XAPP989 (v1.0) April 2, 2008

Correcting Single-Event Upsets with a Self-Hosting Configuration Management Core

Authors: Carl Carmichael, and Chen Wei Tseng

Summary

This application note discusses self-hosting configuration management hardware setup for Xilinx® FPGAs for the purpose of detecting and correcting single-event upsets (SEUs) to the configuration memory array. It is essential for the reader to have a firm understanding of the configuration management application as well as configuration and readback operations. An in-depth review of configuration user guides for the appropriate FPGA is strongly encouraged.

Introduction

Developing SEU mitigation schemes can be both time consuming and costly. Traditional mitigation schemes against accumulating SEUs in the Xilinx FPGA are achieved through the use of a configuration management IP core hosted in a radiation-hardened device. Although this is a proven and efficient way to mitigate SEUs, it does have a few drawbacks.

Since radiation-harden devices are often one-time programmable, a design mistake or a desired modification to the configuration management algorithm often equates to an additional device burn or even a board re-spin, resulting in additional system cost and possible delay of product release.

By shifting most or all of the configuration management's functions back into the SRAM-based FPGA, the design process can proceed without the worry of re-spinning the board to accommodate the latest recommended configuration management algorithm.

Before continuing with this application note. it is assumed that the reader has a solid understanding of configuration management. A careful review of the latest Xilinx application notes is recommend:

XAPP779, *Virtex-II Configuration Management*

XAPP988, *Virtex-4 Configuration Management*

A reference design can be obtained by contacting Xilinx technical support.

External Device Hosting Configuration Management Setup

Configuration management typically resides in a radiation-hardened device (Figure 1), which can be a radiation-hardened FPGA, CPU, or ASIC. This device bridges the Xilinx radiation-tolerant FPGA and the device storing its configuration data. Depending on the functions that the configuration management device performs, it can consume a significant amount of logic resources of the radiation-hardened device. Despite the issues mentioned in "Introduction", this solution has proved to be robust, and provides timely recovery from SEU impact.

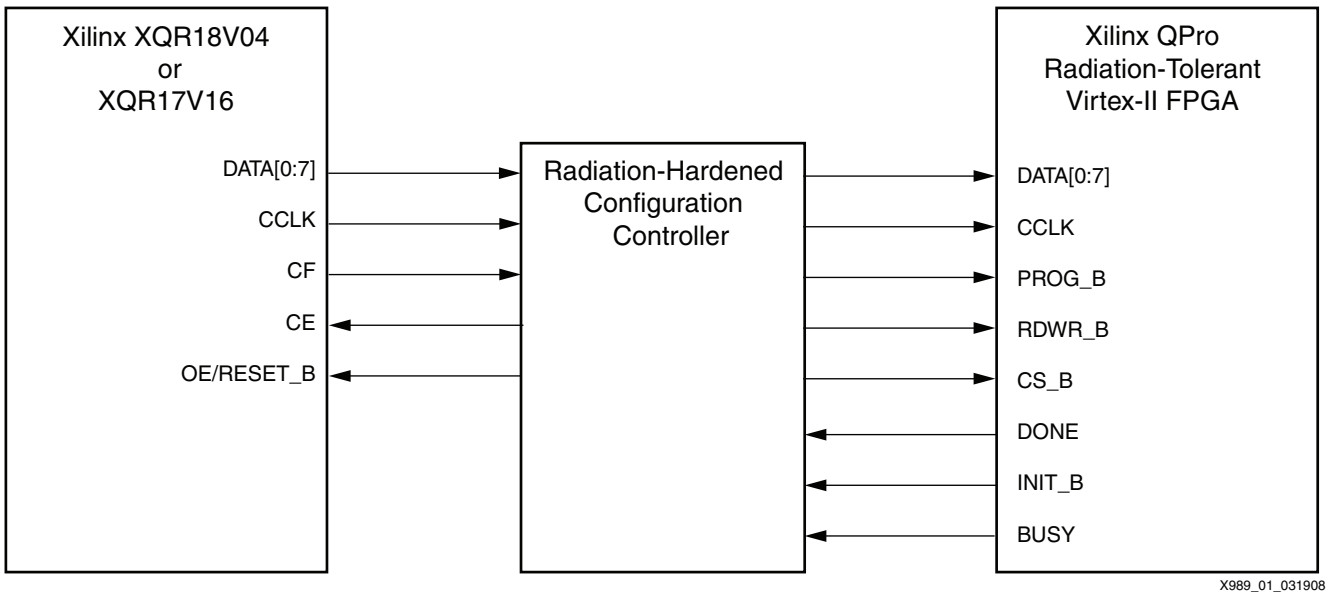

Figure 1: **Overview of External Device Hosting Configuration Management for a Single Xilinx FPGA**

Single FGPA Self-Hosting Configuration Management Setup

A more cost-effective mitigation scheme shifts the configuration management burden back into the SRAM-based FPGA (Figure 1); however, this shift increases the cross section of configuration management failure. To combat this drawback, the configuration management IP can be triplicated with aid of an external watchdog circuit.

External SelectMAP Interface Self-Hosting Configuration Management Core Setup

Figure 2 demonstrates a master SelectMap setup where the majority function of the external radiation-hardened configuration controller in Figure 1 is absorbed into the FPGA. Alternatively, the FPGA can be set to operate in a slave SelectMAP mode by routing the external oscillator's output to clock both the FPGA and the PROM.

Note: The ports shown as [2:0] in Figure 2 indicate triplication of configuration management IP inside the FPGA.

Upon power-up, all FPGA I/Os are 3-stated (assuming HSWAP_EN is pulled up, which is not shown in the diagram) The pull-downs on the CS and RDWR prepare the FPGA for data input while the pull-down on CE of the PROM readies it for data output. The radiation-hardened 3-state buffers need to allow data flow from the PROM to the FPGA. After successful configuration, the self-hosting configuration management core can perform readback, scrub, single-event function interrupt (SEFI) detection, or any other configuration management operation if desired.

The external watchdog and oscillator shown in Figure 3 provide a fallback mechanism in case of configuration failure or SEFI. The watchdog constantly monitors outputs from the configuration management core residing within the FPGA. Two signals in particular should be incorporated into the self-hosting configuration management core: a SEFI flag and RESET.

First, a SEFI flag signal that detects known FPGA SEFIs such as FAR, or SMAP. If a SEFI is detected by the configuration management core, the SEFI flag signal should be asserted for the watchdog to reset the FPGA. A pull-up is added to the output in case of an I/O SEFI that 3-states the output itself.

Second, a RESET signal should periodically pulsed to reset the watchdog's counter circuit. This signal should be closely tied to the configuration management core such that if the core itself is

upset, the RESET signal stays stagnant. A stale RESET signal should cause the watchdog to fully reconfigure the FPGA. A pull-down is attached to the output in case of an I/O SEFI.

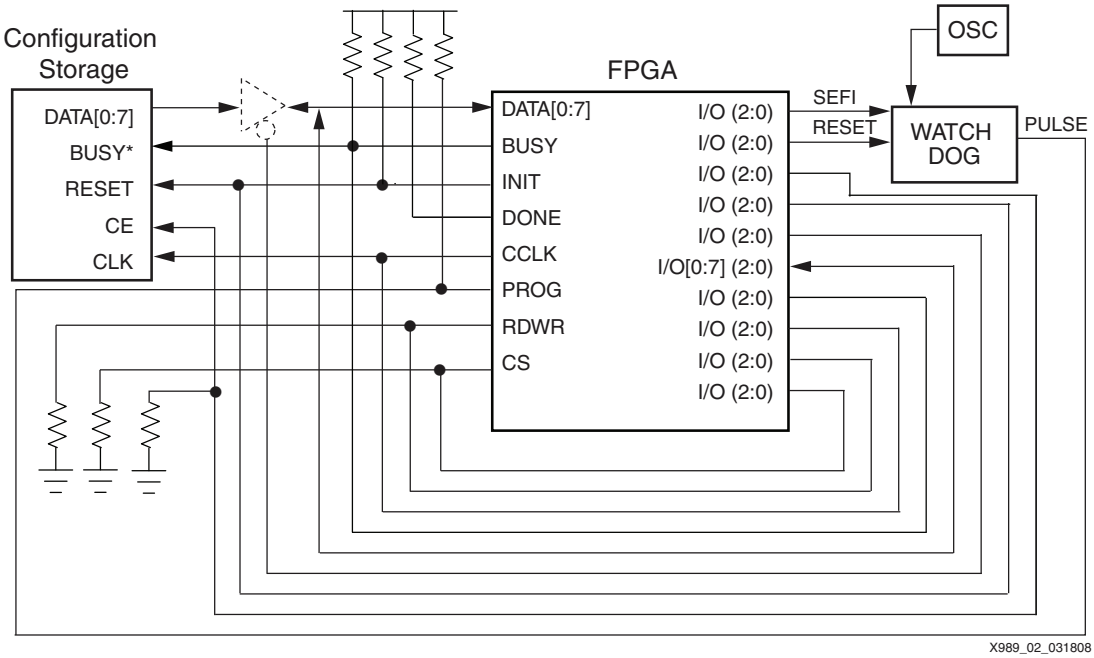

Figure 2: **Overview of single FPGA, in Master SelectMAP Mode, Self-Hosting a Triplicated Configuration Management**

When designing the watchdog's internal counter, the minimum time span for the FPGA to reconfigure and the interval of the reset output from the FPGA's configuration management, whichever has a longer period must be accounted for (preventing premature repulsing of the PROG pin).

Note: When resetting the FPGA, the watchdog needs to fire a one-shot pulse Low for a minimum duration of 300 ns to the FPGA's PROG pin.

Implementation of the watchdog should be in any radiation-hardened devices such as a space PAL. Since the watchdog circuit consume minimal resource, it can also be implemented in any free radiation-hardened resources available.

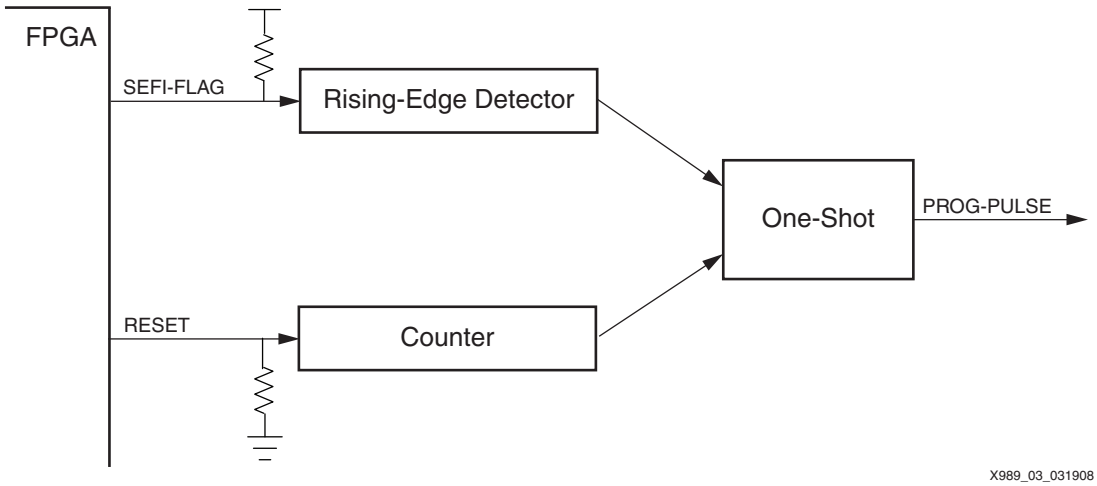

Notes:

1. The one-shot pulses the FPGA PROG pin when SEFI-FLAG is raised and counter is not reset in time.

Figure 3: **Detailed Overview of Watchdog Circuit**

Figure 2 also includes 3-state buffer, optional if the configuration storage device outputs can be placed into high-impedance state. Although Xilinx 18V00 and 17V00 PROMs outputs can be 3-stated by asserting OE/RESET, this input also resets the PROM address. This situation can create a problem if the configuration bitstream (such as for Virtex®-II FPGAs which have the block RAM interconnect bits located after block RAM content bits) is not tailored for scrubbing.

To eliminate the need for the 3-state buffer while using the Xilinx PROMs, the user must either implement the setup discussed in "Internal Configuration Access Port Interface Self-Hosting Configuration Management Core Setup" or rearrange the configuration bitstream to better suit the scrubbing process.

Internal Configuration Access Port Interface Self-Hosting Configuration Management Core Setup

A variance of the setup can eliminate the external watchdog and 3-state buffer circuits at the expense of increased cross section by utilization the internal configuration access port (ICAP) feature of the FPGA. This setup is only supported by Virtex-II or later FPGAs since ICAP does not exist in prior FPGA families.

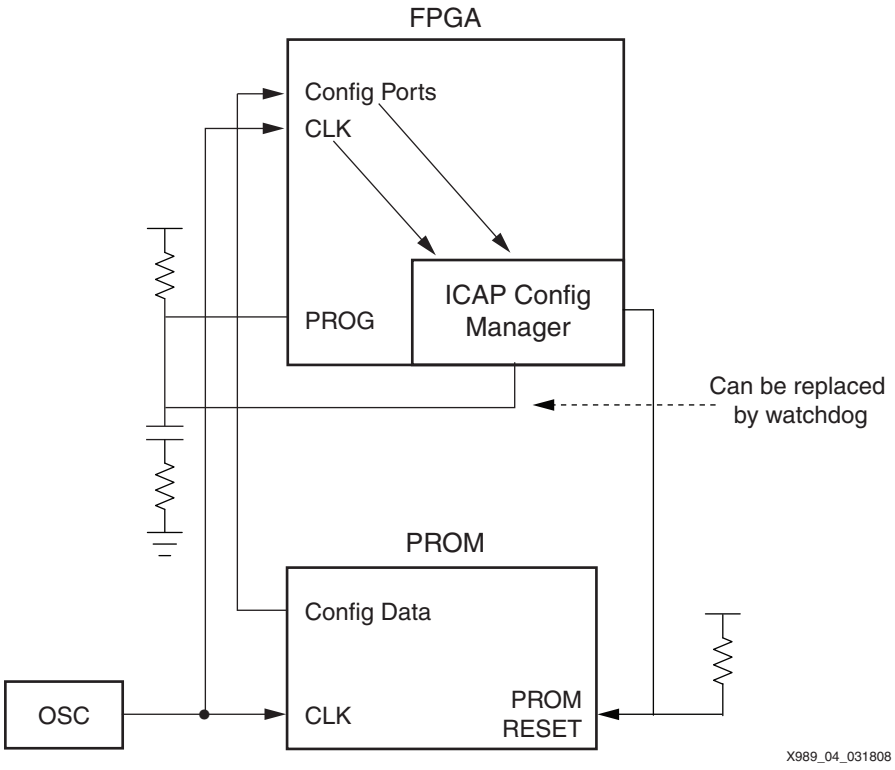

Figure 4: **Single Device Self Configuration with no External Circuitry**

Note: For simplicity purposes, Figure 4 merges the SelectMAP interface into Config Ports. pull-downs on CS and RDWR pins are assumed. It is also assumed that the ports the configuration management core are triplicated as well as.

Upon power-up, the PROM configures the FPGA through the configuration ports. After initial configuration, the ICAP configuration management core can perform optional readback through the ICAP interface and scrub with data arriving through the initial configuration ports. Alternatively, the configuration management core can perform constant scrubbing.

If a SEFI is detected, the ICAP configuration management core fires a Low pulse to the PROG pin of the FPGA. A capacitor should be added to ensure the pulse remains Low for a minimum of 300 ns.

Note: This setup must have the FPGA not in "persist" mode so configuration ports can convert into user I/Os after device configuration, granting configuration data access to user application, and ICAP.

The ICAP configuration management core can control the source (either from the external PROM, or internally via the core) for configuration data and commands (Figure 5). The actual implementation of the core could have both the MUX and the ICAP primitive residing within the ICAP configuration management core itself.

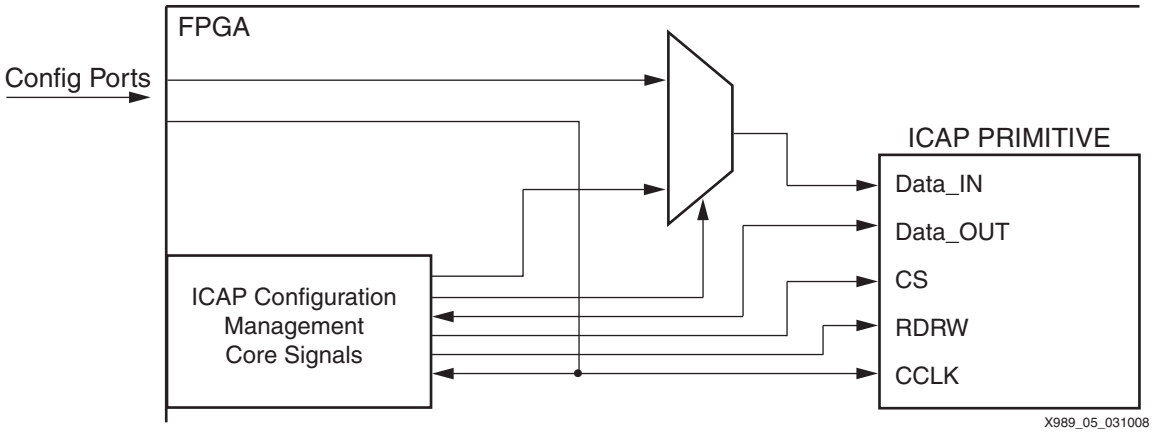

Figure 5: **Interface of the ICAP Configuration Management Core**

Multiple FPGAs Self-Hosting Configuration Managements Setup

Having multiple FPGAs in a system can enable an alternative self-correction technique. In addition to having each FPGA monitoring its own configuration integrity, a chain can be formed for a stronger mitigation scheme. In this chain one FPGA monitors the next FPGA in the chain (Figure 6).

The strength of multiple FPGAs self-hosting a configuration management setup increases as more FPGAs are included in the chain. The only way for a complete system failure to occur is if all of the FPGAs in the chain experience a SEFI within 300 ns of each other (PROG pulse time requirement) plus the time the previous FPGA takes to detect a SEFI. In this case, the previous FPGA monitoring the next FPGA in the chain detects the problem but has insufficient time to completely re-initialize the next FPGA in the chain. However, if all FPGAs in the chain are down due to SEFIs but are reset, the behavior is identical to an initial system boot up. If such a slim window of failure cannot be tolerated, a watchdog circuit can be added to the system to capture a total system failure.

External SelectMAP interface Self-Hosting Configuration Management Core Setup

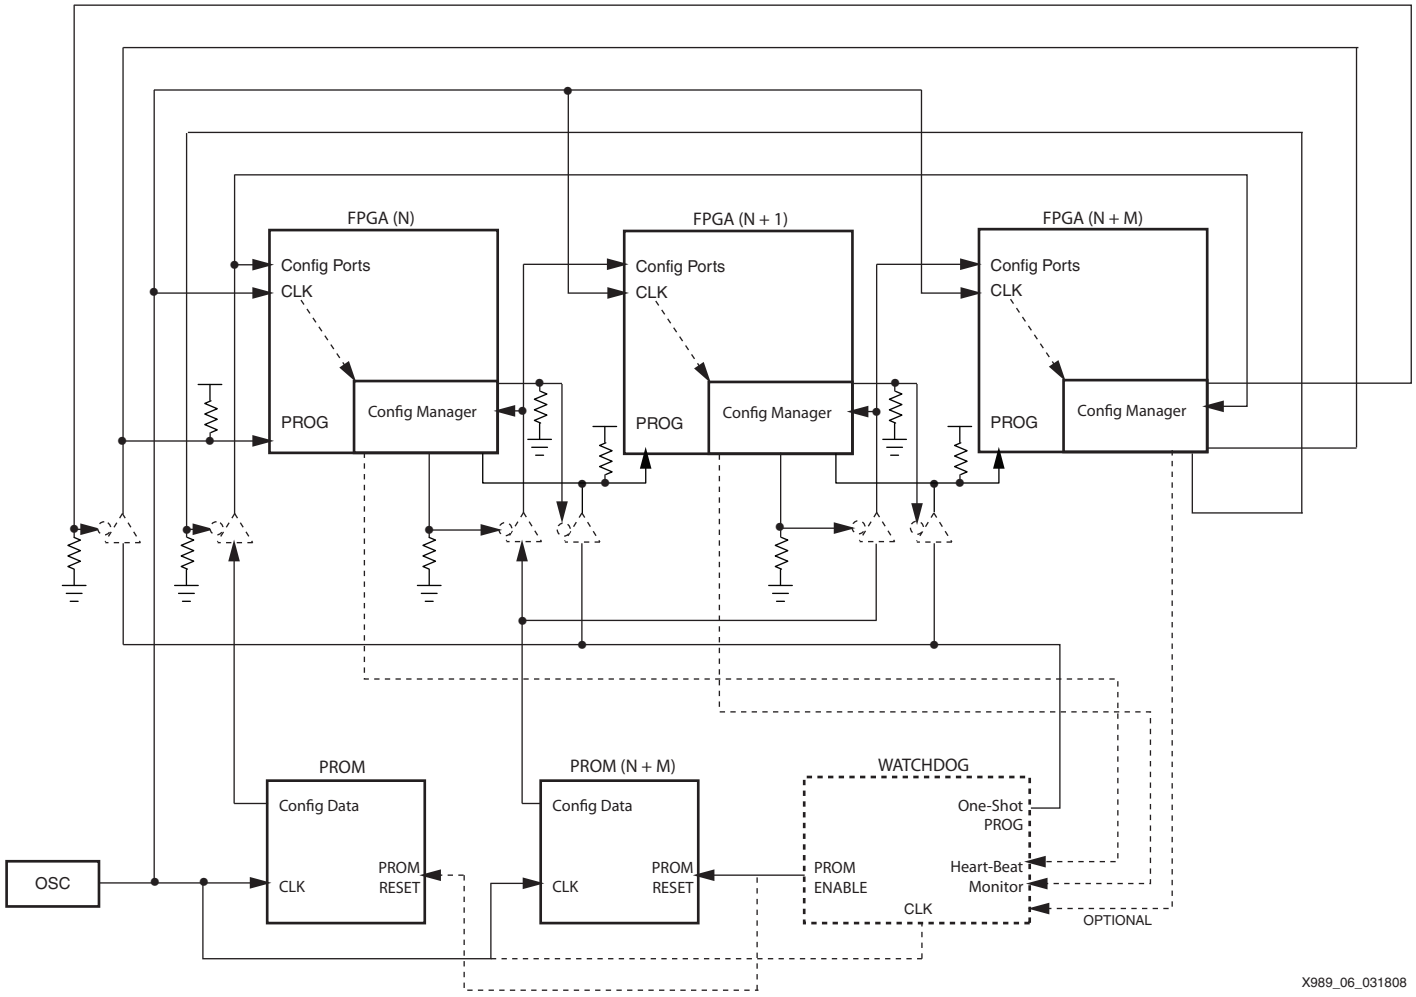

Figure 6: **Overview of Multiple FPGAs, in Slave SelectMAP Mode, Self-Hosting Configuration Managements**

As displayed in Figure 6, a loop of FPGAs is formed with each FPGA hosting the configuration management for the next FPGA in the chain. This setup also allows for heterogeneous FPGA applications with the deployment of multiple sets of PROMs up to N + M.

Triplication of the configuration management can be optional since upsets to the configuration management are now corrected by the previous FPGA in the chain. As long as one FPGA in

the chain functions, all configuration management cores eventually regain their functionality. The robustness of this setup and the need to triplicate the configuration management core are dependent on the number of FPGAs in the chain and the design of the core.

When the system powers up, the FPGAs are configured by the attached PROMs (the configuration ports such as CS and RDWR should be pulled downs). The active-Low 3-state buffers default to pass configuration data from the PROMs to the FPGAs.

After FPGAs are configured, the configuration manager controls the 3-state buffers for the configuration data ports and PROG inputs. The configuration manager can then start scrubbing or readback processes and output an optional heart-beat signal to the watchdog or configuration management in the previous FPGA in the chain. To synchronize the configuration management with the PROM data, the configuration manager can be clocked by the external oscillator clocking the FPGA's CCLK port. Although an FPGA in the chain can be set to master mode and provide the clock, this approach ties potential system failure to just the failure of the single FPGA driving the clock.

The optional watchdog's function is to monitor the configuration management cores in the chain. If implemented, it should at least monitor one heart beat output of the FPGA. When the heart-beat output of the FPGA is lost for the duration of the time it takes to program all but one FPGA in the chain, it can be assumed that all FPGAs in the chain have failed due to SEFI and a one-shot Low pulse to all FPGA's PROG pins should be triggered. Monitoring more heart-beat outputs can accelerate determination of when the entire chain of FPGAs is down.

The PROM enable output of the watchdog needs to account for the minimum FPGA initialization time before configuration can commence. After the FPGAs are operational, the PROM enable output could periodically enable PROM data output for the purpose of scrubbing if needed. If no scrubbing is required, the configuration manager can simply ignore the incoming data from the PROM by 3-stating the data buffer. The watchdog can also monitor optional PROM data output requests from the FPGA's configuration management for immediate scrubbing operation (not shown in the Figure 6)

Note: If the watchdog is not implemented, the control of the PROM's reset should be granted to the configuration management core and default to enable PROM data output.

The optional 3-state buffer shown in Figure 6 can be eliminated if the configuration storage device outputs can be placed into an high-impedance state. Although Xilinx 18V00 and 17V00 PROMs outputs can be 3-stated when OE/RESET is asserted, this input also resets the PROM address. This situation can create a problem if the configuration bitstream (such as for Virtex-II FPGAs which have the block RAM interconnect bits located after block RAM content bits) is not tailored for scrubbing.

To eliminate the need for the 3-state buffer while using the Xilinx PROMs, the user must either implement the setup discussed in "Internal Configuration Access Port Interface Self-Hosting Configuration Management Core Setup" or rearrange the configuration bitstream to better suit the scrubbing process.

Other variations are available to speed up the SEU correction process. For example, by adding a connection from the FPGA's INIT line to the PROM's reset line can bring the system online sooner. Since the FPGA need to be initialized upon power-up before it is able to accept configuration data, INIT indicates when the FPGA is ready for configuration. Without additional monitoring on the INIT line, the PROM enable output from the watchdog needs to account for the minimum FPGA initialization time (T_{POR}) before enabling PROM data output.

Internal Configuration Access Port Interface Self-Hosting Configuration Management Core Setup

Similar to the single-device self-hosting configuration management setup without external circuitries, the multiple-device self-hosting configuration management setup can eliminate the need for any external component. The setup again relies on ICAP to accommodate

configuration storage devices that cannot 3-state their outputs and eliminate the need of external 3-state buffers (Figure 7).

Upon power-up, the FPGAs are configured through the configuration ports. The FPGAs configuration pins can not be in "persist" mode in order to pass of PROM data for scrubbing purposes.

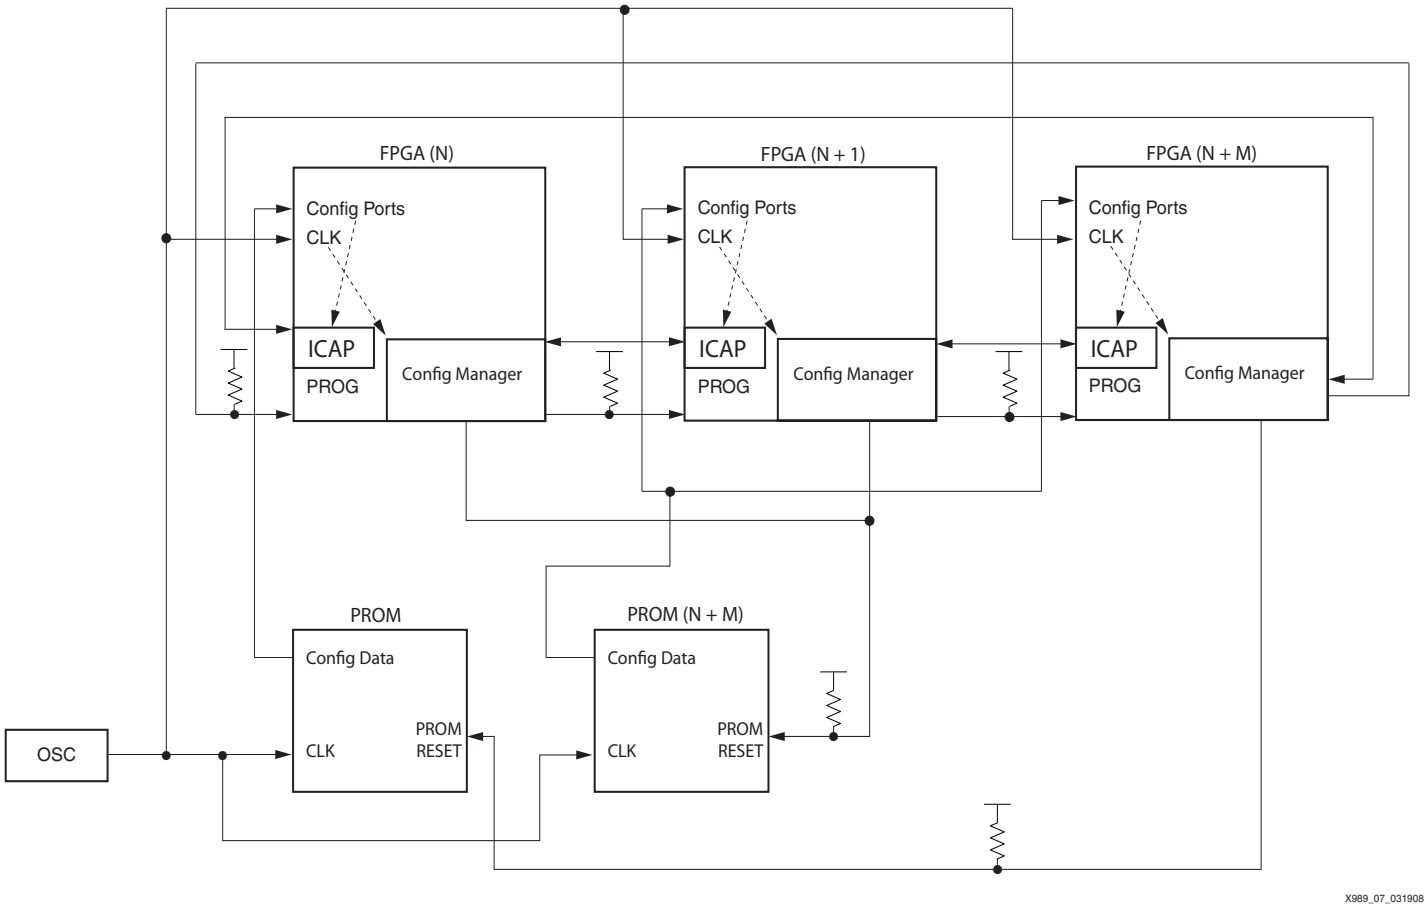

Figure 7: **Setup Overview for Multiple FPGAs Self Configuration**

The setup now involves two parts: first, an ICAP setup that allows data flow selection between the configuration management core and the PROM as shown in Figure 8; second, the configuration management core controlling the next FPGA in the chain.

The ICAP setup should be more than a mere instantiation. It should include data selection and control to accept configuration commands and data from either the PROM or the upstream configuration management core. Additionally, since ICAP has DATA_IN and DATA_OUT separated, the ICAP setup could include bidirectional buffering (not shown in Figure 8) to save I/O usage.

The configuration management core should vary slightly from the single-device core setup since it now should control the configuration command and data inputs from either the PROM or the core to the ICAP in the next device. The core could also control the bidirectional I/O access to the ICAP primitive for the downstream monitored FPGA.

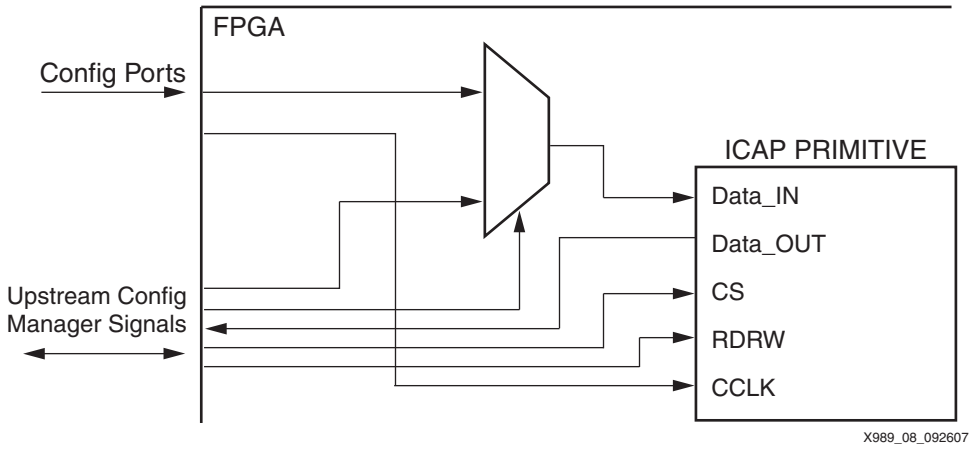

Figure 8: **Interface of the ICAP Configuration Management Core**

Self-Hosting Configuration Management Core Design Considerations

Implementation of self-hosting configuration management has the drawback of slightly increased SEFI cross section. To minimize the impact of this issue, the configuration management should differ slightly from what is typically implemented in external radiation-hardened devices:

- Readback is strongly recommended and scrubbing should be initiated only if readback has errors, minimizing the risk of the upset configuration management core inadvertently corrupting the FPGA.
- The SEFI declaration and PROG output control should be triplicated to reduce accidental reconfiguration of the FPGAs.
- To further reduce configuration management failure, the configuration management core should be reset periodically to ensure stored data integrity.

Conclusion

Configuration mitigation is necessary for applications operating in harsh environments. However, mitigation schemes can quickly mount up to be both a design and cost burden. Self-hosting configuration management can reduce system cost, accelerate system delivery, and assure access to the most up-to-date configuration management algorithm.

Revision History

The following table shows the revision history for this document.

Date	Version	Revision
04/02/08	1.0	Initial Xilinx release.

Notice of Disclaimer