



# RISC-V: #AlphanumericShellcoding

Hadrien Barral<sup>1</sup>, Rémi Géraud-Stewart<sup>1</sup>, Georges-Axel Jaloyan<sup>1,2</sup>, and David Naccache<sup>1</sup>

<sup>1</sup>DIENS, École normale supérieure, CNRS, PSL University, Paris, France

<sup>2</sup>CEA, DAM, DIF, F-91297 Arpajon, France

## Abstract

We explain how to design RISC-V shellcodes capable of running arbitrary code, whose ASCII binary representation use only letters `a-zA-Z`, digits `0-9`, and either of the three characters: `#`, `/`, `'`.

## 1 Introduction

RISC-V [22] is a new *Instruction Set Architecture* (ISA) which development began in 2010. It is based on the concept of *Reduced Instruction Set Computer* (RISC) [17], 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 [1], Atmel AVR [2], MIPS [16], Power [11], and SPARC [21].

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-bit. While only test boards feature RISC-V processors for now, many companies including Western Digital or Nvidia have announced the use of RISC-V chips in their future products [19].

This increasing popularity makes RISC-V an attractive target for low-level attackers, whose tools are traditionally honed chiefly against x86 platforms. This is exacerbated by the widespread adoption of smartphones and their use for almost any task from payment to dating to unlocking one's car, making successful attacks especially profitable. At the same time, mobile environments are improving their security, in part to address such threats.

Nevertheless, mobile applications are still occasionally vulnerable to memory safety vulnerabilities, and the need for performance or obfuscation often drives developers

to implement low-level (e.g., JNI) segments which are particularly susceptible to the usual techniques of buffer overflow exploitation. Unlike traditional local or network scenarios however, the attacker has only limited ways to transmit a payload.

We claim that a reasonable vector is text-based applications, which includes SMS, social networks, chat applications (in a remote context), password entry, note taking, or QR code scanning (in a local context). This being said, the attacker's payload has to be treated by this application as text such as a hashtag, a URL, a sentence, in the most restrictive sense, hence the most widely applicable. We therefore consider alphanumeric programs whose binary representation use only the alphanumeric ASCII characters: the 52 lowercase and uppercase letters of the English alphabet and the 10 digits. As we will discuss, it is only possible to achieve *arbitrary* code execution at the cost of allowing one additional character: either `#`, `/`, or `'`, each being compatible with the use-cases discussed above.

### 1.1 Prior and related work

This work follows a trend initiated in the early 2000s to evade buffer overflow protections (Eller [10] and RIX [20] on x86) and intrusion detection systems [12]. Tools to generate alphanumeric shellcodes on the x86 platform appeared [4] and are now a standard component of attack frameworks including Metasploit (`msfvenom`) and UPX1. The x86 is particularly well suited to this exercise as many letters materialize into `mov` instructions, which form a Turing-complete subset of operations [8]. To this day however neither of these tools are able to generate alphanumeric shellcodes on ARM platforms, such as ARMv8 and RISC-V.

The first automated tool was provided by Younan et al. in 2011 for the ARMv5 platform, relying on an BF interpreter and bytecode [23]. The technique however does not carry over to more recent implementations.

In 2016, Barral et al. introduced the first tool capable of compiling arbitrary ARMv8 code into alphanumeric executable code [3]. This is a *tour de force* but also and most importantly it introduces a generic approach to design such tools.

To the best of the authors’ knowledge, none of the currently available approaches works on RISC-V.

## 1.2 Our contribution

We provide, as far as the authors know, the first analysis of alphanumeric code on RISC-V, as well as a complete framework for automatically generating alphanumeric (+1 character) shellcodes. Through a three-staged modular design, these shellcodes achieve arbitrary code execution on this platform.

This is the second architecture which can be addressed using the methodology from [3], which is an argument in favor of such generic approaches (rather than *ad hoc* ones). Our approach differs on the fact that we do not manually assemble available instructions into higher level constructs for building the unpacker in a bottom-up fashion and instead opt for a partially automated strategy to generate the required alphanumeric instruction sequences to achieve the desired results.

We provide three different constructions, corresponding to each choice of an additional character. All our programs are given in appendix, being to the best of the authors’ knowledge the first automated tool of this kind for RISC-V, as well as the first examples of such shellcodes for each construction.

## 2 Background

### 2.1 Shellcodes and exploitation

In a typical arbitrary code execution (ACE) scenario, attackers can run a relatively short program of their choosing. It is called a *shellcode*, as it can start a shell session, which in turn allows attackers to download and run additional programs.

For instance, a stack overflow ACE can happen when an application allows writing in an array beyond the allocated space for this array, resulting in overwriting stack frame data. In platforms such as x86 the stack frame stores information about the instruction pointer before a call; by overwriting this information an attacker can control the instruction pointer and send it back to the array’s address. The array’s contents are then executed as if they were the vulnerable program’s own instructions: this is where the shellcode is written.

Since a typical array is relatively short, shellcodes must accordingly be concise. Similarly, an application may

restrict what data it manipulates (e.g., strings) and shellcodes must be written to comply with such constraints. Additional protections make shellcode design trickier: *Address Space Layout Randomization* (ASLR), stack-smashing protections, or non-executable stack space for instance. Detection mechanisms may furthermore identify characteristic aspects of a shellcode and prevent the attack from reaching the target application. For all these reasons the modern shellcode designer has to navigate around layers of obstacles.

This difficulty is somewhat offset on embedded and mobile devices, where many protections are only partially implemented, if at all. Such platforms also host to many third-party applications, that can be developed without strictly adhering to secure coding practices, using memory-unsafe languages (sometimes due to performance or obfuscation constraints) and not necessarily updated in a timely fashion.

### 2.2 The RISC-V instruction set

RISC-V splits its instruction set between a mandatory core set (RV64I) and different optional extensions, each of which is designated by a string (a single letter for the most common ones). 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), vector operations (V), ...

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*” [7].

The RV64GC ISA features 32-bit and 16-bit instructions, aligned on 16 bits. 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.

We adopt for the rest of this paper some terminology defined by the RISC-V Instruction Set Manual, Version 1.10 [22]. Assembly instructions are written in the format `add x1,x2,x3`, where `add` is called the *opcode*, and `x1`, `x2`, `x3` are the *operands*. Precisely, `x1` is the *destination register*, `x2` is the *first source register* and `x3` is the *second source register*. When one of the source registers is replaced by a constant, it is called an *immediate*. To those conventions, let  $K$  be a register, we add our slicing notation as  $K[y : x]$  (with  $x < y$ ), meaning we take a slice of bits  $x$  to  $y$  of  $K$ , with the lowest bit denoted as the bit 0.

RISC-V ELF psABI specification [6] provides a register naming convention, reproduced in Table 1.

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

Table 1: Naming convention for registers, per psABI [6].

### 3 Alphanumeric RISC-V

The first step towards building an alphanumeric shellcode for RV64GC consists in generating the subset of alphanumeric valid instructions, which we denote by  $\alpha$ RV64GC. For this purpose, we generated every 16-bit and 32-bit alphanumeric sequence, and tentatively disassembled it using `objdump`. Per RISC-V Instruction Set Manual, 16-bit instructions must have their two least significant bits set to 00, 01 or 10. Similarly, 32-bit instructions must have their five least significant bits set to `bbb11`, with `bbb` different from `111`.

Furthermore, some opcodes may encode invalid or unimplemented instructions. For instance, the little-endian word `700T` corresponds to a load upper immediate (`lui`), whereas `W00T` does not correspond to any valid RV64GC instruction, although its least significant bits are those of a valid 32-bit instruction:

```
700T    0x374f4f54    lui t5,0x544f4
W00T    0x574f4f54    undefined
```

After filtering out all invalid sequences, we regroup the remaining instructions according to their opcode, providing an overview of the available instructions for which there are some operands making them alphanumeric.

The internal structure of the instruction defines the main constraints on the alphanumeric language subset. Each 32-bit instruction has its opcode encoded in the first 7 bits of the first byte. Requiring the first byte to be alphanumeric will therefore greatly reduce the available opcodes, while providing a wide range of operands for each opcode. On the contrary, 16-bit instructions are more entropic in their spread. Henceforth, more opcodes are available, with less operands for each opcode. Consequently, the expressiveness of  $\alpha$ RV64GC relies on the intelligent combination of instructions of various lengths.

Hereafter, we provide a review of those instructions, by explaining their semantics and some insight on the avail-

able operands. For simplicity and following the methodology introduced by Barral et al. in [3], we cluster instructions as *control-flow*, *data processing*, and *memory manipulation* instructions.

#### 3.1 Data processing

Data processing includes every instruction that does not modify the memory or the program counter. Two variants may be available for each instruction, either operating on the usual 64-bit registers or performing the operation on 32 bits and sign-extending the result to the 64-bit register. Using 32-bit variants for pointer manipulation prevents from reaching addresses ranging from `0x8000 0000` to `0xFFFF FFFE FFFF FFFF`. This is a serious caveat for bare-metal shellcodes — as existing boards often have the DRAM start at `0x8000 0000` — forcing us to use the 64-bit variant. Hereafter, we only present the most useful ones, omitting instructions which may have odd effects (like micro-architectural hints for branch predictors):

- The addition `addi` instruction enables adding or removing only some specific values multiple of 16 to `sp`. Its 32-bit signed variant `addiw` is also available, and allows increasing or decreasing registers `a0`, `a2`, `a4`, `a6`, `s0`, `s2`, `sp`, `t1`, and `tp` by a value ranging from  $-20$  to  $30$ .
- The instruction `li`, allows loading signed integers (ranging from  $-20$  to  $30$ ) into registers `sp`, `tp`, `s0`, `s2`, `s4`, `s6`, `s8`, `s10`, `t1`, `t3`, `t5`, `a0`, `a2`, `a4`, and `a6`. We use the letter  $\mathcal{S}$  to designate this set of registers.

Loading immediates to registers may also be done with the `lui` instruction (load upper immediate), which loads a 20-bit signed immediate into the bits 31-12 of a register in  $\mathcal{S}$ . The lowest bits are all set to zero, while the 32 highest bits are computed as the sign-extension of the immediate. We counted 238,791 alphanumeric `lui` instructions, with a large choice of immediates.

- Bitwise manipulation: only the `sra` (shift right arithmetical) instruction is available, with all registers of  $\mathcal{S}$  as source and destination, and registers `s3-s7` as shift amount.
- Floating-point operations: many useful floating-point operations are available in  $\alpha$ RV64GC, in simple, double and quad precision. Among them we find sign manipulation like `fabs` (absolute value) or multiply-accumulate `fmadd` and its variants ( $r \leftarrow \pm a \times b \pm c$ ).
- Control-status register manipulation: many instructions available, such as `csrc`, `csrci`, `csrrc`, `csrrci`, `csrrsi`, `csrrwi`, `csrwi`, not detailed here. As privileged access may be required, we preferred not using

them and instead use the other available data processing instructions.

## 3.2 Control-flow instruction

Both conditional and unconditional jump instructions are available. For unconditional branching, we have both `j` (jump) and `jal` (jump and link) available, with the possibility to link any register of  $\mathcal{S}$ . Conditional branches are also available, with a wide variety of branching conditions: `bgtz`, `ble`, `bleu`, `blez`, `blt`, `bltu`. No backward jump is available, as the immediate offset has its sign set on the highest bit of a byte (hence always equal to zero when alphanumeric). This may prevent the Turing completeness of  $\alpha\text{RV64GC}$ , as no unbounded computation mechanism is available without additional assumptions, such as code-reuse or self-modifying code.

## 3.3 Memory processing

We have both 32-bit `lw` and 64-bit `ld` loads as well as double-precision floating-point `fld` loads. However, no stores are available, which makes it impossible to write arbitrary shellcodes: we are only able to modify the registers and not the machine's memory state.

This turns out to be a strong limitation as for instance the shellcode designer cannot build paths (such as `"/bin/sh"`) in memory (this is not an alphanumeric string). Thus, additional assumptions must be made, either by finding gadgets able to write to memory or by reusing memory previously set to the desired value at a known position (e.g., an environment variable). Either option seems unsuitable in the context of a self-contained shellcode.

We therefore consider the possibility of allowing one non-alphanumeric character — a choice which may be governed by operational constraints as well. Among all ASCII-printable instructions modifying memory, only three non alphanumeric characters stand out: slash `/`, hash `#`, and tick `'`.

- Adding the hash character `#` gives standard 32-bit `sw` and 64-bit `sd` store instructions. The 32-bit store `sw` provides the ability to store almost any variable to various addresses with offsets multiple of 32. Given that there is no possibility to increment a 64-bit register by less than 16 (using `addi`), many memory areas are out of reach. The 64-bit variant `sd` seems more promising: indeed, the available offsets are only 2 bytes apart. Using this, we can efficiently store data by using `addi` increments for coarse-grained pointer manipulation and reaching the exact store address (up to a precision of 2 bytes) by tweaking the offset of `sd`.

- Adding the slash character `/` provides some atomic instructions, such as 32-bit and 64-bit atomic read-modify-write variants of binary conjunction `amoand` and disjunction `amoor`. As an example, `amoor.d t1,s5,(sp)` loads 64 bits from the address in `sp` into `t1`, and stores in the same address the disjunction of `t1` and `s5`. Note that the addresses passed to atomic operations must be naturally aligned, which adds further complexity when designing our shellcode.
- Adding the tick character `'` provides floating-point store instructions `fsd`, `fsq`, `fsw`. Controlling the stored values requires deep technical knowledge of floating-point binary representation, as the associated data manipulation operations are of the form  $\pm a \times b \pm c$  (e.g., `fmadd`, `fmsub`).

For each of these three characters, we define a new subset of  $\text{RV64GC}$ , denoted respectively `#RV64IC`, `/RV64IAC` and `'RV64IDC`. The following section details how we can achieve ACE in `#RV64IC`, setting up the stage and much of the machinery for shellcodes in `/RV64IAC` and `'RV64IDC` as well. Since these require additional work, they are discussed later on.

## 4 High-level design

Several approaches can be used to run arbitrary code from an instruction-limited shellcode. The main techniques available are: virtualization, compilation, and packing.

Virtualization, as used by Younan et al. for 32-bit ARMv7 alphanumeric shellcoding [23], requires the design of a bytecode and an interpreter, both compatible with the limited instruction set, and powerful enough to mount a realistic attack — beyond Turing-completeness, we need to perform system calls or other mechanisms to evade the virtual environment. Virtualization presents a huge runtime overhead as well as a committed engineering effort.

Compilation, when applicable, is very efficient: compilers such as *movfuscator* [8, 9] and *higher subleq* [14] have been provided for *one instruction set computers*, reduced ISA subsets made of only one instruction. However, such methods are not applicable to  $\alpha\text{RV64GC}$  as they often rely on *syntax-directed translation schemes*. Here, the heavy constraints in  $\alpha\text{RV64GC}$  on the instruction operands hinders such methods that systematically translate each grammar symbol into the target language. Furthermore, writing compilers is in itself a daunting task. Perhaps for these reasons, to the best of our knowledge, no work on compilation for alphanumeric shellcoding has been published.

Packing is the third method, and by far the most common approach in shellcoding. This typically results in



multi-staged shellcodes, where one stage decodes a second stage which is then executed. Packers can provide additional functionalities such as encryption, which we do not explore here. However, this technique requires the ability to execute self-modifying code, which may be hindered by the presence of *executable space protection* mechanisms like DEP [15], PaX [18] or NX-bit [13]. Moreover, self-modifying code raises cache issues which need to be handled on a target-specific basis.

We decided to follow this third approach: it is conceptually simpler, much easier to check for correctness, and well suited to our target platform.




Figure 1: General structure of stage 1: an initialization section, with a forward jump over the data-pool that contains the encoded final payload  $\mathcal{P}_{enc}$ , and the unpacker  $\mathcal{U}_1$ . The location at which the stage 2 is unpacked is highlighted in grey.

## 5 Detailed construction

In this section we show how to achieve arbitrary code execution, by detailing each step of the #RV64IC version of the shellcode. Building on the foundations laid with #RV64IC, we achieve similar results in /RV64IAC and 'RV64IDC.

As explained in Section 4, we use a packing multi-staged design. We present a three-stage approach:

- The first stage is a *specific* unpacker written in #RV64IC;
- The second is a *general* unpacker written in a slightly larger subset of RV64IC;
- The third is our arbitrary payload.




Figure 2: General structure of stage 2: an initialization section, with a loop decoding at each iteration one byte of the final payload  $\mathcal{P}$  using two bytes of the encoded payload  $\mathcal{P}_{enc}$ . It finally jumps to the decoded payload, highlighted in grey.

The rationale for using three stages is governed by #RV64IC not containing backjumps, therefore forcing us to unroll the decoding logic. This would result in unwieldy large shellcodes if there were only two stages. Instead, we use the first unpacker  $\mathcal{U}_1$ , whose structure is shown in Figure 1, to unpack a minimal program  $\mathcal{U}_2$  shown in Figure 2. The program  $\mathcal{U}_2$  has backward jumps and can therefore efficiently implement a decoder using a loop.  $\mathcal{U}_2$  unpacks and execute the third stage, which is the payload  $\mathcal{P}$ .

### 5.1 Stage 1

$\mathcal{U}_1$  is an unpacker for the next stage. It is fully written in #RV64IC. As no backward jumps are available, the unpacker is written as a *straight-line program*.

Specifically,  $\mathcal{U}_1$  must: (1) locate the shellcode and jump over the encoded payload; (2) fix-up the store pointer; (3) unpack stage 2; (4) jump to the decoded stage 2.

We achieve (4) simply by placing the decoded stage 2 immediately after  $\mathcal{U}_1$ 's last instruction. The other steps are detailed below:

### 5.1.1 Locating the shellcode and jump over the encoded payload

To make the shellcode position independent, we find its absolute position in memory using the jump and link (`jal`) instruction which stores the program counter to a user-specified register. This instruction consequently increases the shellcode’s size by jumping over a large memory region. Yet, this area is not entirely wasted, as we repurpose it to store our packed payload  $\mathcal{P}$ .

### 5.1.2 Fixing-up the store pointer

The next step consists in setting up the register `XI` containing the address at which we will write stage 2. For this purpose we use the absolute address obtained in Section 5.1.1, to which we add a constant using several `addi` instructions. We must not forget the additional offset required when using the `sd` store instruction in the decoder. Consequently, stage 2 will be unpacked immediately after the shellcode.

The biggest immediate available for the `addi` instruction in `RV64GC` is 464. Since the shellcode is much longer, we use the following trick: we first append several `addi XI, XI, 464` instructions until we exceed the desired value. Then we replace some immediates in the sequence by the second greatest available immediate, i.e. 448, which reduces the total sum, until the desired value is reached.<sup>1</sup> In this way, we are guaranteed to use the least amount of `addi` instructions possible.

### 5.1.3 Unpacking stage 2

We then unpack stage 2 starting at `XI + store_offset`, where `XI` is the register we set previously. This is done sequentially, using the `sd` instruction with carefully chosen offsets. Indeed, we have many offsets only 2 bytes apart. In our case, we chose a long chain of offsets (available from our constrained instruction set), each exactly 2 bytes apart, 1920, 1922, ..., 1938. This allows storing at most 20 consecutive bytes by first loading 2 bytes into a register and then storing them into memory. We use a precomputed table providing for each immediate the minimal sequence of instructions needed for loading it to a given register. We explain below how to compute this table. To store more than 20 bytes, we increment `XI` (using the `addi XI, XI, 16` instruction) between each batch of 16 bytes, and continue with offsets 1924, ..., 1938. The whole stage 2 is 40 bytes long, unpacked in 3 batches of 20, 16, and 4 bytes.

The above strategy relies on a precomputed table of sequences for achieving arbitrary 2-byte loads. We generate this table using a depth-first search strategy, by

<sup>1</sup>A small `NOP` sled of at most 16 bytes may be required for getting an exact match.

iterating over `#RV64IC` instruction sequences and storing the reached values. This approach yields for each 2-byte immediate the shortest sequence required to load it into a register.

More precisely, the first instruction of the sequence is a `lui` (loads an immediate in bits 12 to 31 of the destination register). It is followed by an arithmetical right-shift `sra` instruction (unless the shift amount is null). By intersecting the set of possible registers which may be used both as destination register for `sd` and `lui`, we end up with registers `s4`, `s6`, `t1`, `tp`. As `sra` requires a register as a shift amount, we also iterate over all possible load immediate `li` and `addiw` subsequences to get the desired shift amount.

The next instructions of the sequence are made of `addiw` instructions, with immediates ranging from -20 to 30. We limit the exploration of the instruction sequence space to at most 4 `addiw` instructions, to keep  $\mathcal{U}_1$  compact. This limitation still grants the possibility to load 63448 out of the 65536 possible values (or 96 %) into `s4`, `s6`, `t1`, or `tp`.

In this way, we can design our stage 2 with a substantially expanded set of available instructions. Indeed, we merely require to make sure every pair of bytes in stage 2 could be loaded from an instruction sequence in the table.

## 5.2 Stage 2

Stage 2 ( $\mathcal{U}_2$ ) is more straightforward. It consists of some initialization code followed by a loop whose body decodes two consecutive bytes of  $\mathcal{P}_{enc}$ , the encoded payload. The full implementation can be found in Appendix B. The initialization code sets three registers — the *reading pointer* `XP` pointing to the encoded payload, the *writing pointer* `XQ` pointing to the start of the decoded payload, and the *end pointer* `XS` pointing to the end of the decoded payload. For simplicity,  $\mathcal{U}_2$  performs in-place decoding, meaning that `XP` is initially equal to `XQ`.

We also flush the instruction cache with a `fence.i` instruction, which is required as we modify executable memory. We discuss later in Section 6 the assumption that the first `fence.i` is not shadowed in the instruction cache.

Since 63 characters are available, it is theoretically possible to encode almost 6 bits of the payload in a single alphanumeric byte of the shellcode. However, to keep  $\mathcal{U}_2$  short, we decided to encode only 4 bits per alphanumeric byte. This spreads each byte of the payload over 2 consecutive alphanumeric characters. As stage 2 is unpacked sequentially by the first stage, we need to make stage 2 the shortest possible, even if this makes the encoder more complex. Indeed, any additional length here would lead to a significant increase in stage 1 size.

Let  $K$  be the byte stored at  $XP + 1$ ,  $L$  the byte stored at  $XP$  and  $A$  the byte written at address  $XQ$  by the store instruction. The decoding algorithm we devised only requires 5 instructions in the body of its loop.

```
lw    XS, 4(XP) # Load K and L bytes
# XS == 0x????K[4:7]K[0:3]L[4:7]L[0:3]
mv    XT, XS    # Duplicate value
srli  XT, XT, 4 # Shift right by 4
# XT == 0x????K[4:7]K[0:3]L[4:7]
xor   XS, XS, XT # XS := XS ⊕ XT
# XS == 0x????A[4:7]A[0:3]
sw    XS, 0(XQ) # Store decoded byte A
```

Hereafter, we find the encoding formulae by solving the decoding equations. Henceforth, when encoding byte  $A$ , the encoder must find values for  $K$  and  $L$  so that:

$$\begin{aligned} K \text{ and } L &\text{ are alphanumeric} \\ L[0:3] \oplus L[4:7] &= A[0:3] \\ K[0:3] \oplus L[4:7] &= A[4:7] \end{aligned}$$

One should remark that every byte of the form  $0x4*$  or  $0x6*$  for  $*$  non null is alphanumeric. This simplifies the resolution of the previously given constraints. The following solution can be checked to give an alphanumeric encoding for any input byte.

$$\begin{aligned} L[4:7] &= 0x4 \text{ if } A[0:3] \neq 0x4 \text{ else } 0x6 \\ L[0:3] &= A[0:3] \oplus L[4:7] \\ K[0:3] &= A[4:7] \oplus L[4:7] \\ K[4:7] &= 0x4 \text{ if } A[0:3] \neq 0x0 \text{ else } 0x5 \end{aligned}$$

Finally, as executable memory modifications occurred, we flush the instruction cache again using a `fence.i` instruction, and jump to the decoded payload  $\mathcal{P}$ .

### 5.3 Payload

Stage 3 is the payload  $\mathcal{P}$ , containing arbitrary binary code. We generate this code directly from a C source payload compiled with a standard `gcc`. The resulting binary code is then encoded as described in Section 5.2 with a lightweight PHP script.

The size of the payload is upper bounded by the offset chosen for the forward jump is Figure 1. For our needs, we deemed 1024 bytes to be sufficient, allowing us a decoded payload of 512 bytes. Note that with some minor engineering work, this maximum size can be increased. In the context of usual shellcoding attacks, the payload almost always fits into this limit. As a proof-of-concept, we test in Section 6 three different payloads for a standard Linux: a `printf("Hello world")` shellcode, an `execve("/bin/sh")` shellcode, and one that leaks the contents of `/etc/shadow`.

## 5.4 Integration/Linking

All in all, the complete shellcode is built in the following order:

1. We compute the table of minimal instruction sequences (Section 5.1.3).
2. We build the final payload  $\mathcal{P}$ , and compute its length (Section 5.3).
3. We generate stage 2, with the appropriate values for the reading pointer, the writing pointer, and the end pointer (Section 5.2 and fig. 2).
4. We generate the unpacker for stage 2, and compute its length (Section 5.1.3 and fig. 1).
5. We generate the code for fixing-up the store pointer (Section 5.1.2).
6. We then build the whole shellcode, without its encoded stage 3 payload, for which we allocated the necessary space.
7. We finally insert the encoded payload  $\mathcal{P}$  at the appropriate location in the shellcode.

## 5.5 Shellcoding in /RV64IAC

We have also created a version of the shellcode in `/RV64IAC`, using atomic store instructions instead of regular stores for unpacking in stage 1. Data is stored with the `amoord` instruction which operates on 8 naturally aligned bytes. By opposition to the previous implementation in `#RV64IC`, we do not have offsets for stores, hence we need to modify the store pointer using available `addi` instances, which can only increase a register by a multiple of 16. We thus store our decoded stage 2 in blocks of 16 bytes. As we have control over only the 8 first bytes, we decided to split them into two parts, the first four bytes containing the decoded instruction, whereas the next two bytes contain a jump instruction to the next block (`j .+0xc`, or `0x31A0` in hexadecimal). The structure of the block is shown in Figure 3.




Figure 3: Diagram of a 16-byte block. Our stage 2 instructions are located in the first two bytes, while the next two contain a NOP like instruction followed by a jump to the next block. The last 10 bytes are unused.

As required by the sequences for 2-byte arbitrary load computed in Section 5.1.3, we wrote stage 2 using only compressed instructions. The only exception is `fence.i`, which is unavoidable and does not have a compressed

version. In this case, we use a custom sequence to store its value (0x0000100F in hexadecimal). We would like to particularly thank the authors of RISC-V for the fact that the 16 highest bits of `fence.i` are all zeros, which keeps our sequence of instructions really short. Otherwise we would have required chaining many `addi` instructions, making the shellcode too long to be used in practice.

The sequences used for loading the 2-byte instructions are computed using a table similar to that of Section 5.1.3. By opposition to `#RV64IC`, here the word’s two highest bytes will be executed as an instruction. We make sure that these two bytes do not modify the high-level semantics of the program. Altogether, the table allows loading 58174 possible 16-bit values, out of 65536 (or 88 %) which still allows encoding our stage 2 with only minor modifications, at the expense of a slight size increase of only 2 bytes. The payload and the way it is encoded remains identical.

## 5.6 Shellcoding in `'RV64IDC`

Shellcoding in `'RV64IDC` is more tricky. First of all, it requires the *floating-point unit* (FPU) to be activated, which is always done by the operating system when working in a hosted environment. In the context of the bare-metal examples presented in this paper, we use a small additional piece of non-alphanumeric code, whose sole purpose consists in activating the FPU (in big-endian hexadecimal: 0x896373900330).

Similarly to `/RV64IAC`, the main difference lies in the way stage 2 is unpacked by  $\mathcal{U}_1$ . This time, we store in the data-pool some floating-point values which are used by  $\mathcal{U}_1$  during unpacking. The most general floating-point data manipulation instruction available is `fmadd r, a, b, c` (fused-multiply add): it computes  $r = a \times b + c$ . The store operation `fsd` then stores  $r$  at the desired memory location. We thus have to solve equations of the form  $r_i = a_i \times b_i + c_i$ , where  $r_i$  is a small part of the decoded stage 2, under the constraint that each  $a_i$ ,  $b_i$  and  $c_i$  need to be loaded from the data pool. To keep our data pool as small as possible, we need to share values between different equations. As this increases the mathematical complexity of solving floating-point equations, we decided to work on a simplified version of the problem, in which we only encode 6 bytes of stage 2 into  $r_i$ . Indeed, in this way, the constraint lies only in the mantissa of the floating-point. Furthermore, we fixed the two remaining bytes of  $a_i$ ,  $b_i$  and  $c_i$  to alphanumeric constants which do not propagate carries to the exponent when performing the operation. These simplifications turned out to be sufficient for solving the equations while reducing the total number of different floating-point constants.

Hereafter, we present some of the methods we used for reducing the number of different floating-point constants

used. The first, consists in using the same  $b_i$  for all equations, as, without loss of generality, this does not impede finding a solution by just modifying  $a_i$  and  $c_i$ . This simplification allows solving the equation by testing random alphanumeric values for  $a_i$ , computing the adequate  $c_i$  then checking both  $a_i$  and  $c_i$  are alphanumeric. A simple combinatorial analysis gives us the approximate probability that a randomly chosen alphanumeric  $a_i$  gives an alphanumeric solution for  $c_i$  as:  $(\frac{62}{256})^6 \simeq \frac{1}{50000}$ .

The second method consists in using the same  $a_k$  for two consecutive equations. Formally, we require to find solutions  $a_k, c_{2k}, c_{2k+1}$  for the following set of equations:

$$\begin{aligned} r_{2k} &= a_k \times b + c_{2k} \\ r_{2k+1} &= a_k \times b + c_{2k+1} \end{aligned}$$

Unfortunately, these equations are not guaranteed to always have solutions. Indeed, let  $r_{2k}$  and  $r_{2k+1}$  differ in their highest bit. This means the highest bits of  $c_{2k}$  and  $c_{2k+1}$  are different.<sup>2</sup> Hence one of them is non-alphanumeric. The solution we came up with consists in making stage 2 polymorphic, and trying to solve those equations for all instances of stage 2, hoping to find one for which all the equations have a solution. The different stage 2 instances are generated by either modifying the registers (150k variants), reordering some initialization instructions for the loop (6 variants), or reordering the pointer increment instruction in the loop’s body (7 variants); yielding a total of about 6 millions stage 2 instances.

Algorithm 1 uses memoization to speed up the resolution of equations. In the worst case, the first loop has 12 million iterations (which can be executed in parallel), the second has 4 iterations while the last has 2 million iterations. In practice, when taking into account memoization, we counted  $2.3 \times 10^{11}$  iterations, requiring 1.5 execution hours on a 4-core Atom 2 GHz CPU. Eventually, we found several instances for which all equations had a solution. The rest of the shellcode is built in the same fashion as the previous versions presented in the previous sections.

## 6 Evaluation

### 6.1 QEMU

We initially tested our 3 shellcodes on QEMU [5], a widespread open-source emulator. It emulates a HiFive Unleashed RV64GC development board, without some of its micro-architectural features like caches or timings. The payload is expected to print “*Hello world!*”

<sup>2</sup>We omit the rare and lucky case where carry propagation still provides a solution to the equation.



**Input:**  $b$ , a 64-bit floating-point value  
**Input:**  $s_0, \dots, s_{2\ell+1}$ , the stage 2  
**Result:** a list of 64-bit floating-point values  
 $mem := \text{Array}(\text{None})$  ;  
 $P := \text{Polymorphism}(s_0, \dots, s_{2\ell+1})$  ;  
**foreach**  $r_0, \dots, r_{2\ell+1}$  **in**  $P$  **do**  
  **for**  $k = 0$  **to**  $\ell$  **do**  
    **if**  $mem[r_{2k}][r_{2k+1}]$  **is not**  $\text{None}$  **then**  
      | **continue**  
    **end**  
    **for**  $i = 0$  **to** 2000000 **do**  
       $a := \text{RandAlphanumFloatingPoint}()$   
      Solve  $c_{2k}$  in  
       $r_{2k} = a \times b + c_{2k}$   
      Solve  $c_{2k+1}$  in  
       $r_{2k+1} = a \times b + c_{2k+1}$   
      **if**  $c_{2k}$  **and**  $c_{2k+1}$  **are** alphanumeric  
      **then**  
      |  $mem[r_{2k}][r_{2k+1}] := a$   
      | **break**  
      **end**  
    **end**  
    **if**  $mem[r_{2k}][r_{2k+1}]$  **is**  $\text{None}$  **then**  
      |  $mem[r_{2k}][r_{2k+1}] := \text{NotFound}$   
    **end**  
  **end**  
**if**  $\nexists k, mem[r_{2k}][r_{2k+1}]$  **is**  $\text{NotFound}$  **then**  
  | **return**  $(mem[r_{2k}][r_{2k+1}])_{(k=0..\ell)}$   
**end**  
**end**

**Algorithm 1:** Automated testing of the existence of a solution to the sets of equations induced by a specific stage 2 encoding. The outer loop is parallelized, testing several stage 2 instances concurrently.

on the serial device mapped at address  $0x10013000$ . After generating the corresponding shellcodes for #RV64IC, /RV64IAC and 'RV64IDC, we successfully managed to execute them on QEMU. We provide in Appendix A the generated shellcodes, as well as some instructions to easily reproduce this experiment.

## 6.2 HiFive Unleashed

Subsequently, we moved to a more realistic environment, including a Linux operating system on a HiFive Unleashed board powered by a quad-core Freedom U540 RV64GC processor. It features an off-the-shelf Fedora 28 stage 4 disk image in a buildroot chrooted environment, for which we created a purposely vulnerable application executing its input data.

The first payload uses the `write` system call to print “Hello world!” on the standard output. As previously, we

generate the three different versions of our shellcode, and successfully manage to execute them on the vulnerable application. We successfully test the three shellcodes with two other payloads, one that spawns a shell using the `execve` system call, and one that prints on the standard output the contents of `/etc/shadow` file, using the `openat`, `read` and `write` system calls.

As a side note, as the floating-point unit is activated by the operating system, our 'RV64IDC shellcode does not require the non-alphanumeric previously described gadget anymore. Furthermore, we did not observe any instruction cache issue, as one could dread when using self-modifying code. This can be explained by the use of `fence.i` instructions that synchronize the instruction cache.

## 7 Conclusion and future work

We described a methodology for writing arbitrary alphanumeric (+1) RISC-V shellcodes. This method relies on unpacking, in which a program written in a very constrained instruction set stores into the memory another program written in a less constrained instruction set. Here, we required two unpackers in a three-staged shellcode to achieve arbitrary code execution. As a proof-of-concept, we showed examples of such shellcodes for the HiFive Unleashed board, featuring a standard Linux operating system. These positive results validate our choice for unpacking methods as the most suitable solution to the problem of writing executable code in a very constrained ISA subset.

Besides, the shellcodes provided in this paper only show proof-of-concept attacks. With the wide adoption of RISC-V based devices, we expect the attack surface to widen as new applications are published. Thereupon, we hope to see adequate defense mechanisms being deployed on those platforms, preventing such an attack. On the attacking side, automation seems the most promising way of improvement. Indeed, shellcodes tend to be handwritten or automated using *ad hoc* algorithms. We believe that a more general approach based on a higher-level semantic representation of the available instructions may be able to comprehensively solve the problem of writing code in a constrained ISA subset.

## References

- [1] 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).
- [2] Atmel Corporation. AVR Instruction Set Manual, 2016. <https://ww1.microchip.com/downloads/en/>

- devicedoc / atmel - 0856 - avr - instruction - set - manual.pdf.
- [3] Hadrien Barral, Houda Ferradi, Rémi Géraud, Georges-Axel Jaloyan, and David Naccache. ARMv8 Shellcodes from ‘A’ to ‘Z’. In *Proceedings of the 12th International Conference on Information Security Practice and Experience*, pages 354–377, Berlin, Heidelberg, 2016. Springer-Verlag. [https://link.springer.com/chapter/10.1007/978-3-319-49151-6\\_25](https://link.springer.com/chapter/10.1007/978-3-319-49151-6_25).
- [4] Aditya Basu, Anish Mathuria, and Nagendra Chowdary. Automatic Generation of Compact Alphanumeric Shellcodes for x86. In *Proceedings of the 10th International Conference on Information Systems Security*, pages 399–410, Berlin, Heidelberg, 2014. Springer-Verlag. [https://doi.org/10.1007/978-3-319-13841-1\\_22](https://doi.org/10.1007/978-3-319-13841-1_22).
- [5] Fabrice Bellard. QEMU, a Fast and Portable Dynamic Translator. In *Proceedings of the 2005 USENIX Annual Technical Conference*, pages 41–46, Berkeley, CA, USA, 2005. USENIX Association. <https://www.cse.iitd.ernet.in/~sbansal/cs1862-virt/2010/readings/bellard.pdf>.
- [6] Palmer Dabbelt, Stefan O’Rear, Kito Cheng, Andrew Waterman, Michael Clark, Alex Bradbury, David Horner, max Nordlund, and Karsten Merker. RISC-V ELF psABI Specification, 2016. <https://github.com/riscv/riscv-elf-psabi-doc/>.
- [7] Debian Wiki. RISC-V Port Information, 2019. <https://wiki.debian.org/RISC-V>.
- [8] Stephen Dolan. mov is Turing-Complete, 2013. <https://www.cl.cam.ac.uk/~sd601/papers/mov.pdf>.
- [9] Christopher Domas. The M/o/Vfuscator, 2015. <https://recon.cx/2015/slides/recon2015-14-christopher-domas-The-movfuscator.pdf>.
- [10] Riley Eller. Bypassing MSB data filters for buffer overflow exploits on Intel platforms, 2000. <https://web.archive.org/web/20070221035114/community.core-sdi.com/~juliano/bypass-msb.txt>.
- [11] 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).
- [12] Joshua Mason, Sam Small, Fabian Monrose, and Greg MacManus. English Shellcode. In *Proceedings of the 16th ACM Conference on Computer and Communications Security*, pages 524–533, New York, NY, 2009. ACM. <https://web.cs.jhu.edu/~sam/ccs243-mason.pdf>.
- [13] David Maynor. NX: How well does it say No to attacker’s eXecution attempts?, 2005. <https://www.blackhat.com/presentations/bh-usa-05/bh-us-05-maynor.pdf>.
- [14] Oleg Mazonka. Higher Subleq: Compiler into OISC language, 2009. <http://mazonka.com/subleq/hsq.html>.
- [15] 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/>.
- [16] 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).
- [17] 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. IEEE Computer Society Press, 1981.
- [18] PaX Team. PaX: Twelve Years of Securing Linux, 2012. <https://pax.grsecurity.net/docs/PaXTeam-LATINOWARE12-PaX-linux-security.pdf>.
- [19] Tiernan Ray. Nvidia, Western Digital at Chips’ Frontier, 2018. <https://www.barrons.com/articles/nvidia-western-digital-at-chips-frontier-1516640945>.
- [20] RIX. Writing IA32 alphanumeric shellcodes. *Phrack*, 57, 2001. <http://phrack.org/issues/57/15.html>.
- [21] SPARC International. The SPARC Architecture Manual, Version 8, 1991. <https://www.gaisler.com/doc/sparcv8.pdf>.
- [22] 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>.
- [23] Yves Younan, Pieter Philippaerts, Frank Piessens, Wouter Joosen, Sven Lachmund, and Thomas Walter. Filter-resistant Code Injection on ARM. *Journal in Computer Virology*, 7(3):173–188, 2011. <http://www.fort-knox.org/files/virology.pdf>.

## A Hello World Shellcodes

We provide ready-to-use demo shellcodes, written respectively in #RV64IC, /RV64IAC and 'RV64IDC. They print "Hello world!" on the serial output, when executed on QEMU with the following command:

```
qemu-system-riscv64 -nographic -machine sifive_u
-device loader,file=shellcode.bin,addr=0x80000000
```

The notation (X){Y} means that X is repeated Y times.

Colors have been added to each shellcode, with each color describing a specific high-level operation described in section 5. The instructions that jump over the encoded payload and put the location of the shellcode in sp as described in section 5.1.1 are colored in red. The encoded payload is in blue. Fixing-up the store pointer (section 5.1.2) is in cyan. Unpacking the stage 2 (section 5.1.3) is in purple. The final nopsled is in brown. For /RV64IAC and 'RV64IDC, additional data stored in the data pool is shown in green. In /RV64IAC, some additional code is required to first store the jump instruction (as shown in section 5.5) which is here in orange. Unused parts of the shellcode are in black. Given that colors cannot be reproduced in print, we refer the reader to the Arxiv version of this paper to that end.

### A.1 #RV64IC QEMU Hello World

```
o#0# (BBBB){1304} CGEDEDDOEEDEEDDGEED
ECEDGEEDLAKJDDDBDEDDNCMCDDDDDGMCLCFD
COBGEDDEGDCHCDDDALCDLMFHGDCHCDDDACOKEDAP
FLDLDDDDDDDLAPABHBKBFHFDCKBFBbPEFND
DDDBB (BBBB){1377} 3ZOa3QcAyyayayaya
yayayayayayayayayayayayaEcY3e##0ax
Aj#1Ay75v71J3SEAi##2ax7Eo91J3SEAY##3ax75
#zMJ3SEAM#y#y##4axQcY3E##5ax7ER81J3SEAY#
#6ax7Ej81J3SEAY##7ax75PP9J3ZEA#8Ay7#z81I
3Z#A#9AyAa75r05J3ZEA#2Ay7EBA9J3ZEA#3Ay7#
F#1Im93S#Au3#4ax7Ea85J3SEAY3#5ax7Up01J3Z
EA#6Ay759M5J3SEAi##7axAcy3e3#8axEcY3e##9
axAaAj#2Ay7#h91I3Z#A#3AyySySySs0A4
```

### A.2 /RV64IAC QEMU Hello World

```
ySySo/0/BBBBB03JBbBBBBBBBBBBPCJ
(BBBB){1955}
CGEDEDDOEEDEEDDGEEDDGEEDDGEEDDGEEDDGEEDD
EEDNCMCDDDDDGMCLCFDCOBGEDDEGDCHCDDDALCD
```

```
LMFHGDCHCDDDACOKEDAPFLDLDDDDDDDLAPABHBK
KBHFDCKBFBbPEFNDDBB (BBBB){751}
3YOa3Q/ABj/8Aa/8Aa1J3RHA3ZOa/OAc/8AD//Aa
/2AA9a9a9a9a9a9a9a9a9a9a9a9a9a9a9a9a9a
9a9a9a9a9a9a9a9a9a9a9a9a9a9a9a9a9a9a9a
1313//aDAA3ZOa/OAc/8AD75xG1J3SEAi3//aDAA
3ZOa/OAc/8AD7EqI1J3SEAY3//aDAA3ZOa/OAc/8
AD7EpQ9J3ZEA//AAAa3ZOa/OAc/8AD75gA9J3SEA
y3//aDAA3ZOa/OAc/8AD7ETH1J3SEAY3//aDAA3Z
Oa/OAc/8AD7E1H1J3SEAY3//aDAA3ZOa/OAc/8AD
75PP9J3ZEA//AAAa3ZOa/OAc/8AD7/zH1I3Z/A//
AAAa3ZOa/OAc/8AD75r05J3ZEA//AAAa3ZOa/OAc
/8AD7EBA9J3ZEA//AAAa3ZOa/OAc/8AD7/F/1Im9
3S/Au3//aDAA3ZOa/OAc/8AD7Ea85J3SEAY3//aD
Aa3ZOa/OAc/8AD7UpP1J3ZEA//AAAa3ZOa/OAc/8
AD75aA1J3SEAY3A3//aDAA3ZOa/OAc/8AD7///1I
a93S/AY3M31313//aDAA3ZOa/OAc/8AD75/AIJ3S
EA131313//aDAA3ZOa/OAc/8AD7/h91I3Z/A//AA
AaySySySySySySs0A4
```

### A.3 'RV64IDC QEMU Hello World

```
\89\63\73\90\03\30
o'0'BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
BBBBBBBBBBBBB3B1ozDaBBZzqspbBBBBBBBBBBBBB
BBBB64cinpaBBBBBBBBBBBBBBBBBBBug51zDaBVIQn
4f1A1nKj52aBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
BBBBphYdop1A9R1Yo3aBPtIx'51AMKqGzV1ABBBB
BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
BBBBUUUUUU1ALR5eFXcB (BBBB){1177}
CGEDEDDOEEDEEDDGEEDDGEEDDGEEDDGEEDDGEEDD
EDDNCMCDDDDDGMCLCFDCOBGEDDEGDCHCDDDALCD
LMFHGDCHCDDDACOKEDAPFLDLDDDDDDDLAPABHBK
KBHFDCKBFBbPEFNDDBB (BBBB){1438}
3ZOa3QGAB5b6F'F8f4J9j1N2n3yayayayaya9a9a
9a9a9a9a9a9a9a9a9a9a9a9a9a9a9a9a9a9a9a
9a9a9a9a9a9a9a9a9a9a9a9a9a9a9a9a9a9a9a
azG3Hr'6azG3HB'9azAa07X3L7G3IR'4azG3Ib'7
azG3GZ'9azs0A4
```

## B Source code

The full source code used for this paper is available at: <https://xn--fda.fr/riscv-alphanumeric-shellcoding> and <https://github.com/RischarV/riscv-alphanumeric-shellcoding>. It contains all demos and tools used for this paper.