

Signature-Based Protection from Code Reuse Attacks

Mehmet Kayaalp*, *Student Member, IEEE*, Timothy Schmitt*, Junaid Nomani*, Dmitry Ponomarev*, *Member, IEEE*, Nael Abu Ghazaleh*, *Member, IEEE*

Abstract—Code Reuse Attacks (CRAs) recently emerged as a new class of security exploits. CRAs construct malicious programs out of small fragments (gadgets) of existing code, thus eliminating the need for code injection. Existing defenses against CRAs often incur large performance overheads or require extensive binary rewriting and other changes to the system software. In this paper, we examine a signature-based detection of CRAs, where the attack is detected by observing the behavior of programs and detecting the gadget execution patterns. We first demonstrate that naive signature-based defenses can be defeated by introducing special “delay gadgets” as part of the attack. We then show how a software-configurable signature-based approach can be designed to defend against such stealth CRAs, including the attacks that manage to use longer-length gadgets. The proposed defense (called SCRAP) can be implemented entirely in hardware using simple logic at the commit stage of the pipeline. SCRAP is realized with minimal performance cost, no changes to the software layers and no implications on binary compatibility. Finally, we show that SCRAP generates no false alarms on a wide range of applications.

Index Terms—Processor Architectures, Support for Security, Code Reuse Attacks.

1 INTRODUCTION

Exploits targeting software vulnerabilities remain one of the primary security threats to computer systems, with costs estimated in the 100s of billions of dollars [1]. The NIST national vulnerability database includes tens of thousands of vulnerabilities, with an average reporting rate of 10 new vulnerabilities per day [2]. Thus, it is critical to build systems that make exploits difficult to launch and that detect and limit their effect quickly.

Most current attacks start by exploiting a buffer overflow vulnerability. Despite significant efforts in devising solutions that prevent buffer overflows [3]–[6], they remain prevalent. Early code injection attacks overwrite the buffer with the malicious code on the stack and simultaneously overwrite the return address to point at the start of the exploit code [7], [8]. A number of software and hardware approaches to protect against such attacks were devised [9]–[12]. These efforts have culminated in the recent deployment of hardware memory protection mechanisms that do not allow a memory page to be both writable and executable at the same time (the so called $W \oplus X$ protection). These hardware extensions are supported by both AMD and Intel processors and deployed in both Linux and Windows operating systems [13], [14].

1.1 Code Reuse Attacks: Bypassing $W \oplus X$

In response to these defenses new Code Reuse Attacks (CRAs) emerged that construct a malicious program by stitching together carefully selected fragments of the existing library code; these snippets are called *gadgets* [15]. One example of a CRA is the return-oriented programming (ROP) attack, where each gadget ends with a return instruction to trigger the execution of the next gadget

pointed to by the next return address on the stack. All the attacker has to do is to inject a proper sequence of return addresses onto the stack to point to the needed gadgets. ROP was shown to be Turing-complete on a variety of platforms [16]–[20]. Automated tools have been developed that allow unsophisticated attackers to construct arbitrary malicious programs using ROP [21]–[24].

Several defense mechanisms against ROP have been recently proposed [25]–[30]. Perhaps the simplest of these solutions are the ones that utilize a shadow call/return stack, where the return instructions are matched against the corresponding calls using protected memory space [28]–[30]. We assume that such an enforcement of call-return pairs is already in place and therefore simple ROP-based attacks are defeated.

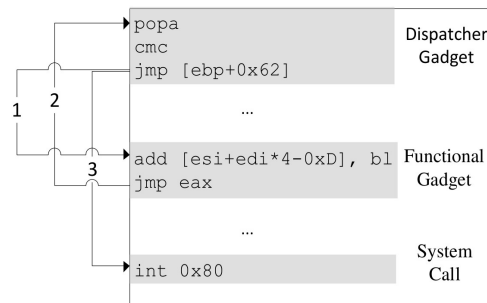


Fig. 1. Example of a simple JOP attack.

Unfortunately, a new form of CRA was developed that does not rely on return instructions [31]–[33]. In this jump-oriented programming (JOP) model, the attacker chains the gadgets using a sequence of indirect jump instructions, rather than return instructions. A special dispatcher gadget is used to orchestrate the control flow among the gadgets. A high level example of the

*Computer Science Department, Binghamton University, Binghamton, NY 13902. Email: {mkayaalp,tschmitt1,jnomani1,dima,nael}@cs.binghamton.edu

JOP attack model is shown in Figure 1. This diagram shows how the attack will jump from the dispatcher gadget to functional gadgets which will then return the control back to the dispatcher gadget. The jump locations change based on the addresses popped off the stack by the dispatcher gadget, and will ultimately result in the execution of a system call.

1.2 Proposed Solution: Signature-based CRA Detection

Although it may appear that CRAs are a narrow form of attack, they represent a wide-open vulnerability that is increasingly used to exploit common buffer overflows. For example, Apple’s operating system for mobile devices (iOS) employs a secure boot chain and code signing to prevent any untrusted code from executing [34]. However recent approaches to jailbreak and software unlock such devices are CRA-based [35] and are able to bypass all these security measures. Thus, it is critical to develop solutions that protect against this major vulnerability, preferably in a way that protects legacy binaries.

In this paper, we propose Signature-based CRA Protection (SCRAP): a simple and low-overhead hardware scheme to protect against JOP attacks that are based on the dynamic detection of attack signatures, or the patterns of executed instructions that are indicative of the JOP attack. SCRAP works because the attack patterns are significantly different from those of the regular programs as they execute frequent indirect jump (or call) instructions to jump from gadget to gadget. Previous work [26] investigated this type of defense for ROP attacks and showed that it has promise. They implemented a defense mechanism called DROP in software using Valgrind tool to detect the ROP pattern. Because it is implemented in software, DROP incurs over 5X performance loss on the average across simulated workloads, mainly due to the overhead of Valgrind.

Starting from DROP, we made several observations about existing signature-based detection that motivated this work. First, the ideas of signature-based detection can be extended to protect against the JOP attacks if one uses the indirect jumps as the gadget boundaries. Second, the high performance overhead of DROP (appropriately adapted to protect against JOP attacks) can be avoided by implementing the checking logic in hardware, placing this hardware off-the-critical path in the commit stage of the pipeline, and performing simple checks during instruction commitment. If successful, this approach can provide protection with much lower overhead and complexity compared to the previous solutions and can naturally protect the existing binaries. Third, and most important, the naive implementations of the signature-based detection along the lines of DROP can be bypassed because of the strong assumptions it makes about usable gadget lengths. For example, we demonstrate an attack that uses a delaying gadget through a function call in the middle of the attack with the only purpose to distort the attack signatures expected by a DROP-like signature-based defense. Finally, the thresholds on the length of gadgets assumed by DROP

are not absolute: although difficult, it is possible to find longer gadgets and integrate them into an attack, avoiding detection. In this paper, we present a complete working example of such a stealth JOP attacks integrating delay gadgets, and using gadgets longer than the DROP thresholds.

Motivated by these observations, we propose an attack signature detection logic that protects against such stealth JOP attacks by filtering out the spurious function calls in the middle of the attack from the attack signature. We develop a language for the possible attack sequences and derive from it a state machine implementation of the detection logic. We show that the proposed mechanism generates no false alarms in any of the regular workloads that we considered and successfully detects CRAs, even when delay gadgets are used, for a large number of shellcodes. Finally, we extend the detectors to tolerate infrequent use of longer gadgets.

We propose implementing the detector in hardware both for performance and legacy binary support reasons, but the main reason is that hardware solutions are able to detect even when *unintended instructions* (see Section 2.2) are used by the attacker. Software solutions such as CFI [36], CFL [37], Google NaCl [38] try to prevent the attacker from ever using an unintended instruction. But if only one control flow change could be executed, then the attacker could bypass all the checking instructions by only using unintended instructions. Such a starting point might be due to a bug in the verifier/binary rewriter or due to a portion of the code that is not checked. In hardware based solutions, even unintended instructions are subject to checks.

SCRAP has the following key characteristics:

- It successfully detects all JOP attacks that we were able to generate, while resulting in zero false alarms across regular code base.
- It incurs minimal performance cost (less than 2%) and only requires simple hardware at the commit stage of the pipeline. There is also no impact on the processor cycle time.
- It does not require complex binary rewriting, binary annotation, or construction of a full control flow graph of a program. It also does not require compiler or ISA support and can be used to protect legacy binaries.
- With a simple hardware support, it performs checks for unintended jumps (in variable instruction-length architectures, such as x86) thus closing the potential security vulnerability of purely software-based solutions.

This submission is an extended version of the paper that appeared in HPCA-2013 conference [39]. The conference paper has been significantly extended in the following ways:

- In the conference version, we only evaluated the impact of SCRAP on SPEC 2006 benchmarks. In this submission, we extend the study of false alarms due to SCRAP to a number of other applications, including Adobe Flash Player, Apache web Server and Mozilla Firefox web browser and Xpdf PDF

viewer. We present including Xpdf, Adobe Flash Player, Apache2 Web Server and Mozilla Firefox web browser. We present detailed results for these applications and demonstrate that no false positive alarms occur during their execution.

- Figure 12 and Figure 14, showing the false positive rates, have been significantly improved to clarify the specific benchmarks that have at least one false positive for varying SCRAP parameters. These figures also include the statistics from the newly evaluated applications listed above.
- To analyze the impact on the critical path delay and dynamic power consumption, we implemented the proposed SCRAP detector in Verilog HDL on an FPGA with a 90nm process. We evaluated both designs: a vanilla SCRAP presented in Section 7 and the two-threshold SCRAP variation discussed in Section 10. For comparison, we also evaluated 8-, 16-, 32- and 64-bit counters in the same technology. Our results are presented in Section 11, showing that the overhead of SCRAP from both timing and power standpoints is negligible. Specifically, a simple $G_{7,4}$ SCRAP logic, has a shorter delay than an 8-bit counter and consumes as much power as a 16-bit counter.
- To estimate the memory overhead of SCRAP, we added Figure 11, which shows the memory footprint of the Secure Call Stack using different alignments for SCRAP counters. Specifically, we evaluated the overhead of adding byte- and word-long SCRAP counters to each Secure Call Stack entry, and we show that the additional memory overhead due to SCRAP counters is negligible.

2 CRA MECHANICS AND EXAMPLE

In this section, we overview a fully functional example of a JOP attack. We follow by discussing how variable length ISAs such as x86 and x86-64 significantly increase the number of gadgets available for attacks.

2.1 Functional JOP Attack Example

Figure 2 shows an example of the malicious shell code to be executed by the attacker. The purpose of this simple code is to execute a system call that starts a new shell. For this example, we use the standard C library (`libc 2.11.3`) as the code base for the gadget composition. Table 1 shows the gadgets that we found in `libc` to carry out the functionality of the attack from Figure 2. Finally, we show the dynamic sequence of the discovered gadgets to execute this attack and explain the functionality and purpose of each dynamic gadget invocation.

In order to launch a shell using the gadgets in Table 1, this type of attack has to accomplish two things: the correct parameters for a system call must be placed in the argument registers and a system call must be made. To launch a shell, our example attack makes a system call to `execve`. When the system call is made, registers `ecx` and `edx` must point to a null word, `0x00000000`,

```

; Load the syscall number for execve to eax
xor eax, eax          ; Set eax to 0
mov al, 0x0b         ; Set eax to 0x0b

; Point ecx and edx to a null word
mov ecx, NULL        ; NULL points to 0x00000000
mov edx, NULL

; Point ebx to the executable path
mov ebx, SH          ; SH points to "/bin/sh"

int 0x80             ; Make the system call

```

Fig. 2. Example shellcode in assembly.

Gadget	Gadget Function
g0	Dispatcher
g1	Null-Writer
g2	System Call

TABLE 1
Gadgets used in example attack.

and `ebx` must point to the string `"/bin/sh"`. Both null words and the string `"/bin/sh"` can be found in memory; we can place their addresses onto the stack and let the JOP attack pop them into the appropriate registers. The remaining step in the attack is to initialize the value of the `eax` register.

When the system call is made, `eax` must contain `0x0000000b`, indicating a call to `execve`. However, a JOP attack typically depends on exploiting a buffer overflow; these attacks typically rely on a buffer overflow which is exploited by the attacker to place data on the stack. The buffer is typically a string buffer, so a `0x00` byte causes the system to terminate reading the string; the attacker cannot use null values in the initial overflow. If the attack needs any null values, such as those in the word `0x0000000b`, the attack must generate them itself.

We make use of a *null-writer* gadget to create null values on the stack that will eventually be popped into `eax`. Our null-writer is constructed with an `add` instruction, adding the byte held in `b1` to the byte on the stack pointed to by `esi+edi*4-0xD`. If we place bytes holding `0xff` on the stack as part of the initial overflow attack and ensure that `b1` contains `0x01`, we can add `0x01` to `0xff` on the stack, overflowing to a `0x00`. Using this method, our attack creates the word `0x0000000b` on the stack where it can be popped into `eax` as the final step before the system call gadget is used.

In the remainder of this section we show how the attack executes using the gadgets described in Table 1. We assume the attacker has exploited a buffer overflow to place data on the stack and redirect control flow to the dispatcher gadget (`g0`). From the dispatcher gadget, the attack proceeds to execute the null-writer gadget (`g1`), then `g0`, `g1`, `g0`, `g1`, `g0`, and finally, the system call (`g2`). Below, each step starts with the gadget number followed

by an explanation of how it advances the attack.

Step 1 - g0 The dispatcher gadget initiates the attack with a `popa` instruction. This instruction populates the registers with useful values the attacker has placed on the stack. The second instruction, `cmc`, has no meaningful effect on this attack. After initializing the registers with values necessary for an attack, the dispatcher jumps to the null-writer gadget.

Step 2 - g1 The null-writer gadget adds the byte held in `bl` to the byte that `esi+edi*4-0xD` points to. In Step 1, the dispatcher gadget populated the registers so that `bl` contains `0x01` and `esi+edi*4-0xD` points to the value `0xff` in the future value of `eax` on the stack.

Step 3 - g0 Populate the registers with the values necessary to perform the null-writer a second time.

Step 4 - g1 Write `0x00` to a second byte in the future value of `eax`.

Step 5 - g0 Populate the registers with values for a third and final execution of the null-writer.

Step 6 - g1 Write the final null value onto the stack where `eax` is popped from.

Step 7 - g0 Populate the registers with the appropriate values for a system call. The value that is popped from the stack to `eax` is `0x0000000b`.

Step 8 - g2 Make a system call to `execve()`, launching a new shell.

2.2 Gadgets and Unintended Instructions

For ISAs such as x86 with variable size instructions, the attackers can find gadgets that are unintended by the programmer. Specifically, these are instructions that start at a byte in the middle of a multi-byte instruction. These instructions account for a large number of the gadgets exploitable by attackers [31].

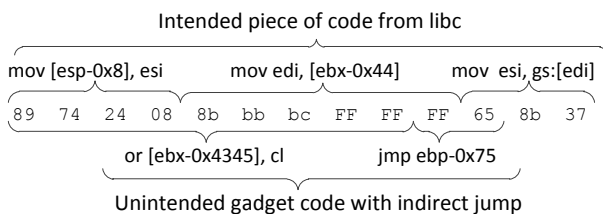


Fig. 3. Example gadget with unintended jump.

To illustrate the concept of unintended branches, we show a sequence of bytes from the `libc` library in the top part of Figure 3. If the decoding starts after skipping the first four bytes, a different instruction sequence can be decoded as shown at the bottom of Figure 3, containing an indirect jump that the programmer did not intend to execute. Although the unintended gadgets far exceed intended gadgets in number, they are often harder to utilize because they can include rarely-used instructions with complicated addressing modes and constants. Thus, only short unintended gadgets are typically usable.

3 UNDERSTANDING SIGNATURES OF JOP ATTACKS

Signature based defenses can only work if the instruction patterns exhibited by the attack code can be distinguished from those of normal programs. The JOP attack patterns (the number and the length of gadgets used) are different from the patterns of ROP attacks examined in [26] because of two factors: 1) the reliance on indirect jumps instead of returns; and 2) the need to execute the dispatcher gadget to orchestrate the gadget-level control flow, thus requiring more gadgets for an attack.

In terms of the number of gadgets, Chen et al. [26] reported that at least three consecutive gadgets are needed to carry out even a simple ROP attack. For JOP, the number of gadgets needed is higher because of the need to call the dispatcher gadget after every functional gadget. In addition, it is much easier to compose an attack using short-length gadgets to limit the undesirable side effects on the program state.

All existing tools for automatic gadget discovery [15], [31] therefore limit the gadget size to at most five instructions and only consider usable the gadgets that perform one operation (and one state update). The work of [26] also used gadget sizes of at most five instructions for implementing the shellcodes in ROP-style attack. Signature based detection relies critically on these threshold values, so it is important to verify that they hold.

3.1 Gadget Analysis for JOP Attack

The size of a usable gadget is limited by the side-effects that the gadget has on the program state (including memory locations and registers). Large gadgets typically overwrite many registers and/or memory locations, thus corrupting the state and making attack continuation very difficult or impossible. This is especially true for the gadgets that are comprised of unintended instructions.

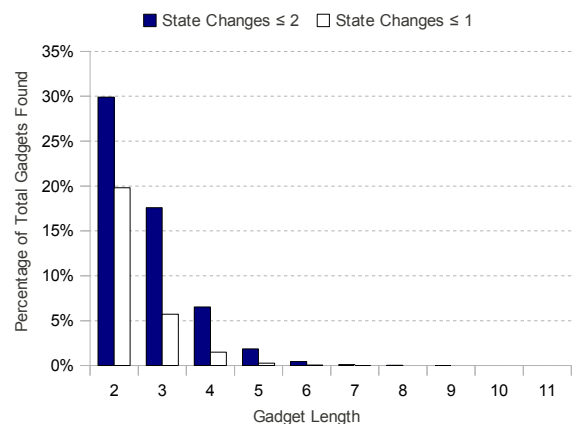


Fig. 4. Gadget length and state changes statistics for standard C library.

To understand the side-effect properties of the JOP gadgets, we performed extensive gadget analysis within the code base of several libraries. Our gadget discovery

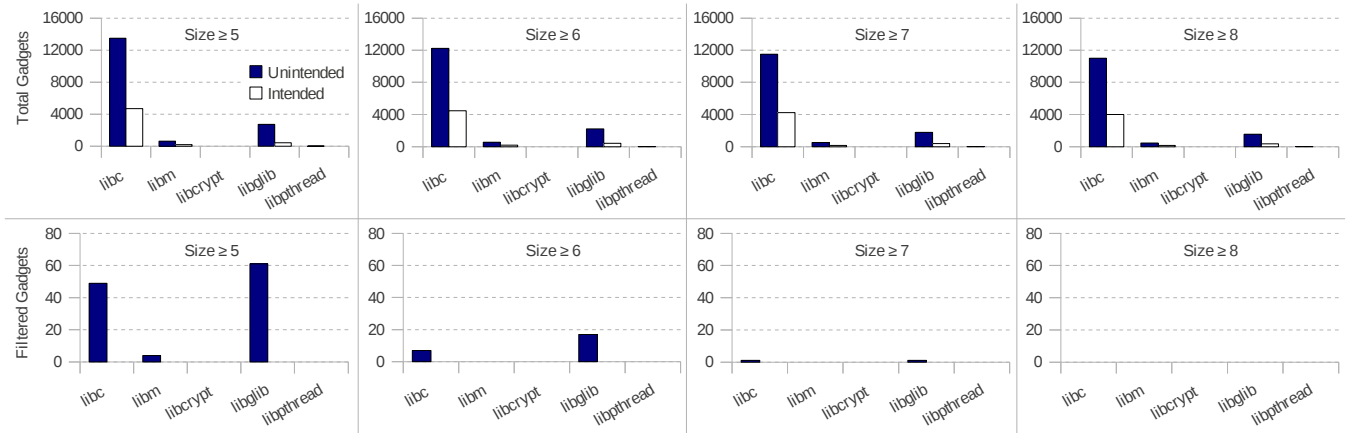


Fig. 5. Gadget length and side effect analysis: top figures show the total number of gadgets of a given length while the bottom figure shows the gadgets for the same length with the shown number of side effects.

algorithm starts with building the gadget trie as described by Shacham et al. [15]. In a gadget trie, indirect jump instructions are represented as nodes immediately under a dummy root node. A child node under an indirect jump represents a possible decoding of an instruction preceding the parent instruction. Since multiple possible instructions (all but one unintended) can precede an indirect branch, the trie can branch leading to multiple gadgets ending at the same indirect branch. Once the trie is constructed, the algorithm traverses the nodes starting with an indirect branch toward its children, and every path along this traversal represents a possible gadget.

Signature detection relies critically on the observation that usable gadgets are short allowing us to distinguish attacks from normal programs where the distance between indirect branches are significantly longer. We base our approach to the usability of gadgets on the number of state updates that a gadget performs. State updates are register limiting instructions such as register writes or indirect memory accesses (which force registers to be a specific value in order to prevent illegal accesses). We contend that longer gadgets that make multiple state updates are difficult to use without destroying the attack state.

Figure 4 shows the total number of gadgets discovered by the algorithm in the standard C library (libc), as well as the number of gadgets that remain after we remove the gadgets that do more state changes than each given threshold. Figure 5 shows the same gadget statistics for other common libraries. The top part of the figure shows the total number of gadgets of a given length (each length is a separate figure). The bottom part shows the number of gadgets present (of the same length as the corresponding top figure) with at most one state update. While a significant number of gadgets of various sizes obviously exist in the libraries, there are no gadgets of size eight instructions or more that perform less than two state updates (to memory or registers).

Figure 6 shows the average number of side effects as the gadget length increases. It also shows the minimum

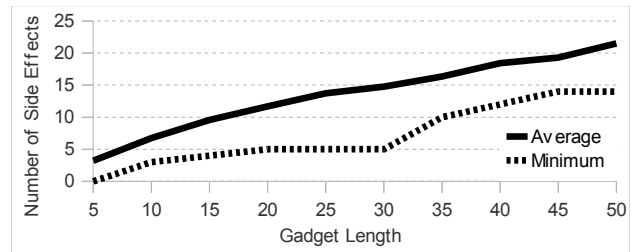


Fig. 6. Number of side effects as gadget length increases.

number of side effects in gadgets of that length found across all the libraries we studied. As the gadget length grows the number of side effects grows linearly making them increasingly more difficult to use.

Even at a threshold of 7, there exists only one gadget with a single state update in libc, and another one in libglib-2.0. Upon further examination, we found both of these gadgets not to be usable because they use unintended instructions that cannot be used. Since no suitable gadgets of seven instructions or more were found in multiple libraries, a threshold of seven instructions can be used by SCRAP to identify a gadget. However, using this length as a hard threshold represents a strong assumption: the attacker may be able to tolerate some of the side-effects in a long gadget, allowing her to use it as a delay gadget and bypass the detection. We later relax this assumption to build signature detectors that are resilient to the presence of some longer gadgets.

4 STEALTH JOP ATTACKS: CONCEALING ATTACK PATTERNS WITH DELAY GADGETS

From the discussion in the previous section, it appears that simple signature-based detection can be effectively applied to protect against JOP attacks. However, when designing security solutions it is important to assume that the attacker is aware of the particular defense that is implemented and consider possible attack modifications that would bypass this protection.

All JOP and ROP variations developed to date only considered the functional requirements of the attack. Therefore, all gadgets used by the attackers were performing some useful part of the attack code. In addition, to avoid the necessity of dealing with gadget side-effects, the existing automatic tools for generating JOP and ROP attacks only consider small gadget sizes. Signature-based approaches are effective under these assumptions, as shown in [26] and also by the analysis in the previous section.

However, what if the attacker is aware of the signature-based protection and modifies the attack to distort its execution patterns from those expected by the defense? One approach for accomplishing this is to introduce a delay gadget in the middle of the attack. The purpose of a delay gadget is not to execute any part of the attack code, but rather perform some spurious computations in a way that would not corrupt the machine state needed by the attack. At the same time, the gadget would be long enough to reset the gadget count used by the signature detector, before an attack is detected. In this section, we introduce such delay gadgets and demonstrate how the attack shown in the background section can be modified to incorporate it.

The analysis in the previous section showed that long gadgets have too many side effects to be usable; however, it is possible to create a small sized delay gadgets by using a call to a function. Since most functions have no side effects, they represent an ideal vehicle for implementing delay gadgets without destroying the program state. If a function call results in executing a larger number of instructions the signature based attack detector will reset (assuming that this is a valid program), allowing the attacker to continue the attack. In the remainder of this section, we demonstrate how to implement a delay gadget using a function call (using `atoi()`).

	Gadget	Gadget Function
	<code>call, [ecx-0x56000a00]</code>	
	<code>add bl, bh</code>	
g3	<code>inc ebx</code>	Delay
	<code>add dh, bh</code>	
	<code>jmp edi</code>	

TABLE 2
Delay gadget used in stealth JOP attack.

An example of a delay gadget that makes a call to the `atoi()` function is shown in Table 2, this gadget was found in the `libc` library. `atoi()` executes many more instructions than the typical JOP gadgets, bypassing signature based detection. When `atoi()` returns, some registers such as `eax`, `ecx`, and `edx` may have been altered and do not contain data that is meaningful to the attack. However, by convention, other registers such as `ebx`, `esi`, `edi`, `esp`, and `ebp` are saved. As long as the delay gadget ends with an indirect jump based on one of these saved registers, the attack can return to the dispatcher gadget which can recover from any side effects caused by the delay.

This new attack, which we call Stealth-JOP, is mounted using the same series of gadgets as our previous example, but with delay gadgets called periodically to avoid detection. Our previous JOP attack jumped from the dispatcher gadget to a functional gadget, and then back to the dispatcher. The Stealth-JOP attack example jumps from the dispatcher to a functional gadget, and then to the delay gadget. After the delay gadget has executed, the control returns to the dispatcher. Thus, there is no sequence in the code with multiple consecutive short gadgets, making DROP-like signature detection fail. At the same time, the attacker is able to execute arbitrary code using the short functional gadgets.

In addition to considering delay gadgets through function calls, it is important to note that if even one gadget of length higher than the detection threshold in DROP can be used (or at least tolerated) in an attack, then an attacker can exploit this gadget to bypass signature detection. We build the basic SCRAP detectors first assuming that the gadget lengths derived in Section 3 represent hard limits; that is, every gadget that makes 2 side effects or more is not usable. However, it is highly likely that a motivated attacker will be able to find at least some longer gadgets whose side effects can be tolerated; we were able to identify multiple such gadgets in constructing our attacks. We then relax this assumption and develop more sophisticated signature detectors in Section 10, that are able to tolerate the presence of some longer gadgets and still detect an attack.

5 THREAT MODEL, ASSUMPTIONS AND LIMITATIONS

We use standard CRA assumptions on the attacker’s access to memory; this could be obtained using a buffer overflow, a string formatting attack, or a non-local jump buffer (using `setjmp` and `longjmp` [40]). We assume that the system has NX support for writable memory such that code injection attacks are not possible.

We assume that the attacker can find arbitrary gadgets limited only by the attack lengths as per the analysis we showed in Section 3. Later we relax this assumption by allowing the use of longer gadgets. Throughout the paper, we present real attacks constructed from existing library code. However, rather than assume security due to our inability to find gadgets in the current version of the libraries, we make the assumption of the existence of arbitrary gadgets such that the defense works with any future code base, and not just the ones we used for the analysis.

We assume that the vulnerability exploited to initiate the attack does not lead to a privilege escalation. If privilege escalation is achieved from the initial vulnerability, then a CRA attack is not necessary. The attacker may seek to obtain privilege escalation through the CRA.

The new stealth JOP attack proposed in this paper uses delay gadgets to obfuscate the JOP execution pattern. We explored the use of function calls as delay gadgets because of the limited side-effects that they generate. Our analysis also showed traditional gadgets are ineffective beyond a certain length because of the presence of state

updates. However, there is a possibility that additional patterns of generating delay gadgets may exist (e.g., a loop gadget), although we have not been able to find and exploit such gadgets. We believe that the detection logic can be extended to capture such delay patterns as well.

6 EXPRESSING ATTACK SIGNATURES IN FORMAL LANGUAGE

In this section, we formalize the attack pattern as a context-free grammar. This formal description is used as the basis for the hardware implementation of SCRAP logic. We encode executions of instructions as strings of symbols denoting types of instructions, called signatures. The attacks are then formalized as formal languages of signatures. The alphabet used in this section is given in Table 3.

Symbol	Instruction
w	Indirect Jump
x	Indirect Call
y	Call
z	Return
a	All Other

TABLE 3
Signature alphabet.

6.1 Expressing Attacks Without Delay Gadgets

We observe that basic CRAs, such as ROP and JOP attacks, can be expressed as a formal language defining an attack as the following regular expression that uses POSIX Extended Regular Expressions:

$$R_{N,S} = (a \{0, N\} (w|x)) \{S, \}$$

Here, w denotes an indirect jump and x denotes an indirect call, while a denotes any other type of instruction. N is a parameter that specifies the number of instructions that a gadget can have, while S specifies the number of consecutive gadgets considered as an attack. For example, in $R_{5,3}$ case, three consecutive gadgets each having no more than five instructions form an attack.

6.2 Expressing Attacks with Delay Gadgets

With the inclusion of function calls as delays, the formal language defining the attack becomes a context free language, formalized as the context-free grammar $G_{N,S}$, where again N is the number of instructions that a gadget can have and S is the number of consecutive gadgets considered as an attack. The definition of $G_{5,3} = (V, \Sigma, Rules, Attack)$ is given in Figure 7.

The grammar starts with *Attack* which is expanded to $S = 3$ phases, each including a gadget and an unbounded number of delays. A gadget is the same as the $G_{N,S}$ regular expression defined above in Section 6.1. A delay starts with a *Call* and ends with a *Return* and a *Body* between them which we further define to capture

$$\begin{aligned}
 V &= \{Attack, P, Gadget, Delays, Delay, \\
 &\quad Call, Body, Return, Gadget, Indirect, \\
 &\quad NotGadget, NotAttack\} \\
 \Sigma &= \{w, x, y, z, a\} \\
 Rules &= \{ \\
 Attack &\rightarrow P P P \\
 P &\rightarrow Gadget Delays \mid Delays Gadget \\
 Gadget &\rightarrow Indirect \mid a Indirect \mid a a Indirect \mid \\
 &\quad a a a Indirect \mid a a a a Indirect \mid \\
 &\quad a a a a a Indirect \\
 Indirect &\rightarrow w \mid x \\
 Delays &\rightarrow Delay Delays \mid \varepsilon \\
 Delay &\rightarrow Call Body Return \\
 Call &\rightarrow x \mid y \\
 Return &\rightarrow z \\
 Body &\rightarrow Delays Body \mid Body Delays \\
 Body &\rightarrow a Body \mid Body a \mid \varepsilon \\
 Body &\rightarrow NotGadget NotAttack \\
 NotGadget &\rightarrow a a a a a Indirect \mid a NotGadget \\
 NotAttack &\rightarrow \varepsilon \mid P \mid P P \}
 \end{aligned}$$

Fig. 7. Definition of $G_{5,3} = (V, \Sigma, Rules, Attack)$.

complex delay gadgets consisting of nested function calls. Specifically, the delay gadget can have any number of delay function calls, and any number of unimportant instructions. It can also include less than S gadgets in it as long as there is a *NotGadget* sequence before it. A *NotGadget* has more than N instructions before the *Indirect* instruction.

The grammar is given for specific N and S values, but it can be reformulated for any N and S value by simply changing some of the production rules. *Attack* has S number of P expansions and *Gadget* allows N many a 's before *Indirect*. *NotGadget* and *NotAttack* would also have to be changed accordingly.

Signature	$\in R_{5,3}$?	$\in G_{5,3}$?
aaawaawaaw	Yes	Yes
awaxaaaaaw	Yes	Yes
awaxaaaaazaxaw	No	Yes
awaxaayaazaxaw	No	Yes

TABLE 4
Example attack signatures.

Table 4 shows example attack signatures and whether they are considered as an attack under prior approaches described in Section 6.1 and under the grammar that excludes delays. The parts of the signature that are matched as delays under $G_{5,3}$ are highlighted.

7 SCRAP: HARDWARE-BASED SIGNATURE DETECTION

In this section, we demonstrate an efficient hardware implementation to recognize the formal grammar that

expresses the attack signatures shown in the previous section. The proposed logic required by SCRAP is located at the commit stage of the pipeline off of the critical timing path. In the subsections below, we describe the components of SCRAP, building from a single state machine towards developing the complete solution. This is a standard exercise of translating the language grammar into the hardware implementation; however, because up to four instructions commit every cycle, we introduce an optimization that significantly simplifies the logic without having any adverse impact on the performance.

7.1 The SCRAP State Machine

The SCRAP state machine is shown in Figure 8. We use a counter to keep track of the current gadget length, and a comparator to decide whether the counter is above the gadget length threshold. When a gadget end is detected (w or x event in the language), the gadget length is used to transition through the shown finite state machine. The remaining step to implement the push down automata is to note that when a call instruction is encountered, we push the current state number to the shadow stack. This number is restored when a return instruction is encountered.

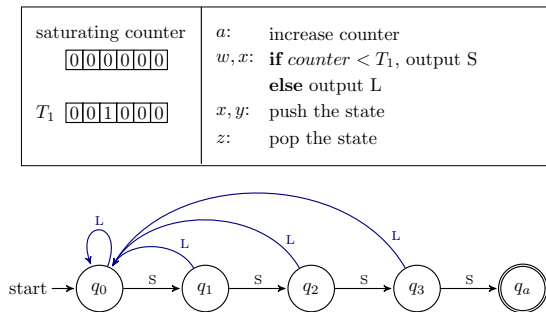


Fig. 8. The state machine for SCRAP.

7.2 Integrating State Counters into Secure Call Stack

As we discussed previously, a shadow call stack is a mechanism that has been proposed to defend against simple ROP attacks [28]–[30], [41]. SCRAP relies on a hardware implementation of the call stack, which is backed up by a larger software stack. In our design, each entry of the hardware stack is augmented with the counter that keeps track of the number of potential attack gadgets that executed consecutively. This makes it possible to track the information about the state of the attack even across function calls, eliminating their use as delay gadgets.

7.3 The SCRAP Microarchitecture

We now describe the microarchitectural changes needed for an out-of-order superscalar processor to implement

SCRAP. First, as the instructions are decoded, the information about the relevant instruction types is extracted and placed in the Reorder Buffer (ROB) entries allocated for the instructions. For this purpose, all instructions are classified into five types, as defined by the attack grammar in Section 6, thus requiring a new 3-bit wide field within each ROB entry to carry this information. When the instructions reach the commit stage of the pipeline, this information is used to update the SCRAP state machine counters.

The complexity of the counter update logic depends on the superscalar width (i.e. how many instructions commit per cycle) and also on the thresholds on the gadget length and the number of consecutive gadgets used by SCRAP. To simplify the logic, to ensure that only one counter update can be performed per cycle, and also to ensure that in a single cycle we operate on the counters within a single entry of the secure stack, we propose a technique called Commit Throttling.

7.3.1 Simplifying SCRAP through Commit Throttling

To simplify the SCRAP state machine counter update logic, we propose Commit Throttling (CT), which allows only one of the following instructions to be committed in a single cycle: CALL, indirect CALL, indirect jump, and RET. The number of these instructions in typical programs is small (less than 5% combined according to our analysis based on the binary instrumentation of SPEC 2006 benchmarks). When encountering the second instruction from this list in the co-committing group in the same cycle, the commit logic blocks and delaying the commit the second instruction to the next cycle. An additional requirement that we impose is that whenever a return instruction is encountered, the commit process also stops to ensure that we always operate on the counters within the same stack entry in each cycle. The impact of CT optimizations on the performance is negligible (less than 0.03% on the average for SPEC 2006 benchmarks), but it allows us to significantly limit the number of different instruction patterns coming out of the commit stage in a single cycle in terms of their impact on the SCRAP detection state. We present details of our implementation in Section 11.

7.4 Allowing software configuration of SCRAP

We allow the SCRAP detector thresholds to be configurable using a privileged system call that sets the detection machine state. We build large detector allowing up to 10 gadgets in a row to be detected. The configuration can be changed to $G_{N,S}$ by changing the T_1 threshold register to N and by marking the S^{th} state in the detector to be the finish state detecting the presence of an attack.

The choice of software configurability is made for two reasons. First we observed significant divergence in application behavior. Without software configurability, we are forced to use the worst case thresholds that do not generate false positives across any applications. Many applications do not use indirect branch and call instructions frequently, and can benefit from lower thresholds which further increase the difficulty of attacks. At the

same time, we want to protect against the potential of an application that does generate false positives against our thresholds. If the thresholds are fixed in hardware, then such an application cannot be supported.

8 PERFORMANCE EVALUATION OF SCRAP

For evaluating the performance impact of SCRAP, we used PTLsim [42] - a cycle-accurate x86 processor simulator. We simulated a 4-wide issue out-of-order core with 64KB L1 data and instruction caches, 512KB L2 cache and 2 MB L3 cache. Memory latency was assumed to be 100 cycles. We used 17 C and C++ SPEC CPU2006 [43] benchmarks for our experiments. The benchmarks were compiled using GCC-4.2 compiler on a x86 machine running Ubuntu with kernel version 2.6.24.

Each benchmark was simulated for 2 billion committed instructions after fast-forwarding for the first 100 million instructions.

First, we studied the impact of the Commit Throttling optimization. We discovered that there was negligible slowdown due to CT (less than 0.1% on average). To explain this slowdown, we show in Figure 9 the percentage of cycles where CT initiated a commit block. The cost of most of these stalls is hidden by out-of-order execution, resulting in the observed low impact on overall performance.

For a 4-entry hardware buffer of the secure call stack, the performance overhead of SCRAP is just over 1% on the average and it is less than 6% for all benchmarks as shown in Figure 10. This includes the overhead of stalls due to CT cycles as well as the overhead of the overflow of the hardware call stack buffer.

The additional memory requirement for SCRAP (and also the secure call stack) is shown in Figure 11. SCRAP uses small counters that easily fits in a byte, but using word-long counters is preferable for alignment purposes. The results show that, even with longer counters, memory footprint of SCRAP is less than a memory page of 4 KBytes.

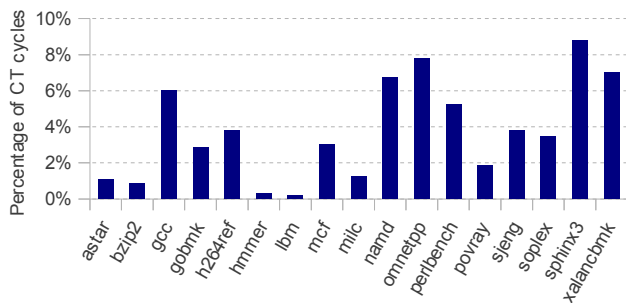


Fig. 9. Percentage of cycles where commit is blocked by CT.

9 SECURITY ANALYSIS OF SCRAP

In this section, we analyze the SCRAP detection effectiveness. We first demonstrate that it results in no false positives for normal programs and then analyze detection of actual shellcodes.

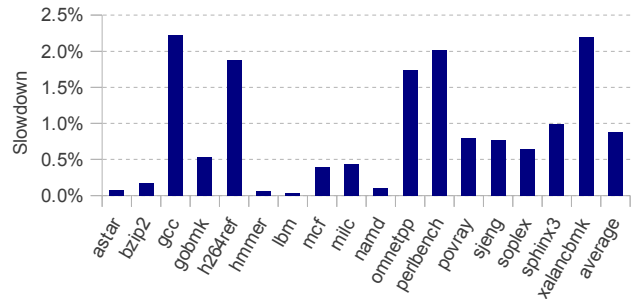


Fig. 10. Performance slowdown of SCRAP.

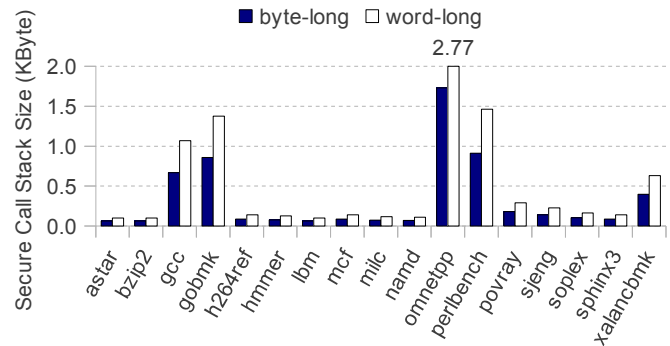


Fig. 11. Secure call stack size when using byte- and word-long SCRAP counters.

9.1 False Positives in Regular Codes

Next, we examine the impact of SCRAP on the execution of real programs to determine if SCRAP generates any false alarms during legal program execution. We used Pin tool [44] to instrument 18 C/C++ SPEC 2006 benchmarks and Apache 2.4.3 Web Server, Firefox 19.0.2 Web Browser, Adobe Flash Player 11.2 and Xpdf 3.03 PDF viewer, for one billion instructions. For instrumenting Apache, we used Apache benchmarking tool ab to remotely send thousands of requests to the web server which serves a static version of the Wikipedia entry¹ with a size of about 65KBs. Firefox benchmark is instrumented by accessing the same Wikipedia entry online and Xpdf is instrumented using a PDF version of the same page. For Flash Player, we used a standalone version of a Flash cartoon called “The Badger Song”².

The instrumentation results are presented in Figure 12, which shows the set of benchmarks that have at least one false positive for given values of N and S . As seen from these figures, for the thresholds with four consecutive gadgets and at most seven instructions in each gadget, none of our benchmarks generated false positives; i.e., a SCRAP detector $G_{7,4}$ generates no false positives for the above applications.

The selection of a SCRAP configuration is a tradeoff between security (the ability of detecting attacks), and false positives (flagging legitimate code as an attack).

1. <http://en.wikipedia.org/wiki/Scrap>

2. <http://weebly-stuff.com/songs/badgers/>

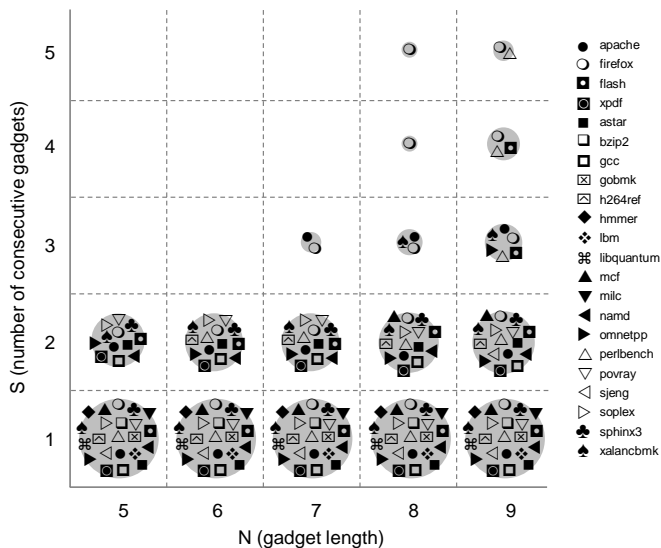


Fig. 12. List of benchmarks with non-zero false positive rates for $G_{N,S}$ for different values of N and S .

Since we only evaluated a subset of possible applications, it is impossible to claim that false negatives will never occur. Instead, our results demonstrate that for some SCRAP configurations that detect all known and even hypothetical attacks, the rate of false positives is likely to be very small (even if they exist at all), such that these false positives can be addressed individually. This, for example, can be achieved by making exceptions, or creating the whitelists. For example, one exception could be made for jump targets that are loaded from read-only memory. Our preliminary experiments on a Windows platform show that there is such necessity for Import Address Tables on Portable Executable format (please refer to the supplementary document for details).

9.2 Detecting JOP Attacks

With a SCRAP detector $G_{7,4}$, SCRAP is capable of detecting any JOP attack that does not use a gadget longer than 7 instructions (8 including the ending control flow instruction). Thusfar, every published attack, and every attack automation tool uses gadgets of size 5 or less [15], [24], [26]. As seen in Section 4, gadgets that call functions can be used in an attack because they preserve half of the registers due to assembly convention. However, SCRAP is capable of detecting attacks that implement these gadgets, while a JOP version of DROP would fail. As discussed in Section 3, in general long gadgets that do not use function calls have too many side effects to be used in an attack. Therefore, all currently published attacks would be detected by SCRAP. However, if an attacker is aware of the SCRAP protection, they may be able to find longer gadgets whose side effects can be tolerated or repaired by a subsequent gadget. Thus, we extend SCRAP in Section 10 to defend against such possible JOP attacks that manage to use an occasional long gadget in the middle of the attack to avoid detection.

To further assess SCRAP detection capabilities, we implemented 140 shell code attacks available from the Shell-Storm Linux shellcode repository [45]. These shellcodes ranged in complexity from simple single system calls, to attacks with multiple system calls, conditional branches, and loops. Even the most basic attack required at least 6 gadgets, which is greater than the minimum number of consecutive gadgets necessary to be detected by SCRAP. Gadgets longer than 6 instructions were extremely difficult to incorporate due to side effects. However, we were able to include a small number of gadgets of intermediate length, a few instructions longer. Attacks that use these longer gadgets are defeated by the improved detector presented in Section 10.

10 TOLERATING LONGER GADGETS

Thusfar, we have assumed that the length of the gadgets usable by attackers is limited to a hard threshold chosen in a way that makes false positives impossible. This assumption is based on the analysis in Section 3 where we showed that longer gadgets create too many state updates, making them difficult to use (e.g., Figure 6). However, it may be possible for attackers to identify some longer gadgets whose side effects do not completely destroy the attack state. Such gadgets can be used as a delay gadget to avoid detection by the basic SCRAP detector. In our own implementation of shellcodes, although it was difficult, we were able to identify a few such gadgets that are longer than the detection threshold and could be integrated into an attack successfully, avoiding detection by the basic SCRAP. These gadgets, for example, updated registers that were not needed for the attack, modified a non-critical memory location while being able to avoid illegal accesses, or had a side-effect that could be undone by another gadget. Thus, for practical signature based detection, it is imperative that we detect attacks even in the presence of some of these longer gadgets.

In the remainder of this section we propose a new multi-threshold detector that is able to detect CRAs quickly, while tolerating the use of longer gadgets. Intuitively, the detector assumes that attackers may be able to find some gadgets longer than the SCRAP threshold whose side-effects can be tolerated or undone by subsequent gadgets. These intermediate gadgets are not easy to find or use constructively in an attack since the number of side effects made by a gadget grows quickly with the length of the gadget. Side effects also increase the number of gadgets necessary for an attack; a repair gadget must be called in order to correct state changes and a dispatcher gadget must be called in order to reach the repair gadget.

The new detector detects attacks as a sequence of gadgets of length T_1 or shorter, while allowing the use of intermediate gadgets (IGs) of length T_2 or shorter such that $T_2 > T_1$. Since IGs typically do not advance the attack but are used only to avoid detection, we do not advance the gadget count (move closer to detection) like we do with short gadgets. At the same time we only reset to the initial state with gadgets of length greater

than T_2 . Now, for every other IG the gadget counter is reduced by one to take advantage of the additional gadgets necessary to repair side effects. To detect an attack, we still need S short gadgets ($\leq T_1$) before a very long gadget ($> T_2$). The state machine for the multi-threshold detector is shown in Figure 13. We call a detector of this type $G_{T_1, T_2, S}$ where S is the gadget count that is needed to detect an attack. Note that all three thresholds are software configurable in privilege mode.

The false positive rate is increased by this new multi-threshold detector. Previously, medium length gadgets reset SCRAP to its initial state, setting all counters to 0, making it more difficult to detect an attack (but making it possible for attackers to avoid detection). Figure 14 shows the benchmarks with false positives for all benchmarks we evaluated. The results show that T_1 can be set to 7, and T_2 can be set to a very high length of 25 without any false positives with gadget count, S , of 4. Gadgets of length 25 in the libraries we examined have a minimum of 5 side effects and an average of 14 side effects (Figure 6)—it is extremely improbable that they can be used without destroying the critical attack state.

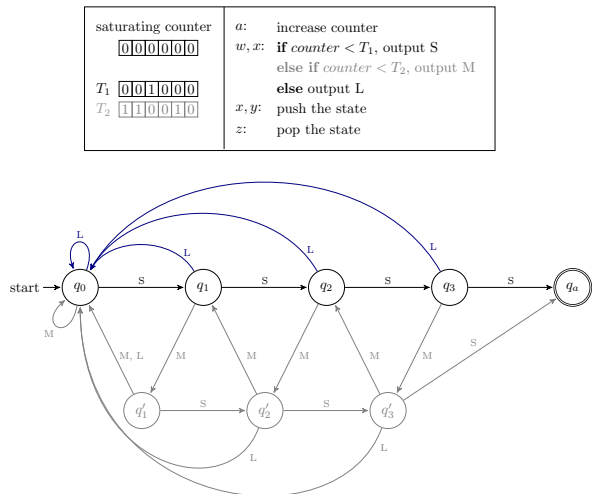


Fig. 13. State machine for the two-threshold detector.

As a further enhancement, a simple $G_{7,4}$ SCRAP module, as discussed in Section 9, could be used concurrently with this multi-threshold detector to catch attacks that use three short gadgets in a row. The overhead of this approach is linear in the number of detectors since a new state machine has to be implemented for each detector, and a space on the stack is needed to save each detector's state upon a function call.

11 FPGA IMPLEMENTATION

We implemented the proposed detectors in Verilog HDL on a Xilinx Spartan-3E XC3S100E FPGA with a 90nm process, using Xilinx ISE WebPACK 14.1. We evaluated both designs; baseline SCRAP presented in Section 7 and

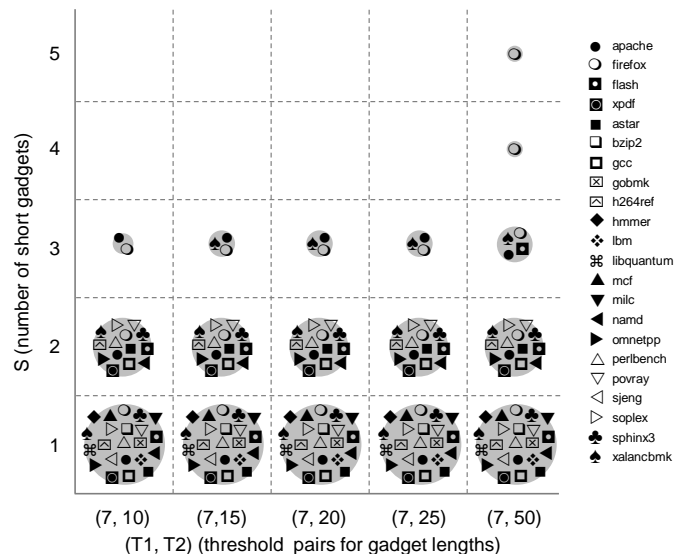


Fig. 14. List of benchmarks with non-zero false positive rates for two-threshold detector $G_{T_1, T_2, S}$ for different values of (T_1, T_2) and S .

two-threshold SCRAP presented in the previous section. In order to allow comparison, we also evaluated 8-, 16-, 32- and 64-bit counters on the same technology.

Figure 15 shows the critical path delays of both designs for varying widths. A baseline SCRAP design shown as (n, s) means it is able to detect $G_{N, S}$ attack language where N and S are encoded using n and s bits respectively. Similarly a two-threshold SCRAP shown as (t_1, t_2, s) , uses t_1 and t_2 bits for the two threshold values and s bits for gadget count.

Results for our unoptimized implementation show that the delay of SCRAP state machine is well under the cycle period of a superscalar processor. With a timing oriented design, it can be implemented with a shorter critical path.

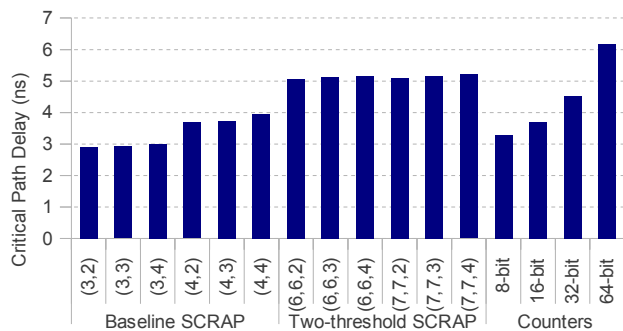


Fig. 15. Critical path delays for two SCRAP designs with different widths and various counters for reference.

We further evaluated the dynamic power dissipation of our FPGA designs, using Xilinx Power Estimator 11.1 for Spartan-3E FPGA Family. We set the ambient temperature to 65°C, toggle rate to 0.5 and clock rate to 256MHz. Results are shown in Figure 16. Again, same x-

axis labels are used as in Figure 15 and also counters of various widths are presented to allow comparison. Using the HDL Synthesis Report, we estimated the transistor count of the SCRAP logic. The largest baseline SCRAP design (4, 4) has as many transistors as a 32-bit up counter and the largest two-threshold design (7, 7, 4) has little less transistors than a 64-bit up counter.

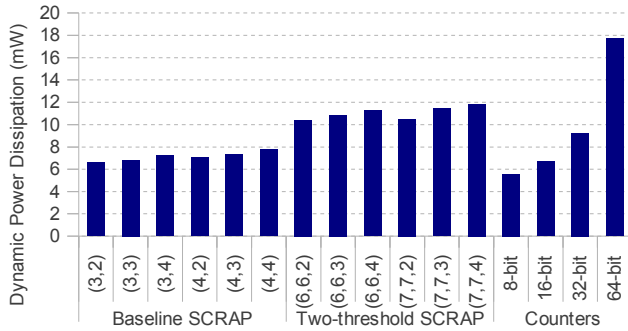


Fig. 16. Dynamic power dissipations for two SCRAP designs with different widths and various counters for reference.

12 RELATED WORK

In this section, we overview different approaches to protecting against CRA attacks. The related work is organized into three parts: (1) defenses against buffer overflow attacks; (2) comprehensive defenses; and (3) defenses specific to Code Reuse Attacks (CRAs).

12.1 Defenses against Buffer Overflows

Several approaches were developed to defeat buffer overflows which are necessary to initiate a CRA attack [9]–[11], [46]–[48]. Stackguard [9] and ProPolice [47] are GCC extensions that use canaries. StackShield separates return addresses into a separate stack at compile time making it impossible for stack buffer overflows to overwrite the return address [48]; similar works save a copy of the return address and validate it before a function return [10], [11]. StackGhost uses the register window feature of the Sun Sparc architecture to verify that return addresses have not been overwritten [49]. Recently, the advent of the NoExecute (NX) bit and its support by mainstream operating systems have made code injection attacks ineffective [13], [14].

12.2 Comprehensive Defenses

Memory bounds checking (MBC) annotate pointers with their legal address range and check every memory access against the base and bound of the associated data structure [3], [4], [50], [51]. However, the overhead of MBC is substantial. MBC cannot prevent all memory exploits: it cannot protect legacy binaries and externally linked or loaded components. Dynamic Information Flow Tracking (DIFT) taints the information coming from insecure sources, and dynamically tracks and propagates the taint

through processor registers and memory locations. The drawback is that DIFT is a heavy-weight approach that entails a significant redesign of the processor datapath and memory system if implemented in hardware [5], [6], [52], or incurs a substantial performance overhead if implemented in software [53], [54]. Data flow integrity [55] derives the data flow graph during compile-time and instrument the program to enforce conformance with the flow in the graph; note that this is a dual approach to control flow integrity.

12.3 CRA Attacks and Defenses

The first CRA attack proposed was the *return-into-libc (RILC) attack* [56], where the attacker subverts the control flow to call a function in the standard C library. Extensions to basic RILC have been proposed to allow a static chain of libc functions to be called [57] and recently to allow a general data-dependent form of chaining of libc functions [58].

Return-oriented Programming (ROP) attacks were recently proposed to execute arbitrary code [15], and the number of solutions to them were introduced [26]–[30]. We discussed those solutions in detail in earlier sections of this paper.

The newer defenses against ROP attacks also attempt to address JOPs. For example, Onarlioglu et al. first use binary rewriting to remove unintended branches and returns [59]. To protect intended branches, they use function-specific markers on each stack frame; they call these markers stack cookies. They also insert checks after every branch to check the stack cookie.

Kayaalp et al [41] propose branch regulation, a hardware supported techniques to protect against JOPs. Using binary rewriting, they insert markers at the beginning of every function, which include a magic number to mark a legal function entry, as well as the length of the function. Control flow integrity [36] is an approach to enforce legal control flow inside of programs; CFI would identify the illegal control flow necessary for code reuse attacks. Control Flow Locking [37] lazily enforces the same property and achieves smaller performance overhead.

Address space layout randomization (ASLR) [60] randomly offsets the program location in memory. ASLR and other optimized heap allocation models [61], [62] hide the correct address of the malicious code hiding the location of the gadgets. Unfortunately, exploits against ASLR are known; for example, a a format string attack can expose the stack location to an attacker allowing the random offset to be derived [63]. Schwartz et al show that even a small part of the code being unrandomized is sufficient to construct CRA attacks [24].

13 CONCLUDING REMARKS

In this paper, we presented SCRAP, a new hardware-based architecture for protecting against the emerging class of code reuse attacks (CRAs). We demonstrated that the latest incarnation of CRAs - jump oriented programming (JOP) attacks - have execution patterns that

are clearly distinguishable from the patterns exhibited by regular programs. However, we also demonstrated a new attack that renders previously proposed signature-based approaches ineffective by introducing delay gadgets. Delay gadgets are gadgets whose only purpose is to obfuscate the execution patterns of the attack without performing any useful computation. We developed a complete working JOP attack that incorporates delay gadgets. We then proposed and developed the SCRAP architecture for efficiently detecting such stealth JOP attacks. Our design started from the development of the formal language describing the stealth JOP attack signature and then subsequent demonstration of the hardware implementation. We also proposed a new microarchitecture optimization to simplify the SCRAP logic without encountering any performance loss.

In summary, SCRAP architecture protects unmodified legacy binaries, involves no changes to the software layers and incurs very small performance degradation: less than 2% on the average across the SPEC 2006 benchmarks. We also show that with appropriate selection of the detection thresholds, SCRAP successfully detects all JOP attacks used to implement many existing shellcodes, but at the same time results in no false alarms for the regular applications. It is therefore an effective and low-overhead protection, which in the very least increases the attack complexity dramatically, if not making it completely impossible. SCRAP can be configured in software to allow the thresholds to adapt to applications or to disable protection if it is not desired.

We extended the basic SCRAP detector to allow it to tolerate the presence of longer gadgets using a two threshold detector. The new detector can tolerate the presence of intermediate gadgets of length up to 50 instructions, without generating any false positives on the SPEC 2006 benchmark. We implemented both SCRAP and the two-threshold SCRAP in a hardware description language; the required hardware is small and fast. It also resides at the commit stage and does not affect any of the critical pipeline stages.

14 ACKNOWLEDGEMENTS

This material is based on research sponsored by Air Force Research Laboratory under agreement number FA8750-09-1-0137 and by National Science Foundation grants CNS-1018496 and CNS-0958501. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies and endorsements, either expressed or implied, of Air Force Research Laboratory, National Science Foundation, or the U.S. Government.

REFERENCES

- [1] W. Baer and A. Parkinson, "Cyberinsurance in IT Security Management," *IEEE Security and Privacy*, vol. 5, no. 3, pp. 50–56, may/june 2007.
- [2] "NIST national vulnerability database," 2012, available online at <http://nvd.nist.gov>.
- [3] D. Dhurjati and V. Adve, "Backwards-compatible array bounds checking for C with very low overhead," in *Proceedings of ICSE*. ACM, 2006, pp. 162–171. [Online]. Available: <http://doi.acm.org/10.1145/1134285.1134309>
- [4] J. Devietti, C. Blundell, M. M. K. Martin, and S. Zdancewic, "Hardbound: architectural support for spatial safety of the c programming language," in *Proceedings of the ASPLOS*. New York, NY, USA: ACM, 2008, pp. 103–114. [Online]. Available: <http://doi.acm.org/10.1145/1346281.1346295>
- [5] G. E. Suh, J. W. Lee, D. Zhang, and S. Devadas, "Secure program execution via dynamic information flow tracking," in *Proceedings of ASPLOS*. ACM, 2004, pp. 85–96. [Online]. Available: <http://doi.acm.org/10.1145/1024393.1024404>
- [6] M. Dalton, H. Kannan, and C. Kozyrakis, "Raksha: a flexible information flow architecture for software security," in *Proceedings of ISCA*. ACM, 2007, pp. 482–493. [Online]. Available: <http://doi.acm.org/10.1145/1250662.1250722>
- [7] Aleph One, "Smashing the stack for fun and profit," Nov. 1996.
- [8] J. Pincus and B. Baker, "Beyond stack smashing: Recent advances in exploiting buffer overruns," *IEEE Security and Privacy*, vol. 2, pp. 20–27, July 2004. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1018027.1018271>
- [9] C. Cowan, C. Pu, D. Maier, H. Hintony, J. Walpole, P. Bakke, S. Beattie, A. Grier, P. Wagle, and Q. Zhang, "StackGuard: Automatic adaptive detection and prevention of buffer-overflow attacks," in *Proceedings of USENIX Security*, vol. 7, 1998.
- [10] A. Baratloo, N. Singh, and T. Tsai, "Transparent run-time defense against stack smashing attacks," in *Proceedings of the USENIX Annual Technical Conf.*, 2000, pp. 251–262.
- [11] T. cker Chiueh and F.-H. Hsu, "RAD: A compile-time solution to buffer overflow attacks," in *ICDCS'01*, 2001.
- [12] M. Prasad and T. cker Chiueh, "A binary rewriting defense against stack based overflow attacks," in *Proceedings of the USENIX Annual Technical Conf.*, 2003, pp. 211–224.
- [13] P. Team, "Pax non-executable pages design & implementation," <http://pax.grsecurity.net/docs/noexec.txt>.
- [14] S. Andersen, "Part 3: Memory protection technologies," in *Changes to Functionality in Microsoft Windows XP Service Pack 2*. Microsoft Corp., 2004, <http://technet.microsoft.com/en-us/library/bb457155.aspx>.
- [15] H. Shacham, "The geometry of innocent flesh on the bone: Return-into-libc without function calls (on the x86)," in *Proceedings of CCS*. ACM Press, Oct. 2007, pp. 552–61.
- [16] E. Buchanan, R. Roemer, H. Shacham, and S. Savage, "When good instructions go bad: generalizing return-oriented programming to risc," in *Proceedings of CCS*. ACM, 2008, pp. 27–38. [Online]. Available: <http://doi.acm.org/10.1145/1455770.1455776>
- [17] A. Francillon and C. Castelluccia, "Code injection attacks on harvard-architecture devices," in *Proceedings of CCS*, 2008.
- [18] F. Lindner, "Developments in cisco ios forensics. confidence 2.0. presentation," 2009, http://www.recurity-labs.com/content/pub/FX_Router_Exploitation.pdf.
- [19] S. Checkoway, A. J. Feldman, B. Kantor, J. A. Halderman, E. W. Felten, and H. Shacham, "Can DREs provide long-lasting security? The case of return-oriented programming and the avc advantage," in *Proceedings of EVT/WOTE*. USENIX/ACCURATE/IAVoSS, aug 2009. [Online]. Available: <https://cs.ucsd.edu/~scheckow/papers/evt2009.html>
- [20] L. Davi, A. Dmitrienko, A.-R. Sadeghi, and M. Winandy, "Return-Oriented Programming without returns on ARM," System Security Lab - Ruhr University Bochum, Tech. Rep., 2010. [Online]. Available: <http://www.ei.rub.de/media/trust/veroeffentlichungen/2010/07/21/ROP-without>Returns-on-ARM.pdf>
- [21] R. G. Roemer, "Finding the bad in good code: Automated return-oriented programming exploit discovery," Master's thesis, University of California, San Diego, 2009. [Online]. Available: <https://cseweb.ucsd.edu/~rroemer/doc/thesis.pdf>
- [22] R. Hund, T. Holz, and F. C. Freiling, "Returnoriented rootkits: Bypassing kernel code integrity protection mechanisms," in *Proceedings of Usenix Security*, 2009.
- [23] T. Dullien and T. Kornau, "A framework for automated architecture-independent gadget search," 2010.
- [24] T. A. Edward J. Schwartz and D. Brumle, "Q: Exploit hardening made easy," in *Proceedings of USENIX Security*, 2011.
- [25] L. Davi, A.-R. Sadeghi, and M. Winandy, "Dynamic integrity measurement and attestation: towards defense against return-oriented programming attacks," in *Proceedings of ACM STC*. ACM, 2009, pp. 49–54. [Online]. Available: <http://doi.acm.org/10.1145/1655108.1655117>
- [26] P. Chen, H. Xiao, X. Shen, X. Yin, B. Mao, and L. Xie, "Drop: Detecting return-oriented programming malicious code," in *Proceedings of ICISS*. Springer-Verlag, 2009, pp. 163–177. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-10772-6_13

- [27] J. Li, Z. Wang, X. Jiang, M. Grace, and S. Bahram, "Defeating return-oriented rootkits with "return-less" kernels," in *Proceedings of EuroSys*. New York, NY, USA: ACM, 2010, pp. 195–208. [Online]. Available: <http://doi.acm.org/10.1145/1755913.1755934>
- [28] J. Xu, Z. Kalbarczyk, S. Patel, and R. K. Iyer, "Architecture support for defending against buffer overflow attacks," in *Proceedings of Workshop on Evaluating and Architecting Systems for Dependability*, 2002.
- [29] J. McGregor, D. Karig, Z. Shi, and R. Lee, "A processor architecture defense against buffer overflow attacks," in *Proceedings of ITRE*, aug. 2003, pp. 243 – 250.
- [30] S. Sinnadurai, Q. Zhao, and W. fai Wong, "Transparent runtime shadow stack: Protection against malicious return address modifications," 2008.
- [31] T. Bletsch, X. Jiang, V. W. Freeh, and Z. Liang, "Jump-oriented programming: a new class of code-reuse attack," in *Proceedings of ASIACCS*. ACM, 2011, pp. 30–40. [Online]. Available: <http://doi.acm.org/10.1145/1966913.1966919>
- [32] S. Checkoway, L. Davi, A. Dmitrienko, A.-R. Sadeghi, H. Shacham, and M. Winandy, "Return-oriented programming without returns," in *Proceedings of CCS*. ACM Press, oct 2010, pp. 559–72.
- [33] P. Chen, X. Xing, B. Mao, L. Xie, X. Shen, and X. Yin, "Automatic construction of jump-oriented programming shellcode (on the x86)," in *Proceedings of ASIACCS*. ACM, 2011, pp. 20–29. [Online]. Available: <http://doi.acm.org/10.1145/1966913.1966918>
- [34] Apple Inc., "iOS security," 2012, retrieved October 2013 from http://www.apple.com/ipad/business/docs/iOS_Security_Oct12.pdf.
- [35] SOGETI ESEC R&D Lab, "Analysis of the jailbreakme v3 font exploit," 2012, retrieved September 2012 from <http://esec-lab.sogeti.com/post/Analysis-of-the-jailbreakme-v3-font-exploit>.
- [36] M. Abadi, M. Budi, U. Erlingsson, and J. Ligatti, "Control-flow integrity," in *Proceedings of CCS*. ACM, 2005, pp. 340–353.
- [37] T. Bletsch, X. Jiang, and V. Freeh, "Mitigating code-reuse attacks with control-flow locking," in *Proceedings of the 27th Annual Computer Security Applications Conference*, ser. ACSAC '11. New York, NY, USA: ACM, 2011, pp. 353–362. [Online]. Available: <http://doi.acm.org/10.1145/2076732.2076783>
- [38] B. Yee, D. Sehr, G. Dardyk, J. Chen, R. Muth, T. Ormandy, S. Okasaka, N. Narula, and N. Fullagar, "Native client: A sandbox for portable, untrusted x86 native code," in *Security and Privacy, 2009 30th IEEE Symposium on*, 2009, pp. 79–93.
- [39] M. Kayaalp, T. Schmitt, J. Nomani, D. Ponomarev, and N. Abu-Ghazaleh, "SCRAP: Architecture for signature-based protection from code reuse attacks," in *Proceedings of HPCA*, 2013.
- [40] The Open Group, "IEEE Std 1003.1, 2004," <http://pubs.opengroup.org/onlinepubs/009695399/functions/setjmp.html>.
- [41] M. Kayaalp, M. Ozsoy, N. Abu-Ghazaleh, and D. Ponomarev, "Branch regulation: Low overhead mitigation of code reuse attacks," in *Proceedings of ISCA*, 2012.
- [42] M. T. Yourst, "PtlSim: A cycle accurate full system x86-64 microarchitectural simulator," in *Proceedings of ISPASS*, 2007, pp. 23–34.
- [43] C. D. Spradling, "Spec cpu2006 benchmark tools," *SIGARCH Comput. Archit. News*, vol. 35, no. 1, pp. 130–134, 2007.
- [44] C. Luk, R. Cohn, R. Muth, H. Patil, A. Klauser, G. Lowney, S. Wallace, V. Reddi, and K. Hazelwood, "Pin: building customized program analysis tools with dynamic instrumentation," in *Proceedings of the 2005 ACM SIGPLAN conference on Programming language design and implementation*, ser. PLDI '05, 2005, pp. 190–200.
- [45] "The shell storm linux shellcode repository," 2012, accessed Sept. 2012 at <http://www.shell-storm.org/shellcode/shellcode-linux.php>.
- [46] C. Cowan, S. Beattie, J. Johansen, and P. Wagle, "Pointguardtm: protecting pointers from buffer overflow vulnerabilities," in *Proceedings of USENIX Security*. USENIX Association, 2003, pp. 7–7. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1251353.1251360>
- [47] H. Etoh and K. Yoda, "Propolice: Improved stack-smashing attack detection," *IPSI SIG notes on computer security*, Oct 2001.
- [48] Vindicator, "Stack shield technical info file v0.7," January 2001, <http://www.angelfire.com/sk/stackshield/>.
- [49] M. Frantzen and M. Shuey, "StackGhost: Hardware facilitated stack protection," in *Proceedings of USENIX Security*. USENIX Association, 2001, pp. 5–5. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1251327.1251332>
- [50] S. Ghose, L. Gilgeous, P. Dudnik, A. Aggarwal, and C. Waxman, "Architectural support for low overhead detection of memory violations," in *Proceedings of DATE*, 2009.
- [51] S. Nagarakatte, M. Martin, and S. Zdancewic, "Watchdog: Hardware for safe and secure manual memory management and full memory safety," in *Proceedings of ISCA*, 2012.
- [52] M. Ozsoy, D. Ponomarev, N. Abu-Ghazaleh, and T. Suri, "SIFT: A low-overhead dynamic information flow tracking architecture for smt processors," in *Proceedings of CF*, May 2011.
- [53] J. Newsome and D. Song, "Dynamic taint analysis for automatic detection, analysis, and signature generation of exploits on commodity software," in *Proceedings of NDSS*, feb 2005.
- [54] F. Qin, C. Wang, Z. Li, H.-s. Kim, Y. Zhou, and Y. Wu, "Lift: A low-overhead practical information flow tracking system for detecting security attacks," in *Proceedings of MICRO*. IEEE Computer Society, 2006, pp. 135–148. [Online]. Available: <http://dx.doi.org/10.1109/MICRO.2006.29>
- [55] M. Castro, M. Costa, and T. Harris, "Securing software by enforcing data-flow integrity," in *Proceedings of OSDI*. USENIX Association, 2006, pp. 147–160. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1298455.1298470>
- [56] S. Designer, "'return-to-libc' attack," 1997, bugtraq.
- [57] Negral, "The advanced return-into-lib(c) attacks," 2001, <http://www.phrack.org/issues.html?issue=58&id=4>. Retrieved June 2012.
- [58] M. Tran, M. Etheridge, T. Bletsch, X. Jiang, V. Freeh, and P. Ning, "On the expressiveness of return-into-libc attacks," in *Proceedings of RAID*, Sep. 2011.
- [59] K. Onarlioglu, L. Bilge, A. Lanzi, D. Balzarotti, and E. Kirda, "Gfree: Defeating return-oriented programming through gadget-less binaries," in *Proceedings of ACSAC*, 2010, pp. 49–58.
- [60] P. Team, "Pax address space layout randomization (aslr)," <http://pax.grsecurity.net/docs/aslr.txt>.
- [61] E. D. Berger and B. G. Zorn, "Diehard: probabilistic memory safety for unsafe languages," in *Proceedings of PLDI*. ACM, 2006, pp. 158–168. [Online]. Available: <http://doi.acm.org/10.1145/1133981.1134000>
- [62] O. Whitehouse, "An analysis of address space layout randomization on windows vista," 2007.
- [63] T. Newsham, "Format string attacks," September 2000, <http://julianor.tripod.com/bc/tn-usfs.pdf>.

Mehmet Kayaalp is a PhD student in the Department of Computer Science at SUNY Binghamton. His research interests are in the areas of computer architecture and secure system design.

Timothy Schmitt is an undergraduate student in the Department of Computer Science at SUNY Binghamton. His research interests are in the areas of security and computer architecture.

Junaid Nomani is a PhD student in the Computer Science Department at Yale University. His research interests are in the areas of security and computer architecture. This work was performed while Junaid was an undergraduate student at SUNY Binghamton.

Nael Abu-Ghazaleh is an Associate Professor in the Department of Computer Science at SUNY Binghamton. His research interests are in the areas of secure system design, parallel discrete event simulation, networking and mobile computing. He received his PhD from the University of Cincinnati in 1997.

Dmitry Ponomarev is an Associate Professor in the Department of Computer Science at SUNY Binghamton. His research interests are in the areas of computer architecture, secure and power-aware systems and parallel discrete event simulation. He received his PhD from SUNY Binghamton in 2003.