

FPGA-BASED SINGLE CHIP CRYPTOGRAPHIC SOLUTION (U)

By Mark McLean and Jason Moore

(U) ABSTRACT

(U) The use of Field Programmable Gate Arrays (FPGAs) in Type I Cryptographic equipment has historically been limited. While FPGA use is allowed, restrictions on how they are used can result in inefficient processing and an increase in system size, weight and power. For example, redundancy and isolation of functionality is required through physically separate devices. This paper introduces new technology that will provide the industry with an FPGA-based single chip cryptographic solution.

(U) The National Security Agency (NSA) and Xilinx have been working together to bring the advantages of FPGA technology to the High Assurance industry. Utilizing the Xilinx Virtex-4 FPGA, the NSA and Xilinx have developed a design flow and verification process based on NSA requirements for high-grade cryptographic processing. This paper will outline the design flow process and summarize the results of the evaluation effort.

(U) INTRODUCTION

(U) The flexibility of Programmable logic affords the Aerospace and Defense Industry many advantages. However, in certain applications like Information Assurance, government contractors and agencies have not been able to realize the full capability of programmable logic due to isolation, reliability and security concerns.

(U) In an effort to address these concerns, the NSA and Xilinx worked together to analyze and evaluate the Xilinx Virtex™-4 series of FPGAs. The goal of this analysis was to understand the design, construction, and functionality of the Virtex-4 such that a solution could be developed that would allow independent functions to operate on a single chip. Examples of such applications include, but are not limited to: Redundant Type I encryptors operating on a single chip, Red and Black data resident on a single chip and functionality operating on multiple independent levels of security on a single chip.

(U) The successful completion of the isolation analysis against the NSA Fail Safe Design Assurance (FSDA) specifications, and an exhaustive vulnerability assessment, has allowed the NSA and Xilinx to develop new technology for the industry. This paper will summarize the analysis completed and then describe, in detail, the design flow required to realize the solution.



(U) VIRTEX-4 FPGA Overview

(U) The Virtex™-4 family is Xilinx's 90nm Field Programmable Gate Array (FPGA) technology. It is based on an innovative Advanced Silicon Modular Block or ASMBL™ column-based architecture. Virtex-4 FPGAs contain three families or platforms: LX, SX and FX. The Virtex-4 LX devices are logic-centric with a high ratio of logic and I/O to feature. The Virtex-4 SX devices are signal processing-centric with a high ratio of embedded signal processing elements and internal memory to logic. The Virtex-4 FX devices are System-on-a-chip (SoC) solutions that include embedded PowerPC™ processors, Tri-Mode Ethernet MACs and 6.5Gb/s serial transceives. Virtex-4 devices are produced on a 90-nm copper process using 300mm (12 inch) wafer technology.

(U) SINGLE-CHIP ISOLATION ANALYSIS

(U) The level of logical and functional isolation that can be achieved within a single Virtex-4 FPGA was one of the critical analyses performed. The complexity, and sheer density, of the Virtex-4 device would seem to preclude any efficient isolation analysis. In addition, the difficulty in performing such an analysis is assumed to be magnified considering there are seventeen FPGAs total in the Virtex-4 family. However, Xilinx FPGAs are very modular devices. No matter which device is chosen, they consist of the same basic building blocks tiled over and over again. In the Virtex-4 FPGA, the basic building block is a Configurable Logic Block (CLB). As shown in Figure 1, the CLB includes a Global Switch Matrix (GSM) used for interconnect that is common for all features in the FPGA. In other words, the exact same GSM is used to provide interconnect from CLB to CLB, or from any CLB to dedicated functions within the FPGA like embedded PowerPC™ processors, Digital Clock Managers (DCM), DSP48 Blocks, etc. Therefore, the isolation analysis focused on an exhaustive review of the CLB and GSM.

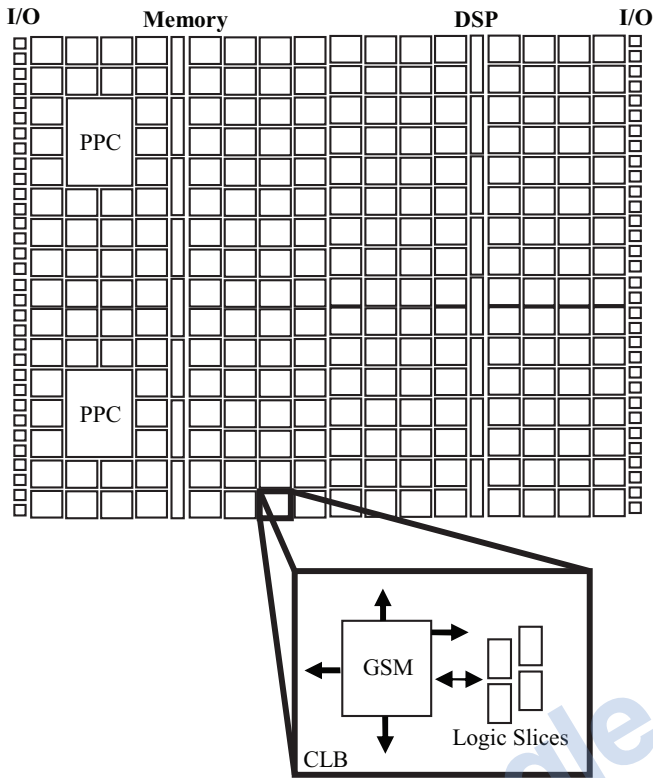


Figure 1. (U) The GSM is common to all features in the FPGA and hence was the focus of the isolation analysis.

(U) In order to achieve isolation within a single device the concept of a *fence* is introduced. The *fence* is nothing more than a set of unused CLB(s). In the *fence*, no routing or logic may be present. The result of the analysis showed that a single unused CLB between isolated regions provides, at a minimum, three physical failures of isolation, and hence exceeds the FSDA requirements. A hypothetical FPGA design with three isolated regions, separated by the fence, is shown in Figure 2.

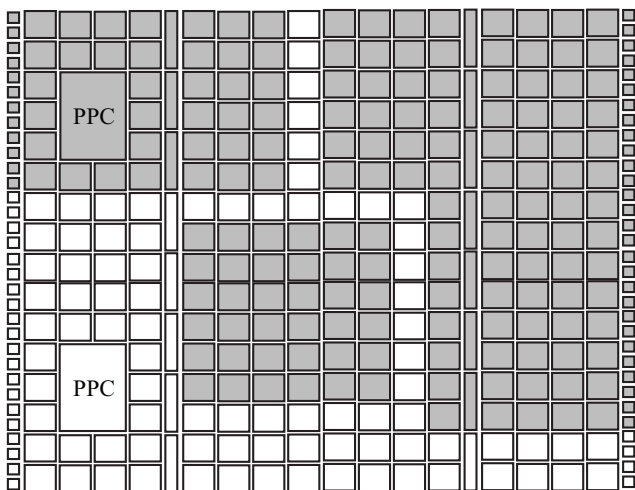


Figure 2. (U) User-defined isolated regions must be separated by one unused CLB to meet FSDA requirements.

(U) The determination of the number of unused CLBs to constitute a reliable fence meeting the FSDA requirements is but one part of the solution. This assumes that the FPGA development tools can then control the placement of logic and routing to keep the CLB completely unused. In addition to controlling the logic and routing, verification of the implementation is also required.

(U) Historically, Xilinx development tools have provided FPGA designers very granular logic placement. However, the same level of routing control has either been non-existent or left to only those with an in-depth understanding of very complicated routing constraints. Recently, Xilinx has invested in the maturation of its Partial Reconfiguration (PR) technology. This technology allows for a portion of the FPGA to be re-configured, while the rest of the FPGA operates normally. This technology is applicable to Information Assurance applications for two reasons. First, it will provide the ability to support algorithm agility – the ability of adding and tearing down multiple algorithms in a single FPGA without disrupting those that are actually processing data. Second, in order to reconfigure a portion of the FPGA, without disrupting the remaining circuits, a level of isolation must exist, and hence routing must be controlled. Therefore, by using the PR Design Flow and Toolset, both logic and routing can be controlled to create a *fence* between isolated regions in a single chip.

(U) SINGLE CHIP Design Flow

(U) Leveraging the Partial Reconfiguration capability and design tools for Virtex-4, a user can develop isolated regions of logic within a single FPGA. As mentioned earlier, the PR toolset provides the user the ability to contain routing within a specific region. In addition to the PR toolset, a couple of unique capabilities have been added.

(U) A new design constraint, called NOBOUNDARYCROSS, instructs the Place and Route (PAR) tools to keep the entirety of a route within a specified area. This is important as the tools now see the entire interconnect as a utilized route, whether it is or not. For example, a certain type of route in the FPGA, known as a HEX route shown in Figure 3, traverses from a single source to two destinations that are three and six CLBs away.

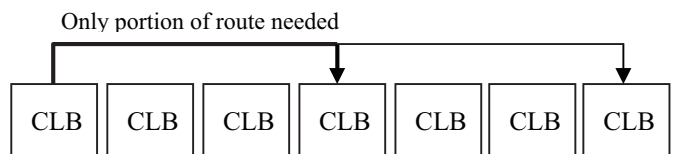


Figure 3. (U) PAR Tools understand that the entirety of a given route is used even when only a portion is needed.

(U) If a signal placed on this route gets off at the destination only three CLBs away, the router understands that the entirety of the route is actually used. Therefore the placement is constrained such that the entirety of the route is contained within the isolated

region and no part of the HEX line is allowed to cross the user-defined boundary.

(U) Another significant addition to the PR toolset was the ability to instantiate user inputs and outputs (I/O) in lower level Hardware Descriptive Language (HDL) modules. Isolated logic is defined in the PR flow through hierarchy and lower level HDL modules. For isolation to be maintained to the boundary of the FPGA, the I/Os must be contained within the isolated regions. The only way to include the I/Os in the isolated regions is to allow the user to infer I/Os in lower level HDL modules. This is done with a simple attribute statement to allow the FPGA designer the ability to create generic HDL and simulate the design as normal.

(U) Single Chip-Isolation Verification

(U) Once the design running on the FPGA has been completed it is necessary to ensure that the isolation designed into the solution still exists and has not been compromised. Xilinx's FPGAEEditor tool provides a compelling visual representation of the placement and routing of the design.

(U) However, FPGAEEditor presents an abstract view of the actual silicon and does not display all the information needed for isolation verification. Therefore, Xilinx has developed a verification tool for NSA known as IVT (Isolation Verification Tool) to address this issue.

(U) The IVT was developed independently from the Xilinx® ISE™ software but utilizes the same FPGA design database and code libraries. IVT can be used in two different ways. First, IVT is used to analyze the final FPGA design to ensure the isolation designed into the FPGA, did not get compromised due to errors in the design flow. Second, IVT can be used early in the design process to identify potential mistakes before releasing a Printed Wiring Board (PWB) to fabrication and working through the FPGA implementation process.

(U) When analyzing the final design, the placed and routed designs from the Partial Reconfiguration Design Flow are inputs into the IVT. In this design flow, a design database, also known as an .ncd file, will exist for each isolated region, and the complete merged design. Knowing what functionality exists in each isolated region, and how the isolated regions have been merged to form one final design provides all the information that is required by the IVT to verify isolation. IVT performs an exhaustive analysis tracing every potential interconnect path that could be created to determine if the hypothetical interconnect would cause the isolation to be compromised. If a path between regions could exist within less than a user defined number of failures, IVT will generate an error and report the hypothetical failing path to the designer.

(U) In addition to the logic and routing isolation analysis, the IVT also analyzes device pin and package locations to ensure that if red and black data exist in the same FPGA, red I/Os and

black I/Os do not share the same bank and are not physically adjacent in both the silicon and at the device package (i.e. the ball of the FPGA Ball Grid Array (BGA) package).

(U) As mentioned earlier, the IVT can also be used to identify potential problems early in the design flow process. IVT can analyze an FPGA Designers User Constraint File (.ucf). The constraints file is how the system designers creates isolated regions in the FPGA with proper logic and routing placement. The analysis can be done early in the design process, which will help eliminate costly PWB layout mistakes.

(U) Communication between isolated regions

(U) Communication between isolated regions is expected for applications such as the transfer of keying material or the comparison of redundant encryptors output. In these scenarios the use of Bus Macros are required.

(U) Bus Macros, as shown in Figure 4, are small pieces of IP that are pre-designed by Xilinx, analyzed by NSA, and then used by the FPGA designer. Since they are designed in advance and provided in the form of an .nmc file, they are considered hard macros since the logic and routing used is pre-defined. The value of the Bus Macro is that it eliminates un-deterministic routing behavior when providing a communication path between isolated regions. Bus Macros are commonly used in commercial Partial Reconfiguration applications. A set of special Bus Macros have been designed by Xilinx to meet the FSDA requirements.

(U) There are a number of different bus macro types varying in direction, data width, asynchronous vs synchronous and with and without enables.

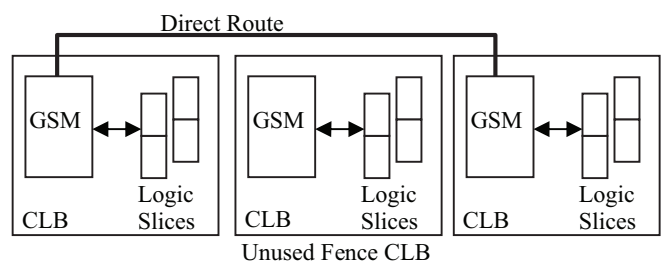


Figure 4. (U) Logical representation of a Bus Macro

(U) SECURITY EVALUATION

(U) Before a device can be used in a U.S. cryptosystem, a thorough security analysis must be completed. This analysis is conducted to determine the level-of-effort required by an adversary to exploit the device. As a result of the analysis, strategies can be developed and deployed to strengthen system security. FPGA technology is critical to future crypto system development and the Xilinx Virtex-4 (V4) was the most comprehensive FPGA evaluation done to date. The evaluation results and test setups are available from NSA with the proper clearance and a valid need-to-know.

(U) The Virtex-4 was tested against publicly known vulnerabilities such as Differential Powerline Analysis, Transient Fault Analysis and JTAG attacks. The device was also tested to ascertain remanence issues within the V4's SRAM cells that could potentially be exploited by an adversary. The remanence evaluation is focusing mainly on the Battery Backed RAM (BBR), and will include a cursory evaluation of the configuration RAM. Additionally, NSA applied classified in-house developed evaluation techniques in an attempt to better quantify the overall security of the V4.

(U) The evaluation also included analysis of the Verilog code used to synthesize the V4's control, security and AES key programming functions. These functions are critical to the secure operation of the device.

(U) Due to the high security requirements of U.S. crypto equipment, a security monitor was developed to be implemented within the V4 fabric to further enhance security. The monitor will always be vigilant and handle all security related functions for the device. This IP core will remove the burden of physical security requirements from the system designers and allow them to concentrate more on their particular application. The security monitor will be embedded into all designs processing sensitive data. The security monitor will be ITAR controlled and will be available from Xilinx.

(U) When used in conjunction with the security monitor, NSA found the V4 to be a robust architecture capable of processing classified information and maintaining a very high level of security.

(U) SUMMARY

(U) Working together, NSA and Xilinx have developed technology that will enable the Information Assurance industry to maximize the advantages of programmable logic. Using published design guidelines, verification tools and the Security Monitor IP, Xilinx Virtex-4 FPGAs can be used to support redundancy, red/black data, and multiple levels of security on a single chip.



Mark McLean is a technical director for the National Security Agency. He is responsible for microelectronic evaluations of devices that are used in secure systems, and his guidance not only supports the Information Assurance Directorate at NSA, but the entire U.S. government and its NATO allies. He has more than 10 years of experience in his present field and more than 23 years of experience as a design engineer. He has a Bachelor's in Computer Science from National University, a Master's in Computer Engineering from Loyola, and is currently working on his PhD in Computer Science, focusing on neural networks, at the University of Maryland at College Park.

National Security Agency

Fort George G. Meade, MD 20755
 410-854-6561
 mrmclean@restarea.ncsc.mil
 www.nsa.gov



Jason Moore is a program manager in the Aerospace and Defense Division of Xilinx, responsible for FPGA technology development in support of major DoD Transformation programs (primarily the Joint Tactical Radio System and crypto modernization). Previous to Xilinx, Jason was responsible for the development of FPGA-based communication security equipment in a wide range of avionics and ground-based platforms at the Motorola Government Group. He was awarded two U.S. patents on cryptographic design and has a patent application pending for novel approaches to address logical and functional isolation within a single FPGA. Jason has 14 years of experience in the aerospace and defense field and holds a BSEE from New Mexico State University.

Xilinx
 7801 Jefferson St
 Albuquerque, NM 87109
 505-292-0750
 jmoore@xilinx.com
 www.xilinx.com