETs, to each primary output of a combinational circuit. The SET suppressor circuit consists of three components; an AND gate to suppress an SET reaching the primary output, when the primary output is logic '0', and an OR gate when the primary output is logic '1'. The third component is a simple two input multiplexer with its output connected to its own select line such that it will select the AND gate output when the combinational circuit primary output is logic '0' and the OR gate output when the primary output is logic '1'. A delay element is used to split each primary output of the combinational circuit into two signals. The two signals, one being the original primary output and the other a delayed copy of it, is sent to input one and input two of the SET suppressor. An alternative embodiment of the SET suppressor circuit is to use Double Modular Redundancy (DMR) instead of the delay element implementation.

The SET Suppressor method is thoroughly tested on MCNC'91 benchmarks using the ModelSim simulator. The SET Suppressor circuit provides total immunity against SETs, however it does so with an area savings of 11.6–62.2% with respect to TMR when the delay element technique is use. When the DMR SET Suppressor technique is used, the area savings with respect TMR is between 16.1% and 31.9%.

1. Introduction

SETs are caused by charged particles depositing charge on circuit elements through ionization. These deposited charges causes elevated local voltage levels in the circuit elements, which leads to incorrect logic values [1].

In a combinational logic element, the charge will leak away (over several hundreds of picoseconds) and the element, and consequently the system, will return to a consistent state. However when synchronous logic is disturbed by a SET on a clock edge, the temporarily incorrect logic value is latched into the register. This incorrect value can then propagate through the rest of the system compromising its functional integrity. SETs that are erroneously latched by a register are called Single Event Upsets (SEUs).

It is apparent that some kind of single event transient mitigation scheme is crucial for the successful deployment of FPGA's (and even ASICs) for space-based applications.

2. Background and related work

2.1. Single event transients soft errors

For a single event transient to result in a soft error in a sequential circuit, three conditions have to be satisfied: (1) an active path must exist between the afflicted node and the output of the circuit; (2) the pulse must be wide enough to avoid inertial delay filtration through subsequent gates and survive electrical attenuation along the active path; and (3) the pulse should arrive within the setup and hold time of a latch element, such as a flip-flop, to be captured and cause a soft fault [2,3].

2.2. SET and SEU mitigation in FPGAs

The most common mitigation scheme for correcting SET and SEU errors in sequential circuits in orbit is TMR [4]. The TMR technique is a spatial redundancy technique that compares three signal values by means of a voting circuit, where the output is equal to the two inputs that agree.

The main disadvantage of TMR is the excessive area overhead. The hardened design has at least three times more area and power consumption than the original circuit, excluding TMR overhead.

☆ International Patent Pending, Int Pat Appl No PCT/IB2012/055738

* Tel.: +27 832819015.
E-mail address: farouk.smith@nmmu.ac.za

0141-9331/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.micpro.2012.12.002

2.3. SET's in non-volatile FPGA's

The configuration memory of non-volatile FPGA's have been shown to be resistant to SET's and SEU's [5–7]. However, when combinational circuits are mapped to non-volatile FPGA's and Application Specific Application Circuits (ASICs) the user logic is susceptible to SET's [8]. When such circuits are part of sequential logic, the SET's can be latched if it arrives within the setup and hold time of a memory element and propagate throughout the system. However, if an SET appears in the user logic but is not latched by a memory element, its effect will be transient [8]. Therefore, in such a case, redundancy to protect against SET's is not required for the user logic in non-volatile FPGAs or ASICs, but only in the memory elements. A scheme to filter SETs was previously proposed by [8], where a SET filter using Guard gates [9] was used before the input of every memory cell in a sequential circuit that was mapped to a Flash FPGA.

2.4. SET filtering techniques

A SET filtering technique for combinational logic was previously proposed using Guard Gates (GGs) [8–12]. The guard-gate is a circuit consisting of a combination of four Field Effect Transistors (FETs) with two inputs and one output. If the two inputs differ, the output floats in high impedance. In this case, the output voltage will maintain its value until leakage current degrades it. However, with the two inputs identical, the GG acts like an inverter.

The method proposed in [9] was to use the GG with a delay in the signal issued from the combinational logic cells with a delay higher than the SET pulse width. Thus, if an SET occurs at the output of the combinational circuit, it will be available immediately at the one input of the GG, but will be delayed by the delay element before it reaches the other input. The delay results in a differing value between the two GG inputs, and hence the output will maintain its value. This will only be valid if the delay in the GG input path is longer than the SET pulse width. If the SET pulse width is longer than the delay, there will be an overlap of the erroneous signals at the inputs of GG resulting in an erroneous output.

The efficiency of using the delay element will depend on the maximum SET pulse width, since the wider the allowable SET pulse, the lower the maximum allowed frequency of the mitigated design.

Instead of using an SET filter with a delay element, the combinational logic between two sequential elements could be duplicated and the outputs of these two combinational logic paths will be the GG-inputs [8]. The time performance of using a dual redundancy implementation is much improved when compared to the delay element solution. However, the area overhead is twice that of the delay element solution.

A novel double-modular-redundancy design technique for hardening combinational circuits against SET in non-volatile FPGAs was proposed in [13]. DMR is used to detect the presence of a SET in a sequential circuit. The central idea of the method proposed in [13] is to "freeze" the sequential circuit at a particular state when a SET is detected. As soon as the SET dissipates, the circuit is "unfrozen" so that it can continue with normal operation. Due to the short SET lifetime versus much longer circuit clock periods, the "frozen" state will normally not last more than one clock period. The scheme in [13] is suitable for delay-insensitive applications requiring minimal hardware overhead.

Prior circuits for removing SETs and glitches in digital circuits during high and low transitions were previously proposed using AND, NAND, OR and NOR gates [14,15]. NAND gates and NOR gates have characteristics similar to AND and OR gates that inhibit transitions on any one input from affecting the output node if the second input is set correctly. For example, when both the input signals to an AND gate or a NOR gate is low, and a change occur in one input, the output will remain low. Similarly, for a NAND gate or an OR gate, if the two inputs is high, a change in one input will not affect the output, which remains high. By using a delay in the signal issued from the combinational logic cells with a delay higher than the SET pulse width, a SET occurring when the input is low can be removed by an AND gate or similarly a NOR gate, and by a NAND gate or an OR gate if the input is high.

These circuits are able to handle SETs (or glitches) occurring either during the high or low states of a circuit line, but not both simultaneously. It would be desirable to have a single circuit that can handle a SET occurring during both the high and low states of a circuit signal. The proposed method in this paper combines the AND–OR gate circuits in parallel with a two input multiplexer to provide a single circuit that can dissipate a SET irrespective of whether the input state is high or low.

3. Detailed description of the proposed set suppressor method in FPGAs

3.1. Delay element SET suppressor

SETs occurring during signal transitions can be categorized into two types. The first type is when the SET occur while the signal is in the low voltage state and the second type occurs while the signal is in the high state.

A prior circuit, as mentioned above, for removing glitch signals occurring following the falling edge of a clock pulse is illustrated in Fig. 1. For all practical purposes, a SET can be thought of as being a glitch, and has the same effect on circuits. Therefore, a circuit that removes a glitch can theoretically also remove a SET.

Although the circuit in Fig. 1 was designed to remove glitches (or SETs) following the falling edge of a clock pulse, i.e. when the clock pulse is in the low logic state, it can equally be applied to any circuit signal, whether it is a clock signal or some other internal logic signal. The circuit consists of an AND gate and a delay element. The primary input signal is branched off and one of the branches is passed through a delay element. The delay elements produce a delayed signal in the one branch greater than the maximum time width (or lifetime) of the anticipated SETs. The input signal and the delayed input signal are input into a two input AND gate. It may happen that a SET occurs in the primary input that propagates to the AND gate. If the current value of the input signal is logic 0, as illustrated in Fig. 1, an SET will temporarily cause the input signal (and hence, the delayed signal to become logic 1). However, because SETs have a duration shorter than the delay introduced by the delay element, the SET in the primary input will arrive immediately at the one input of the AND gate and will have dissipated by the time the SET in the delayed signal reaches the other input of the AND gate. The output of the AND gate will thus, remain at logic 0, since "1 AND 0" = "0". Hence, when the primary input signal of the circuit in Fig. 1 is at logic 0, the output is insensitive to SETs with a shorter lifetime than the delay introduced by the delay element. However, the opposite is true if the input signal is at logic 1. In order to eliminate SETs occurring when the primary input logic level is high, the circuit in Fig. 2 is needed.

Fig. 2 is similar to Fig. 1 except that the AND gate has been replaced with an OR gate. With both inputs to the OR gate at logic 1, an SET which results in a temporary bit flip to logic 0 at the pri-

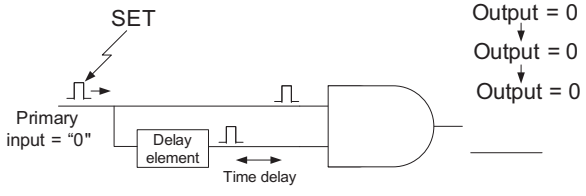

Fig. 1. AND gate SET suppressor for primary inputs at logic 0.