

# Return-Oriented Programming on RISC-V

Georges-Axel Jaloyan  
DIENS, École normale supérieure,  
CNRS, PSL University, Paris, France  
CEA, DAM, DIF  
F-91297, Arpajon, France

Konstantinos Markantonakis  
Royal Holloway, University of London  
London, United-Kingdom

Raja Naeem Akram  
Royal Holloway, University of London  
London, United-Kingdom

David Robin  
DIENS, École normale supérieure,  
CNRS, PSL University, Paris, France

Keith Mayes  
Royal Holloway, University of London  
London, United-Kingdom

David Naccache  
DIENS, École normale supérieure,  
CNRS, PSL University, Paris, France

## ABSTRACT

This paper provides the first analysis on the feasibility of Return-Oriented programming (ROP) on RISC-V, a new instruction set architecture targeting embedded systems. We show the existence of a new class of gadgets, using several Linear Code Sequences And Jumps (LCSAJ), undetected by current Galileo-based ROP gadget searching tools.

We argue that this class of gadgets is rich enough on RISC-V to mount complex ROP attacks, bypassing traditional mitigation like DEP, ASLR, stack canaries, G-Free and some compiler-based backward-edge CFI, by jumping over any guard inserted by a compiler to protect indirect jump instructions.

We provide examples of such gadgets, as well as a proof-of-concept ROP chain, using C code injection to leverage a privilege escalation attack on two standard Linux operating systems. Additionally, we discuss some of the required mitigations to prevent such attacks and provide a new ROP gadget finder algorithm that handles this new class of gadgets.

## CCS CONCEPTS

• **Security and privacy** → **Embedded systems security**; *Malware and its mitigation*; *Operating systems security*.

## KEYWORDS

RISC-V, return-oriented programming, code overlap, Galileo algorithm

## ACM Reference Format:

Georges-Axel Jaloyan, Konstantinos Markantonakis, Raja Naeem Akram, David Robin, Keith Mayes, and David Naccache. 2020. Return-Oriented Programming on RISC-V. In *Proceedings of AsiaCCS '20: ACM ASIA Conference on Computer and Communications Security (AsiaCCS '20)*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*AsiaCCS '20, June 01–05, 2020, Taipei, Taiwan*

© 2020 Association for Computing Machinery.  
ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00  
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

## 1 INTRODUCTION

Memory corruption vulnerabilities are one of the most popular entry points for hackers to hijack a program. Amongst them, stack overflow attacks have been popular since 1996 [2]. It was for long thought that the hacker would always inject some standalone payload, that could be detected as malicious as such, using methods such as executable space protection [29]. This assumption has been invalidated by *Return-Oriented Programming* (ROP), introduced on par with the Galileo detection algorithm by Shacham in 2007 [43], proving, as formulated by Dino Dai Zovi in 2010, that “*preventing the introduction of malicious code is not enough to prevent the execution of malicious computations*” [51].

Since then, many countermeasures have been developed against ROP attacks [17, 19, 34, 35]. Each time, the publication of new ROP variants, such as JOP, SROP, SOP, or even JIT-spray [9, 10, 12, 38] bypassed those stopgap mitigations. At the same time, these attacks have been extended to many architectures, including much simpler *Reduced Instruction Set Computer* (RISC) [13], confirming that those design flaws are widespread among all architectures. State-of-the-art mitigation methods such as gcc’s `-mmitigate-rop` option or G-Free [34], tend to uproot such attacks by detecting and eliminating any code section that could be reused by an attacker, in the hope that the remaining gadgets would not be sufficient to mount complex attacks. Other even more radical methods like *Control-Flow Integrity* (CFI) try preventing arbitrary control-flow transfers by validating the target of indirect jumps [1, 27, 39], often at the cost of performance, thus reducing their usability [14, 16].

Likewise, these methods do hardly more than increasing the cost of such attacks, as it may be sufficient to find new unexpected gadgets to get back to step one of stack overflow exploitation. In this paper, we show once again, how to challenge the existing security mechanisms using a new class of gadgets that are undetected by the vast majority of published methods, based on the well-known Galileo algorithm. We explain how to produce such gadgets in RISC-V [49], a new *Instruction Set Architecture* (ISA) which development began in 2010. Consequently, an attacker may be able to insert such gadgets in an open source program and exploit them unnoticed.

RISC-V is based on the concept of RISC [37], targeting simplicity by providing few and limited computer instructions. RISC ISAs have become increasingly popular with the wide adoption of embedded devices such as smartphones, tablets, or other Internet of Things devices. The most popular RISC ISAs are currently ARM [6], Atmel AVR [7], MIPS [30], Power [25], and SPARC [45].

RISC-V is the fifth RISC ISA published by UC Berkeley. It is completely free and open-source, with its User-Level ISA published in May 2017 in version 2.2. It features 32-bit and 64-bit little-endian variants (designated as RV32 and RV64), with a future extension to 128 bits. While only expensive test boards feature RISC-V processors currently, many companies including Western Digital or Nvidia have announced the use of RISC-V chips in their future products [41]. Hence, this architecture is of particular interest for such attacks, as most programs are in the process of being ported to this architecture, leaving the insertion of backdoors easy for an ill-intentioned programmer.

We summarize our contributions as follows.

- (1) We provide the first analysis on the feasibility of ROP attacks on the new RISC-V architecture.
- (2) We introduce a new and stealthy class of ROP gadgets, undetected by all previously published methods based on the Galileo algorithm.
- (3) We show the achievability of complex ROP attacks using this class of gadgets on RISC-V ISA, under the assumption of malicious C source code insertion generating such gadgets.
- (4) We implement a proof-of-concept backdoored SUID program allowing privilege escalation on two standard Linux operating systems running on RISC-V, with every available ROP mitigation mechanism enabled.
- (5) We present a new algorithm able to find ROP gadgets of this class and discuss the plausibility of their presence in existing RISC-V binaries.

## 2 BACKGROUND

In this section, we briefly introduce the key concepts related to this paper's scope-of-work and contributions. More particularly, we describe the memory corruption exploitation technique known as Return-Oriented Programming and detail some RISC-V features, later used in the paper.

### 2.1 Return-Oriented Programming

The first methods aiming at exploiting memory corruption bugs were as simple as a straightforward data injection into the program, which would end up being executed by the processor [2]. The introduction of *Data Execution Prevention* (DEP) [29] made those attacks almost impossible, as injected data could not be executed anymore. In this battle between the shield and the sword, *Return-Oriented Programming* (ROP) has been the answer from malware developers. The first ROP attack was publicly presented in 2001 by Nergal in Phrack [32].

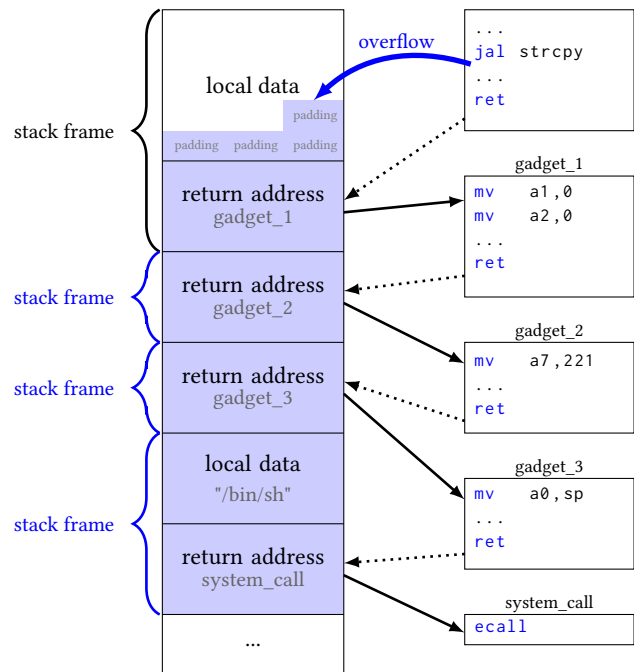
As shown by Fig. 1, it bypasses DEP by injecting in the stack a succession of call frames. Each call frame will result in the execution of a *gadget*: a small snippet of legitimate code containing a small number of instructions ended by a `ret`. When the `ret` instruction is reached, the address of the next gadget is popped from the stack into the program counter. Provided that enough different gadgets are available in the executable, arbitrary code may be executed by chaining those gadgets.

Two categories of gadgets can be distinguished. The first one using only legitimate code written by the programmer, also called

the *Main Execution Path* (MEP). The second category uses overlapping code, called *Hidden Execution Path* (HEP), *i.e.* code sections that have another interpretation by the CPU depending on its internal status (32 or 64 bits, Thumb mode, or on the offset at which the execution has started). The latter has the advantage of bypassing any compiler-added stack protection mechanism, presenting a wider variety of side-effects and being undetectable by traditional linear or recursive disassemblers, which only handle the MEP of a program.

The first academic paper studying this technique was published in 2007 by Shacham [43], in which he presents ROP on x86 and the Galileo algorithm, which allows the detection of gadgets in any executable memory region. It is based on a backward disassembly method, starting from every return instruction, and then trying to recursively bruteforce the length of the previous instruction. This provides a tree of possible gadgets all ending with a return.

The most common attack scheme consists in scanning the executable sections of the binary with Galileo based [15, 26, 50] or with other *ad hoc* algorithms [42] to find gadgets which are thereupon used to devise a ROP chain performing the required computation. Intermediate languages are sometimes used, allowing the design of higher-level ROP chains that are then compiled to the gadget language [44, 50]. Finally, the payload is adapted to the injection method, with techniques like padding, NUL bytes removal, or even alphanumeric conversion, which are not within the scope of this study.



**Figure 1: General principle of Return-Oriented Programming attacks. The vulnerability shown here consists in a buffer overflow from an unchecked strcpy allowing the user to smash the contents of the stack.**

By design, the Galileo algorithm is only able to find gadgets made of a straight-line instruction sequence, with no jumps except for the last instruction. Such a sequence is called a *Linear Code Sequence And Jump* (LCSAJ). Gadgets spanning over several LCSAJs are thus undetected by Galileo, and, to the best of our knowledge, have never been subject to study in the context of ROP attacks.

## 2.2 RISC-V

RISC-V splits its instruction set between a mandatory core set (RV64I) and different optional extensions, each of which is designated by a letter. The defined extensions include integer multiplication and division (M), atomic operations (A), single-, double- or quad-precision (F, D, Q) floating-point operations, decimal floating-point operations (L), compressed instructions (C), bit-manipulation (B), just-in-time (J), transactional memory (T), packed-SIMD instructions (P), vector operations (V), and user-level interrupts (N). The general purpose ISA, which includes IMAFD, is designated by the letter G. In what follows, we focus on the RV64GC ISA, which is the one agreed on by Debian and Fedora porters, as well as members of the RISC-V Foundation. On top of that, the Foundation intends to provide “a profile for standard RISC-V Unix platforms that will include C extension as mandatory”.<sup>1</sup>


There are 31 general purpose 64-bit registers (x1-x31), 32 floating-point registers (f0-f31), a program counter (pc), as well as various control-and-status registers. The pseudo-register x0 designates the zero constant. RISC-V provides a standard ELF *Application Binary Interface* (ABI), called psABI [18], with the naming convention provided in Fig. 2.

| Register | ABI Mnemonic | Meaning                |
|----------|--------------|------------------------|
| x0       | zero         | Zero                   |
| x1       | ra           | Return address         |
| x2       | sp           | Stack pointer          |
| x3       | gp           | Global pointer         |
| x4       | tp           | Thread pointer         |
| x5-x7    | t0-t2        | Temporary registers    |
| x8-x9    | s0-s1        | Callee-saved registers |
| x10-x17  | a0-a7        | Argument registers     |
| x18-x27  | s2-s11       | Callee-saved registers |
| x28-x31  | t3-t6        | Temporary registers    |


**Figure 2: Naming convention for registers, per RISC-V ELF psABI.**

While most RISC ISAs require naturally aligned instructions, RV64GC features 32-bit and 16-bit instructions, aligned on 16 bits, like in Thumb-2 extension introduced with ARMv6T2 [5]. Instruction length is encoded in the least-significant byte (hence with the lowest address as RISC-V is little-endian): 16-bit instructions require the last two bits to be different from 11 whereas 32-bit instructions have their last two bits equal to 11 with the three previous bits different from 111.

Combining those two peculiarities of RV64GC opens the door to overlapping instructions, that can be obtained by either using two 32-bit instructions 2 bytes apart (Fig. 3), or by using a 32-bit



**Figure 3: Two 32-bit overlapping instructions of  $I_1$  (little-endian representation). Instruction length encoding for each instruction is emphasized in blue.**



**Figure 4: A 32-bit instruction of  $I_2$  whose last 2 bytes are also a 16-bit valid instruction (little-endian representation)**

instruction whose last 2 bytes are also a valid 16-bit compressed instruction (Fig. 4). In what follows, we use  $I_1$  to designate the set of 32-bit instructions allowing overlapping sequences, whereas the set of 32-bit instructions whose last 2 bytes are valid 16-bit instruction will be denoted by  $I_2$ . Examples of overlapping for both sets  $I_1$  and  $I_2$  are given in Fig. 3 and 4. Typically, an overlapping sequence consists of several instructions of  $I_1$  chained together, optionally ending with an instruction of  $I_2$ .

## 3 THREAT MODEL AND ATTACK OVERVIEW

In this section, we explicit our target platforms, aiming run-of-the-mill RISC-V systems with off-the-shelf ROP mitigations deployed. We also present two attack scenarios taking advantage of our new class of gadgets for improved concealment.

Our target platform features a standard Linux operating system, such as Debian or Fedora, with two levels of privilege, that we call user and root. Standard protections are deployed, such as *Address Space Layout Randomization* and *Data Execution Prevention*, that prevent common stack overflow exploits. Programs are compiled with the standard gcc provided by the operating system, adding gcc’s ROP mitigation mechanism using compiler flag `-fstack-protector-strong`. Note that some other mitigation specific to x86 are not available on RISC-V, like gcc’s `-mmi tigate-rop` option or clang’s CFI. In Section 7, we discuss the ability of such mitigations, if ported to RISC-V, to hamper attacks using this new class of gadgets.

### 3.1 Closing (stealthily) the gap between vulnerability and exploitation

The first attack scenario focuses on adding a backdoor to a program leading to a ROP attack. Backdoors allow any person aware of their existence to reach a *privileged state* upon a specific *input*. In order to create a backdoor, two distinct elements must be stealthily inserted by an attacker: a *trigger* and a *payload* [46]. In our scenario,

<sup>1</sup><https://wiki.debian.org/RISC-V>

we assume the attacker already managed to insert a trigger (or find an existing one), in the form of a *ROP exec vulnerability*: a one-time memory write like a buffer overflow combined with an arbitrary control-flow redirect, such as a return at the end of the function, use-after-free, type confusion, or even corrupted instruction through fault injection [47]. Such vulnerabilities are pretty common in programs, and are often rendered non exploitable by reducing the number of available gadgets and by deploying ROP mitigations, such as ASLR, stack canaries, backward-edge CFI, or G-Free.

To lower the bar of exploitability, the attacker must embed gadgets in the payload of his backdoor, aiming at preventing any unaware outsider from stumbling upon those gadgets. As a stepping-stone for future elaboration, we consider generic C code injection through traditional backdooring, as we believe that one variant of this scenario may target C++ Just-in-Time compilers (like Cling [48] or ClangJIT [21], once they get ported to RISC-V) to mount JIT-spraying attacks [23]. Indeed, identical assumptions are required for the latter: code injection and ROP exec vulnerability.

As an illustration, we consider the case where the attacker has a user privileged level access to the system, including shell, ability to run programs, read access to binaries and libraries. The goal of the attacker is to increase his privilege level to root, which in practice thoroughly compromises the system by granting a read-write access to the whole target. Such an attack is called a *privilege escalation attack*, and is at the core of highly publicized attacks such as iOS jailbreaking [20]. To that end, the attacker will use a program that can be executed by the user, but running at a root privilege level. Those programs are called SUID (*Set owner User ID up on execution*), and are abundant on any system. Indeed, actions as simple as changing a password, plugging a USB key or granting root privilege for an authorized user require the execution of SUID programs.

To backdoor such programs, the attacker may upstream under-handed C code in some open-source project. Details on how to achieve this have been provided by Gilbertson [24] and thoroughly studied by Prati [40], with some examples provided in the *Under-handed C Contest* and DEF CON's *Hiding backdoor in plain sight* contest. Here, the payload consists in a set of ROP gadgets that span over several LCSAsJs. Furthermore, those gadgets are using overlapping techniques, so that only the last LCSAJ is in the MEP, whereas all the previous ones are in the HEP, thus hiding them to currently available ROP gadget searchers. To trigger the exploit and gain root access, the attacker only has to execute the SUID program with the adequate user input.

### 3.2 Creating a (concealed) persistent backdoor on a compromised system

The second attack scenario leverages privilege escalation through SUID to build a persistent backdoor in a compromised target. Persistence is considered as a key step in a complex attack chain to maintain access into compromised systems upon slight environment changes (reboot, updates, password change). This attack is much easier to implement than inserting backdoors in highly scrutinized SUID programs, as it requires the attacker to only get a one-time root access, and grant SUID permission to a program for

which he has knowledge of the existence of a privilege escalation exploit. Such backdoors are quite common,<sup>2</sup> as they involve modifying the permissions of only one file, which is not monitored by default on popular intrusion detection systems such as rkhunter, chkrootkit, or samhain.

For better chances of success, this can be combined with the first attack scenario, by inserting hidden gadgets in a non SUID open-source program, which is much easier to achieve. This backdoored program, embedding the hidden gadgets and a ROP exec vulnerability, will be legitimately deployed on the target. Should a security analyst audit the program before the attack, he will wrongly conclude that the vulnerability is not exploitable, hence not requiring an urgent patch. After the attack has been discovered, even if a forensics analyst comes across the program granted with SUID permission, without the knowledge of the ROP-chain, he will waste precious time and effort trying in vain to identify the mechanism allowing privilege escalation.

## 4 INSERTING HIDDEN GADGETS

For the sake of realism, we intend to use code created by a standard C compiler like gcc. We create exactly one function per gadget (named `function1, ...`), each ending with a C return instruction. For each function, the compiler may add assembly code at the beginning and the end of the function whose purpose is to respectively insert (*save sequence*) and remove (*restore sequence*) the call frame from the stack, depending on whether a callee-saved register is modified by the function. Inserting a nested call in the function is an easy way to be sure that the compiler will emit these save and restore sequences.

Indeed, the presence of a restore sequence is crucial for mounting a ROP attack, as we need to tamper with the return address register `ra`, which is callee-saved. Inserting malicious call frames into the stack hence grants control over the program counter through `ra`. In practice, a vast majority of functions do call other functions, either in the program, or in any library. In our proof-of-concept attack, we purposely added a call to a dummy function in every gadget function. Other ROP variants using alternative control-flow instructions such as indirect jumps or exceptions are beyond the scope of this paper.

The malicious gadget is made of two LCSAsJs, the first being hidden with code overlapping and the last being the legitimate restore sequence. A detailed example for one of the gadgets is provided in Fig. 5. The C code (using `gcc -Os -fstack-protector-strong`) used to generate it is:

```
long long function15c(){
    dummy();
    dummy4((signed) 0x9932000,
           0,
           (signed) 0xa0212000,
           (signed) 0x23371000);
    return 0;}
```

The hidden instructions are directly written in C code, and feature one or two instructions followed by a jump to a relative offset.

<sup>2</sup><https://attack.mitre.org/techniques/T1166/>

|                   | MEP                      | FUNCTION15c      | HEP  |
|-------------------|--------------------------|------------------|--|
| <code>addi</code> | <code>sp, sp, -16</code> | save sequence    |  |
| <code>sd</code>   | <code>ra, 8(sp)</code>   |                  |  |
| <code>jal</code>  | <code>ra, dummy</code>   | dummy call       |  |
| <code>lui</code>  | <code>a0, 0x9932</code>  | overlapping code | <code>addi</code> <code>s3, a4, 363</code> |
| <code>lui</code>  | <code>a3, 0x23371</code> |                  | <code>lui</code> <code>t1, 0x26372</code>  |
| <code>lui</code>  | <code>a2, 0xa0212</code> |                  | <code>jmp</code> <code>0x8</code>          |
| <code>mv</code>   | <code>a1, zero</code>    | instructions     |  |
| <code>jal</code>  | <code>ra, dummy4</code>  |                  |  |
| <code>ld</code>   | <code>ra, 8(sp)</code>   | restore sequence |  |
| <code>mv</code>   | <code>a0, zero</code>    |                  |  |
| <code>addi</code> | <code>sp, sp, 16</code>  |                  |  |
| <code>ret</code>  |                          |                  |  |

Figure 5: Segmentation of the different code sequences present in function15c. The gadget is highlighted in gray.

In the example of Fig. 5, the MEP consists of two 32-bit  $I_1$  instructions followed by one  $I_2$  instruction, whereas the HEP comprises two 32-bit  $I_1$  instruction followed by one 16-bit jump instruction. Here, the jump is only 8 bytes off its target, but it is definitely possible to modify this value to hide the overlapping LCSAJ anywhere, even in other functions. In this gadget, magic constants are loaded into the arguments of a function. The other gadgets use a mix of arithmetical and floating-point operations, as well as load and stores instructions. To have a consistent output among different compiler versions and environments, we forced register allocation (using the `register` keyword), and prevented instruction reordering in the overlapping sequence. Magic constants as arguments of the function cannot be prevented, as the opcode of a HEP instruction lies in the operand of the MEP instruction. However, many source code obfuscation techniques may come to help here, such as C-preprocessor [28] or lightweight constant blinding, hiding the magic constants respectively until the preprocessing and constant folding passes of the compiler.

## 5 CHAINING THE GADGETS

In the previous section, we described our method to build one gadget hiding some  $I_1$  instructions. In our full privilege escalation attack, we need to chain several of such gadgets together. We will aim at spawning a root shell, by invoking two system calls, the first being `setuid(0)` and the second `execve("/bin/sh", 0, 0)`.

In RISC-V, each syscall requires the execution of a special instruction named `ecall`, with register `a7` set to a value encoding the call.<sup>3</sup> For each call, one or several arguments may be passed, in registers `a0`, `a1`, `a2`, ... The `setuid` syscall requires `a7` to be set to 146, and `a0` to the desired `userid`, in our case zero. The `execve` syscall requires register `a7` to be set to 221 (0xdd), `a1` and `a2` to zero, and `a0` to point to the address of the string `/bin/sh`. The next paragraphs explain how to achieve this result by using only  $I_1$  instructions. We summarize the high-level overview of the ROP

chain in an assembly-like pseudocode in Fig. 6. The link to the full source code is available in Appendix A.

Let us start by zeroing (resetting) a register. For this purpose, we use the `slti` instruction (store less than immediate), that compares its source register to a constant, and if lower resets the destination register, else sets it to 1. By performing two `slti` instructions with a negative immediate and with same source and destination register, we are guaranteed to reset the register. In Fig. 6, this happens at offset 88. We can then reset other registers by just performing an `slti` with a zeroed source register and a negative immediate (offset 104).

The execution of an `ecall` instruction is trickier, as `ecall`  $\notin I_{1,2}$ . Hence, we must find an existing `ecall` and insert its location into the stack, so that the program counter points to it after the execution of the last gadget. If the program is statically compiled, this does not raise any issue. However, in most operating systems, the program is compiled dynamically, which results in every `ecall` instructions to be located in the libraries. Consequently, in order to find the address of such an instruction, we must outsmart the *Address Space Layout Randomization* (ASLR), which loads the linked libraries at random addresses. Randomized libraries are then linked to the program through the *Procedure Linkage Table* (PLT), in which the dynamic loader (`ld.so`) stores the randomized addresses of each external function called by the program. The PLT itself is always stored in the same memory area, statically known to the attacker (offset 48). Programs compiled as *Position Independent Executable* with `-fPIE` require an information leak to locate the PLT. By reading into the PLT, we compute the address of our `ecall` instruction and write it into the stack, so that the last gadget before the `ecall` will pop its address and jump on it, triggering the syscall.

If a program uses the standard C library, then an initialization function called `__libc_start_main` is systematically included in the PLT. In version 2.27 of the library, there is an `ecall` at offset 220, making a perfect candidate for the `execve` syscall. However, this instruction is not satisfactory enough for our `setuid` syscall, as

<sup>3</sup><https://www.lurklurk.org/syscalls.html>

```

.base
8   slti   t2, t2, 225
24  slti   t2, t2, 225    //t2:=1
40  slti   t2, t2, 225    //NOP
48  .plt_address+1823
56  slti   a1, t2, -1999  //a1:=0
72  mul    a4, t2, sp     //a4:=.base+80
88  slti   t2, t2, -1999  //t2:=0
104 slti   a2, t2, -1999  //a2:=0
120 addi   a4, a4, -1278
136 addi   a4, a4, 1275   //a4:=.base+77
152 addi   t2, t2, -31    //t2:=-31
168 ld     s6, -29(a4)    //s6:=.plt+1823
184 ld     s6, -1823(s6)
    //s6:=.__libc_start_main@libc
200 addi   t1, s6, -1823
208 .ecall1_offset+1823
216 addi   s11, t1, s2
    //s11:=.setuid@libc:34
232 sd     s11, 315(a4)   //.base+392<-s11
248 addi   s3, a4, 363   //s3:=.base+440
264 sd     s3, 307(a4)
    //.base+384<-.base+440
280 sd     s3, 363(a4)
    //.base+440<-.base+440
296 addi   t1, s6, -1823
304 .ecall2_offset+1823
312 addi   s11, t1, s2
    //s11:=.setuid@libc:38
328 sd     s11, 411(a4)  //.base+488<-s11
344 addi   t2, t2, -31   //t2:=-62
360 addi   t2, t2, -31   //t2:=-93
376 sltiu  a0, t2, 2017  //a0:=0
384 0     // .base+440
392 0     //ecall1 at .setuid@libc:34
440 0     // stack canary
456 addi   a7, t2, 314   //a7:=221
472 addi   a0, a4, 67    //a0:=.base+507
488 0     //ecall2 at .setuid@libc:38
507 "/bin/sh"

```

**Figure 6: High level assembly description of the ROP chain. The first column describes the offset (in bytes) relative to the beginning of the ROP chain. Each instruction is hidden inside a gadget. As call-frames are at least 16 bytes long, we omit the offsets of null value. The notation with a leading dot .xxx@yyy:off designates the address of xxx in yyy at offset off. The notation <- designates a memory store, and := an assignment. The .ecall1\_offset+1823 indicates the location where we put the offset of the ecall instruction in the setuid function of the C library relative to the \_\_libc\_start\_main function. Similarly, the .plt\_address+1823 indicates the location where the PLT address should be inserted.**

we need to continue the execution of our ROP chain after invoking the syscall. Here, the candidate is part of an infinite loop.

One may think that jumping at the beginning of the `setuid@libc` function of the C standard library may be a good idea. This is definitely not, as the function inserts its own call frame into the stack, based on the value of `ra` at its entry. Since we already use `ra` to hijack the control flow with `ret` instructions, the function would return at its beginning, causing an infinite loop. Jump and link instructions that could modify `ra` are inadequate as well, inasmuch as they are detectable by Galileo.

Our solution involves jumping directly into the middle of the `setuid@libc` function, making use of the instruction that sets register `a7` to 146 immediately followed by the `ecall`. As a downside, we now must bypass gcc's *stack protector* (SP), that enforces backward-edge control-flow integrity, obliging the function to return to its caller. Concretely, it checks whether the call frames have been tampered with by generating a random number, the *canary*, at the beginning of the function, and storing it in two different locations. During the restore sequence, the two values are compared, and, if different, the program aborts.

Howbeit, the other location at which the canary is stored is pointed to by `s0`, which happens to be a callee-saved register, also used by gcc as a frame pointer. Hence it may be possible to obtain a gadget whose restore sequence pops `s0` from the stack, which allows hijacking the canary. We do so by writing at offset 384, which smashes the value of `s0`, thence pointing both copies of the canary to the same memory area. In this way, the canary test will always pass, as both pointers are now aliased. Finally, the gadget at offset 232 inserts into the stack the address of the `ecall` in `setuid@libc` using the location of `__libc_start_main` obtained through the PLT.

The `execve` syscall is easier to prepare. We reset `a2`, and straightforwardly set `a7` to 221. The gadget at offset 328 inserts into the stack the address of the `ecall` candidate, also in `setuid@libc`. Note that we do not need to bypass SP this time, as the `execve` syscall will spawn a new process. Finally, we take advantage of the previously leaked stack pointer (at offset 72) to set `a0` to the address of the string `/bin/sh`, located after the last call frame of our ROP chain.

## 6 ATTACK PROOF-OF-CONCEPT ON DIFFERENT PLATFORMS

In this section, we experiment our attack on two Linux operating systems, Debian and Fedora, running as a chroot environment on a HiFive Unleashed development board, featuring a quad-core Freedom U540 RV64GC processor.

### 6.1 Debian chroot on HiFive Unleashed

We first try our attack on the HiFive Unleashed board with a reduced Linux buildroot system shipped with the board. We add a Debian chroot, allowing the access of Debian features within the minimal operating system. Additionally, we create an unprivileged user, setting up the stage for our attack. Given that there is no gcc available on Debian RISC-V, we statically cross-compile the binary from another host computer. Static compilation greatly simplifies our attack, as all the libraries are now included within the



If we consider compiler-based backward-edge CFI variants like LLVM-CFI,<sup>5</sup> MCFI or Picon [14, 22, 33], the restore sequence may be hardened in a way that may not allow reusing any part of it, *e.g.* by putting the target validation guard between the return and the assignment to *ra* from the stack. This leaves us with only the last return instruction that can be jumped to from the HEP. Although we hypothesize it may be possible to assign any value to *ra* directly from the HEP, it is actually much easier to fall back on the restore sequence of another function that is not protected by compiler-based CFI, like glue-code. For the C standard library, the `__libc_csu_init` function of `crt1.o` inserted by `gcc` and `clang` is a perfect candidate, as it contains an unprotected restore sequence, even when compiled with SP (`-fstack-protector-all`) and LLVM-CFI (`-fsanitize=cfi` on `clang`).

OpenBSD has its own SP version called *RetGuard* [31], running on par with gadget reduction techniques, with the same shortcomings as `gcc`'s SP. More generally, gadget reduction techniques like G-Free [34] or code randomization [36] intend to eliminate any unaligned indirect jump, relying on canaries or backward-edge CFI to prevent malicious use of aligned branches, which is effective only against gadgets having one LCSAJ. The new gadgets presented in this paper fall out of reach of those mitigations.

To include this new class of gadgets in existing mitigation, we would have combine them with a static analysis pass verifying that every main and hidden execution path ending with an indirect jump does go through the canary check (SP) or reaches target validation (backward-edge CFI). For this purpose, we provide Algorithm 1 finding each and every execution path in a program. Its source code is available in Appendix A. It tentatively disassembles one instruction at every program byte, and checks whether it yields a valid instruction. It then inserts these valid instructions into a graph, whose nodes are defined by their addresses and the outgoing edges by the values that the program counter might take after the execution of the instruction. For example, conditional jumps may have two outgoing edges, while data processing instructions may only have one outgoing edge to the immediately following instruction in the program.

Indirect jumps (like `ret`) do not have outgoing edges as the value of the program counter may not be known statically. We mark such instructions as *Points of Interest* (or PoI, term coined in [50]), to keep only the instructions that can reach one of those PoI, loop indefinitely or trigger an invalid instruction causing the program to crash. This can equivalently be rephrased as keeping only the subgraph coreachable from those PoI. Additional work can be performed on this graph, like merging chains of nodes, yielding a *control-flow graph* (CFG) showing both the MEP and HEP. We show in Fig. 8 an example of such CFG.

We used this algorithm to find such gadgets in the C standard library. Out of the 1957 unaligned sequences ending with a fixed jump offset, only one can realistically be used as a gadget in a traditional ROP attack. The scarcity of such gadgets on RISC-V architecture confirms our need for magic constants when encoding the gadgets in Section 4. Indeed the opcode of a HEP instruction lies in the operand of the MEP instruction.

Some more radical solutions would consist in trying to prevent overlapping code in RISC-V, either by deleting the compressed instruction C extension, or by requiring 32-bit instructions to be naturally aligned, or by changing the ISA so that the length of the instruction is encoded in first bit of every half-word. Though, we may lose one bit per half-word, hampering with the range of opcodes, *i.e.* less immediates, or less registers. Furthermore, this requires extensive changes to the instruction set, and we believe that such a solution could only be implemented on next generation ISAs.

**Input:**  $B_0, \dots, B_n$ , a binary program

**Result:**  $G$ , a directed graph of all execution paths

$G \stackrel{\text{def}}{=} (V, E)$ ;

$End \stackrel{\text{def}}{=} \emptyset$ ;

**for**  $pc \stackrel{\text{def}}{=} 0$  **to**  $n$  **do**

$I := \text{Disasm\_one\_inst}(B_{pc}, \dots)$ ;

**if**  $I$  **is not** a valid instruction **then**  
        | **continue**

**end**

$V.\text{insert}(pc)$ ;

**foreach**  $pc'$  **in**  $I.\text{get\_next\_pc}()$  **do**

        |  $E.\text{insert}(pc, pc')$

**end**

**if**  $I$  **is** an indirect jump **then**

        |  $End.\text{insert}(pc)$

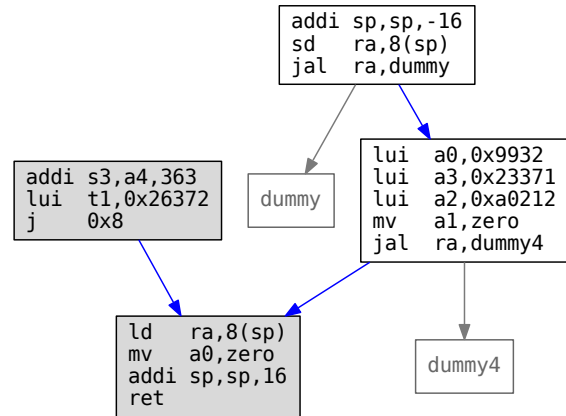
**end**

**end**

$G' \stackrel{\text{def}}{=} \text{coreachable}(G, End)$ ;

**return**  $G'$ ;

**Algorithm 1:** Disassembly algorithm finding all execution paths in a binary.



**Figure 8:** The `function15c` (first presented in Fig. 5) as shown by our disassembler. Unnecessary details such as instruction addresses or hexadecimal representations have been deleted. The gadget is highlighted in gray, and the dummy functions are shown in light-gray.

<sup>5</sup><https://clang.llvm.org/docs/ControlFlowIntegrityDesign.html>



## 8 RELATED WORK

Andriese et al. [4] have shown a method to hide malicious code using overlapping instructions in x86. It splits the code into smaller fragments and bruteforces a prefix and a suffix, for which the code fragment becomes a valid x86 MEP. This bruteforce method relies on the high density of the x86 instruction set, although it still sometimes requires manual intervention to conceal the fragments. The resulting hidden fragments are only one LCSAJ long, and always end by an indirect jump, hence easily caught by any ROP gadget searcher. Our approach allows better stealth by splitting the hidden code over several LCSAJs, for which the bruteforce method may not work anymore. We also apply our method to a RISC architecture, which does not benefit from the same code density.

ROP attacks have been subject to many academic studies since their first publication in 2007 [43] introducing the Galileo algorithm. Many variants based on the same algorithm have been published, like gadgets ending with indirect jumps [10], gadgets popping signal-contexts from the stack instead of call-frames [12], or attacks using format string vulnerabilities [38]. Amongst popular ROP gadget searchers, only two have added support for RISC-V - *xrop* and *Radare2* [3, 15], both of them implementing the Galileo algorithm, falling short of detecting this new class of gadgets. The closest to our work could be *ROPgadget* [42], which tentatively disassembles a fixed number of instructions starting from each byte of the program. This method is particularly inefficient compared to our algorithm and to Galileo, but it could find some gadgets spanning over several LCSAJs, if they are shorter than a given threshold (by default 10 instructions). Quite surprisingly, after finding them, ROPGadget discards those gadgets by default, unless passed the option `--multibr`. The algorithm that we provided comprehensively solves this aspect of gadget detection by revealing any gadget, whatever their length or number of LCSAJs is.

More recently, Borrello et al. [11] published a method to insert backdoors in programs with encrypted ROP gadgets and a small decryption procedure. While encryption methods provide a definitive proof that the malicious behavior will indeed be hidden to static analysis, this does not address the problem of detection, as the decryption procedure is not concealed, and thus may be detected by static analysis. In this paper, we provided another method for adding such backdoors, without having any unconcealed element in the program. To achieve this result, we rely on a fine understanding of how current detectors work, exploiting their inability to find gadgets spanning over more than one LCSAJ.

## 9 CONCLUSION AND FUTURE WORK

ROP attacks still pose a threat, although despite the wide deployment of dedicated countermeasures. Those protections fail to provide a satisfactory solution to these attacks, as we managed to design a new type of gadget on RISC-V, undetectable by existing tools, made of several linear-code sequences and jumps, that bypasses ASLR, DEP, stack canaries, G-Free and some compiler-based backward-edge CFI. We showed how to use such gadgets in two different attack schemes concealing a backdoor to perform privilege escalation attack on two standard Linux operating systems. Although the gadgets are written in C, we believe that it can generalize to other languages, such as JIT compilers once they become

available on RISC-V, as well as other architectures featuring code overlap.

We provided a new algorithm aiming to replace previous Galileo based algorithms, that manages to find all the hidden execution paths of a program, and not just the last LCSAJ. This algorithm may be used both for offensive and defensive purposes. However, we believe that its defensive usage is only provisional, as a definitive solution to prevent code overlap requires thorough changes in the ISA, which may only be implemented on next-generation architectures.

## REFERENCES

- [1] Martin Abadi, Mihai Budiu, Ulfar Erlingsson, and Jay Ligatti. Control-flow integrity. In *Proceedings of the 12th ACM Conference on Computer and Communications Security*, pages 340–353, New-York, NY, 2005. ACM. [http://www.cs.columbia.edu/~suman/secure\\_sw\\_devel/p340-abadi.pdf](http://www.cs.columbia.edu/~suman/secure_sw_devel/p340-abadi.pdf).
- [2] Aleph One. Smashing the stack for fun and profit. *Phrack*, 7(49), 1996. <http://phrack.org/issues/49/14.html>.
- [3] Sergi Alvarez. Radare2, 2008. <https://rada.re/n/radare2.html>.
- [4] Dennis Andriese and Herbert Bos. Instruction-level steganography for covert trigger-based malware. In *Detection of Intrusions and Malware, and Vulnerability Assessment*, pages 41–50, Berlin, Heidelberg, 2014. Springer-Verlag. [https://www.cs.vu.nl/~herbertb/papers/stega\\_dimva14.pdf](https://www.cs.vu.nl/~herbertb/papers/stega_dimva14.pdf).
- [5] ARM Limited. ARM1156T2-S technical reference manual, 2005. [http://infocenter.arm.com/help/topic/com.arm.doc.ddi0338g/DDI0338G\\_arm1156t2s\\_r0p4\\_trm.pdf](http://infocenter.arm.com/help/topic/com.arm.doc.ddi0338g/DDI0338G_arm1156t2s_r0p4_trm.pdf).
- [6] ARM Limited. ARM architecture reference manual. ARMv8, for ARMv8-A architecture profile, 2013. [https://static.docs.arm.com/ddi0487/da/DDI0487D\\_a\\_armv8\\_arm.pdf](https://static.docs.arm.com/ddi0487/da/DDI0487D_a_armv8_arm.pdf).
- [7] Atmel. AVR instruction set manual, 2016. <http://ww1.microchip.com/downloads/en/devicedoc/atmel-0856-avr-instruction-set-manual.pdf>.
- [8] Fabrice Bellard. QEMU, a fast and portable dynamic translator. In *Proceedings of the Annual Conference on USENIX Annual Technical Conference*, Berkeley, CA, 2005. USENIX Association. <https://www.cse.iitd.ernet.in/~sbansal/csl862-virt/2010/readings/bellard.pdf>.
- [9] Dion Blazakis. Interpreter exploitation: Pointer inference and JIT spraying, 2010. <http://www.semanticscope.com/research/BHDC2010/BHDC-2010-Paper.pdf>.
- [10] Tyler Bletsch, Xuxian Jiang, Vince W. Freeh, and Zhenkai Liang. Jump-oriented programming: A new class of code-reuse attack. In *Proceedings of the 6th ACM Symposium on Information, Computer and Communications Security*, pages 30–40, New York, NY, 2011. ACM. <https://www.comp.nus.edu.sg/~liangzk/papers/asiaccs11.pdf>.
- [11] Pietro Borrello, Emilio Coppa, Daniele Cono D’Elia, and Camil Demetrescu. The ROP needle: Hiding trigger-based injection vectors via code reuse. In *Proceedings of the 34th ACM SIGAPP Symposium on Applied Computing*, New York, NY, 2019. ACM. [https://www.researchgate.net/publication/329538169\\_The\\_ROP\\_Needle\\_Hiding\\_Trigger-based\\_Injection\\_Vectors\\_via\\_Code\\_Reuse](https://www.researchgate.net/publication/329538169_The_ROP_Needle_Hiding_Trigger-based_Injection_Vectors_via_Code_Reuse).
- [12] Erik Bosman and Herbert Bos. Framing signals - a return to portable shellcode. In *Proceedings of the 2014 IEEE Symposium on Security and Privacy*, pages 243–258, Washington, DC, 2014. IEEE Computer Society. [https://www.cs.vu.nl/~herbertb/papers/srop\\_sp14.pdf](https://www.cs.vu.nl/~herbertb/papers/srop_sp14.pdf).
- [13] Erik Buchanan, Ryan Roemer, Hovav Shacham, and Stefan Savage. When good instructions go bad: Generalizing return-oriented programming to RISC. In *Proceedings of the 15th ACM Conference on Computer and Communications Security*, pages 27–38, New York, NY, 2008. ACM. <https://hovav.net/ucsd/dist/sparc.pdf>.
- [14] Nathan Burov, Scott A. Carr, Joseph Nash, Per Larsen, Michael Franz, Stefan Brunthaler, and Mathias Payer. Control-flow integrity: Precision, security, and performance. *ACM Computer Surveys*, 50(1):16:1–16:33, 2017. <http://hexhive.epfl.ch/publications/files/17CSUR.pdf>.
- [15] Amat Cama. xrop: Tool to generate ROP gadgets for ARM, AARCH64, x86, MIPS, PPC, RISC-V, SH4 and SPARC, 2017. <https://github.com/acama/xrop>.
- [16] Nicholas Carlini, Antonio Barresi, Mathias Payer, David Wagner, and Thomas R. Gross. Control-flow bending: On the effectiveness of control-flow integrity. In *Proceedings of the 24th USENIX Security Symposium*, pages 161–176, Washington, DC, 2015. USENIX Association. <https://www.usenix.org/system/files/conference/usenixsecurity15/sec15-paper-carlini.pdf>.
- [17] Ping Chen, Hai Xiao, Xiaobin Shen, Xinchun Yin, Bing Mao, and Li Xie. DROP: Detecting return-oriented programming malicious code. In *Proceedings of the 5th International Conference on Information Systems Security*, pages 163–177, Berlin,

- Heidelberg, 2009. Springer-Verlag. [https://link.springer.com/chapter/10.1007%2F978-3-642-10772-6\\_13](https://link.springer.com/chapter/10.1007%2F978-3-642-10772-6_13).
- [18] Palmer Dabbelt, Stefan O'Rear, Kito Cheng, Andrew Waterman, Michael Clark, Alex Bradbury, Davod Horner, Max Nordlund, and Karsten Merker. RISC-V ELF psABI specification, 2016. <https://github.com/riscv/riscv-elf-psabi-doc/>.
- [19] Lucas Davi, Ahmed-Reza Sadeghi, and Marcel Winnandy. ROPdefender: A detection tool to defend against return-oriented programming attacks. In *Proceedings of the 6th ACM Symposium on Information, Computer and Communications Security*, pages 40–51, New York, NY, 2011. ACM. <https://people.csail.mit.edu/hes/ROP/Readings/ROPdefender.pdf>.
- [20] Stephan Esser. iOS kernel exploitation, 2011. [https://media.blackhat.com/bh-us-11/Esser/BH\\_US\\_11\\_Esser\\_Exploiting\\_The\\_iOS\\_Kernel\\_Slides.pdf](https://media.blackhat.com/bh-us-11/Esser/BH_US_11_Esser_Exploiting_The_iOS_Kernel_Slides.pdf).
- [21] Hal Finkel, David Poliakoff, and David F. Richards. Clangjit: Enhancing c++ with just-in-time compilation, 2019. <https://arxiv.org/pdf/1904.08555.pdf>.
- [22] Arnaud Fontaine, Pierre Chifflier, and Thomas Coudray. PICON: Control Flow Integrity on LLVM IR. In *Symposium sur la Sécurité des Technologies de l'Information et des Communications*, Rennes, 2015. [https://www.sstic.org/2015/presentation/control\\_flow\\_integrity\\_on\\_llvm\\_ir/](https://www.sstic.org/2015/presentation/control_flow_integrity_on_llvm_ir/).
- [23] Robert Gawlik and Thorsten Holz. Sok: Make jit-spray great again. In *Proceedings of the 12th USENIX Workshop on Offensive Technologies*, Baltimore, MD, 2018. USENIX Association. <https://www.usenix.org/system/files/conference/woot18/woot18-paper-gawlik.pdf>.
- [24] David Gilbertson. I'm harvesting credit card numbers and passwords from your site. Here's how., 2018. <https://hackernoon.com/im-harvesting-credit-card-numbers-and-passwords-from-your-site-here-s-how-9a8cb347c5b5>.
- [25] IBM. Power ISA version 3.0B, 2017. [https://openpowerfoundation.org/?resource\\_lib=power-isa-version-3-0](https://openpowerfoundation.org/?resource_lib=power-isa-version-3-0).
- [26] Peter Kacherginsky. Ida sploiter, 2017. <https://github.com/iphelix/ida-sploiter/>.
- [27] Jinku Li, Zhu Wang, Xuxian Jiang, Michael Grace, and Sina Bahram. Defeating return-oriented rootkits with "return-less" kernels. In *Proceedings of the 5th European Conference on Computer Systems*, pages 195–208, New York, NY, 2010. ACM. <http://www.cs.fsu.edu/~zwang/files/eurosys10.pdf>.
- [28] Flávio Medeiros, Christian Kästner, Márcio Ribeiro, Sarah Nadi, and Rohiy Gheyi. The love/hate relationship with the c preprocessor: An interview study. In *Proceedings of the 29th European Conference on Object-Oriented Programming*, pages 999–1022, Germany, 2015. Dagstuhl Publishing. <https://www.cs.cmu.edu/~ckaestne/pdf/ecoop15.pdf>.
- [29] Microsoft. Microsoft security toolkit delivers new bluehat prize defensive technology, 2012. <https://news.microsoft.com/2012/07/25/microsoft-security-toolkit-delivers-new-bluehat-prize-defensive-technology/>.
- [30] Inc MIPS Technologies. MIPS32 architecture for programmers. Volume II: The MIPS32 instruction set, 2001. [https://www.cs.cornell.edu/courses/cs3410/2008fa/MIPS\\_Vol2.pdf](https://www.cs.cornell.edu/courses/cs3410/2008fa/MIPS_Vol2.pdf).
- [31] Todd Mortimer. Removing ROP gadgets from OpenBSD, 2018. <https://www.openbsd.org/papers/eurobsdcon2018-rop.pdf>.
- [32] Nergal. Advanced return-into-lib(c) exploits (PaX case study). *Phrack*, 11(58), 2001. <http://phrack.org/issues/58/4.html>.
- [33] Ben Niu and Gang Tan. Modular control-flow integrity. In *Proceedings of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 577–587, New York, NY, 2014. ACM. <http://www.cse.psu.edu/~gxt29/papers/mcfi.pdf>.
- [34] Kaan Onarligolu, Leyla Bilge, Andrea Lanzi, Davide Balzarotti, and Engin Kirda. G-Free: defeating return-oriented programming through gadget-less binaries. In *Proceedings of the 2010 Annual Computer Security Applications Conference*, pages 49–58, New York, NY, 2010. ACM. <http://www.eurecom.fr/fr/publication/3235/download/rs-publi-3235.pdf>.
- [35] Vasileios Pappas. Defending against return-oriented programming, 2015. <https://academiccommons.columbia.edu/doi/10.7916/D8CZ35VH>.
- [36] Vasilis Pappas, Michalis Polychronakis, and Angelos D. Keromytis. Smashing the gadgets: Hindering return-oriented programming using in-place code randomization. In *Proceedings of the 2012 IEEE Symposium on Security and Privacy*, pages 601–615, Washington, DC, 2012. IEEE Computer Society. <https://www.doc.ic.ac.uk/~livshits/classes/CO445H/reading/smashing-rop.pdf>.
- [37] David A. Patterson and Carlo H. Sequin. RISC I: A reduced instruction set VLSI computer. In *Proceedings of the 8th Annual Symposium on Computer Architecture*, pages 443–457, Los Alamitos, CA, 1981. IEEE Computer Society. [http://www.cecs.pdx.edu/~alaa/ece587/papers/patterson\\_isca\\_1981.pdf](http://www.cecs.pdx.edu/~alaa/ece587/papers/patterson_isca_1981.pdf).
- [38] Mathias Payer and Thomas R. Gross. String oriented programming: When ASLR is not enough. In *Proceedings of the 2nd ACM SIGPLAN Program Protection and Reverse Engineering Workshop*, pages 2:1–2:9, New York, NY, 2013. ACM. <https://nebelwelt.net/publications/files/13PPREW.pdf>.
- [39] Manish Prasad and Tzi cker Chiuieh. A binary rewriting defense against stack based buffer overflow attacks. In *USENIX Annual Technical Conference*, pages 211–224, San Antonio, TX, 2003. USENIX Association. [https://www.usenix.org/event/usenix03/tech/full\\_papers/prasad/prasad.ps](https://www.usenix.org/event/usenix03/tech/full_papers/prasad/prasad.ps).
- [40] Marco Prati. Rop gadgets hiding techniques in open source projects, 2012. <http://ams.laurea.unibo.it/4682/>.
- [41] Tiernan Ray. Nvidia, Western Digital at Chips' Frontier, 2018. <https://www.barrons.com/articles/nvidia-western-digital-at-chips-frontier-1516640945>.
- [42] Jonathan Salwan. ROPgadget: Gadgets finder and auto-roper, 2011. <https://github.com/JonathanSalwan/ROPgadget/>.
- [43] Hovav Shacham. 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*, pages 552–561, New York, NY, 2007. ACM. <https://hovav.net/ucsd/dist/geometry.pdf>.
- [44] Yan Shoshitaishvili, Chris Salls, and Nick Stephens. Angrop, 2014. <https://github.com/salls/angrop>.
- [45] SPARC International, Inc. The SPARC architecture manual, version 8, 1991. <https://www.gaisler.com/doc/sparcv8.pdf>.
- [46] Sam L. Thomas and Aurélien Francillon. Backdoors: Definition, deniability and detection. In *Proceedings of the 21st International Symposium on Research in Attacks, Intrusions and Defenses*, pages 92–113, Berlin, Heidelberg, 2018. Springer-Verlag. [http://s3.eurecom.fr/docs/raid18\\_thomas.pdf](http://s3.eurecom.fr/docs/raid18_thomas.pdf).
- [47] Niek Timmers, Albert Spruyt, and Marc Witteman. Controlling PC on ARM using fault injection. In *Proceedings of the 2016 Workshop on Fault Diagnosis and Tolerance in Cryptography*, pages 25–35, Santa Barbara, CA, 2016. IEEE Computer Society. <https://www.riscure.com/uploads/2017/09/Controlling-PC-on-ARM-using-Fault-Injection.pdf>.
- [48] Vasil Vasilev, Philippe Canal, Axel Naumann, and Paul Russo. Cling – the new interactive interpreter for ROOT 6. *Journal of Physics: Conference Series*, 396(5):052071, 2012. <https://iopscience.iop.org/article/10.1088/1742-6596/396/5/052071>.
- [49] Andrew Waterman and Krste Asanović. The RISC-V instruction set manual, volume I: User-level ISA, document version 2.2, 2017. <https://content.riscv.org/wp-content/uploads/2017/05/riscv-spec-v2.2.pdf>.
- [50] Patrick Wollgast, Robert Gawlik, Behrad Garmany, Benjamin Kollenda, and Thorsten Holz. Automated multi-architectural discovery of CFI-resistant code gadgets. In *Proceedings of the 21st European Symposium on Research in Computer Security*, pages 602–620, Heraklion, 2016. IEEE Computer Security. [https://www.syssec.ruhr-uni-bochum.de/media/emma/veroeffentlichungen/2016/10/04/esorics2016\\_XOP.pdf](https://www.syssec.ruhr-uni-bochum.de/media/emma/veroeffentlichungen/2016/10/04/esorics2016_XOP.pdf).
- [51] Dino Dai Zovi. Practical return-oriented programming, 2010. <http://repository.root-me.org/Exploitation%20-%20Sys%20C%3A8me/Microsoft/EN%20-%20Practical%20Return%20Oriented%20Programming.pdf>.

## A SOURCE CODE AND ARTIFACT

The C source code used for generating the gadgets, as well as the images of the Fedora and Debian virtual machines are available on the following link: [ec2-13-58-4-171.us-east-2.compute.amazonaws.com](https://ec2-13-58-4-171.us-east-2.compute.amazonaws.com) using password "AsiaCCS".