

ROPecker: A Generic and Practical Approach For Defending Against ROP Attacks

Yueqiang Cheng[†], Zongwei Zhou[§], Miao Yu[§], Xuhua Ding[†], Robert H. Deng[†]

[†] School of Information Systems, Singapore Management University

[†] {yqcheng.2008, xhding, robertdeng}@smu.edu.sg

[§] Cylab, Carnegie Mellon University

[§] {stephenzhou, superymk}@cmu.edu

Abstract

Return-Oriented Programming (ROP) is a sophisticated exploitation technique that is able to drive target applications to perform arbitrary unintended operations by constructing a gadget chain reusing existing small code sequences (gadgets). Existing defense mechanisms either only handle specific types of gadgets, require access to source code and/or a customized compiler, break the integrity of the binary code, or suffer from high performance overhead.

In this paper, we present a novel system, ROPecker, to efficiently and effectively defend against ROP attacks without relying on any other side information (e.g., source code and compiler support) or binary rewriting. ROPecker detects an ROP attack at run-time by checking the presence of a sufficiently long chain of gadgets in past and future execution flow, with the assistance of the taken branches recorded in the Last Branch Record (LBR) registers and an efficient technique combining offline analysis with run-time emulation. We also design a sliding window mechanism to invoke the detection logic in proper timings, which achieves both high detection accuracy and efficiency.

We build an ROPecker prototype on x86-based Linux computers and evaluate its security effectiveness and performance overhead. In our experiment, ROPecker can detect all ROP attacks from real-world examples and generated by the general-purpose ROP compiler *Q*. It only incurs acceptable performance overhead on CPU computation, disk I/O and network I/O.

I. INTRODUCTION

Return-Oriented Programming (ROP) is a code-reuse security exploitation technique introduced by Shacham et al. [26], [7], [8]. By chaining together existing small instruction sequences (gadgets) from target programs, ROP empowers a remote adversary to perform Turing-complete computation without injecting any malicious code. Due to its great threat, recent years have witnessed many proposed methods (Table I) to defend against ROP attacks [18], [13], [9], [15], [6], [20], [28], [34], [32], [14], [33].

The approaches, such as ROPDefender [13], DROP [9], Return-less kernel [18] and ROPGuard [15], only focus on the ROP gadgets ended with return instruction, allowing the

adversary to use other gadgets (e.g., jmp-based). In addition, the first two schemes [13], [9] also incur high overhead and degrade the run-time performance of the protected application.

Defense mechanisms, such as CFLocking [6] and G-Free [20], aim to defend against all types of ROP attacks, but they require the knowledge of side information (e.g., source code and/or customized compiler tool chain). In fact, this side information is often unavailable to the end users in the real world.

Recent proposals, such as Binary stirring [34], ILR [14], IPR [32], CCFIR [36] and KBouncer [33], achieves good attack-type coverage and run-time efficiency, and requires no side information. However, they all leverage binary rewriting technique to instrument the code, in the purposes of randomly shuffling the instructions, forcing control flow integrity, or monitoring abnormal control transfers. Binary instrumentation breaks the integrity of the binary code, which raises compatibility issues against security mechanisms, such as Window 7 system library protection, Integrity Measurement Architecture (IMA) [23], and remote attestation. Note that KBouncer [33] only monitors the application execution flow on selected critical paths, e.g., system APIs. It inevitably misses the ROP attacks that do not use those paths.

This paper presents the first generic and practical ROP countermeasure, called *ROPecker*, that effectively and efficiently defends against all types of ROP attacks without requiring source code access, customized compiler support and binary rewriting, as summarized in Table I. We observe that the distinguishing feature of an ROP attack is that *the execution flow of a victim application consists of a sufficiently long sequence of gadgets chained together by branching instructions*. Thus, ROPecker analyzes all gadgets located in the target application binary and shared libraries in offline pre-processing. During run-time check points, ROPecker identifies the gadget chain in the *past* execution flow, using the history of taken branches recorded in the Last Branch Record (LBR) registers, and also inspects the *future* execution flow to detect any ROP gadget chain, leveraging the information from offline analysis and occasional instruction emulation.

We also propose a novel *sliding window* mechanism to

	ROP Types	No Source Code	No Binary Rewriting	Run-time Efficiency
DROP [9]	Ret-based	✓	X	X
ROPDefender [13]	Ret-based	✓	X	X
ROPGuard [15]	Ret-based	✓	X	✓
Return-less Kernel [18]	Ret-based	X	✓	✓
CFLocking [6]	ALL	X	✓	✓
G-Free [20]	ALL	X	✓	✓
ILR [14]	ALL	✓	X	✓
Binary Stirring [34]	ALL	✓	X	✓
IPR [32]	ALL	✓	X	✓
CCFIR [36]	ALL	✓	X	✓
KBouncer [33]	ALL	✓	X	✓
ROPecker	ALL	✓	✓	✓

Table I
THE COMPARISON OF SEVERAL TYPICAL ROP DEFENDING APPROACHES.

decide run-time check points, which sets the most recent visited code regions of the protected application as executable, and leaves the application code outside of the window as non-executable. Any attempt to execute the code beyond the window boundary automatically triggers the ROP checking logic, while the application execution within the window remains unaffected. The sliding window design takes advantage of the temporal and spatial locality of application code and the sparsely distribution of meaningful gadgets across the application code base, to achieve both run-time efficiency and detection accuracy.

We have built an ROPecker prototype on x86-based Linux platform, though our design can be extended to other commodity operating systems. We have experimented ROPecker with real world ROP attacks and those generated by the ROP compiler Q [25]. The results demonstrate that ROPecker successfully detects all of them. We also evaluate the performance by running several macro-benchmark (e.g., SPEC INT2006, bonnie++, and Apache server httpd) and micro-benchmark tools. The results show that ROPecker introduces reasonable performance overhead on CPU computation, disk I/O, and network I/O. Specifically, ROPecker incurs only 2.60% overhead on average on CPU computation, 1.56% overhead on disk I/O, and 0.08% overhead on typical (4KB) HTTP communications.

Contributions. In specific, we make the following contributions:

- Design the first generic and practical ROP countermeasure to protect legacy applications from all types of ROP attacks without side information and binary rewriting.
- Propose novel techniques combining sliding window, offline gadget analysis with run-time instruction emulation to achieve high efficiency without compromising detection accuracy.
- Implement an ROPecker prototype on x86-based Linux platform and evaluate its security effectiveness and run-

time performance.

Organization. The rest of the paper is structured as follows. In Section II and Section III, we briefly describe the background knowledge, and highlight our system goals, threat model and assumptions. We introduce in detail the design rationale, system architecture and implementation of ROPecker in Sections IV, V, and VI, respectively. Section VII presents the security and performance evaluation of ROPecker. In Sections VIII and Section IX, we discuss several subtle attacks and compare our system with the existing work. At last, we conclude this paper and discuss the future work in Section X.

II. BACKGROUND

A. Return-Oriented Programming

The main idea of ROP attack is to reuse instructions from existing code space to perform malicious operations. There are two major steps to launch an ROP attack: (1) to identify a set of useful instruction sequences, called *gadgets*, from the entire code segment, e.g., the application code and the shared libraries; (2) to link the selected gadgets into a gadget chain through a crafted payload. Note that the gadgets are not limited to using aligned instructions, e.g., on x86 platform, a sequence of unaligned instructions may also be converted to a valid gadget.

A typical gadget has a code section for computation operations (e.g., assigning a value to a general register), and a link section manipulating the control flow to link gadgets. The control flow manipulation is achieved through the *indirect branch* instructions such as *ret* and *indirect jmp/call* instructions¹. According to the difference of the link section, ROP attacks are classified into ret-based ROP and jmp/call-based ROP or JOP. In a real-life ROP attack,

¹Traditionally, the direct branch instructions can not be used as the link section, since their destinations are fixed. For completeness, in Section VIII-B, we discuss extensions of our ROP defense mechanism against gadgets with direct branch instructions.

the adversary may mix both types of gadgets. The gadgets used in ROP attacks typically have the following features.

Small Size. A gadget’s code section is usually small, e.g., consisting of 2 to 5 instructions [31], which leads to the lack of the functionality of a single gadget. Though the gadgets with large code sections can perform more operations, they inevitably lead to more side effects and some of them may conflict with each other, e.g., a gadget accidentally changes the stack pointer, which may lead to the failure of the execution of the next gadget. In fact, the adversary usually prefers to collect the gadgets only with the intended operations, instead of using long gadgets². As illustrated in [8], a jump-oriented Turing-complete JOP needs up to 34 gadgets.

Sparsely Distributed. Although the gadgets are distributed across the entire code space, the ones meeting for the adversary’s needs are not guaranteed to exist due to the sparse density. To have higher success probability, the adversary usually needs a large code base to collect enough gadgets to perform the malicious operations. **The experiment results of Q [25] imply that the adversary has low possibility to launch a meaningful ROP attack, if we can limit the scope of the executable code within 20KB at any time. If we can further reduce the scope, the possibility will consequently go down.**

B. Last Branch Record

In our solution, we leverage Last Branch Record (LBR) registers to provide reliable information about the execution trace of the protected application. LBR are dedicated CPU registers widely available on modern Intel and AMD processors. LBR provides a looped buffer to store the sources and destinations of the most recently executed branch instructions. The length of the buffer is limited, e.g., the Intel i5 only has 16 register pairs for recording the branches. It leads to that the new records inevitably override the old ones when the LBR buffer gets full.

The LBR functionality is disabled by default, and can only be enabled/disabled through certain Model Specific Registers (MSRs). Without the kernel privilege (i.e., ring-0), the user space applications cannot modify the value in the LBR MSRs. The LBR can be configured to only record the branches taken in user space. However, when recording user-space branches, the LBR does not distinguish branches in different processes. Thus, we have to filter out the unavoidable noise records to get the ones relevant to a specific process. Note that the branch recording is performed by the hardware processor and it introduces almost *zero* overhead for application executions.

²Our mechanism can be tuned to detect gadgets of different length, as discussed in Section ??

III. PROBLEM DEFINITION

A. System Goals

Our goal is to design a security system to *detect* and *prevent* ROP attacks at run-time with the following features.

G1: Generic. We aim to protect binary applications against all types of user space ROP attacks. The ROP gadget chain can be constructed by ret-based gadgets, jmp-based gadgets, or both of them.

G2: Transparent. Our system should transparently work for the legacy binary applications. In addition, 1) it does *not* rely on source code or customized compiler tools; 2) it does *not* instrument the binary code.

G3: Efficient. We aim to minimize the performance overhead. Our defense mechanisms should not inflict the protected applications and the whole system with high performance degradation.

B. Threat Model

In this paper, we focus on defending against application-level ROP attacks. We consider a remote adversary attacking a target application by manipulating inputs in order to launch an ROP attack. We suppose that the adversary knows all implementation details of the target application and can send arbitrary inputs. Nonetheless, it cannot subvert the target platform’s hardware and the operating system (OS) to their favor.

C. Assumptions

We assume that both the processor and the operating system enable the Data Execution Prevention (DEP) mechanism. In fact, the DEP mechanism is supported by default in modern operating systems. We do not assume that the Address Space Layout Randomization (ASLR) mechanism is enabled. In essence, our attack detection mechanism does not rely on ASLR. The application is not created with malicious purpose, e.g., it is not maliciously compiled to contain abundant gadgets within a small code base. Furthermore, we do not assume that the application is released and distributed with side information (e.g., source code and debugging information). We do not attempt to protect any self-modifying applications, because they require writable code segment, conflicting with the DEP mechanism.

IV. DESIGN RATIONALE

Essentially, our approach is to capture the application execution features at proper moments and then to identify the existence of the ROP payload. The detection accuracy and efficiency are affected by two critical design issues. The first issue is the types of run-time features to capture. Ideally, the features should be *reliable* in the sense that the adversary can not modify them to evade detection. In addition, they should be *sound* in the sense that they provide a *solid* evidence to infer the presence of ROP attacks. The second is the

timing of detection. A poorly designed timing may either miss critical information which leads to a low detection success rate or introduces unnecessary checks with a high performance loss.

Hallmark of ROP Attack. Most existing ROP countermeasures are based on catching run-time abnormalities of the ROP-infected application, e.g., call-ret violation [9], [13] and deviation from the control flow graph (CFG) or call graph (CG) [4], [6]. *None* of these approaches can achieve all our goals listed in Section III. For instance, the method of catching any broken call-ret pairing is not applicable to jump based ROP attacks. It is extremely difficult to extract control flow information *completely* and *accurately*, without source code access. In Section IX, we elaborate more on the disadvantages of using those abnormalities.

In this paper, we observe that the hallmark of ROP attacks is that *the execution flow of a victim application consists of a long sequence of gadgets chained together by branching instructions*. Because all gadgets can be located from the constant binary code [26], [7], [8], the history of branches in an execution flow, enhanced by prediction of future branches, can be the *solid* evidence to decide whether the execution comprises a gadget chain.

Based on this observation, we propose a novel ROP defense mechanism, which relies on a payload detection algorithm: 1) identifying past executed gadget chain using LBR information, and 2) searching potential gadget chain in future execution flow, assisted with information from offline gadget analysis, and/or occasional run-time instruction emulation (Section V-C). If the total number of gadgets chained together in both past and future detection exceeds a chosen threshold, ROPEcker reports the ROP attack and terminates the application. Note that the adversary can not tamper with the LBR values due to the hardware and the OS protection, i.e., the accessing of LBR registers needs kernel privilege. Furthermore, when the detection algorithm is executing, the monitored application is suspended and cannot change its future execution flow. Thus, all information collected is *reliable* for the detection.

An offline pre-processor is introduced to analyze the instruction and gadget information. The results cover the most frequently used gadget cases and leaves only minor cases to run-time analysis/emulation, which greatly reduces the ROPEcker run-time overhead without sacrificing detection accuracy (Sections V-A and VI-A).

Timing of Checking. To meet the goal of G2, we refrain from inserting code in critical execution paths of the protected application [15], [33] to trigger our ROP checking. In addition, it is obviously inefficient to monitor every branch instruction of the application. Periodical sampling may incur less cost, but it is prone to miss the ROP attacks taking place within the sampling interval.

In this paper, we propose a *sliding window* technique

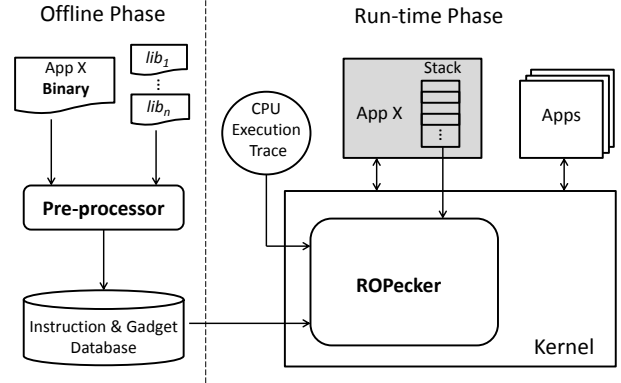


Figure 1. **The architecture of ROPEcker.** The shaded square represents a protected application.

which can catch an ROP attack in a timely fashion without a heavy performance toll (Section V-B). Shifting along with the execution flow, a sliding window refers to the portion of the application code region which are set executable by ROPEcker, while the instructions outside of the window are non-executable. Therefore, attempting to execute an instruction out of the window triggers an exception where the ROP checking logic is invoked. This design is advantageous for the following two reasons. Firstly, the window size is small comparing with the large code base that is needed in launching a meaningful ROP attack. As there are not enough gadgets for an ROP attack in the sliding window, an attack always triggers our ROP checking. Secondly, it is highly efficient since there is no intervention for executions within the sliding window. In addition, the window covers the most recent, non-contiguous code pages visited by the monitored application, meaning that we can take full advantage of the temporal and spatial locality feature of the application execution [11]. This guarantees that the ROP checking logic is not frequently invoked.

A prerequisite of using the sliding window technique is that the adversary cannot bypass the checking or disable it. Despite of lacking the kernel privilege, the adversary may mislead the kernel through certain special system calls, e.g., *mprotect* and *mmap2*. The adversary may leverage ROP to invoke those system calls with malicious inputs, which leads the kernel to set the code pages of the monitored application that contains gadgets to be executable, evading the ROPEcker’s detection, or even to disable the *DEP* mechanism and allow code injection. Thus, ROPEcker have to intercept those system calls and check the existence of the ROP payload before passing them to the kernel.

V. SYSTEM OVERVIEW

In this section, we elaborate on the design of the ROPEcker system consisting of a pre-processor, a kernel module and the Instruction & Gadget (IG) database as depicted in Figure 1. The general work-flow of ROPEcker proceeds in

two phases: an offline pre-processing phase followed by a run-time detection phase. In the offline pre-processing phase, the ROPEcker pre-processor extracts the instruction and gadget information from the application binary and system libraries it depends on, and saves the results in the IG database. During the run-time detection phase, the ROPEcker kernel module is invoked due to the attempted execution outside of the sliding window or the invocation of risky system calls. The kernel module suspends the protected application (i.e., the shaded box in Figure 1); acquires information from the CPU execution trace, the stack and code segments of the protected application; and analyzes the information to detect ROP attacks by querying the IG database and invoking instruction emulation when necessary. Note that ROPEcker can protect multiple applications in parallel.

A. Offline Pre-processing

In this phase, the ROPEcker pre-processor extracts all instruction information (e.g., offsets, types, and alignment) of the target application as well as the dependent shared libraries. The pre-processor also identifies the potential gadgets from the application binary and libraries. **In order to eliminate the dependency on a perfect disassembler, we start from the first byte of the code segment to identify gadget, and then redo this step from the next byte. If an indirect branch is found within a predefined length (e.g., 6 instructions), we will treat this short code sequence as a gadget. Thus, we only require the disassembler correctly analyze a short sequence of code, instead of the whole application code all in once. For that, we are able to identify all possible gadgets, while using existing disassemblers, such as diStorm [10].** The pre-processor analyzes gadget types and their impacts on the stack (e.g., pop and push instructions can change stack pointers). The pre-processor constructs the IG database to store all the above instruction and gadget information. Note that the pre-processing of shared libraries (e.g., libc) is a one-time effort, as their results can be re-used among different protected applications. We will introduce the implementation of the pre-processor and the detailed format of the IG database in Section VI-A.

B. ROP Detection Triggers

Before describing the run-time detection phase, we first introduce two types of events that trigger the detection logic: execution out of the sliding window and invocation of the risky system calls.

1) *Sliding Window*: As explained in previous sections, application code pages outside of the sliding window are **all** set as non-executable by ROPEcker kernel module. Consequently, any attempt to execute an instruction out of the window triggers a page fault exception, and the exception handler invokes the ROPEcker checking logic. If no ROP attack is detected, ROPEcker slides the window by replacing

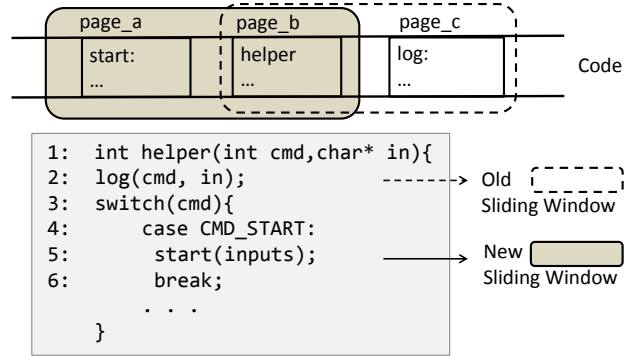


Figure 2. **The sliding window and its update.** In this example, the size of the sliding window is 2 pages. When the execution flow reaches to *log* function (line 2), the sliding window is (*page_b*, *page_c*). Later the window is updated to (*page_a*, *page_b*) when the *start* function (line 5) is executed. Note that the code pages in the sliding window can be non-contiguous.

the oldest page with the newly accessed code page. In this way, the window always covers the most recent *active* pages. For instance, the sliding window in Figure 2 consists of *page_b* and *page_c* when the execution flow reaches to the *log* function (line 2). When the *start* function (line 5) is executed, *page_a* is added to the sliding window and *page_c* is deleted.

Due to the temporal and spatial locality of application-s [11], the usage of the sliding window significantly reduces the number of times of the ROP checking, and thereby reduces the performance overhead. The window size is a critical parameter to keep the balance between the detection accuracy and performance. Increasing the window size can decrease the frequency of the invocation of the ROP checking and thus reduce the system performance overhead, but at the same time, it increases the possibility for the adversary to launch an ROP attack only using the gadgets within certain sliding window. **According to the experiment results of Q [25], we recommend using 16KB or 8KB as the window size. Note that it is safe to use a relatively larger window under certain cases. For instance, the target application has been processed by a gadget-elimination tool, such as G-Free [20]. In that case, the gadgets in the target application are expected to be extremely sparse.**

2) *Risky System Calls*: Risky system calls such as *mprotect* and *mmap2* allow the ROP adversary to disable the *DEP* mechanism so that no exception is raised when executing outside of the sliding window. To intercept those security system calls, we modify the system call table such that the kernel invokes the ROP checking logic prior to serving the system call request. If an ROP payload is detected in the application context, ROPEcker rejects the request and terminates the application execution; otherwise it transfers the execution flow back to the requested system call handler.

C. Run-time ROP Detection

The ROPEcker kernel module implements the ROP checking logic which detects the existence of an ROP gadget chain. Figure 3 depicts the flow chart of the detection logic. It begins with condition filtering, which filters out the irrelevant events that are not triggered by the monitored application. The second step is the past payload detection, which identifies the gadget chain from the branching history saved in LBRs. The last step is the future payload detection to identify a potential gadget chain that will be executed in the near future. If the total number of the chained gadgets exceeds a priorly chosen threshold, ROPEcker concludes that an ROP attack is caught and terminates the application. Otherwise it updates the sliding window and resumes the original execution flow.

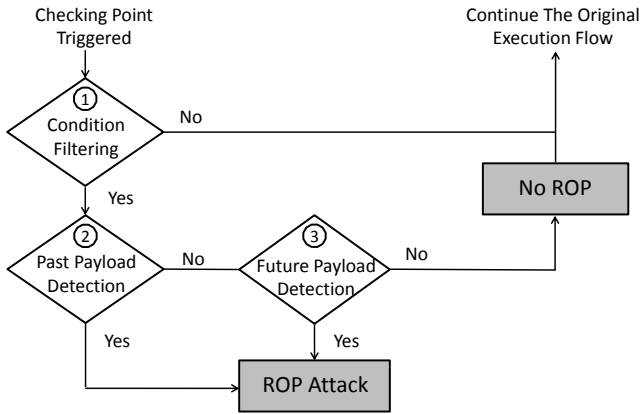


Figure 3. The ROP checking logic.

1) *Condition Filtering*: In this step, the ROP module checks whether the triggering event is due to the target application. For page fault exceptions, ROPEcker can distinguish the exceptions triggered by the sliding window from others, because the page fault error codes relevant to the sliding window indicates executing instructions in a non-executable code page, whereas the normal execution of applications never trigger page faults with this kind of error code. For system call invocation, the kernel module compares the process ID of the invoker against the IDs of the target applications to select relevant system calls.

2) *Past Payload Detection*: As introduced earlier, the recent branches are recorded in the LBR registers. Each LBR entry records the source and destination address of an executed branch instruction. To identify a gadget chain, ROPEcker verifies the following two requirement on each LBR entry, starting from the most recent one: 1) the instruction at the source address is an indirect branch instruction, such as *ret* or *jmp %eax*, and 2) the destination address points to a gadget. Once a branch fail to satisfy *either* requirement, ROPEcker stops LBR checking and outputs the length of the identified gadget chain. If the identified gadget

chain length exceeds the threshold, ROPEcker reports the ROP attack and terminates the application. Otherwise, it moves to the future payload detection with the length of identified chain as input parameters. Note that the instruction and gadget information have been extracted and saved in the table in the offline processing phase. Thus, ROPEcker can simply query the table to get the answer for 1) and 2), which saves the checking cost and thereby reduces the performance overhead.

Recall that the LBR value is not process specific. Thus, we face the challenge of sifting out those records belonging to the target process, since those noise records will contribute to the false negative rate of ROPEcker. Our solution is to traverse the LBR queue backwards starting from the most recent one, until reaching one entry that indicates a context switch. The context switch entry can be recognized by its unique feature. Namely, its source address is in kernel space and its destination address is in user space. Other types of branches must have the source and destination addresses in the same space. The ROPEcker module does not use LBR entries older than the context switch one since it is not certain they belong to the target application. Interestingly, if the context switch entry with its source address in the ROPEcker module, it is exactly the mark left by the previous detection. In this case, the following LBR entries are still in our monitored application.

Another challenge is from the limitations of the LBR mechanism. In LBR, the new records will override the old ones when the LBR buffer gets full, and the current implementation of LBR does not provide any mechanism to intercept the override event or to backup the previous ones. Thus, ROPEcker can not get the whole historical taken branches. In addition, the noise records may occupy several entries in the limited LBR registers. Thus, in certain cases, the identified gadget chain is not long enough to confirm an ROP attack, necessitating additional information collecting from the future payload detection.

3) *Future Payload Detection*: This step predicts any incoming gadget along the execution flow. According to the gadget types, i.e., *ret*-based and *jmp*-based gadgets, we apply two methods to handle them respectively.

It is relatively easy to handle *ret*-based gadgets, because their destinations are stored on the application stack and the relative positions can be calculated according to the gadget instructions. For instance, in the gadget (*pop %eax, pop %ebp; ret*), the two *pop* instructions move the stack pointer 8 bytes towards the stack base, while the instruction *ret* retrieves the top integer of the application stack as its next destination and increases the stack pointer 4 bytes at the same time. Following their stack operations, we can correctly re-position the stack pointer and get the destination of the next gadget. The ROPEcker pre-processor performs this costly instruction analysis and stack pointer calculation offline and stores the results in the IG database (Sec-

tion VI-A). At run-time, the ROPEcker kernel module simply queries the database and retrieves the results efficiently.

Dealing with the jmp-based gadgets is more challenging, as their destinations are dependent on the semantic of the gadget instructions. Moreover, it is rather difficult to enumerate all possible combinations in the offline processing phase due to its enormous storage cost. We propose to emulate such jmp-based gadgets at run-time to reveal the destinations of the jump instructions. Since most gadgets only have a few instructions, and the number of jmp-based gadgets are relatively smaller comparing with ret-based ones in a normal application, the overhead of gadget emulation is limited. **Moreover, because the offline databases cover the most gadget cases, ROPEcker seldom triggers its emulation step. The performance evaluation results also demonstrate this point.**

To avoid any side effect on the actual context due to the emulation, we build a new instruction emulator that emulates the instructions in a temporally created environment with an initial context identical to the present one. The implementation details of the emulator is provided in Section VI-B4.

VI. IMPLEMENTATION

To evaluate the effectiveness and performance of our approach, we have implemented a prototype ROPEcker for x86 32bit Ubuntu 12.04 with kernel 3.2.0-29-general-pae. The ROPEcker system consists of two key components: 1) an offline pre-processor which aims to generate a database accelerating run-time checking, and 2) a loadable kernel module which responds to the run-time events and invokes the ROP checking algorithm.

A. Offline Pre-processor

Given a target application, the ROPEcker pre-processor collects the application binaries and their dependent shared libraries. For each binary, the pre-processor extracts the size of the executable code segments from the binary header. For a code segment of n bytes, the pre-processor allocates a bit-vector of $n/2$ bytes, such that each byte in the executable code segment corresponds to a 4-bit slot in the bit-vector. The slots are indexed by the byte offset in the executable code segment. As shown in Table II, the slot values are assigned according to: 1) the type and **optional** alignment information of the instruction starting from this byte; and 2) the type and stack modification behavior of the gadget (in our case, a maximum of six contiguous instructions that ended with an indirect branch) starting from this byte. During emulation, instruction information is used to decide whether the end of the emulated gadget is reached, which is an indirect branch instruction. During run-time gadget chain detection, we use the gadget information to identify gadgets and their chaining points (e.g., a gadget manipulates the stack pointers to point to the next gadget). The bit-vector tables from all binaries form the IG database. The IG

database covers the most frequently used instruction types, gadget types and their stack manipulation behavior, and leaves only the rare cases to the run-time emulator. This design minimizes the invocations of our emulator and thus reduces the run-time detection overhead.

To build the IG database, we develop a tool to analyze the gadget information of a given binary, based on the disassembly library diStorm [10] (around 11K SLOC). Specifically, for each byte in the executable code segment, we run the gadget-analysis tool to disassemble six instructions starting at that byte. If an indirect branch instruction is found, we further analyze the stack manipulation behavior of the gadget, and store the result in the IG database. The alignment information in our implementation is extracted from *readelf* and *objdump* outputs. Note that we could use other disassembler tools (e.t., IDA Pro or the tool in [17]) to verify the aligned information.

The alignment information is **optional** since it is platform-dependent and the alignment analysis may not be completely accurate for certain applications. Thus, end users can selectively or completely disable the alignment checking. The main benefit of the alignment checking is increasing the bar for launching ROP attacks on x86 platform, because the adversary has to give up using unaligned gadgets. By doing so, the adversary generally needs more gadgets to launch an ROP attack. According to the analysis on the Q data set, the adversary generally needs two more gadgets for an ROP attack. **Moreover, the adversary also needs more code base. For example, Q generally requires 100KB or more code base for constructing an ROP payload with only aligned gadgets.**

We choose bit-vectors as our pre-processing database instead of hash tables, because the bit-vector indexing/query is more efficient and stable. For the bit-vector, the time cost of a query is only one memory access without false positive, while for a hash table, the time cost of a query is the computation cost of the hash function plus one or more memory access if there is a collision. **In addition, the bit-vector is also space efficient, which is also indicated in our experiment (Section VII-B).**

B. Kernel Module

We build the ROPEcker module as a loadable kernel module, which can be automatically installed when the system boots up. The module consists of 7K SLOC, where a large portion (around 4.4K) is attributed to the instruction emulator [5]. The ROPEcker module undertakes four main tasks: 1) to set up ROP checking triggers; 2) to collect the address locations of the shared libraries, 3) to check the branching history; and 4) to emulate instructions.

1) ROP Checking Triggers: The module inserts hooks in the Interrupt Descriptor Table (IDT) and the system call table to intercept the page fault exceptions and sensitive system calls, respectively.

Value	Instruction Information		Gadget Information	
	Instruction Type	Emulation Decision	Gadget Type	Payload Detection
0000	Aligned non-branch	Continue emulation.	Not a gadget	Stop gadget chaining
0001	Aligned direct branch	Stop emulation	Not a gadget	Stop gadget chaining
0010	Aligned indirect branch (ret)	Stop emulation	Ret-based gadget	Stack offset 4
0011	Aligned indirect branch (call *.jmp *)	Stop emulation	Jmp-based gadget	Need emulation
0100	Aligned non-branch	Continue emulation	1) Stack pivoting	Need emulation
			2) Jmp-based gadget	
			3) Stack offset too large	
0101 to 1110	Aligned non-branch	Continue emulation	Ret-based gadget	Stack offset 4*(Value-4)
1111	Unaligned	Stop Emulation	Unaligned gadget	ROP Found

Table II
THE FORMAT OF THE BIT-VECTOR TABLE.

Sliding Window Setup. We use the NX (Never eXecute) page table permission in any DEP-capable modern processor, to set up the sliding window. ROPEcker initially sets the NX bits in the loaded virtual memory pages of the application and library code. During execution, the application will trigger page faults by trying to execute code in non-executable pages. Once the page fault is captured, the ROPEcker module identifies the relevant faults by using process ID and page fault error code. The error code is a special combination (*PF_INSTR|PF_USER|PF_PROT*), which means that the page fault exception is triggered due to the protection violation during instruction fetching from user space. In the normal application execution, the code (NX pages) without execution right will never be executed. Thus, this error code confirms that the exception is triggered by the sliding window mechanism.

Sensitive System Call Interception. The system calls to intercept are *mprotect*, *mmap2* and *execve*. For the *mprotect* and *mmap2* calls, the module checks the request before passing it to the kernel. Any request to change a read-only code region to writable is rejected. Any request to change the data regions to executable is rejected. Then, the module invokes the ROP checking algorithm to ensure that there is no ROP gadget chain in the current stack. Otherwise ROPEcker will return the system call with an error code to indicate the failure of the request.

The system call *execve* is able to start a new process with some prepared input parameters. For instance, the adversary may use the libc function *system* or directly invoke the system call *execve* to open a shell, in which the adversary is able to execute arbitrary commands. The end user can create a simple policy: ROPEcker directly rejects all such requests for all selected applications. In most cases, it works well. However, for some legitimate applications, they may also send such requests for certain purposes. The first policy may lead to the non-coexistence of ROP with these applications. To support such applications, ROPEcker can be configured to launch the ROP checking algorithm to verify if there is an ROP attack. If not, ROPEcker passes the request to the

kernel. By doing so, ROPEcker and such applications can coexist, but the performance overhead may slightly increase due to the extra checking.

2) *Memory Mapping Acquisition:* ROPEcker should acquire the virtual memory mapping of the protected application and its shared libraries. If the Address Space Layout Randomization (ASLR) mechanism is enabled, the shared libraries are loaded to random addresses in different application instances. In the commodity OSes, there is *no* exported interface for kernel modules to get the mappings of a particular application. Thus, in order to locate the exact memory mappings for the target application, the ROPEcker module has to intercept the mapping-manipulation operations or analyze the corresponding kernel data structures to construct the mappings.

Operation Interception. For the memory mapping manipulation operations, they are driven by the application requests through certain system calls. In a Linux system, the kernel reserves the mapping region for the application binary in the system call *execve*, while in the system call *mmap2/mummap*, the kernel reserves/releases the mapping region for shared libraries. By intercepting such system calls, ROPEcker can get the mapping information from the parameters and the return values. For instance, in the system call *mmap2*, the base address of a shared library can be obtained either from the return value or the first parameter, and the length of the occupied memory region can be obtained from the second parameter. In addition, we need to intercept the *open* and *close* system calls, as they provide the name of the shared libraries. The names will be used in the database installation phase.

Data Structure Analysis. Certain shared libraries (e.g., the *library loader*) are loaded by the kernel by default, rather than being driven by the application request. Thus, we can not find their mapping information by intercepting system calls. To handle such cases, we can analyze the corresponding data structures to get the needed information. In a commodity operating system, the kernel usually has one or more dedicated data structure to maintain the memory

mapping information. The Virtual Memory Area (VMA) is the data structure in Linux that maintains the start and end of memory mappings/segments. Each library is represented by a VMA structure (i.e., memory segment). The name of the shared library is also linked to the VMA structure. In addition, all VMAs are linked together and the list header can be easily found following the task structure. Thus, by traversing the VMA list, we can easily locate the corresponding VMA and get the memory mapping information. Note that the traversing is only done once, since the library locations are fixed when the application starts to run.

3) *LBR Access*: The `IA32_DEBUGCTL` MSR is the control register in the LBR mechanism. Through this MSR, ROPEcker can enable, disable and manipulate the recording behaviors, e.g., it allows the LBR to only record user space branches. The LBR values are stored in the MSR registers, from `MSR_LASTBRANCH_k_FROM_IP` and `MSR_LASTBRANCH_k_TO_IP`, where k is the number of the branch record, e.g., the range of k is 0-16 on the Intel i5 processors. To read and write such LBR MSRs, ROPEcker must leverage the privileged instructions such as `rdmsr` and `wrmsr`.

4) *Instruction Emulation*: The ROPEcker module creates a shadow environment for emulating instructions. In the shadow environment, the initial context is initialized using the context of the interrupted application. It is challenging to create the shadow virtual address space. The virtual address space is quite large, e.g., `3GB` in an x86-32bit system. To get a complete copy, the time and space cost will be extremely high. In addition, the shadow environment will be immediately dropped when the current round of the checking finishes. To save cost, we borrow the *copy-on-write* idea from the Linux kernel. Specifically, we make the virtual address space read-only for the emulator, meaning that the emulator can freely *read* any address but can not write. When a write operation is needed, ROPEcker creates a mapping table which records the destination address together with the new value. Later, any reading or writing to this address will be redirected to the corresponding table entry. After the checking algorithm, the mapping table is cleared for avoiding polluting the next round emulation.

VII. EVALUATION

We evaluate both security and performance of ROPEcker. For the security evaluation, we first calculate the threshold for ROP gadget chain detection, and verify our algorithm with real ROP attacks and those generated by Q. For the performance evaluation, we use the SPEC CPU2006 benchmark suite [3], which is processor, memory and compiler stressing, Bonnie++ [22] which is disk I/O stressing and Apache server `httpd-2.4.6` [1] which is network stressing. The source code of the SPEC CPU2006 benchmark suite is compiled with gcc, g++ version 4.6.3 with the default Makefile. We run

Configurations	Descriptions
CPU	Intel i5 M540 with 2.53GHZ
Memory	4GB DDR3 1333MHZ
Network Card	Intel 82577LM Gigabit
Disk	320G ATA 7200RPM

Table III
THE CONFIGURATIONS OF THE EXPERIMENT MACHINE.

our experiments on a system with the configurations shown in Table III.

A. Parameter Effects

In this section, we discuss and evaluate the effects of ROPEcker parameter tuning, in terms of performance and security (accuracy).

1) *Gadget Chain Threshold*: The gadget chain threshold affects the performance of the monitored processes. Bigger threshold generally produces higher performance overhead. However, the degree of the performance decrease is very limited according to our experiment results (Figure 4), which indicates that ROPEcker stops the ROP detection after one or two attempts in up to 96.8% cases. The gadget chain threshold also affects the detection accuracy. An ideal threshold should be smaller than the minimum length min_{rop} of all ROP gadget chains, and at the same time, be larger than the maximum length max_{nor} of the gadget chains identified from normal execution flows. The threshold can be any number between max_{nor} and min_{rop} , and the choice does not affect the performance, because the ROP detection algorithm always stops before it reaches the threshold. If max_{nor} is larger than min_{rop} for certain applications, there is no ideal threshold and ROPEcker inevitably produces false positives and/or false negatives.




Figure 4. **The length of payload measurement results.** The number (0 to 5) is the payload length. In the major cases (84.37%), there is no gadget identified at all.

To find max_{nor} , we measured the lengths of gadget chains from many normal applications, which include 17 popular

Linux tools (e.g., *grep*, *find*, *less* and *netstat*) under directory `/bin` and `/user/bin`, 12 benchmark tools of SPEC INT2006, and 3 *large* binaries (the video processing tool *ffmpeg* 2.2.0, the graphic processing tool *graphics-magick*-1.5.1, and the Apache web server *httpd*-2.4.6). In our experiments, ROPEcker counts the lengths of the identified gadget chain in each invocation. As shown in Figure 4, in 84.37% of the ROPEcker invocation there is no gadget detected. The counts of detected gadget chain drops when gadget chain length increases. The detected chains with 6, 7, 8, 9 and 10 gadgets together only occupy 0.00006% in the total measurements. 10 is the largest length of gadget chain we have detected in the experiments. Note that the percentage results also reflect the efficiency of our ROP detection algorithm on normal applications, i.e., in the major cases (84.37%) the future payload detection stops the analysis after one attempt. To decide min_{rop} , we also measured the ROP payload collected from the real world ROP attacks and generated by the general-purpose ROP compiler Q [25]. The shortest aligned gadget chain contains 17 gadgets and the longest one has 30 gadgets. Based on the above results, we can safely choose a number from 11 to 16 as the threshold for major applications. Given a specific application, we can even choose a smaller threshold for better run-time detection performance. For instance, the longest gadget chain in Apache server has only 4 gadgets, so the threshold for Apache can be 5.

2) *Gadget Length*: In our implementation, we assume the gadget contains no more than 6 instructions, which is true for most of gadgets in the real-world ROP attacks. However, a smart adversary may attempt to use *long gadgets* (with more than 6 instructions) to bypass our ROP detection. Fortunately, long gadgets usually introduce more side effects, and thus the adversary has less opportunity to collect enough *suitable* long gadgets to form a valid chain for a meaningful ROP attack. If the adversary has to use consecutive short gadgets, it is likely to be detected by ROPEcker. However, a smarter adversary may carefully insert long gadgets into consecutive short gadgets to make the length of each segmented gadget chain not exceed the gadget-chain threshold, e.g., use a long gadget for every 10 gadgets. In this way, the current version of ROPEcker will miss this ROP attack.

To *mitigate* long gadget attacks, we propose a solution to extend ROPEcker. Specifically, we observe that there are still many segmented gadget chains although each of them is shorter than the threshold. If we accumulate their lengths in several consecutive sliding windows, the accumulated length may longer than the maximum one collected from the normal execution flows. We evaluated this extension in the experiment, where we choose to measure *three* consecutive sliding windows and count the total number of gadgets. According to the measurement results, the minimum accumulated length of the tested ROP attacks collected from real

world samples and generated by Q is 43, while the maximum accumulated length of the major tested applications is 13. The gap still tolerates a certain degree of long-gadget attack, allowing our algorithm to distinguish ROP attacks from normal executions. For a particular application, the gap is even larger, e.g., the maximum length of *htedit* is 9. Even so, if the adversary inserts too many long gadgets, our detection still fails. The *extra* space cost for the multiple-window accumulation extension is small, since ROPEcker only needs a short cycle buffer (i.e., the length is 3 in our example) to maintain the gadget lengths. In addition, the *extra* time cost is also small, because the checking only needs one integer comparison.

3) *Sliding Window Size*: The sliding window size is also a critical parameter for ROP detection. Larger window size offers better performance, but it may give more opportunities for the adversary to collect enough gadgets within the sliding window to launch a meaningful ROP attack. Thus, we should select a proper window size to balance the performance and security. According to the statistical results from Q [25], the adversary has low possibility to launch a meaningful ROP attack, if we can limit the scope of the executable code at any time within $20KB$. To be safety, we recommend setting the window size smaller than $20KB$. In our experiment we set the window size as $8KB$ (2 pages) and $16KB$ (4 pages). According to the performance results in Section VII-D, the performances with larger window size are generally better, but the performance differences are not that big. Even in some cases, they can achieve almost the same performance.

B. Space Evaluation

The space cost includes disk and memory costs, which are mainly from the generated databases. In our experiment, the databases for all 2393 shared libraries under `/lib` and `/usr/lib` of the Ubuntu Linux 12.04 distribution is about $210MB$. On average, a database is $90KB$. For example, the database of *libc*-2.15 is $832KB$, and the *ld*-2.15 database is $68KB$. Note that the size of a database is only related to the size of the original binary file, rather than the detection parameters (i.e., gadget length, gadget chain threshold and sliding window size). To further reduce the space cost on the disk, the database can be compressed. The experiment results show that all databases of shared libraries can be compressed to about $19MB$ using *bzip2*. The compression rate reaches up to 91%, which also indicates that the potential gadgets are sparse in the binary files.

At run-time, the databases are dynamically loaded into memory according to the demands of the monitored processes. To further reduce memory cost, we allow the loaded databases to be shared by all processes. For instance, if one process has loaded the *libc* database, all other monitored processes can share it without reloading. Thus, the memory

cost of ROPEcker is not large. Even if ROPEcker loads all databases of the shared libraries into memory, the memory cost is still acceptable. The memory space effects from gadget chain threshold, gadget length and sliding window size are quite small and limited. To maintain n sliding windows for a process, ROPEcker only needs an n -length *integrity* list recording the base addresses of the code pages within the sliding windows. Due to the smallness of gadget length and gadget chain threshold, they only need *limited* data structures *temporal* available in the gadget chain identification procedure.

C. Security Evaluation

To verify the strength of ROPEcker, we conduct several tests on two real-world examples and a number of payloads generated from Q. The experiment results show that ROPEcker can detect all these attacks with zero false negatives.

In the first test, we use an artificial program (demonstrated in ROPEME [19]) that has a simple stack buffer overflow triggered by a long input parameter. We use ROPEME to analyze the program and generate usable gadgets. We manually chain them together to craft an ROP attack that causes the program to start a shell which can facilitate the following attacks. We then verify our system with a realistic program: a Linux Hex-editor (*htediter*) (2.0.20). The ROP example on the *htediter* can be found on the web site [2]. An appropriately chosen long input can trigger a stack overflow. The template of the ROP payload is available on the website. We only replace the gadget pointers according to the compiled binary in our system.

At last, we test ROPEcker against the gadget chains generated by Q. Q uses semantic program verification techniques to identify the functionality of gadgets and generates a valid ROP payload. Specifically, we generate the payloads for the applications under directory */bin/* and */usr/bin/*. **To avoid manually exploiting each application, we build a new tool as a simulator to simulate the application execution environment. We assume that the adversary has compromised the process and successfully make the execution flow start from the gadget chain. Once a gadget jumps out of the simulated sliding windows, an algorithm will count the number of gadgets in the rest of chain. If the number is larger than the threshold, an ROP attack is identified.** Specifically, the detection algorithm itself can successfully detect 100% all the payloads generated from 253 applications under directory */bin* and directory */usr/bin/*.

D. Performance Evaluation

In this experiment, we use micro- and macro-benchmark to evaluate the performance of ROPEcker. We also evaluate its performance impact to the target applications and other applications in the same platform.

Operation	open	close	execve	mmap	mprotect
Time (μs)	0.03	0.04	0.86	0.05	0.03
Operation	munmap	pre-exception	post-exception		
Time (μs)	0.04	0.03	0.01		

Table IV
THE MICRO-BENCHMARK RESULTS FOR ONE INTERCEPTION.

Operations		Time (μs)
Past Gadget Chain Detection		0.07
Future Gadget Chain Detection	W/o Emulation	0.91
	W/ Emulation	2.61

Table V
THE TIME COST IN ROP CHECKING.

1) *Micro-Benchmark*: The micro-benchmark is used to evaluate the time cost of each operation introduced by ROPEcker. The time cost introduced by system call interception may affect the performance of other applications, but other operation costs (e.g., database installation) are confined to the protected application.

To measure the time cost on the other applications, we run ROPEcker without any target application. The system call interception code in the *execve* needs to check the application name, while others only check the process ID (PID). The checking costs are listed in Table IV. The experiment results show that the time overhead introduced to other applications is quite small.

We then use the *htediter* as an example to measure the time cost of loading the IG database. Note that the time is dependent on the size of the database. For instance, it takes $283\mu s$ to load the database for *libc-2.15* ($849.93KB$), and $51.07\mu s$ for *ld-2.15* ($65.55kB$). Although the loading is relatively high, it is only performed once in the whole life cycle of the protected application. In addition, the database installations are mostly done in the application initialization stage, before the application starts execution. Note that the databases of the shared libraries will not be re-loaded if they have already been loaded by other protected applications.

We also measured each step of the ROP checking algorithm. As shown in Table V, the performance overhead is quite low. Note that the total time cost for a particular checking is not the sum of all these steps. The reason is that some branches (e.g., the payload checking with emulation) may not be taken in that round. In fact, payload detection (with emulation) is rarely performed. In our test cases, the emulation is only invoked in about 1.7% of all detection invocations.

2) *Macro-Benchmark*: The macro-benchmark tests the overall performance impacts of ROPEcker on the target application and the system. We choose three aspects to evaluate: CPU computation, disk I/O and network I/O. Considering accuracy, we set the size of the sliding window always under the upper bound ($20KB$). Specifically, we

choose 8KB (2 pages) and 16KB (4 pages) as the size of the sliding window in our experiments.

SPEC CPU Benchmark. We choose the benchmark tool SPEC CPU2006 benchmark suite to evaluate the computation performance. Specifically, we run the testing suits with and without ROPEcker. The results are illustrated in Figure 5, which shows that ROPEcker only introduces 2.60% performance loss on average.

Disk I/O Performance Evaluation. We choose the benchmark tool Bonnie++ (version 1.96) to evaluate disk I/O performance. The tool sequentially reads/writes data from/to a particular file in different ways. The read/write granularity varies from a character to a block (i.e., 8192 Bytes). Furthermore, we also test the time cost of the random seeking. Figure 6 shows the disk I/O measurement results which show the performance overhead on the disk I/O is quite low, i.e., 1.56% overhead on average.




Figure 6. The disk performance results.

Network I/O Performance Evaluation. We choose the Apache web server httpd-2.4.6 to evaluate the network I/O. In our experiments, two machines are directly connected through a network cable. The tool *ab* on the client machine sends network requests (e.g., *ab -n 100000 -c 50 http://192.168.1.11/index.html*) to retrieve a web page on the server running with ROPEcker. The Apache web server is configured to work in mpm-worker mode. It has one worker process with 40 threads. Figure 7 shows the experiment results of multiple test scenarios: “HTTP” and “HTTP-4k” represent that the server running HTTP protocol serves the default work page (45 bytes) and a 4K bytes web page, respectively. “HTTPS” and “HTTPS-4k” denote that the server running HTTPS protocol with a 1024 bit RSA key. The test results indicate that a larger sliding window achieves better performance, and the performance overhead on network I/O is acceptable. For example, the performance overhead is 6.33% when the Apache server delivers the default 45-byte HTTP web pages, and is only 0.08% when serving 4K byte HTTP web page. Even in more computation-bound cases, e.g., HTTPS, the ROPEcker overhead is only 9.72%, when choosing the 4-page sliding window.



Figure 7. The network performance results.

VIII. DISCUSSIONS

A. Stack Pivoting Attack

The adversary can launch the *stack pivoting* attack [29] to evade ROPEcker detection. In this attack, the adversary puts the payload into another memory region (e.g., a heap buffer or the global data region), and manipulates the stack pointer to point to that region. Fortunately, the future payload checking algorithm is able to defeat such attacks, since ROPEcker always follows the execution flow to identify the potential gadget chain, without relying on the trustworthiness of the stack information.

B. Gadget Gluing Attack

Our payload detection algorithm is designed based on the assumption that a gadget does not contain direct branch instructions, which is also used in the many previous work [26], [8], [25], [32], [14], [34]. Therefore, the gadget chain detection stops when a direct branch instruction is encountered. However, the ROP attack in theory can foil our detection by constructing a special gadget which consists of two short code sequences glued together by a direct branch instruction. We call this type of attack *Gadget Gluing Attack*, which has not been found in real-life to the best of our knowledge.

Gadget gluing attacks are very powerful, as they lead to high ambiguousness between ROP attacks and normal executions. The current version of ROPEcker can not well defend against them. To mitigate such attacks, in the future, one possible extension of ROPEcker is to follow the direct branch at run-time. If its destination is a gadget, ROPEcker treats the whole as a glued gadget. The following step can be repeated one, two or more times (according to the configuration) until getting a gadget or reaching the times limit. This approach may introduce more false positives and/or false negatives.

C. Short Gadget Chain

In certain extreme cases, the gadget chain may only contain one or two gadgets. For such special ROP attacks, our scheme can not detect or prevent them because the length

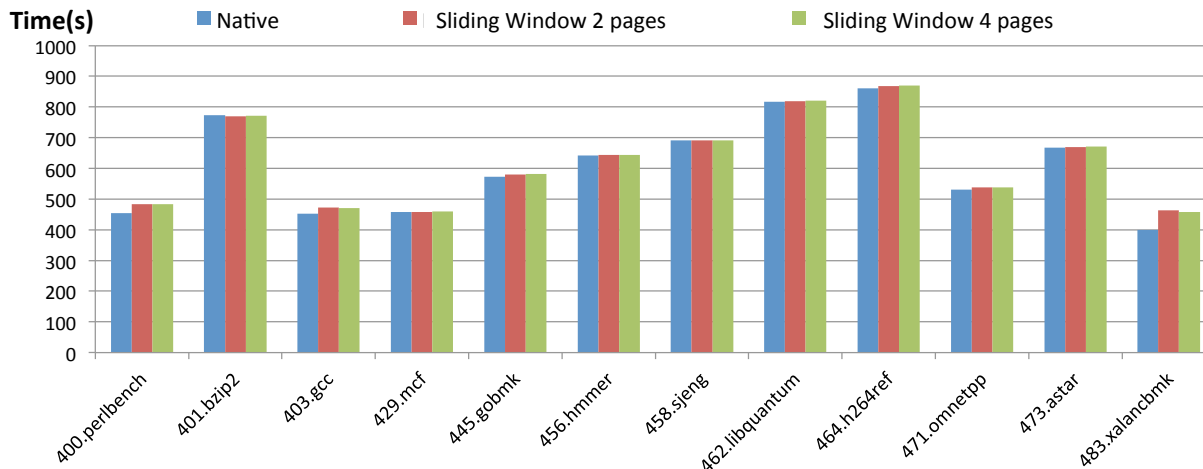


Figure 5. The SPEC INT2006 Benchmark Results.

of the gadget chain does not exceed the threshold. However, due to the limited power of the adversary in such attacks, we guess that their goals are not to directly launch malicious behaviors, while they aim to open doors for facilitating the following ROP attacks. In that way, our scheme is likely to detect the following ROP attacks.

D. Gadgets Within Sliding Windows

In most cases, the gadgets within the sliding windows are not enough for an ROP attack. However, in the whole life-cycle of a process, we can not eliminate the possibility for the adversary to find one sliding window where an ROP attack is possible. For such cases, the current version of ROPEcker cannot detect them. One possible solution is to dynamically reduce the window size to lower the possibility of ROP attacks. Specifically, we can observe all the sliding windows that ever occurred during normal executions, analyze each of them to evaluate the possibility of launching an ROP attack, and record the highly possible ones. According to those recorded sliding windows, ROPEcker could choose a dynamic way to temporally reduce the window size when the execution flow is running in one of them, and immediately restore the size when the execution flow moves out.

IX. RELATED WORK

A. Randomization

Address Space Layout Randomization (ASLR) is proposed to prevent ROP attack by randomizing base addresses of code segments. Given that the adversary has to know the locations of the gadgets to launch an ROP attack, the ASLR technique seems to effectively prevent ROP attacks. However, it has been shown that ASLR can be bypassed by leveraging brute-force attacks [27] or information leakage attacks [30]. In addition, some libraries or applications may

not be ASLR-compatible, which allows the adversary to find useful gadgets to circumvent ASLR mechanism.

Facing such difficulties, researchers proposed the *binary stirring* technique [34], which imbues x86 native code with the ability to self-randomize its instruction addresses each time it is launched. Note that the size of the modified binary file increases on average by 73%. The ILR [14] technique is proposed to randomize the location of every instruction in a program for thwarting an attacker’s ability to re-use program functionality. The new generated ILR-programs have to be executed on a dedicated Virtual Machine (VM). The runtime performance overhead is relatively low, but the rule files are quite large and their in-memory size are even worse, e.g., the on-disk size of the rule file for the benchmark tool 481.wrf is about 264MB, while its in-memory size reaches to 345MB. Note that the *whole* database in our system is only 210MB..

B. Compiler-Based Approaches

The control flow integrity [4] is a typical technique to defend against the code reuse attack. However, the traditional control flow works in function level, which leads to the failure of the detection of the ROP attack since its disorder of control flow is on instruction level. CFLocking [6] aims to limit/lock the number of abnormal control flow transfer by recompiling a program. Essentially this technique can not handle the ROP attacks that use unaligned gadgets.

The return-less kernel [18] is a compiler-based approach that aims to remove the *ret* opcode from the kernel image by placing the control data into a dedicated buffer instead of the stack. Obviously this technique only defends against ret-based ROP attack. Following the same idea, Shuo et al. [28] propose a virtualization based approach, which requires the source code and compiler to insert control-data-integrity-checking code into function prologue and epilogue. Again,

this approach can not defend against the ROP attack that use jump/call instructions.

G-free [20] is a compiler-based approach, which aims to 1) eliminate all unaligned indirect branch instructions with *aligned sled* [20] and 2) protect the aligned indirect branch instructions to prevent them from being misused. **G-free requires side information, i.e., the source code. Moreover, the inserted code may introduce new gadgets, which can be potentially used by the adversary.**

C. Instrumentation-Based Approaches

TRUSS [24], ROPDefender [13], DROP [9] and TaintCheck [16] use code instrumentation technique to insert checking code into binary code to detect ROP attack. Such approaches not only break the binary integrity, but also suffer high performance overhead. For instance, the preliminary performance measurements for DROP range from 1.9X to 21X, and the performance overhead for TaintCheck is over 20X. In addition, some of them, such as ROPDefender and DROP only focus on ret-based ROP, rather than handling all types of ROP attack.

To overcome the high performance overhead issue, several new approaches are proposed. Specifically, the IPR [32] technique is proposed, which aims to smash the gadgets in place without changing the code size. However, many gadgets can not be removed using the in-place technique. In addition, the in-place smashing technique may not always smash a significant part of the executable address space [32], and it is hard to give a definitive answer on whether the remaining un-modifiable gadgets would be sufficient for constructing a meaningful ROP attack. The ROPGuard [15] and KBouncer [33] approaches only add the checking points in the *selective* critical functions (e.g., Windows APIs) by instrumenting the binary code on-the-fly. **Although the infrequent invocations of the checking algorithm lead to low performance overhead, they inevitably miss the ROP attacks that do not use those paths. Note that the ROPGuard only works the non-JOP code and KBouncer completely relies on non-fully reliable LBR records, which could be overflowed or polluted during context switches.** CCFIR [36] randomly inserts all legal targets of indirect control-transfer instructions into a dedicated *Springboard board*, and then instruments the binaries to limit indirect transfers to flow only to the board. CCFIR suffers from compatibility issues due to the integrity protection of the system shared libraries on certain platforms (e.g., Windows 7). In addition, the compatibility issue may lead to the failure of the ROP defence since the adversary may only use the gadgets from such non-instrumented system libraries.

D. Others

CFIMon [35] is to detect a variety of attacks violating control flow integrity by collecting and analyzing run-time traces on-the-fly. **However it has high detection latency that**

may cause it too late to detect an attack. MoCFI [12] is a framework to mitigate control flow attacks on smart phones. It performs control flow integrity checking on-the-fly without requiring the applications source code, and the experiment results show that it does not induce notable overhead when applied to popular iOS applications. Polychronakis et al. propose a method [21] to identify ROP payloads in arbitrary data. The technique speculatively drives the execution of code that already exists in the address space of a targeted process according to the scanned input data, and identifies the execution of valid ROP code at run-time. The basic idea of the payload detection method has been adopted by our payload checking algorithm.

X. CONCLUSIONS AND FUTURE WORK

We have presented ROPEcker, a novel and universal ROP attack detection and prevention system. We have innovated two techniques as the main building blocks of ROPEcker. One is the gadget chain detection algorithm which detects the chain in the past and future execution flow, and the other is the sliding window mechanism which allows the detection algorithm to be triggered with a proper timing so as to achieve high accuracy and efficiency. **ROPEcker does not require source code, special compiler, or binary instrumentations. In fact, our scheme is complementary to instrumentation- and compiler-based approaches, as well as randomization schemes. Our scheme provides another line to defend against ROP attacks. We evaluated the security of ROPEcker by running experiments with real-life ROP attacks and those generated from Q. We also evaluated its performance by running benchmark tools which show acceptable performance loss.**

Future Work. The current design and implementation of ROPEcker focus on user-space ROP attacks. In the future work, we will extend ROPEcker to defend against kernel-space ROP attacks, because the hallmarks of ROP attacks are still there. However, ROPEcker should not be placed in the kernel space any more, since it may be disabled by the adversary once the kernel is compromised. A promising way is to move ROPEcker into the hypervisor space, where the adversary cannot access. By doing so, the performance penalty will be higher because the host-guest switch cost is higher than the one of kernel-user context switch. The performance loss may also be introduced by interrupts and exceptions, because they may introduce poor temporal and spatial locality for the kernel execution flow, and thereby increase the frequency of the detection algorithm invocation.

In the future work, we also aim to extend ROPEcker to support ARM platform. The ARM-based CPUs might not have LBR registers or the registers with similar functionality. Thus, the past payload detection will be removed in such environment. Fortunately, the ROPEcker future payload detection still works, unlike other schemes [33] that completely rely on LBR. On ARM platforms, instructions are aligned

to 16 bits or 32 bits. This will simplify the disassembler and the gadget analysis of the ROPecker pre-processor. The minor change is that the optional alignment checking does not work on ARM-based platform.

In the future, we plan to publish the code of ROPecker, facilitating other researchers who want to evaluate and/or improve the system.

XI. ACKNOWLEDGMENTS

The authors would like to thank the anonymous reviewers and our shepherd Davide Balzarotti for their numerous, insightful comments that greatly helped improve the presentation of this paper. The authors are also grateful to Virgil D. Gligor, Edward J. Schwartz for useful discussions. This research is supported by the Singapore National Research Foundation under its International Research Centre Singapore Funding Initiative and administered by the IDM Programme Office, Media Development Authority (MDA).

REFERENCES

- [1] The apache software foundation, Apache http server. <http://httpd.apache.org/>.
- [2] HT Editor 2.0.20 Buffer Overflow (ROP PoC). <http://www.exploit-db.com/exploits/22683/>.
- [3] Standard Performance Evaluation Corporation, SPEC CPU2006 Benchmarks. <http://www.spec.org/osg/cpu2006>.
- [4] ABADI, M., BUDI, M., ERLINGSSON, U., AND LIGATTI, J. Control-flow integrity. In *Proceedings of the 12th ACM conference on Computer and communications security* (2005), pp. 340–353.
- [5] BARHAM, P., DRAGOVIC, B., FRASER, K., HAND, S., HARRIS, T., HO, A., NEUGEBAUER, R., PRATT, I., AND WARFIELD, A. Xen and the art of virtualization. In *SOSP '03: Proceedings of the nineteenth ACM symposium on Operating systems principles* (2003), pp. 164–177.
- [6] BLETSCH, T., JIANG, X., AND FREEH, V. Mitigating code-reuse attacks with control-flow locking. In *Proceedings of the 27th Annual Computer Security Applications Conference* (2011), pp. 353–362.
- [7] BUCHANAN, E., ROEMER, R., SHACHAM, H., AND SAVAGE, S. When good instructions go bad: Generalizing return-oriented programming to RISC. In *Proceedings of the 15th ACM conference on Computer and communications security* (2008), pp. 27–38.
- [8] CHECKOWAY, S., DAVI, L., DMITRIENKO, A., SADEGHI, A.-R., SHACHAM, H., AND WINANDY, M. Return-oriented programming without returns. In *Proceedings of the 17th ACM conference on Computer and communications security* (2010), pp. 559–72.
- [9] CHEN, P., XIAO, H., SHEN, X., YIN, X., MAO, B., AND XIE, L. Drop: Detecting return-oriented programming malicious code. In *Proceedings of the 5th International Conference on Information Systems Security* (2009), pp. 163–177.
- [10] DABAH, G. A lightweight, Easy-to-Use and Fast Disassembler/Decompiler Library for x86/AMD64. <http://ragestorm.net/distorm/>.
- [11] DANIEL P., B., AND MARCO, C. Understanding the Linux Kernel, Third Edition, chapter 2 - hardware cache and chapter 9 - demand paging, 2005.
- [12] DAVI, L., DMITRIENKO, A., EGELE, M., FISCHER, T., HOLZ, T., HUND, R., NÜRNBERGER, S., AND SADEGHI, A.-R. Mocfi: A framework to mitigate control-flow attacks on smartphones. In *19th Annual Network & Distributed System Security Symposium (NDSS)* (2012).
- [13] DAVI, L., SADEGHI, A.-R., AND WINANDY, M. Ropdefender: a detection tool to defend against return-oriented programming attacks. In *Proceedings of the 6th ACM Symposium on Information, Computer and Communications Security* (2011), pp. 40–51.
- [14] HISER, J., NGUYEN-TUONG, A., CO, M., HALL, M., AND DAVIDSON, J. W. IIr: Where'd my gadgets go? In *Proceedings of the 2012 IEEE Symposium on Security and Privacy* (2012), pp. 571–585.
- [15] IVAN, F. Runtime prevention of return-oriented programming attacks. <https://code.google.com/p/ropguard/>.
- [16] JAMES, N., AND DAWN, S. Dynamic taint analysis for automatic detection, analysis, and signature generation of exploits on commodity software.
- [17] KRUEGEL, C., ROBERTSON, W., VALEUR, F., AND VIGNA, G. Static disassembly of obfuscated binaries. In *Proceedings of the 13th conference on USENIX Security Symposium - Volume 13* (2004), pp. 18–18.
- [18] LI, J., WANG, Z., JIANG, X., GRACE, M., AND BAHAM, S. Defeating return-oriented rootkits with "return-less" kernels. In *Proceedings of the 5th European conference on Computer systems* (2010), pp. 195–208.
- [19] LONGLD. Payload Already Inside: Data Resue For ROP Exploits. <https://media.blackhat.com/bh-us-10/whitepapers/Le/BlackHat-USA-2010-Le-Paper-Payload-already-inside-data-reuse-for-ROP-exploits-wp.pdf>, Aug. 2010.
- [20] ONARLIOGLU, K., BILGE, L., LANZI, A., BALZAROTTI, D., AND KIRDA, E. G-free: defeating return-oriented programming through gadget-less binaries. In *Proceedings of the 26th Annual Computer Security Applications Conference* (2010), pp. 49–58.
- [21] POLYCHRONAKIS, M., AND KEROMYTIS, A. Rop payload detection using speculative code execution. In *Proceedings of the 6th International Conference on Malicious and Unwanted Software* (2011), pp. 58–65.
- [22] RUSSELL, C. Disk Performance Benchmark Tool - Bonnie. <http://www.coker.com.au/bonnie++/>.
- [23] SAILER, R., ZHANG, X., JAEGER, T., AND VAN DOORN, L. Design and implementation of a tcg-based integrity measurement architecture. In *Proceedings of the 13th conference on USENIX Security Symposium - Volume 13* (2004), pp. 16–16.
- [24] SARAVANAN, S., QIN, Z., AND WENG-FAI, W. Transparent runtime shadow stack: Protection against malicious return address modifications, 2008.

- [25] SCHWARTZ, E. J., AVGERINOS, T., AND BRUMLEY, D. Q. exploit hardening made easy. In *Proceedings of the 20th USENIX conference on Security* (2011), pp. 25–25.
- [26] SHACHAM, H. The geometry of innocent flesh on the bone: Return-into-libc without function calls (on the x86). In *Proceedings of the 14th ACM conference on Computer and communications security* (2007), pp. 552–61.
- [27] SHACHAM, H., PAGE, M., PFAFF, B., GOH, E.-J., MODADUGU, N., AND BONEH, D. On the effectiveness of address-space randomization. In *Proceedings of the 11th ACM conference on Computer and communications security* (2004), pp. 298–307.
- [28] SHUO, T., YEPING, H., AND BAOZENG, D. Prevent kernel return-oriented programming attacks using hardware virtualization. In *Proceedings of the 8th international conference on Information Security Practice and Experience* (2012), pp. 289–300.
- [29] SNOW, K. Z., DAVI, L., DMITRIENKO, A., LIEBCHEN, C., MONROSE, F., AND SADEGHI, A.-R. Just-in-time code reuse: On the effectiveness of fine-grained randomization. In *Proceedings of the 34th IEEE Symposium on Security and Privacy* (2013), pp. 574–588.
- [30] SOTIROV, A., AND DOWD, M. Bypassing browser memory protections. In *In Proceedings of BlackHat* (2008).
- [31] TRAN, M., ETHERIDGE, M., BLETSCH, T., JIANG, X., FREEH, V., AND NING, P. On the expressiveness of return-into-libc attacks. In *Proceedings of the 14th international conference on Recent Advances in Intrusion Detection* (2011), pp. 121–141.
- [32] VASILIS, P., MICHALIS, P., AND ANGELOS, D. K. Smashing the gadgets: Hindering return-oriented programming using in-place code randomization. In *Proceedings of the 2012 IEEE Symposium on Security and Privacy* (2012), pp. 601–615.
- [33] VASILIS, P., MICHALIS, P., AND ANGELOS D., K. Transparent rop exploit mitigation using indirect branch tracing. In *Proceedings of the 22nd USENIX Security Symposium* (2013).
- [34] WARTELL, R., MOHAN, V., HAMLIN, K. W., AND LIN, Z. Binary stirring: Self-randomizing instruction addresses of legacy x86 binary code. In *Proceedings of the 19th ACM Conference on Computer and Communications Security* (2012).
- [35] YUBIN, X., YUTAO, L., HAIBO, C., AND BINYU, Z. C-fimon: Detecting violation of control flow integrity using performance counters. In *Proceedings of the 42nd Annual IEEE/IFIP International Conference on Dependable Systems and Networks* (2012), pp. 1–12.
- [36] ZHANG, C., WEI, T., CHEN, Z., DUAN, L., SZEKERES, L., MCCAMANT, S., SONG, D., AND ZOU, W. Practical control flow integrity and randomization for binary executables. In *Proceedings of the 2013 IEEE Symposium on Security and Privacy* (2013), pp. 559–573.