

available at www.sciencedirect.comjournal homepage: www.elsevier.com/locate/cose

**Computers
&
Security**



Enforcing memory policy specifications in reconfigurable hardware

Ted Huffmire^{a,*}, Timothy Sherwood^b, Ryan Kastner^c, Timothy Levin^a

^aNaval Postgraduate School, Department of Computer Science, Monterey, CA 93943, USA

^bUniversity of California, Santa Barbara, Department of Computer Science, Santa Barbara, CA 93106, USA

^cUniversity of California, San Diego, Department of Computer Science and Engineering, La Jolla, CA 92093, USA

ARTICLE INFO

Article history:

Received 14 February 2008

Accepted 9 May 2008

Keywords:

Reconfigurable hardware

Protection mechanisms

Security and privacy protection

Access controls

Covert channels

ABSTRACT

While general-purpose processor based systems are built to enforce memory protection to prevent the unintended sharing of data between processes, current systems built around reconfigurable hardware typically offer no such protection. Several reconfigurable cores are often integrated onto a single chip where they share external resources such as memory. While this enables small form factor and low cost designs, it opens up the opportunity for modules to intercept or even interfere with the operation of one another. We investigate the design and synthesis of an FPGA memory protection mechanism capable of enforcing access control policies and a methodology for translating formal policy descriptions into FPGA enforcement mechanisms. The efficiency of our access language design flow is evaluated in terms of area and cycle time across a variety of security scenarios. We also describe a technique for ensuring that the internal state of the reference monitor cannot be used as a covert storage channel.

Published by Elsevier Ltd.

1. Introduction

Reconfigurable hardware is at the heart of many high performance embedded systems. Satellites, set-top boxes, electrical power grids, and the Mars Rover all rely on Field Programmable Gate Arrays (FPGAs) to perform their respective functions for everything from encryption to FFT, or even entire customized processors. The bit-level configurability of these devices can be used to implement specific logic circuits that are highly optimized compared to the processing required in a general-purpose CPU. Because the logic of the fabricated device is reconfigurable, special-purpose circuits can be developed and deployed at a fraction of the cost associated with custom fabrication (e.g., ASIC). Furthermore, the logic on an FPGA board can even be changed in the field. These advantages of

reconfigurable devices have resulted in their proliferation into critical systems, yet many of the security primitives which software designers take for granted in general-purpose processors are simply nonexistent.

Due to Moore's law, FPGAs today have enough transistors on a single chip to implement over 200 separate RISC processors. Increased levels of integration are inevitable, and reconfigurable systems are no different. Current reconfigurable systems-on-chip include diverse elements such as specialized multiplier units, integrated memory tiles, multiple fully programmable processor cores, and a sea of reconfigurable gates capable of implementing significant ASIC or custom datapath functionality. The complexity of these systems and the lack of separation between different hardware modules on the FPGA device has increased the possibility that security

* Corresponding author.

E-mail addresses: tdhuffmi@nps.edu, televin@nps.edu (T. Huffmire), sherwood@cs.ucsb.edu (T. Sherwood), kastner@cs.ucsd.edu (R. Kastner).

0167-4048/\$ – see front matter Published by Elsevier Ltd.

doi:10.1016/j.cose.2008.05.002

vulnerabilities may surface in one or more components, which could threaten the entire device. New methods are needed to provide separation and security in these highly integrated reconfigurable devices.

One of the most critical aspects of separation that needs to be addressed is in the management of external resources such as off-chip DRAM. While a general-purpose processor will typically provide virtual memory mapping primitives such as TLBs that are used to enforce some form of memory protection, reconfigurable devices usually operate in a flat physical address space with a flat program structure (e.g., without underlying operating system support). Lacking these mechanisms, the FPGA environment is assumed to be benign, since any hardware module can normally read or write to the memory of any other module at any time. Whether purposefully, accidentally, or maliciously, destructive interference between cores can result. This situation calls for a *memory access policy* and related control mechanisms that all modules on chip must obey. In this paper we present a method that utilizes the reconfigurable nature of field programmable devices to provide a mechanism to enforce such a policy.

In the context of this paper, a *memory access policy* is a description of what accesses to memory are legal and which are not. Our method rests on the ability to formally describe the access policy using a specialized language. The formalism results in two significant capabilities: the ability to reason about policy soundness and the ability to automatically derive refinements to the policy. We present a set of tools through which the policy description can be automatically transformed and *directly synthesized to a circuit*. This circuit, represented as a bit-stream, can then be loaded into a reconfigurable hardware module and used as an execution monitor to analyze memory accesses of individual cores on the FPGA and enforce the memory access policy.

The techniques presented in this paper are steps towards a cohesive methodology for those seeking to build reconfigurable systems that can securely control data at different sensitivity labels and modules acting at different security clearance levels on a single chip (i.e., systems that can provide multi-level security). In order for such a methodology to be accepted by the embedded design community it is critical that the resulting hardware provides both high performance and efficient use of the FPGA fabric. Within the security community, the methods must be formally grounded. Finally, the integration of these requirements must be understandable to those in both communities. Throughout this paper we strive to strike a balance between engineering and formal evaluation; between performance, security, and clarity. Specifically, this paper makes the following contributions:

- We specify a memory access policy language, based on formal regular languages, for expressing the set of legal accesses and allowed policy transitions for stateful policies.
- We demonstrate how our language can express classical security scenarios, such as isolation, controlled sharing, and Chinese wall.
- We present a policy compiler that translates an access policy described in this language into a synthesizable hardware module.

- We evaluate the effectiveness and efficiency of this novel enforcement mechanism by synthesizing several policies down to a modern FPGA and analyzing the area and performance.

In this article, we extend our preliminary work (Huffmire et al., 2006) to incorporate:

- A more thorough discussion of the architecture of reconfigurable systems.
- A motivating example of a reconfigurable system from the field of computer vision.
- Additional example policies and synthesis results, including B&L, Biba, high water mark, and dynamic policies.
- A description of a technique to prevent the internal state of the reference monitor from being used as a covert storage channel.
- Substantial revisions and corrections throughout the paper.

The remainder of the paper is organized as follows: Section 2 provides background on FPGAs and describes the threat model that we are addressing. In Section 3, we explain the algorithms behind our reference monitor design flow. In Section 4, we describe our access policy language including several example policies. We present our reference monitor synthesis results in Section 5. We describe our technique for preventing the reference monitor from being used as a covert storage channel in Section 6. Finally, we conclude in Section 7 and discuss where there is room for future work.

2. Reconfigurable systems

Increasingly we are seeing reconfigurable devices emerge as the flexible and high performance workhorses inside a variety of high performance embedded computing systems (Bondalapati and Prasanna, 2002; Compton and Hauck, 2002; DeHon and Wawrzynek, 1999; Kastner et al., 2004; Mangione-Smith et al., 1997; Schaumont et al., 2001). The power of reconfigurable systems lies in the immense amount of flexibility that is provided. Designs can be customized down to the level of individual bits and logic gates. They combine the post-fabrication programmability of software running on a general-purpose processor with the spatial computational style most commonly employed in hardware designs (DeHon and Wawrzynek, 1999). Reconfigurable systems use programmability and regularity to create a flexible computing fabric that can lower design costs, reduce system complexity, and decrease time to market, while achieving 100× performance gain per unit silicon as compared to a similar microprocessor (Buell and Pock, 1995; DeHon, 1998; Vuillemin et al., 1996). The growing popularity of reconfigurable logic has forced practitioners to start to consider the security implications, yet the resource constrained nature of embedded systems is a challenge to providing a high level of security (Kocher et al., 2004). To provide a security technique that can be used in practice, it must be both robust and efficient. To understand what is a practical design, we must first examine the architecture of a modern reconfigurable system.

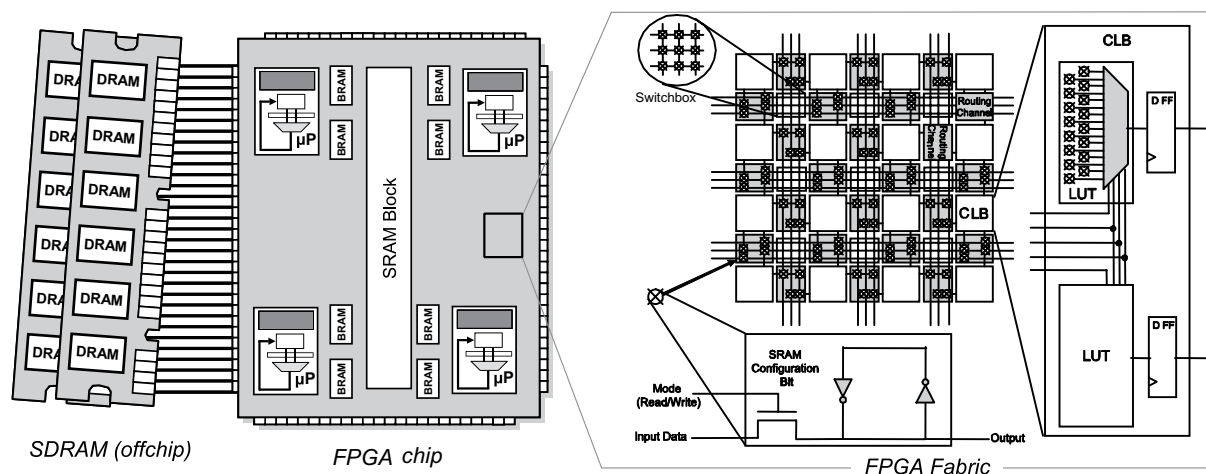


Fig. 1 – A modern FPGA-based embedded system: reconfigurable logic, blocks of SRAM, and hard-wired microprocessors all share the same piece of silicon, and, more importantly, the same off-chip memory. The reconfigurable logic is a fabric of tiny lookup tables and statically scheduled routing hardware that can be configured to emulate almost any possible circuit.

2.1. Architecture of a reconfigurable system

Field Programmable Gate Arrays (FPGAs) are the most common reconfigurable devices. An FPGA is a collection of programmable gates embedded in a flexible interconnect network. FPGAs use truth tables (known as lookup tables or LUTs) to implement logic gates, flip-flops for timing and registers, switchable interconnect to route logic signals between different units, and I/O blocks (IOB) for transferring data into and out of the device. A circuit can be mapped to an FPGA by loading the LUTs and switch-boxes with a *configuration*, a method that is analogous to the way a traditional circuit might be mapped to a set of *and* and *or* gates. Fig. 1 shows a modern FPGA-based embedded system.

LUTs employ static RAM cells as programming bits. A LUT is an extremely generic computational component. It can compute “any” function; i.e., any n -input LUT can be used to compute any n -input function. A LUT requires 2^n bits to describe, but it can implement 2^{2^n} different functions. LUTs are limited to a small number of inputs due to the size of SRAM cells as a programming point. A typical LUT has either 4 or 5 inputs, a number based on extensive empirical work aimed at optimizing physical aspects of the FPGA architecture (Betz et al., 1999). An FPGA is programmed using a *bit-stream*. This binary data is loaded into the FPGA to execute a particular task. The bit-stream contains all the parameters needed such as the configuration interface and the internal clock cycle supported by the device.

2.1.1. Reconfigurable devices and security

FPGAs are a natural platform for performing many cryptographic functions because of the large number of bit-level operations that are required in modern block ciphers. While there is a great deal of work centered around exploiting FPGAs to speed cryptographic or intrusion detection primitives, researchers are now starting to realize the security ramifications of building systems around hardware which is reconfigurable.

One major problem is that hardware, not just software, can now be copied from existing products, and there has been a flurry of research to protect this intellectual property (Bossuet et al., 2004; Kean, 2002; Lach et al., 1999a) and to secure the FPGA’s program logic update channels (Harper et al., 2003; Harper and Athanas, 2004). However, few researchers have begun to consider the security ramifications of compromised hardware (Hadzic et al., 1999).

It is important to understand the different attacks against FPGAs that are possible in order to develop counter measures (Wollinger et al., 2004). In a covert channel attack, an observable property such as power consumption is analyzed by a malicious module in order to steal secrets such as cryptographic keys or the bit-stream contained in the FPGA, which is valuable intellectual property (Standaert et al., 2003). In some systems, the bit-stream can be modified remotely, and authentication mechanisms should be employed to prevent unauthorized users from uploading a malicious design, which could change the intended functionality of the device. Even worse, the malicious design could physically destroy the FPGA by causing the device to short-circuit (Hadzic et al., 1999). Solutions to these problems include encryption (Bossuet et al., 2004; Kean, 2001, 2002), fingerprinting (Lach et al., 1999a), and watermarking (Lach et al., 1999b). While there are a variety of attacks possible, our work is concerned with addressing the problem of memory protection on reconfigurable systems. In particular this paper is concerned with techniques to provide *separation* while allowing controlled interaction between multiple interacting cores and modules with respect to their use of off-chip memory.¹ In our attack model, there may be subverted modules or remote attacks that originate from the network through I/O, but we assume that the attacker cannot physically modify or monitor the device.

¹ The same approach is applicable to on-chip memory, but we leave this to future work.

2.1.2. Protecting memory on an FPGA

A secure runtime management system must protect different logical modules from interfering, intercepting, or corrupting any use of a shared resource without authorization. On an embedded system, the primary resource of concern is memory. Whether it is on-chip block RAM, off-chip DRAM, or backing-store such as Flash, the allocation and sharing of memory must be performed in a way that is efficient, flexible, and protected. On a general-purpose processor, interaction via shared memory can be controlled through the use of page table and associated TLB attributes. Use of Superpages, which are very large memory pages, makes it possible for the TLB to have a lower miss rate (Navarro et al., 2002). Segmented Memory (Saltzer, 1974) and Mondrian Memory Protection (WitchelGates and Asanovic, 2002), a finer-grained scheme, address the inefficiency of providing per-process memory protection via global attributes by associating each process with distinct permissions on the same memory region.

While a TLB may be used to speed up page table accesses, this requires additional associative memory (not available on FPGAs) and greatly decreases the performance of the system in the worst case. Therefore, few embedded processors and even fewer reconfigurable devices support even this most basic method of protection. Instead, reconfigurable architectures on the market today support a simple linear addressing of the physical memory. Hence, on a modern FPGA the memory is essentially flat and unprotected by hardware mechanisms.

Preventing unauthorized accesses to memory is fundamental to both effective debugging, error prevention, and computer security. However, memory management in software is complex and difficult: many of the most insidious bugs are a result of errant memory accesses which affect multiple sub-systems. Ensuring protection and separation of memory when multiple concurrent logic modules are active requires a new approach to ensure that the security properties of the system are enforced.

To provide separation in memory between multiple interacting modules, we adapt some of the key concepts from separation kernels. Rushby originally proposed that a separation kernel (Irvine et al., 2004; Timothy Levin et al., 2004; Rushby, 1984) creates within a single shared machine an environment which supports the various components of the system, and it provides the communication channels between them in such a way that individual components of the system cannot distinguish this shared environment from a physically separated one. A separation kernel partitions all resources under its control into blocks (subsets) such that the actions of a subject in one block are isolated from (viz., cannot be detected by or communicated to) a subject in another block, unless an explicit means for that communication has been established. For a multilevel secure system, each block typically represents a different classification level, and the allowed communications conform to the MLS-label lattice (Denning, 1976).

We propose to treat the separate cores of the FPGA and related memory regions just as blocks of a separation kernel. The cores are isolated through a means we call “moats,” and then we control interaction between cores in a highly

assured manner. By building a specialized circuit that recognizes a *language of legal accesses* between blocks, and then by realizing that circuit directly onto the reconfigurable device as a specialized state machine through which all off-chip memory accesses are routed, every memory access can be checked with only a small additional latency. Although implementing the enforcement module into a separate off-chip hardware module would lessen the impact of covert channel attacks between modules on the chip, this would introduce additional latency. We describe techniques to isolate the enforcement module described by Huffmire et al., 2007.

2.2. Video redaction: a motivating example

The purpose of redaction is to delete sensitive information from a document, audio recording, video feed, or other data stream. For example, a document containing sensitive information would need to have all top secret and secret data removed before a person with only a confidential clearance could read the document. In video redaction, the faces of people in video feeds from surveillance cameras are blurred if the person viewing the video does not have a high enough clearance level. IBM’s PeopleVision project has developed such a video privacy system (Senior et al., 2003). FPGAs are a natural choice for streaming applications because they can provide deep regular pipelines of computation, with no shortage of parallelism.

Consider how such a system might be developed. There would need to be at least three modules on the FPGA: a video interface for decoding the video stream, a redaction mechanism for blurring faces in accordance with a policy, and a network interface for sending the redacted video stream to the security guard’s station. Each of these modules would need buffers of off-chip memory to function, and our enforcement module could prevent sensitive information from being shared between modules improperly (e.g., directly between the video interface and the network). In Section 4.4 we describe how such a situation might be handled using our methods.

3. Policy description and synthesis

While reconfigurable systems typically do not have traditional memory protection enforcement mechanisms, the programmable nature of the devices means that we can build whatever mechanisms we need as long as they can be implemented efficiently. In fact, we exploit the fine grain re-programmability of FPGAs to provide word-level stateful memory protection by implementing a compiler that can translate a memory access policy directly into a circuit. The enforcement mechanisms generated by our compiler will help prevent a corrupted module or processor from compromising other modules on the FPGA with which it shares memory. We have developed a security primitive for providing isolation of cores at the gate level by surrounding each core with a “moat” that blocks wiring connectivity from the outside (Huffmire et al., 2007).

We begin with an explanation of our memory access policies, and we describe how a policy can be expressed and then compiled down to a synthesizable module. In this section

we explain both the high level policy description and the automated sequence of steps, or *design flow*, for converting a memory access policy into a hardware enforcement module. Assurance that the conversion is accurate and complete is discussed as future work.

3.1. Memory access policy

Once a high level policy is developed based on the requirements of the system and the organizational security policy (Stern, 1991), it must be expressed in a precise form to allow engineers to build concrete enforcement mechanisms. In the context of this paper we concentrate on policies as they relate to memory accesses. In particular, the enforcement mechanisms we consider in this paper belong to the Execution Monitoring (EM) class (Schneider, 2000), which monitor the execution of a *target*, which in our case is one or more modules on the FPGA. The enforcement mechanism is also a Reference Validation Mechanism (RVM) (Anderson, 1972), which must be tamper-proof, always invoked, and small enough to be subject to analysis and test, the completeness of which can be assured. We describe techniques for isolating the reference monitor in Ref. by Huffmire et al. (2007).

Although Erlingsson et al. have proposed the idea of merging the reference monitor in-line with the target system (Erlingsson and Fred Schneider, 1999), in a system with multiple interacting cores, this approach has the drawback that the reference monitors are distributed, which is problematic for stateful policies. It may also prohibit the use of third-party bit-streams or require access to source code and the re-compilation of third-party bit-streams. Although there exist security policies that execution monitors are incapable of enforcing, such as *information flow* policies (Sabelfeld and Myers, 2003), we argue that in the future our execution monitors could be combined with static analysis techniques to enforce a more broad range of policies if required. We therefore begin by describing a well defined method for describing memory access policies.

The goal of our memory access policy description is to precisely describe the set of legal memory access patterns, specifically those that can be recognized by an execution monitor capable of tracking address ranges of arbitrary size within an enforcement framework that prohibits all other access. Furthermore, it should be possible to describe complex behaviors such as sharing, exclusivity, and atomicity, in an understandable fashion. An engineer can then write a policy description in our input form (as a series of “re-writing” productions) and have it transformed automatically to an extended type of regular expression. By extending regular languages to fit our needs we can have a human-readable input format, and we can build off of theoretical contributions which have created a refinement path to state machines and hardware (Aho et al., 1988).

There are three pieces of information that we will incorporate into our execution monitor. The *Accessing Modules* (M) are the unique identifiers for a specific principal on the chip, such as a specific intellectual property core or one of the on-chip processors. Throughout this paper we simply refer to these distinct units of activity on the FPGA as “Modules.” The *Access Methods* (A) are typically Read and Write, but may include

special memory operators such as execution, zeroing or incrementing if required. Elements of A are used to describe “permissions.” The set P is a partitioning of physical memory into “ranges.” The *Memory Range Specifier* (R) describes a set of contiguous physical addresses to which a specific permission can be assigned. Our language describes an access policy through a sequence of *productions*, which specify the relationship between principals (M : modules), access rights (A : read, write, etc.), and objects (R : memory ranges²).

The terminals of the language are *memory accesses descriptors* which ascribe a specific right for a specific module to access a specific object until the descriptor is negated or deleted.³ Formally, the terminals of the productions are tuples of the form (M, A, R) , and the universe of tuples forms a power set $\Sigma = M \times A \times R$. Given two sets of tuples, a and b , “ ab ” indicates the *union* of a and b . A memory access policy is precisely defined as a formal language $L \subseteq \Sigma$ which can be either generalized as being infinite or focussed on a fixed number of modules, ranges, and accesses. L needs to satisfy the property that $\forall x, t : \text{tuple set} | t \in \Sigma, xt \in L \rightarrow x \in L$, so that any legal access sequence will be incrementally recognized as legal along the way.

One thing to note is that memory accesses refer to a specific memory address, while memory access descriptors are defined over the set of all memory ranges R (i.e., the power set of addresses). A memory access (M, A, k) , where k is a particular address, is *contained* in a memory access descriptor (M', A', R) iff $M = M'$, $A = A'$, and $R_{\text{low}} \leq k \leq R_{\text{high}}$. A sequence of memory accesses $a = a_0, a_1, \dots, a_n$ is said to be legal iff $\forall_{0 \leq i \leq n} \exists s_i \in L | a_i \in s_i$. In order to enforce this policy during the execution of an FPGA, we need three things.

1. A notation with the details for a specific policy can be precisely defined under L .
2. A method for automatically creating a circuit which recognizes memory access sequences that are legal under L .
3. A method for preventing all accesses that are not legal under L .

We begin with a description of (1) through the use of a simple example. Consider a straightforward isolation policy that simply enforces the separation in memory of two different modules. *Module*₁ is only allowed to access memory in the range of $[0x8e7b008, 0x8e7b00f]$, and *Module*₂ is only allowed to access memory in the range of $[0x8e7b018, 0x8e7b01b]$. In our memory access *policy definition format*, this is coded as the following set of productions:

```
rw → r|w;
Range1 → [0x8e7b008, 0x8e7b00f];
Range2 → [0x8e7b018, 0x8e7b01b];
Access1 → {Module1, rw, Range1};
Access2 → {Module2, rw, Range2};
Policy → (Access1|Access2)*;
```

Each of these productions is a re-writing rule as in a standard grammar. The non-terminal *Policy* is the start symbol of

² An interval of the address space including high (R_{high}) and low (R_{low}) bounds.

³ Details of revocation will be discussed in Section 4.

the grammar that defines the overall access policy (L as described above). Through the use of a grammar we allow the hierarchical composition of more complex policies. In this case $Access_1$ and $Access_2$ are simple access descriptors, but we want to allow more complex sets of memory accesses, such that all sequences of accesses that can be derived from Policy by application of the grammar's productions are legal.

Since we eventually want to transform the access policy to hardware logic in a limited space, we limit our language to sequences that can be described with grammatical constructs no more complex than a regular expression (Linz, 2001), with the added ability to express ranges. Although a regular language is limited to a type-3 regular grammar in the Chomsky hierarchy, it is inconvenient for security administrators to express policies in right-linear or left-linear form, which would not allow "range" expressions. Since a language can be recognized by many grammars, any grammar that can be automatically transformed into type-3 form is acceptable, so we present the end user with an extended regular grammar that is later transformed by extracting first terminals from non-terminals.

Note that the atomic unit of enforcement is an address range, and that the ranges are of arbitrary granularity. The smallest granularity that we currently allow in the policy definition format is at the word boundary, and we can support any sized range from a single word to the entire address space. Also, ranges may be of the same or different size, unlike traditional memory pages. We will later show how this ability can be used to set up special control words that help in securely coordinating between modules.

Although we are restricted to policies that are equivalent to a finite automata with range checking, we have constructed many example policies including isolation and Chinese wall in order to demonstrate the versatility and efficiency of our approach. In Section 4.4 we describe a "redaction policy," in which modules with multiple security clearance levels are interacting within a single embedded system. However, now that we have introduced our memory access policy definition format, we describe how it can be transformed automatically to an efficient circuit for implementation on an FPGA.

3.2. Hardware synthesis

We have developed a policy compiler that converts an access policy, as described above, into a circuit that can be loaded onto an FPGA to serve as the policy enforcement module. At a high level the technique partitions the module into two parts, range discovery and language recognition. Specifically the steps of our design flow are:

- User creates the access policy (described above) and inputs it to the compiler, which:
- Builds a syntax tree from the policy.
- Transforms the syntax tree to an expanded intermediate form.
- Expands Policy to a regular expression defined over the alphabet Σ .
- Converts the regular expression to a non-deterministic finite automaton (NFA).

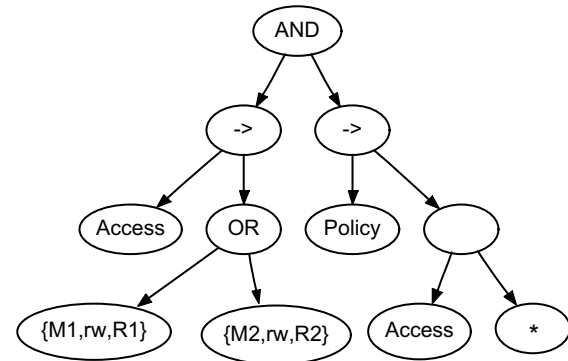


Fig. 2 – Parse tree of the simple access policy.

- Constructs an equivalent minimized state machine from the NFA.
- Factors the ranges into sizes that are a power of two.
- Organizes the set of ranges as a trie,⁴ and creates a logic tree that recognizes them.
- Exports the state machine and range detection logic as Synthesizable Verilog.
- Inputs hardware description expressed in Verilog to Quartus software, which synthesizes, places, and routes circuit.
- Bit-stream loader loads the synthesized bit-stream onto the FPGA.

3.3. Design flow details

3.3.1. Access policy

To describe the process of transforming a policy to a circuit, we again consider a simple isolation policy with two modules, which can only access their own single range:

$Access \rightarrow \{Module_1, rw, Range_1\} \{Module_2, rw, Range_2\};$
 $Policy \rightarrow (Access)^*;$

3.3.2. Building and transforming a parse tree

Next, we use Lex (Lesk and Schmidt, 1975) and Yacc (Johnson, 1975) to build a parse tree from our security policy. Internal nodes represent operators such as concatenation, alternation, and repetition. Fig. 2 shows the parse tree for our example policy.

We must then transform the parse tree into a large single production with no non-terminals on the right hand side, from which we can generate a regular expression. This process of macro expansion requires an iterative replacement of all the non-terminals in the policy. We apply the productions to the parse tree by substituting the left hand side of each production with its right hand side. Fig. 3 shows the transformed parse tree for our policy.

3.3.3. Building the regular expression

Next, we find the subtree corresponding to Policy and traverse this subtree to obtain the regular expression. By this stage we

⁴ An ordered tree data structure for storing lookup tables.

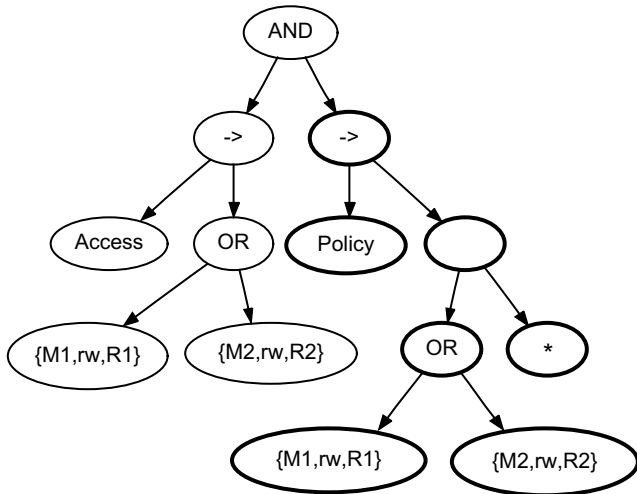


Fig. 3 – Expanded parse tree.

have completely eliminated all of the RHS non-terminals, and we are left with a single regular expression which can then be converted to an NFA. The regular expression for our access policy is: $((\{Module_1,rw,Range_1\})|(\{Module_2,rw,Range_2\}))^*$.

3.3.4. Constructing the NFA

Once the regular expression has been formed, we construct an NFA from this regular expression using Thompson's Algorithm (Aho et al., 1988) as implemented by Gerzic (2003). Fig. 4 shows the NFA for our policy. Notice that the policy transitions can occur in parallel. We will use the FPGA to exploit this for faster processing.




Fig. 4 – NFA derived from the regular expression.




Fig. 5 – NFA converted to a minimized DFA.

3.3.5. Converting the NFA to a DFA

From this NFA we can construct a DFA through subset construction (Aho et al., 1988) as implemented by Gerzic (2003). Following the creation of the DFA, we apply Hopcroft's Partitioning Algorithm (Aho et al., 1988) as implemented by Grail (Raymond and Wood, 1995) to minimize the DFA. Fig. 5 shows the minimized DFA for our policy.

3.3.6. Processing the ranges

Before we can convert the DFA into Verilog, we must perform some processing on the ranges so that the circuit can efficiently determine which range contains a given address. Our system converts the ranges to an internal format using "don't care" bits. For example, 10XX can be 1000, 1001, 1010, or 1011, which is the range (Bossuet et al., 2004; Compton and Hauck, 2002). Hardware can be easily synthesized to check if an address is within a particular range by performing a bit-wise XOR on just the significant bits.⁵ Using this optimization, any aligned power of two range (i.e., the cardinality of the range is a power of two) can be efficiently described, and any non-power of two range can be converted into a covering set of $O(\log_2|range|)$ power of two ranges. For example the range (Bondalapati and Prasanna, 2002; Cormen et al., 1990) (0111, 1000, 1001, 1010, 1011, 1100) is not an aligned power of two range but can be converted to a set of aligned power of two ranges: $\{[7,7],[8,11],[12,12]\}$ (or equivalently $\{0111|10XX|1100\}$).

3.3.7. Converting the DFA to Verilog

Because state machines are a very common hardware primitive, there are well-established methods of translating a description of state transitions into a hardware description language such as Verilog. Fig. 6 shows the hardware decision module we wish to build.

As previously described, an access descriptor specifies the allowed accesses between a module and a range. Each DFA transition represents an access descriptor, consisting of a module ID, an op, and a range ID bit vector. The range ID bit vector contains a bit for each possible range (currently a max of N ranges), and the descriptor's range is indicated by the (one) bit that is set.

⁵ This is equivalent to performing a bit-wise XOR, masking the lower bits, and testing for non-zero except that in hardware the masking is unnecessary.

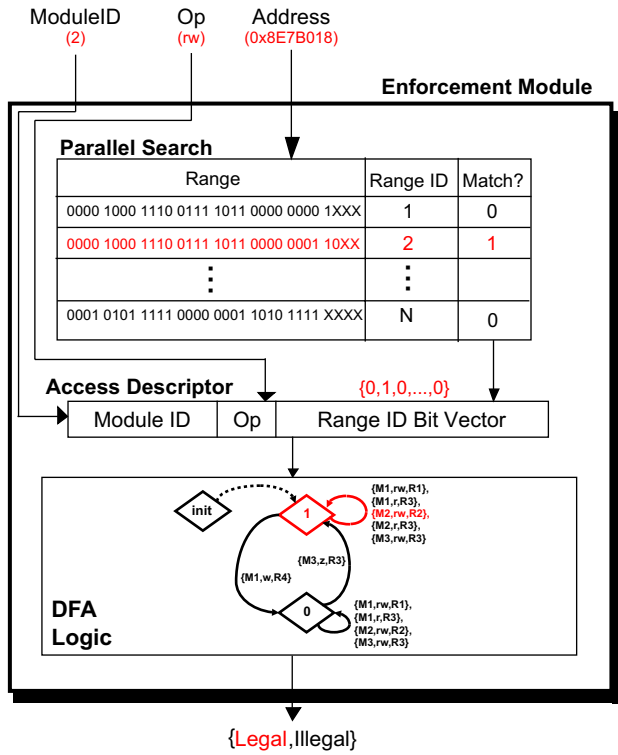


Fig. 6 – The inputs to the enforcement module are the module ID, op, and address. The range ID is determined by performing a parallel search over all ranges, similar to a content addressable memory (CAM). The module ID, op, and range ID together form an access descriptor, which is the input to the state machine logic. The output is a single bit: either grant or deny the access.

A memory access request comprises three inputs: the module ID, the op {read, write, etc.}, and the address. The output is a single bit: 1 for grant and 0 for deny. First, the hardware converts the memory access address to a bit vector. To do this, it checks all the ranges in parallel and sets the bit corresponding to the range ID that contains the input address (if any).

Then the memory access request is processed through the DFA. If an access descriptor matches the access request, the DFA transitions to the accept state and outputs a 1. If there is no transition for an access request, the machine always transitions to the rejecting state, which is a “dummy” sink state. This is important for security because an attacker might try to access an address not covered by the policy or try to insert illegal characters into the input, and results in a “fail secure” machine.

3.3.8. State machine synthesis

The final step in the design flow is the actual conversion of Verilog code to a bit-stream that can be loaded onto an FPGA. Using the Quartus tools from Altera, which does synthesis, optimization, and place-and-route, we turn each machine into an actual implementation. After testing the circuit to verify that it accepts a sample of valid accesses and rejects invalid accesses, we are ready to measure the area and cycle time of our design.

4. Example applications

To further demonstrate the usefulness of our language, we use it to express several different policies. We have already demonstrated an isolation policy, which can be easily extended to include overlapping ranges, shared regions, and most any static policy. The true power of our system comes from the description of *stateful* policies that involve revocation or conditional access or other forms of dynamic policy. Let us first discuss a traditional example: access control lists.

4.1. Access control list

A secure system that employs access control lists will associate every object in the system with a list of principals along with the rights of each principal to access the object. For example, suppose our system has two objects, *Range₁* and *Range₂*. *Class₁* is a class of principals (*Module₁* and *Module₂*), and *Class₂* is another class of principals (*Module₃* and *Module₄*). Either *Class₁* or *Class₂* may access *Range₁*, but only *Class₂* may access *Range₂*. We express such an access control list policy below:

```

Class1 → Module1|Module2;
Class2 → Module3|Module4;
List1 → Class1|Class2;
List2 → Class2;
Access1 → {List1,rw,Range1};
Access2 → {List2,rw,Range2};
Policy → (Access1|Access2)*;
    
```

In general, since access control list policies are stateless, the resulting DFA will have one state, and the number of transitions will be the sum of the number of principals that may access each object. In this example, *Module₁*, *Module₂*, *Module₃*, and *Module₄* may access *Range₁*, and *Module₃* and *Module₄* may access *Range₂*. The total number of transitions in this example is 4 + 2 = 6.

4.2. Controlled sharing

Secure system design requires the prevention of unintended flows of information between principals such as cores, but there are times when cores need to communicate with each other. Our language makes possible the secure transfer of data from one core to another. Rather than requiring large communication buffers or multiple copies of the data, we can simply transfer the control of a specified range of data from one module to the next. For example, suppose *Module₁* wants to securely transfer some data to *Module₂*. Rather than establishing a direct channel between *Module₁* and *Module₂*, an access policy can be created that synchronizes the transition of permissions during the exchange. Using formal languages to express security policies makes such an exchange possible. Consider the example below:

```

Module1|2 → Module1|Module2;
Access1 → {Module1,rw,Range1}{Module1|2,rw,Range2};
Access2 → {Module2,rw,(Range1|Range2)};
Trigger → {Module1,rw,Range2};
Policy → (Access1)* (∈|Trigger (Access2)*);
    
```


Initially, $Module_1$ can access $Range_1$ and $Range_2$, and $Module_2$ can only access $Range_2$. However, the first time $Module_1$ accesses $Range_2$ (signaling to $Module_2$ that $Module_1$ is ready to exchange), $Access_1$ is deactivated by this trigger event, revoking the permissions for $Module_1$ from both Ranges. As a result of the trigger, $Module_2$ has exclusive access to $Range_1$ and $Range_2$.

4.3. Chinese wall

Another security scenario that can be efficiently expressed using our policy language is the Chinese wall (Brewer and Nash, 1989). Consider an example of this scenario, in which a lawyer who looks at the set of documents of $Company_1$ should not view the set of files of $Company_2$ if $Company_1$ and $Company_2$ are in the same conflict-of-interest (COI) class. This lawyer may also view the files of $Company_3$ provided that $Company_3$ belongs to a different COI class than $Company_1$. Fig. 7 shows a Venn Diagram for this situation. We express a Chinese wall policy below, where $Module_1$ corresponds to the lawyer and each range corresponds to a company:

```

Access1 → {Module1,rw,(Range1|Range3)*};
Access2 → {Module1,rw,(Range1|Range4)*};
Access3 → {Module1,rw,(Range2|Range3)*};
Access4 → {Module1,rw,(Range2 | Range4)*};
Policy → Access1|Access2|Access3|Access4;
    
```

In our Chinese wall policy, there are two COI classes. One contains $Range_1$ and $Range_2$, and the other contains $Range_3$ and $Range_4$. For simplicity, we have restricted this policy to one module since with multiple modules, the restrictions to a module are independent of the actions of other modules so each module requires its own state machine. Fig. 7 shows the DFA that recognizes legal accesses for this policy.

In general, for Chinese wall security policies, the number of states scales exponentially to the number of COI classes. Because the number of possible legal accesses is the serial product of the number of ranges (companies) in each separate COI

class. The number of transitions also scales exponentially to the number of COI classes for the same reason. Fortunately, the number of states and the number of transitions both scale linearly to the number of ranges. In addition, the number of transitions scales linearly in the number of ranges.

4.4. Redaction

Our security language can also be used to enforce forms of redaction (Sami Saydjari, 2004), even at very high throughput (such as for video). Military hardware such as avionics (Weissman, 2003) may contain processing components that are “cleared” for different levels of data, and a TS component must not leak sensitive information to a U component (Smith, 2001). However, the TS component may be required to send a document to the U component; a third component does this by redacting TS data from the document. Fig. 8 shows the architecture of a redaction scenario that is based on separation.

A multilevel database contains both top secret (TS) and unclassified (U) data. $Module_1$ has a TS label, and $Module_2$ has a U label. $Module_1$ and $Module_2$ are initially isolated, since they have different labels. Therefore, $Range_1$ belongs to $Module_1$, and $Range_2$ belongs to $Module_2$. $Module_3$ acts as a trusted server of information contained in the database, and this server must have a security label range from U to TS. $Range_3$ is temporary storage used for holding information that has just been retrieved from the database by the trusted server. $Range_4$ (the control word) is used for performing database queries: a module writes to $Range_4$ to request that $Module_3$ retrieve some information from the database and then write the query result to the temporary storage. Any database query requested by $Module_2$ must have all TS data redacted by the trusted server. If a request is made by $Module_1$ for top secret information, it is necessary to revoke $Module_2$'s read access to the temporary storage, and this access must not be reinstated until the trusted server zeroes out the sensitive information contained in the temporary storage. One way of implementing

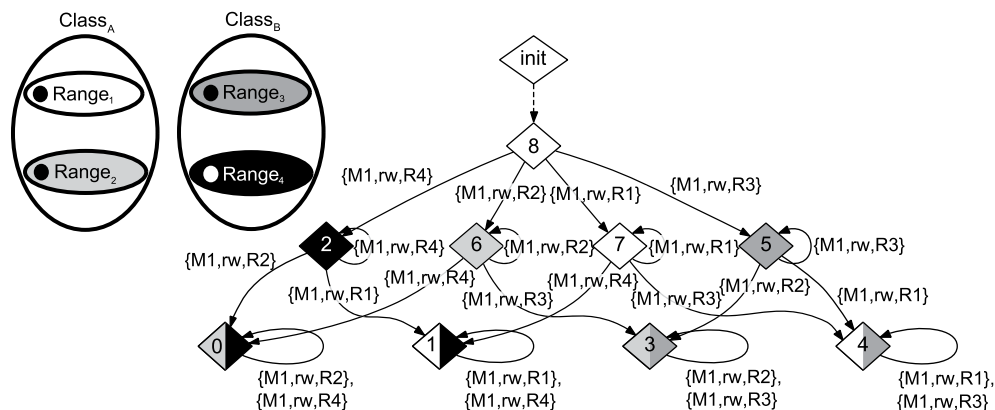


Fig. 7 – A Chinese wall policy. The Venn Diagram shows two conflict-of-interest classes, $Class_A$ and $Class_B$, and the DFA recognizes legal accesses for this Chinese Wall policy. //A principal that accesses $Range_4$ (black) is subsequently prohibited from accessing $Range_3$ (dark gray), but it may access either $Range_1$ (white) or $Range_2$ (light gray), because they are in a different class. An access to $Range_4$ results in a transition to state 2 (black), from which an access to $Range_1$ results in a transition to state 1 (black or white).

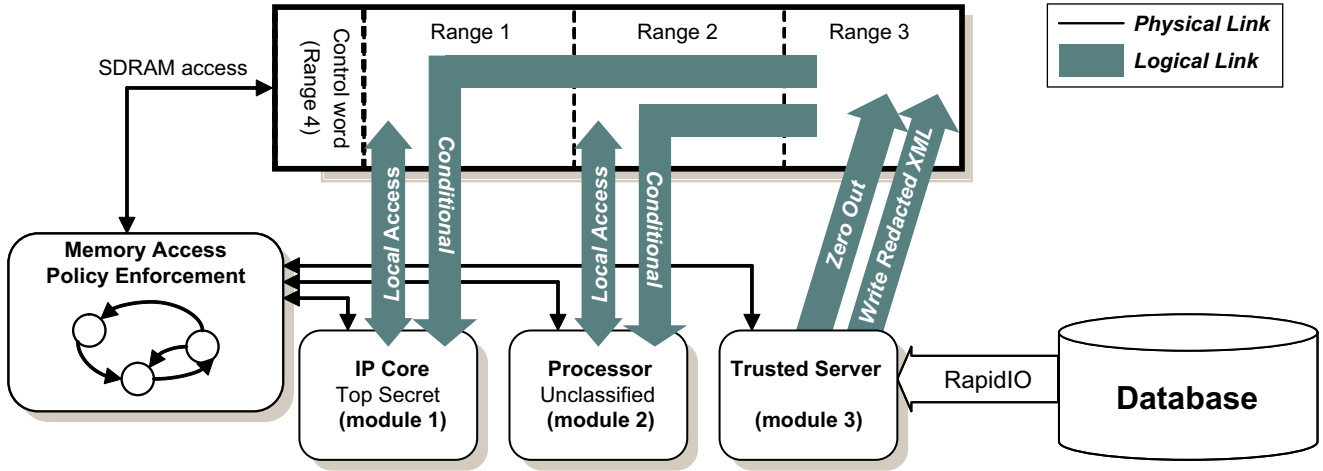


Fig. 8 – A redaction architecture. A database contains both Top Secret and Unclassified data. *Module₁* has a Top Secret (TS) clearance, and *Module₂* has an Unclassified (U) clearance. Any database query requested by *Module₂* must have all TS data redacted by the Trusted Server *Module₃*. Furthermore, *Module₂* must be prevented from accessing the result of a database query performed by *Module₁* because such a query result may contain TS data. This is accomplished by revoking *Module₂*'s permission to access the temporary storage (*Range₃*) where query results are written by the Trusted Server. IP stands for Intellectual Property.

the zeroing out functionality is to use a special access right (*z*) in conjunction with logic that erases the contents of the temporary storage. We express our redaction policy below:

$$\begin{aligned}
 rw &\rightarrow r|w; \\
 Access_2 &\rightarrow \{Module_1, rw, Range_1\} \{Module_1, r, Range_3\} \\
 &\quad \{Module_2, rw, Range_2\} \{Module_2, w, Range_4\} \\
 &\quad \{Module_3, rw, Range_3\}; \\
 Access_1 &\rightarrow \{Module_2, r, Range_3\} | Access_2; \\
 Trigger &\rightarrow \{Module_1, w, Range_4\}; \\
 Clear &\rightarrow \{Module_3, z, Range_3\}; \\
 SteadyState &\rightarrow (Access_2 | Clear Access_1^* Trigger)^*; \\
 Policy &\rightarrow \in | Access_1^* | Access_1^* Trigger SteadyState \\
 &\quad | Access_1^* Trigger SteadyState Clear Access_1^*;
 \end{aligned}$$

$Access_1$ is the less restrictive access mode, and $Access_2$ is the more restrictive access mode. The Trigger event changes the access mode from $Access_1$ to $Access_2$, and the Clear event causes the machine to transition from $Access_2$ back to $Access_1$. In general, the DFA for a redaction policy will have one state for each access mode. Applying our redaction policy to a real-world video privacy system would likely require some additional complexity.

4.5. Bell and LaPadula confidentiality model

The Bell and LaPadula (B&L) Model is a formal model of multi-level security in which a subject may not read an object with a higher security label (no read-up), and a subject may not write to an object with a lower security label (no write-down) (Bell and LaPadula, 1973). This model is designed to protect the confidentiality of classified information. All B&L policies are stateless in that the rules don't change and the labels of individual subjects and objects upon which the rules are based, don't change. We express a B&L policy below:

$$\begin{aligned}
 Access_{B\&L} &\rightarrow \{Module_1, r, Range_1\} \{Module_1, r, Range_2\} \\
 &\quad \{Module_2, r, Range_2\} \{Module_2, w, Range_1\} \\
 &\quad \{Module_2, w, Range_2\}; \\
 Policy &\rightarrow (Access_{B\&L})^*;
 \end{aligned}$$

In our simple example, *Module₁* has a TS label, *Module₂* has a U label, *Range₁* has an S label, and *Range₂* has a U label. We leave to future work the covert channel analysis of these mechanisms.

4.6. High water mark

High water mark is similar to B&L in that no read-up is permitted, but object labels change over time, and write-down is allowed. Following a write-down, the security label of the object written to must change to the label of the subject that performed the write; thus, high water mark policies are stateful. We express our high water mark policy below:

$$\begin{aligned}
 Access_1 &\rightarrow \{Module_1, r, Range_1\} \{Module_1, r, Range_2\} \\
 &\quad \{Module_1, w, Range_2\} \{Module_2, w, Range_1\}; \\
 Access_2 &\rightarrow Access_{B\&L} | \{Module_1, w, Range_1\}; \\
 Access_3 &\rightarrow Access_1 | \{Module_2, w, Range_2\}; \\
 Access_4 &\rightarrow Access_1 | \{Module_1, w, Range_1\} \{Module_2, w, Range_2\}; \\
 Trigger_1 &\rightarrow \{Module_1, w, Range_1\}; \\
 Trigger_2 &\rightarrow \{Module_1, w, Range_2\}; \\
 Path_1 &\rightarrow (\in | Trigger_1 Access_2^* (\in | Trigger_2 Access_4^*)); \\
 Path_2 &\rightarrow (\in | Trigger_2 Access_3^* (\in | Trigger_1 Access_4^*)); \\
 Policy &\rightarrow Access_{B\&L}^* | (\in | Path_1 | Path_2);
 \end{aligned}$$

We use trigger events to express the write-downs. The number of triggers *T* in the high water mark policy is equal to the number of write-downs that would be illegal in the B&L policy. The number of states *S* in the DFA that enforces the high water mark policy is $O(2^T)$, and the number of

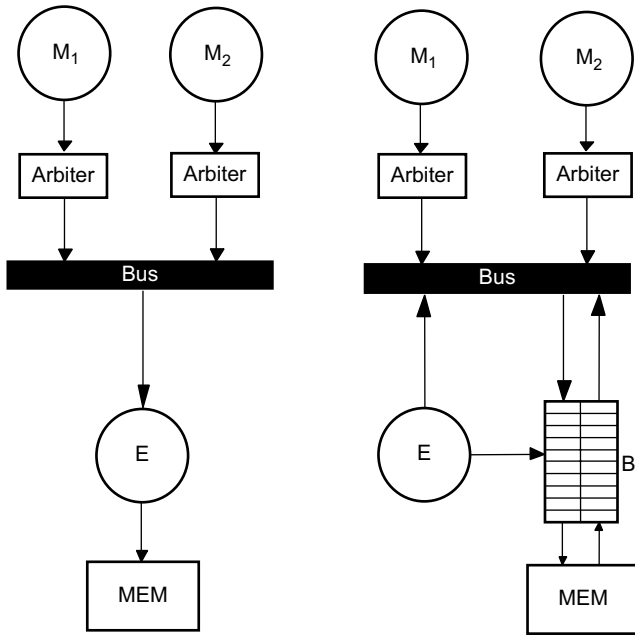


Fig. 9 – Two alternative architectures for the enforcement mechanism. In the figure on the left, a memory access must pass through the enforcement mechanism (E) before going to memory. In the figure on the right, the enforcement mechanism (E) snoops on the bus, and a buffer (B) prevents access to the data until the access is approved. Arbiters prevent the bus from being accessed by more than one module at a time.

transitions in the DFA that are triggers is $O(T!T)$. If N is the number of transitions in the corresponding stateless B&L policy, then the number of transitions in the high water mark DFA that are not triggers is $O((N)(S))$. Therefore, the total number of transitions is $O(T!T + (N)(S))$.

4.7. Biba integrity model

The Biba model is the dual of the Bell-LaPadula model (Biba, 1977), but the label spaces of the policies are distinct. Both read-down and write-up with respect to the ordering of integrity labels are prohibited. Like B&L, all Biba policies are stateless. We express our B&L policy below:

$$\text{Access}_{\text{Biba}} \rightarrow \{ \text{Module}_{1,w,\text{Range}_1} \} \{ \text{Module}_{1,w,\text{Range}_2} \} \\ \{ \text{Module}_{2,r,\text{Range}_1} \} \{ \text{Module}_{2,r,\text{Range}_2} \} \\ \{ \text{Module}_{2,w,\text{Range}_2} \}; \\ \text{Policy} \rightarrow (\text{Access}_{\text{Biba}})^*;$$

Low water mark is to Biba as high water mark is to B&L. Since low water mark is similar to high water mark, we do not discuss it further.

4.8. Dynamic policies

The ability to change the policies in response to external events is useful. For example, if the system comes under attack, it may be necessary to change to a more restrictive policy. We express a dynamic policy below:

$$\text{Policy} \rightarrow \text{Policy}_1 (\in | \text{Trigger}_1 (\text{Policy}_2) (\in | \text{Trigger}_2 (\text{Policy}_3)));$$

Policy_1 , Policy_2 , and Policy_3 can be any three policies. If the policies come from different sources, pre-processing can be used to prevent naming conflicts (e.g., if two policies define Access_1 differently). Trigger events specify the circumstances under which a policy change can occur. Trigger_1 causes the policy to change from Policy_1 to Policy_2 , and Trigger_2 causes the policy to change from Policy_2 to Policy_3 . Every state in Policy_1 has an additional transition (Trigger_1) to the first state of Policy_2 , and every state in Policy_2 has an additional transition (Trigger_2) to the first state of Policy_3 . The number of states in the combined policy is $O((S_1) + (S_2) + (S_3))$, where S_N is the number of states in Policy_N . The number of transitions in the combined policy is $O((T_1) + (T_2) + (T_3))$, where T_N is the number of transitions in Policy_N .

In the above scenario, the system must start in Policy_1 . The system may or may not transition to Policy_2 . If the system transitions to Policy_2 , the system may or may not transition to Policy_3 . Supporting the ability to go in any order requires more complex expressions and more complex DFAs. In addition, the ability to return to an earlier policy has several security implications, especially when stateful policies are involved. Understanding the organizational requirements for dynamic security policies is the topic of related research (Bossuet et al., 2004; Fraser and Badger, 1998).

Although switching back and forth between an arbitrary number of stateful policies would require modifying our compiler, it is possible to use our language to switch back and forth between two stateless policies Policy_1 and Policy_2 using the following expression:

$$\text{SteadyState} \rightarrow (\text{Policy}_2 | \text{Trigger}_2 \text{Policy}_1 \text{Trigger}_1)^*;$$

$$\text{Policy} \rightarrow \text{Policy}_1 | \text{Policy}_1 \text{Trigger}_1 \text{SteadyState} \\ | \text{Policy}_1 \text{Trigger}_1 \text{SteadyState} \text{Trigger}_2 \text{Policy}_1 | \epsilon;$$

Trigger_1 changes the policy from Policy_1 to Policy_2 , and Trigger_2 changes the policy back to Policy_1 .

5. Integration and evaluation

Now that we have described several different memory access policies that could be enforced using a stateful monitor, we need to demonstrate that such systems could be efficiently realized on reconfigurable hardware.

5.1. Enforcement architecture

The placement of the enforcement mechanism can have a significant impact on the performance of the memory system. Fig. 9 shows two architectures for the enforcement mechanism which assumes that modules on the FPGA can only access shared memory via the bus.

In the figure on the left, the enforcement mechanism sits between the memory and the bus, which means that every access must pass through the enforcement mechanism before going to memory. In the case of a read, the request cannot proceed to memory until the enforcement mechanism approves the access. This results in a large delay which is the sum of

the time to determine the legality of the access and the memory latency. We can mitigate this problem by having the enforcement mechanism snoop on the bus or through the use of various caching mechanisms for keeping track of accesses that have already been approved. This scenario is shown in the figure on the right. In the case of a read, the request is sent to memory, and the memory access occurs in parallel with the task of determining the legality of the read. A buffer holds the data until the enforcement mechanism grants approval, at which time E sends the data across the bus. In the case of a write, the data to be written is stored in the buffer until the enforcement mechanism grants approval, at which time E sends the data from the bus to memory. Thus, both architectures provide the isolation and omnipotence required of a reference or execution monitor.

Since a module may be sending sensitive data over the bus, it is necessary to prevent other modules from accessing the bus at the same time. We address this problem by placing an arbiter between each module and the bus. In a system with two modules, for example, the arbiters could allow one module to access the bus on even clock cycles and the other module to access the bus on odd clock cycles. We discuss a secure communication architecture for FPGAs as well as a method of ensuring the isolation of the reference monitor at the gate level in Ref. (Huffmire et al., 2007).

5.2. Evaluation

Of the different policies we discussed in Section 4, we focus primarily on characterizing the isolation policy in order to separate the effect of range detection on system efficiency. Rather than tying our results to the particular reconfigurable system prototype we are developing, we quantify the results of our design flow on a randomly generated set of ranges over which we enforce isolation. The range matching constitutes the majority of the hardware complexity (assuming there are a large number of ranges), and there has already been a great deal of work in the CAD community on efficient state machine synthesis (De Micheli, 1994).

To obtain data detailing the timing and resource usage of our range matching state machines, we ran the memory access policy description through our front-end and synthesized⁶ the results with Quartus II 4.2 (Inc and QuartusManual, 2004). Compilations are optimized for the target FPGA device (Altera Stratix EPS1S10F484C5), which has 10,570 available logic cells, and Quartus will utilize as many of these cells as possible.

5.3. Synthesis results

In general, a DFA for an isolation policy always has exactly one state, and there is one transition for each $\{\text{ModuleID}, \text{op}, \text{RangeID}\}$ tuple. We have determined that for our isolation policy, there is a linear relationship between the number of transitions and the number of ranges. Fig. 10 shows that the area of the resulting circuit scales nearly linearly with the number of ranges for the compartmentalization policy. The slope is approximately four logic cells for every range.

⁶ The back-end handles netlist creation, placement, routing, and optimization for both timing and area.

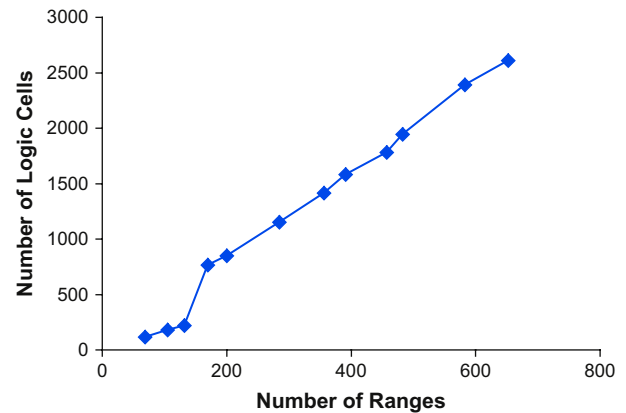


Fig. 10 – Circuit area versus number of ranges. There is a nearly linear relationship between the circuit area and the number of ranges.

Fig. 11 shows the cycle time (T_{clock}) for machines of various sizes. T_{clock} is primarily the time for one DFA transition, and it is very close to the maximum frequency of this particular Altera Stratix device. Fig. 12 shows the setup time (T_{su}), which is primarily the time to determine the range to which the input address belongs. Although T_{clock} remains nearly constant with the number of ranges, T_{su} increases nearly linearly above 170 ranges. Fortunately, T_{su} can be reduced by pipelining the circuitry that determines what range contains the input address.

Fig. 13 shows the area of the circuits resulting from the example policies presented in this paper. These circuits are much smaller in area than the series of isolation circuits above because the example policies have very few ranges, states, and transitions. The complexity of the circuit is a combination of the number of ranges and the number of DFA states and transitions. In our dynamic policy, Policy_1 is our isolation policy, Policy_2 is our Biba policy, and Policy_3 is our controlled sharing policy. Returning to an earlier policy is not allowed since Policy_3 is stateful. As expected, the circuit for the dynamic policy has the greatest area because it consists

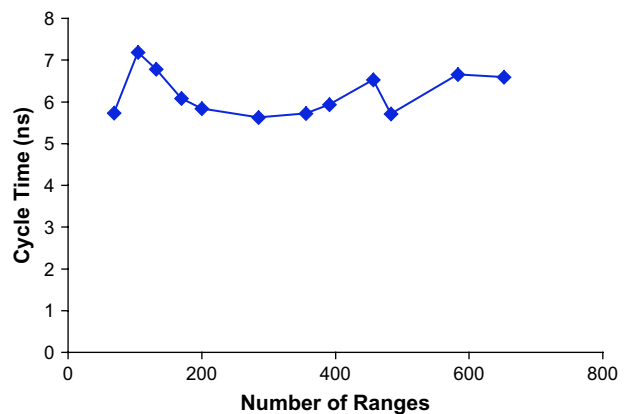


Fig. 11 – Cycle time versus number of ranges. There is a nearly constant relationship between the cycle time and the number of ranges.




Fig. 12 – Setup time versus number of ranges. Above 170 ranges, there is a nearly linear relationship between the setup time and the number of ranges. This time can be reduced with pipelining.

of three policies. The next biggest circuit belongs to Chinese wall, followed by redaction, high water mark, low water mark, and controlled sharing. Fig. 14 shows that the cycle time is greatest for redaction, followed by high water mark, low water mark, B&L, and isolation. Fig. 15 shows that the setup time is greatest for Biba, followed by B&L, dynamic, redaction, and high water mark.

5.4. Impact of the reference monitor on system performance

Since FPGAs do not operate at high frequency, they achieve their performance from computational parallelism. Many FPGA applications such as DSPs, signal processing, and intrusion detection systems are throughput-driven, but do not have rigid latency requirements. Therefore are latency-insensitive. For these applications, we argue that our technique does not impact performance significantly. For example, since




Fig. 13 – Circuit area versus access policy. The area is related to the number of states, transitions, and ranges. The circuit area is greatest for the dynamic policy.

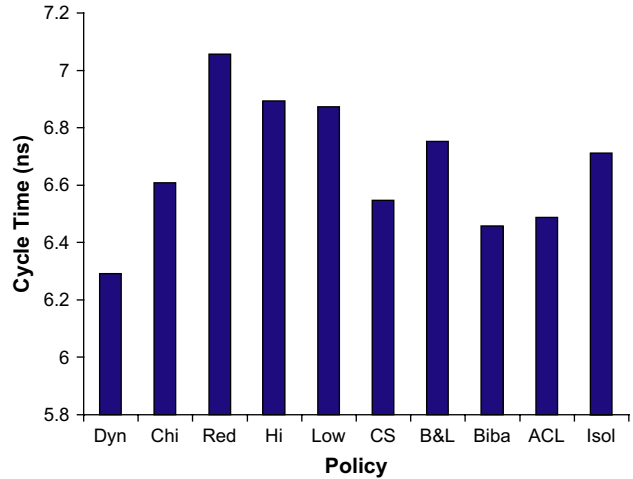


Fig. 14 – Cycle time for each access policy. Cycle time is greatest for redaction, followed by high water mark and low water mark.

an FPGA operating at 200 MHz will have a cycle time of 5ns, our reference monitor only adds at most a two cycle delay in this case.

6. Covert storage channels

In a covert channel, internal state of the enforcement mechanism is used to transmit information in a manner that is contrary to the security policy. Some stateful policies define (and require) internal state changes that, if unchecked, can form the basis of a covert channel. In this section, we describe a method of analyzing stateful policies to detect these “inherent” covert channels. We describe a method for measuring the potential bandwidth of the covert channel at runtime so that corrective action can be taken if the bandwidth exceeds a predetermined threshold value.

A reference monitor makes a binary decision to either grant or deny a particular access based on a policy. As a result,




Fig. 15 – Setup time for each access policy. Setup time is greatest for Biba, followed by B&L and dynamic.

with certain stateful policies, subjects may be able to observe the internal state of the policy by observing the reference monitor's decisions. Subjects can change the state of the policy enforcement mechanism by making access requests that affect subsequent requests. The state of the policy is in effect a shared resource. This ability to observe and modify the policy itself makes it possible for two subjects to illegally communicate.

Some stateful policies only allow a few bits to be leaked, while others allow a high bandwidth of data to be leaked. A stateful policy expressed in a regular language is equivalent to a directed graph, which contains cycles that allow the policy to alternate between two or more states continuously. If certain properties are met, then the cycle represents a possible covert channel. If it is not feasible to revise the policy to eliminate the cycle, we propose a technique for coping with such a storage channel by counting the number of times that a cycle goes around at runtime. If the counter exceeds a threshold, the system can take corrective action. One option for corrective action is to change to a policy without the covert channel. This limits the amount of data that can be leaked to a specific value.

6.1. Storage channels in stateful policy enforcement systems

In both covert storage and timing channels, the sender has a higher security label than the receiver and interferes with the receiver's access to a shared system resource (e.g., the processor or disk) to signal the receiver. In a timing channel, the interference changes the time needed for the receiver to perform a task, and the receiver interprets the delay or lack of delay. On the other hand, in a storage channel, the interference changes the receiver's ability to access the resource. This section will focus on storage channels rather than timing channels. We close timing channels with respect to access decisions but assume that data access times, for example, are not otherwise modified by the actions of subjects. We leave to future work the application of our methods to timing channels. In this section, we will show how our reference monitor can be used as a storage channel if the policy has certain properties.

Kemmerer (1983, 2002) has devised a shared resource matrix method of identifying storage and timing channels in computer systems. Kemmerer identifies four criteria for storage channels: the sender and receiver must have access to the same shared resource attribute, the sender must be able to change the shared attribute, the receiver must be able to detect the change, and the sender and receiver must be able to initiate communication and sequence events (Kemmerer, 1983, 2002). In the case of our embedded system, the sender and the receiver are cores, and the shared resource is the state of the policy.

Since a policy is enforced by a DFA, we can think of policies as directed graphs. Each node of the graph is a state of the policy, and each edge is a transition of the policy. In order for a policy to have a storage channel, there must be a non-trivial cycle in the graph. Of course, stateless policies do not have non-trivial cycles.

The presence of a cycle in the graph allows the internal state of the reference monitor to alternate between two or more states. This property allows a stream of information to be leaked if the sender can cause a DFA transition and the two have different access matrices with respect to the receiver such that the receiver can sense the policy change.

For covert channel analysis, we assume the most conservative assignment of security labels to principals that results in the largest information flow. For example, if a possible covert channel from *Module₁* to *Module₂* is identified, in the absence of any evidence to the contrary, we assume that *Module₁* has a higher security label than *Module₂*.

Fig. 16 shows a DFA that enforces a memory access policy. Each node in the graph is a state in the policy, and we show the access matrix at each node. The columns of the access matrix are the principals (modules), and the rows are the objects (ranges). This DFA contains a non-trivial cycle that satisfies our criteria for a covert channel: (1) *Module₁* has a higher security label than *Module₂*, (2) *Module₁* controls at least one of the transitions within the cycle, and (3) at least two nodes within the cycle have access matrices that differ with respect to *Module₂*.

We now describe how this cycle is used to establish an illegal information flow from *Module₁* to *Module₂*. *Module₂* continually tries to read *Range₁*. If access is granted, *Module₂* knows that the current state is the first state, but if access is denied, *Module₂* knows that the current state is the second state. *Module₁* alternates between the two states by reading *Range₁*. *Module₂* receives a bit of information when the current state remains stable for a fixed number of cycles *T*. Another way of transmitting a bit is to treat one complete cycle as a bit, similar to a Morse code pulse. There are many ways for the sender to encode the data to be leaked. The bandwidth of the information flow can be calculated in terms on the number of




Fig. 16 – A non-trivial cycle. This figure shows the DFA that enforces a security policy. Each node of the graph is a state in the policy, and we show the access matrix at each node. *Module₁* has a higher security label than *Module₂*. Initially, *Module₂* can read *Range₁*. *Module₁* can then change the state by reading *Range₁*. Now, *Module₂* can no longer read *Range₁*. *Module₁* can then change the state again by reading *Range₁*. According to our criteria, there is a possible storage channel from *Module₁* to *Module₂*.

possible encodings of the data and the probability of each symbol (Millen, 1987, 1989), over time.

As mentioned above, one of the criteria for a covert channel is that the sender must be able to change the shared attribute. The sender does not have to be able to cause every DFA transition within the cycle. One is sufficient because the remaining transitions can be caused by the receiver. If neither the sender nor the receiver is able to cause a particular transition within the cycle, the sender can wait a large number of cycles so that such a transition is very likely to occur. This allows the cycle to come around again, but the sender and receiver would need to be able to sync their activities, e.g., by polling the state of the policy. This case is shown in Fig. 17. *Module*₁ would like to send some data to *Module*₂, but *Module*₃ controls one of the transitions. *Module*₁ simply waits a sufficient length of time during which *Module*₃'s transition is likely to occur. In this case, bandwidth would depend on the estimated rate of *Module*₁'s activity.

Since we are using a static technique to detect possible covert channels, not all possible channels identified can be exploited at runtime. Furthermore, analysis of the design may be required to estimate false positives. Another criterion for a storage channel mentioned above is that the receiver must be able to detect the change. Not every node in the cycle must differ with respect to the receiver. In fact, the access matrices of just two states within the cycle must differ with respect to the receiver. Suppose we have a large cycle with many nodes, most of which have access matrices that are identical with respect to the receiver. If just two of them differ, the receiver will be able to detect that the cycle has repeated. Fig. 18 shows this case. In this cycle that contains three nodes, only one of the nodes differs from the other two nodes with respect to *Module*₂. Still, *Module*₂ is able to detect the change, allowing data to leak from *Module*₁, which controls all of the transitions within the cycle, to *Module*₂.




Fig. 17 – Suppose that *Module*₁ would like to leak some data to *Module*₂. In this example, *Module*₁ causes one of the transitions, and *Module*₃ causes the other transition. The two access matrices differ with respect to *Module*₂. Since *Module*₃ is not a party in the exchange, *Module*₁ must wait a sufficient length of time for *Module*₃'s transition to occur, allowing the cycle to come around again.




Fig. 18 – A cycle consisting of three nodes. Two of the nodes are identical with respect to *Module*₂, but one is different from the other two. Since at least one node differs, *Module*₂ can detect the change, allowing data to leak from *Module*₁ to *Module*₂.

6.2. Automatically detecting policy-based covert storage channels

To automatically detect a possible covert channel in a policy, we first determine if its DFA has any cycles. Topological sort is a well-known algorithm for detecting cycles in a directed graph (Cormen et al., 1990). The first step is to select a vertex in the graph with no incoming edges and remove it, repeating this process until there are no more vertices left. If this process does not finish in an amount of time proportional to the size of the graph, then the graph contains a cycle. Identifying the set(s) of vertices that make up the cycle(s) involves tracing the graph recursively. As shown:

```

Procedure Detect Channels (Graph G)
{
    Array of Lists Senders
    Array of Lists Receivers
    If (Topological_Sort(G) == False)
        Output "Graph G Contains No Cycles."
    Return
    C = Recursively_Trace_Graph_to_Find_Cycles (G)
    For (All Cycles C)
        For (All Edges E in C)
            M = Module that causes transition E
            Add M to Senders[C]
        For (All Vertices V in C)
            For (All Vertices V' in C) if V' != V then
                For (All Rows row)
                    For (All Columns col)

```

```
If (Matrix(V)[row][col] != Matrix(V)[row][col])
Add col to Receivers[C]
```

```
Output "Possible Covert Channels:"
For (All Cycles C Found)
    Output Cross_Product(Senders[C], Receivers[C])
Return
}
```

We applied our covert channel detector to several example policies. Fig. 19 shows the redaction policy described in Section 4.4, which alternates between a more restrictive and less restrictive access matrix. Our detector correctly identified four possible covert channels: from $Module_1$ to $Module_2$, from $Module_1$ to $Module_3$, from $Module_3$ to $Module_1$, and from $Module_3$ to $Module_2$.

Dynamic policies that switch back and forth between two or more policies may have covert channels (Woodward, 1987; Zheng and Myers, 2004). One way of dealing with the problem of covert channels in dynamic policies is to have a module with a low security label perform the policy transitions. If this is not possible and a module with a higher security label must be used to perform the policy transitions, then it is essential that this module be “trusted” to not drive the covert channel in violation of the policy. Even if a covert channel exists, if policy switching is infrequent, the bandwidth is low, and the channel might be acceptable.

In Section 4.3, we described a Chinese wall policy shown in Fig. 8 with two conflict-of-interest classes: $\{Range_1, Range_2\}$ and $\{Range_3, Range_4\}$. Although it does not have any cycles, this policy is not completely free of covert channels. $Module_1$ could leak one bit of information to $Module_2$ and one bit to $Module_3$, or $Module_1$ could leak one bit to $Module_2$ and $Module_4$, or $Module_1$ could leak one bit to $Module_3$ and $Module_4$. While two bits does not seem like a lot of information, there may be highly sensitive applications for which even leaking two bits is unacceptable. In a graph that does not have any cycles, the maximum

amount of information that can be leaked is bounded by the longest path length from the initial state to any final state.

6.3. Approaches to covert channel management

Once a possible covert channel has been identified, the system designer, working with the enterprise security manager, can modify the policy in order to eliminate the problematic cycle. If this is not an option, we can close the covert channel or confine the bandwidth within certain limits, restricting the behavior of the system. This approach requires a method of tracking current bandwidth. One way of measuring the usage bandwidth of the covert channels is with counters. A counter keeps track of the number of times the cycle occurs within a sliding window of time relative to the current time, and the system ensures that the covert channel bandwidth stays below a threshold value.

6.3.1. Approaches for measuring covert channels

A cycle can be expressed as a regular expression, and a piece of hardware to recognize this expression can be easily built. For example, a cycle from $State_1$ to $State_2$ to $State_3$ and back to $State_1$ can be expressed as $S_1(S_2S_3S_1)^+$. Regular expressions can even identify large cycles that contain smaller cycles within. This “monitor monitor” can be incorporated into the reference monitor. The price of this measurement mechanism should be balanced against the cost of ensuring that the module with a high security label will not leak secret information. Typically, the cost of such mechanisms is much lower than the price of ensuring that a module is trusted.

6.3.2. Options for closing and throttling covert channels

As we explained in Section 6.1, terminating a core is highly problematical because critical services may be disabled. Rather than terminating the receiver, we propose changing the policy in response to a counter exceeding its threshold. Fig. 20 shows an example of this concept. A stateful policy with two states has a cycle resulting in a possible covert channel from $Module_1$ to $Module_2$. A counter monitors the number of times the cycle completes. If the counter exceeds a threshold value, during the measurement window, the policy changes so that the nodes in the cycle are identical with respect to $Module_2$, thus closing the covert channel. This is accomplished in the example by revoking $Module_2$'s privilege to write to $Range_1$ in the second state of the stateful policy. After a period of time, the policy can revert, if desired. This will be highly desirable in practice. Alternatively, we can add delays to transition between states to throttle the channel within acceptable limits.

Modification to the policy as above adds to the amount of logic that must reside on the FPGA. In a stateful policy with M modules that are receivers in a possible covert channel and S states, the total number of states in this combined policy will be $O(S(2^M))$. If the stateful policy has T transitions, the total number of transitions in the combined policy will be $O(T(M!M + 2^M))$. The cost of this privilege revocation mechanism should be balanced against the cost of ensuring that a module with a high security label is trusted. The cost of ensuring that a module is trusted is usually much higher than the price of building a mechanism. To ensure the greatest

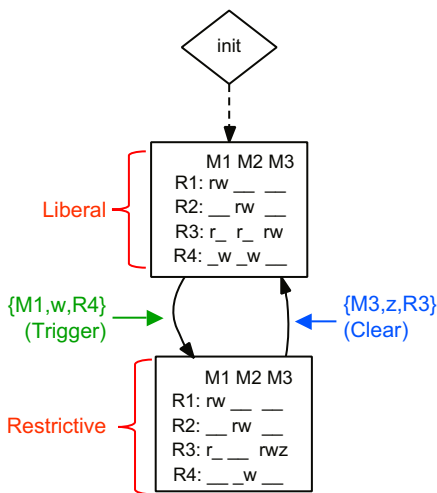


Fig. 19 – This redaction policy has four possible covert channels: from $Module_1$ to $Module_2$, from $Module_1$ to $Module_3$, from $Module_3$ to $Module_1$, and from $Module_3$ to $Module_2$.

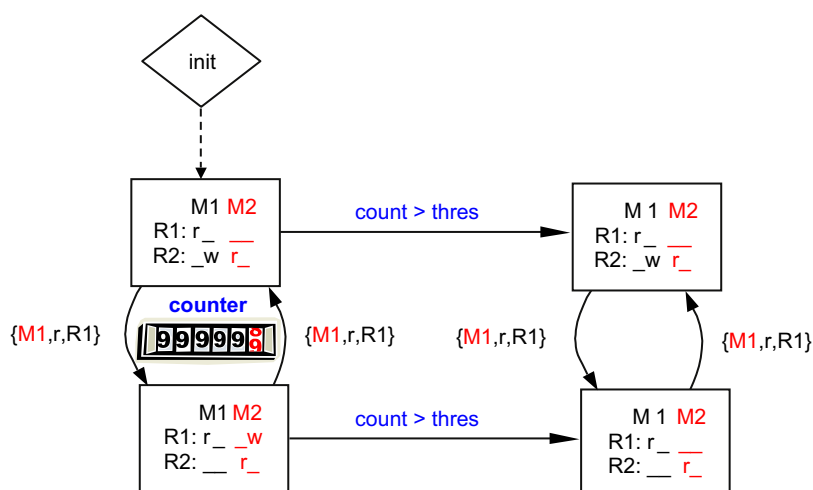


Fig. 20 – Coping with a covert channel in a stateful policy with two states. A counter measures the bandwidth of a covert channel from *Module*₁ to *Module*₂ by counting the number of times that a cycle occurs in the original policy on the left. If this counter exceeds a predetermined threshold, it is necessary to switch to the policy on the right, in which the covert channel has been eliminated by making the access matrices in both nodes of this policy identical with respect to *Module*₂.

likelihood that critical services will be maintained, only those privileges that pertain to the covert channel should be revoked when a core causes a counter to exceed its threshold. In other words, the revocation would be performed on a per-channel basis. To ensure that the combined policy does not introduce any new covert channels, the combined graph should be run through the detector, although the combined policy should not introduce any new covert channels if done correctly.

6.4. Related work

Policy engineering is an extremely important problem because an enforcement mechanism is only as good as the policy it enforces. Correctly designing a system that relies on a set of complex security policies calls for a new set of techniques to make it tractable for a human to correctly formulate policy specifications. Fong has developed a new approach to policy design by constraining the reference monitor to only track a “shallow execution history” of permitted resource access events (Fong, 2004). Although this restriction limits the number of enforceable policies, many classic security policies can still be enforced. Breaking down the class of policies that can be enforced by an execution monitor into subclasses makes the problem of policy design more tractable because specialized policy languages and verification techniques can be tailored to these classes, and they are more easily decomposed into reusable components.

Since Lampson first introduced the concept of covert channels (Lampson, 1973), several techniques for detecting covert channels in policy specifications have been proposed, including shared resource matrix methodology (Kemmerer, 1983), information flow (Millen, 1976), and noninterference (Haigh et al., 1987; Goguen and Meseguer, 1982), although this section focuses on the shared resource matrix method. Tsai et al. developed a static method of identification of covert storage channels in source code by using information flow analysis to identify kernel variables that are visible or can be

altered (Tsai et al., 1990). They observe that not all potential covert channels can be exploited because the conditions that make the channel possible may not exist at runtime. They distinguish between potential covert channels and real covert channels.

There is much prior work in estimating the bandwidth of covert channels. Millen applied information theory to calculate the bandwidth of a covert channel based on the number of possible encodings of the data and the probability of each symbol (Millen, 1987, 1989). Shieh proposes a method of measuring the bandwidth of covert channels in multilevel operating systems (Shieh, 1999). He observes that resource-exhaustion channels can be modeled as finite-state graphs, but event-count channels cannot. Tsai and Gligor developed a Markov model to compute the maximum bandwidth of a covert storage channel under different system loads (Tsai and Gligor, 1988).

7. Conclusions

Reconfigurable systems are blurring the line between hardware and software, and they represent a large and growing market. Due to the increased use of reconfigurable logic in mission-critical applications, a new set of security primitives is needed to prevent improper memory sharing and to contain memory bugs in these physically addressed embedded systems. We have demonstrated a method and language for specifying access policies that can be used as both a description of legal access patterns and as an input specification for direct synthesis to a reconfigurable logic module. Our architecture ensures that the policy module is invoked for every memory access.

Our formal access policy language provides a convenient and precise way to describe coarse or fine-grained computer security policies for modules on an FPGA. We have used our policy compiler to translate a variety of security policies to hardware enforcement modules, and we have analyzed the area requirements and performance of these circuits. Our

synthesis data show that our methods are both efficient and scalable in the number of ranges that must be recognized. In addition to the reconfigurable domain, our methods can be applied to systems-on-a-chip as part of a more general scheme.

We have also developed an automatic method of identifying security policies with inherent covert channels. We have identified a range of corrective actions that can be considered once a possible covert channel is detected. The ideal alternative is to eliminate the covert channel by changing the policy, but in case this option is not available, we have presented a method of coping with the covert channel by closing it when it reaches a designated threshold, as determined by a hardware counter that measures the real-time bandwidth of potential covert channel exploitation.

Since expressing some policies in our language requires complex expressions, we do not expect the typical engineer to work in our language. Because usability is fundamental to system security (Cynthia Irvine et al., 2002; Gutmann and Grigg, 2005), we plan to develop a higher-level language along with a set of tools to assist the engineer in constructing mathematically precise policies. This will build on the policy engineering work of Fong et al. (Fong, 2004). A higher-level language will allow the engineer to express policies in terms of security concepts (e.g., isolation, controlled sharing, etc.) rather than in terms of modules and ranges.

REFERENCES

- Aho A, Sethi R, Ullman J. Compilers: principles, techniques, and tools. Reading, MA: Addison Wesley; 1988.
- Altera Inc. Quartus II manual; 2004.
- Anderson JP. Computer security technology planning study. Technical Report ESD-TR-73-51, ESD/AFSC, Hanscom AFB, Bedford, MA; 1972.
- Bell DE, LaPadula LJ. Secure computer systems: mathematical foundations and model. Bedford, MA, USA: The MITRE Corporation; May 1973.
- Betz Vaughn, Scott Rose Jonathan, Marquardt Alexander. Architecture and CAD for deep-submicron FPGAs. Boston, MA: Kluwer Academic; 1999.
- Biba KJ. Integrity considerations for secure computer systems. Technical Report ESD-TR-76-372, USAF Electronic Systems Division, Bedford, MA; 1977.
- Bondalapati K, Prasanna VK. Reconfigurable computing systems. Proceedings of the IEEE 2002;90(7):1201-17.
- Bossuet L, Gogniat G, Burlinson W. Dynamically configurable security for SRAM FPGA bitstreams. In: Proceedings of the eighteenth international parallel and distributed processing symposium (IPDPS '04); April 2004. Santa Fe, NM.
- Brewer DFC, Nash MJ. The Chinese wall security policy. In: Proceedings of the 1989 IEEE symposium on security and privacy; 1989.
- Buell DA, Pocek KL. Custom computing machines: an introduction. Journal of Supercomputing 1995;9(3):219-29.
- Compton K, Hauck S. Reconfigurable computing: a survey of systems and software. In: ACM Computing Surveys, vol. 34 (2). USA: ACM; 2002. p. 171-210.
- Cormen TH, Leiserson CE, Rivest RL, Stein C. Introduction to algorithms. MIT Press and McGraw-Hill; 1990.
- DeHon A. Comparing computing machines. In: SPIE - int. soc. opt. eng. Proceedings of SPIE - the international society for optical engineering, vol. 3526; 1998. p. 124-33.
- DeHon A, Wawrzynek J. Reconfigurable computing: what, why, and implications for design automation. In: Proceedings of the design automation conference; 1999. p. 610-5. West Point, NY.
- Denning DE. A lattice model of secure information flow. Communications of the ACM May 1976;19(5).
- Erlingsson Ulfar, Schneider Fred B. Sasi enforcement of security policies: a retrospective. In: Proceedings of the 1999 workshop on new security paradigms; 1999.
- Fong Philip WL. Access control by tracking shallow execution history. In: Proceedings of the 2004 IEEE symposium on security and privacy; 2004.
- Fraser T, Badger L. Ensuring continuity during dynamic security policy reconfiguration in dte. In: Proceedings of the 1998 IEEE symposium on security and privacy; 1998. p. 15-26.
- Gerzic Amer. Codeguru: write your own regular expression parser; November 2003.
- Goguen JA, Meseguer J. Security policy and security models. In: Proceedings of the 1982 IEEE symposium on security and privacy; 1982. p. 11-20.
- Gutmann Peter, Grigg Ian. Security usability. IEEE Security and Privacy Magazine July/August 2005.
- Hadzic I, Udani S, Smith J. FPGA viruses. In: Proceedings of the ninth international workshop on field-programmable logic and applications (FPL '99); August 1999. Glasgow, UK.
- Haigh JT, Kemmerer RA, McHugh J, Young WD. An experience using two covert channel analysis techniques on a real system design. IEEE Transactions on Software Engineering February 1987;13(2):157-68.
- Harper Scott, Athanas Peter. A security policy based upon hardware encryption. In: Proceedings of the thirty-seventh Hawaii international conference on system sciences; 2004.
- Harper Scott, Fong Ryan, Athanas Peter. A versatile framework for fpga field updates: an application of partial self-reconfiguration. In: Proceedings of the fourteenth IEEE international workshop on rapid system prototyping; June 2003.
- Huffmire T, Brotherton B, Wang G, Sherwood T, Kastner R. Moats and drawbridges: an isolation primitive for reconfigurable hardware based systems. In: Proceedings of the 2007 IEEE symposium on security and privacy; May 2007. Oakland, CA, USA.
- Huffmire Ted, Prasad Shreyas, Sherwood Tim, Kastner Ryan. Policy-driven memory protection for reconfigurable systems. In: Proceedings of the European symposium on research in computer security (ESORICS); September 2006. Hamburg, Germany.
- Irvine C, Levin T, Nguyen T, Dinolt G. The trusted computing exemplar project. In: Proceedings of the fifth IEEE systems, man and cybernetics information assurance workshop; June 2004. p. 109-15. West Point, NY.
- Irvine Cynthia E, Levin Timothy E, Nguyen Thuy D, Shifflett David, Khosalim Jean, Clark Paul C, et al. Overview of a high assurance architecture for distributed multilevel security. In: Proceedings of the 2002 IEEE workshop on information assurance and security; June 2002. West Point, NY.
- Johnson S. Yacc: yet another compiler-compiler. Technical Report CSTR-32, Bell Laboratories, Murray Hill, NJ; 1975.
- Kastner Ryan, Kaplan Adam, Sarrafzadeh Majid. Synthesis techniques and optimizations for reconfigurable systems. Boston, MA: Kluwer Academic; 2004.
- Kean T. Secure configuration of field programmable gate arrays. In: Proceedings of the eleventh international conference on field programmable logic and applications (FPL '01); August 2001. Belfast, UK.
- Kean T. Cryptographic rights management of FPGA intellectual property cores. In: Tenth ACM international symposium on field-programmable gate arrays (FPGA '02); February 2002. Monterey, CA.
- Kemmerer RA. Shared resource matrix methodology: an approach to identifying storage and timing channels. In: ACM transactions on computer systems; 1983.

- Kemmerer RA. A practical approach to identifying storage and timing channels: twenty years later. In: Proceedings of the eighteenth annual computer security applications conference; December 2002. Las Vegas, Nevada, USA.
- Kocher P, Lee R, McGraw G, Raghunathan A, Ravi S. Security as a new dimension in embedded system design. In: Proceedings of the forty-first design automation conference (DAC '04); June 2004. San Diego, CA.
- Lach J, Mangione-Smith W, Potkonjak M. FPGA fingerprinting techniques for protecting intellectual property. In: Proceedings of the 1999 IEEE custom integrated circuits conference; May 1999a. San Diego, CA.
- Lach J, Mangione-Smith W, Potkonjak M. Robust FPGA intellectual property protection through multiple small watermarks. In: Proceedings of the thirty-sixth ACM/IEEE conference on design automation (DAC '99); June 1999b. New Orleans, LA.
- Lampson BW. A note on the confinement problem. *Communications of the ACM* October 1973;16(10):842–56.
- Lesk M, Schmidt E. Lex: a lexical analyzer generator. Technical Report 39, Bell Laboratories, Murray Hill, NJ; October 1975.
- Timothy E. Levin, Cynthia E. Irvine, Thuy D. Nguyen. A least privilege model for static separation kernels. Technical Report NPS-CS-05-003, Naval Postgraduate School; 2004.
- Linz Peter. An introduction to formal languages and automata. Sudbury, MA: Jones and Bartlett; 2001.
- Mangione-Smith WH, Hutchings B, Andrews D, DeHon A, Ebeling C, Hartenstein R, et al. Seeking solutions in configurable computing. In: *Computer*, vol. 30 (12); 1997. p. 38–43.
- De Micheli Giovanni. Synthesis and optimization of digital circuits. New York: McGraw-Hill; 1994.
- Millen JK. Security kernel validation in practice. *Communications of the ACM* May 1976;19(5):243–50.
- Millen JK. Covert channel capacity. In: Proceedings of the 1987 IEEE symposium on security and privacy; April 1987. Oakland, CA, USA.
- Millen JK. Finite-state noiseless covert channels. In: Proceedings of the computer security foundations workshop II; June 1989. Franconia, NH, USA.
- Navarro J, Iyer S, Druschel P, Cox A. Practical, transparent operating system support for superpages. In: Fifth symposium on operating systems design and implementation (OSDI '02); December 2002. Boston, MA.
- Raymond D, Wood D. Grail: A C++ library for automata and expressions. *Journal of Symbolic Computation* 1995;11:341–50.
- Rushby John. A trusted computing base for embedded systems. In: Proceedings seventh DoD/NBS computer security conference; September 1984. p. 294–311.
- Sabelfeld Andrei, Myers Andrew C. Language-based information-flow security. *IEEE Journal on Selected Areas in Communications* January 2003;21(1).
- Saltzer J. Protection and the control of information sharing in multics. *Communications of the ACM* July 1974;17(7):388–402.
- Sami Saydjari O. Multilevel security: reprise. *IEEE security and privacy magazine*; September/October 2004.
- Schaumont P, Verbauwheide I, Keutzer K, Sarrafzadeh M. A quick safari through the reconfiguration jungle. In: Proceedings of the design automation conference; 2001. p. 172–7.
- Schneider Fred B. Enforceable security policies. *ACM Transactions on Information and System Security* February 2000;3(1).
- Senior AW, Pankanti S, Hampapur A, Brown L, Tian Y-L, Ekin A. Blinkering surveillance: enabling video privacy through computer vision. Technical Report RC22886, IBM; 2003.
- Shieh S. Estimating and measuring covert channel bandwidth in multilevel secure operating systems. *Journal of Information Science and Engineering* 1999;15:91–106.
- Smith Richard E. Cost profile of a highly assured, secure operating system. In: *ACM transactions on information and system security*; 2001.
- Standaert F, Oldenzeel L, Samyde D, Quisquater J. Power analysis of FPGAs: how practical is the attack?. In: *Field-Programmable Logic and Applications*, vol. 2778 (2003); September 2003. p. 701–11.
- Stern DF. On the buzzword “security policy”. In: Proceedings of the 1991 IEEE symposium on security and privacy; 1991. p. 219–30. Oakland, CA.
- Tsai CR, Gligor V. A bandwidth computation model for covert storage channels and its applications. In: Proceedings of the IEEE symposium on security and privacy; 1988. p. 108–21.
- Tsai CR, Gligor V, Chandrasekaran C. On the identification of covert storage channels in secure systems. *IEEE Transactions on Software Engineering* June 1990;16(6).
- Vuillemin JE, Bertin P, Roncin D, Shand M, Touati HH, Boucard P. Programmable active memories: reconfigurable systems come of age. In: *IEEE transactions on very large scale integration (VLSI) systems*, vol. 4 (1); 1996. p. 56–69.
- Weissman Clark. MLS-PCA: a high assurance security architecture for future avionics. In: Proceedings of the annual computer security applications conference. Los Alamitos, CA: IEEE Computer Society; December 2003. p. 2–12.
- Cates Witchel J, Asanovic K. Mondrian memory protection. In: Tenth international conference on architectural support for programming languages and operating systems (ASPLOS-X); October 2002. San Jose, CA.
- Wollinger T, Guajardo J, Paar C. Security on FPGAs: state-of-the-art implementations and attacks. *ACM Transactions on Embedded Computing Systems* August 2004;3(3):534–74.
- Woodward J. Exploiting the dual nature of sensitivity labels. In: *IEEE symposium on security and privacy*; 1987. p. 23–30. Oakland, CA, USA.
- Zheng L, Myers A. Dynamic security labels and noninterference. Technical Report 2004–1924, Cornell University; 2004.

Ted Huffmire is an Assistant Professor of Computer Science at the Naval Postgraduate School. His research focus is hardware-assisted security, especially the development of policy-driven mechanisms for special-purpose devices and next-generation processors. Ted has a PhD in Computer Science from UC Santa Barbara. He is a member of the IEEE.

Timothy Sherwood is an Assistant Professor of Computer Science at UC Santa Barbara. His research is in the area of computer architecture, specifically in the development of novel high throughput methods by which systems can be constructed, monitored, and analyzed. Prior to joining UCSB in 2003, he received his BS in Computer Science and Engineering from UC Davis, and his MS and PhD in the same from UC San Diego. He is a member of both ACM and IEEE.

Ryan Kastner is an Associate Professor in the Department of Computer Science and Engineering at UC San Diego. His research interests include many aspects of embedded computing systems including reconfigurable architectures, digital signal processing, and security. Ran received his PhD in Computer Science from UCLA.

Timothy Levin is an Associate Research Professor at the Naval Postgraduate School. His research interests include design, analysis, and verification of high-assurance security architectures and dynamic security policies. Mr. Levin has a BS in Computer Science from UC Santa Cruz and is a member of both ACM and IEEE.