

Towards Automated Augmentation and Instrumentation of Legacy Cryptographic Executables

Karim Eldefrawy¹, Michael Locasto¹, Norrathep Rattanaivanon^{*2}, and Hassen Saidi¹

¹ SRI International

{karim.eledefrawy, michael.locasto, hassen.saidi}@sri.com

² Prince of Songkla University, Phuket Campus

norrathep.r@phuket.psu.ac.th

Abstract. Implementation flaws in cryptographic libraries, design flaws in underlying cryptographic primitives, and weaknesses in protocols using both, can all lead to exploitable vulnerabilities in software. Manually fixing such issues is challenging and resource consuming, especially when maintaining legacy software that contains broken or outdated cryptography, and for which source code may not be available. While there is existing work on identifying cryptographic primitives (often in the context of malware analysis), none of this prior work has focused on replacing such primitives with stronger (or more secure ones) after they have been identified. This paper explores feasibility of designing and implementing a toolchain for *Augmentation and Legacy-software Instrumentation of Cryptographic Executables (ALICE)*. The key features of ALICE are: (i) automatically detecting and extracting implementations of weak or broken cryptographic primitives from binaries without requiring source code or debugging symbols, (ii) identifying the context and scope in which such primitives are used, and performing program analysis to determine the effects of replacing such implementations with more secure ones, and (iii) replacing implementations of weak primitives with those of stronger or more secure ones. We demonstrate practical feasibility of our approach on cryptographic hash functions with several popular cryptographic libraries and real-world programs of various levels of complexity. Our experimental results show that ALICE can locate and replace insecure hash functions, even in large binaries (we tested ones of size up to 1.5MB), while preserving existing functionality of the original binaries, and while incurring minimal execution-time overhead in the rewritten binaries. We also open source ALICE’s code at <https://github.com/SRI-CSL/ALICE>.

Keywords: Binary analysis · Cryptographic executables · Software instrumentation.

* Work done partially while at SRI International.

1 Introduction

Cryptography is instrumental to implementing security services such as confidentiality, integrity, and authenticity in most software (both, new and legacy). In practice, proper usage and correct implementation of cryptographic primitives are difficult; vulnerabilities often occur due to misuse or erroneous implementations of cryptographic primitives. Example vulnerabilities arising from misuse of cryptography include weak and/or broken random number generators, enabling an adversary to recover servers’ private keys [1]. Cryptographic APIs are sometimes misused by software developers, e.g., causing applications to be insecure against specific attacks, such as the chosen plaintext [2], which a typical software developer may be unaware of.

In addition, incorrect implementations of cryptographic primitives can result in leakage of secrets through side-channels [3] or through “dead memory” [4]. Other vulnerabilities in software for embedded and generic systems include implementation flaws in cryptographic libraries (e.g., the HeartBleed [5] and Poodle [6] vulnerabilities in the OpenSSL library), weaknesses in protocol suites (e.g., cryptographic weakness in HTTPS implementations [7,8]), and algorithmic vulnerabilities in cryptographic primitives (e.g., a chosen-prefix collision attack on the MD5 hash function [9] or on the SHA1 hash function [10]). Even after such vulnerabilities are discovered, it may take a while before appropriate fixes are applied to existing software as demonstrated by a recent large-scale empirical study [11] that showed many software projects did not patch cryptography-related vulnerabilities for a full year after their public disclosure. This represents a large window for adversaries to exploit such vulnerabilities.

We argue that (automated) tools that assist software and system designers, and developers, in performing identification, analysis, and replacement in binaries (without requiring source code) can help shorten such vulnerability window, especially for legacy software. While there is existing work on identifying cryptographic primitives (often in the context of malware analysis), none of this prior work has focused on replacing such primitives with more secure ones after they have been identified. To address this issue, we explore feasibility of designing and developing a toolchain for Augmentation and Legacy-software Instrumentation of Cryptographic Executables (ALICE).

Contributions: Specifically, our goal is to make the following contributions:

1. We design the ALICE framework to automatically augment and instrument executables with broken or insecure cryptographic primitives.
2. We develop heuristics to identify (binary) code segments implementing cryptographic primitives (see the extended version [12] of this work for a list of such primitives).
3. We develop heuristics to determine the scope of the (binary) code segments requiring augmentation if the cryptographic primitives are replaced with stronger ones.

4. We implement ALICE and experimentally evaluate its performance on several executable open source binaries of varying complexity. We also open source ALICE’s code at <https://github.com/SRI-CSL/ALICE>.

Outline: The rest of this paper is organized as follows: Section 2 discusses related work. Section 3 overviews the ALICE toolchain, while Section 4 contains its design details. Section 5 contains the results of our experimental evaluations. Section 6 discusses ALICE limitations, while Section 7 concludes the paper.

Full Paper: Due to space constraints, we refer to the extended version of this paper [12] for more detailed discussion of: challenges in addressing ALICE limitations, ALICE implementation details and full experimental results, as well as background on our evaluation dataset.

2 Related Work

Identifying Cryptographic Primitives. Several publicly available tools [13–15] utilize static analysis to identify cryptographic primitives by detecting known (large) constants used in their operation. Such constants, for example, can be in the form of look-up tables (e.g., S-Boxes in AES) or a fixed initialization vectors/values (e.g., IV in SHA-128/256). Such tools do not always produce accurate results as the detected algorithm may be another function or another cryptographic primitive that uses the same constant values [16]. They are also ineffective when dealing with obfuscated programs [17].

In terms of academic efforts, Lutz [18] detects block ciphers from execution traces based on three heuristics: the presence of loops, high entropy, and integer arithmetic. Grobert et al. [19] introduce an additional heuristic to extract cryptographic parameters from such execution traces and identify primitives by comparing the input-output relationships with those of known cryptographic functions. Lestringant et al. [16] propose a static method based on data flow graph isomorphism to identify symmetric cryptographic primitives. Recently, the CryptoHunt [20] tool develop a new technique called bit-precise symbolic loop mapping to identify cryptographic primitives in obfuscated binaries.

Our work focuses on non-obfuscated programs as we target common (and possibly legacy) software and not malware. We rely on finding known constants to identify cryptographic primitives as our first step. We then improve the accuracy of detection by applying a heuristic based on input-output relationships, similar to the work in [19]. In contrast to [19], our identification algorithm does not require program execution traces.

While there is existing work on identifying executable segments implementing cryptographic primitives, none of such work investigates the problem of replacing an identified weak primitive with a more secure one. Such replacement requires non-trivial operations, even if one can successfully identify executable segments implementing cryptographic primitives. To accomplish such replacement, one has to perform the following: (1) determining all changes throughout the binary necessary for replacing the identified primitive, and (2) rewriting the binary

to apply all of the determined changes. *To the best of our knowledge, there is no prior work addressing the first task, as a standalone, or in conjunction with the second.* The second task can be tackled using slight modifications of existing binary rewriting techniques. In this paper, we categorize different types of necessary changes one may require when replacing a cryptographic primitive, and then discuss how to locate and rewrite each category of such changes in Section 4.2. In the rest of this section, we overview general binary rewriting techniques.

Rewriting Binaries. Binary rewriting is a technique that transforms a binary executable into another without requiring the original’s source code. Typically, the transformed binary must preserve the functionality of the original one while possibly augmenting it with extra functionalities. There are two main categories of binary rewriting: static and dynamic binary rewriting.

In *static binary rewriting* [21–24], the original binary is modified offline without executing it. Static binary rewriting is typically performed by replacing the original instructions with an unconditional jump that redirects the program control flow to the rewritten instructions, stored in a different area of the binary. This relocation can be done at different levels of granularity such as inserting a jump for each modified instruction, for the entire section or for each routine containing modified instructions. Static binary rewriting often requires disassembling the entire binary and thus incurs high overhead during the rewriting phase, but typically results in small runtime overhead in the rewritten binary. This technique is thus well-suited for scenarios where the runtime performance of the rewritten binary is a primary concern. Another approach for static rewriting is to transform the binary into the relocatable disassembled code and directly rewrite instructions in the transformed code. Doing so completely eliminates runtime and size overhead in the rewritten binary. Nonetheless, this approach relies on many heuristics and assumptions for identifying and recovering all relocating symbols and is still subject to the high overhead during the rewriting phase. Some example tools that are based on this approach are `Uroboros` [25] and `Ramblr` [26].

Dynamic binary rewriting (or dynamic instrumentation) [27–30] modifies the binary’s behaviors during its execution through the injection of instrumentation code. Due to the need to instrument the code at runtime, this technique may result in higher execution time compared to the original binary. The main advantage of dynamic rewriting is its ability to accurately capture information about a program’s states or behaviors, which is much harder when using static rewriting. Example dynamic binary rewriting tools include `Pin` [28] and `DynamoRIO` [29].

In this work, we first leverage the runtime information retrieved from dynamic instrumentation to accurately locate instructions that need to be rewritten. Instruction rewriting is then performed statically in order to minimize the runtime overhead of the rewritten binary.

3 Overview of the ALICE Framework

The most straightforward, and obvious, approach to replace implementations of vulnerable cryptographic primitives requires modifying (and then recompiling) a program’s source code. This takes time and effort, and renders it difficult to fix legacy software for which source code may not be available. Instead, we propose ALICE – a toolchain that automatically augments and replaces weak, vulnerable, and/or broken cryptographic primitives at the binary level.

To better illustrate how ALICE works, we start by presenting a simple representative example shown in Figure 1. This example program first computes an MD5 digest over an input string. The digest is then converted into a human-readable form, which is in turn displayed to the user.

MD5 has been shown to be vulnerable to collision and pre-image attacks [31, 32]. Suppose that a system or software developer would like to manually rewrite parts of the binary in order to support a more secure hash algorithm – e.g., SHA-256. One way to accomplish this task is to perform the following steps:

Step-1: Identify the functions in the binary that implement MD5.

Step-2: Recover the type and order of parameters in the identified functions.

Step-3: Insert an implementation of a SHA-256 function with the same type and order of parameters into the original binary.

Step-4: Redirect all calls to MD5 to the newly added SHA-256 function.

Step-5: Determine all changes throughout the binary affected by an increase in the digest size (MD5’s digest size is 128 bits while that of SHA-256 is 256 bits).

Step-6: Rewrite the binary according to changes discovered in step-5.

Goal & Scope: ALICE is designed to automate the aforementioned steps. It targets ELF-based X86/64 binaries generated by compiling C programs. We do not assume any knowledge of, or require, the corresponding source code or debugging symbols. Since ALICE is built as a defensive tool to work on standard and legacy software, it assumes the target programs are not malicious. *Obfuscated or malware binaries are out of scope in this work.* We demonstrate concrete feasibility on cryptographic hash functions, but the design and ideas behind ALICE are general and can be applied to other primitives too.

NOTE: We use *cryptographic hash functions* (or *hash functions*) to denote the algorithmic details behind the implementations/executables. We denote the functions (or methods) in such implementations/executables that realize such hash function(s) as *hash routines*. We use *target* hash functions/routines to refer to the insecure hash functions/routines that need to be identified and replaced.

4 Design Details of ALICE

This section discusses the design details of the ALICE toolchain. The operation of ALICE consists of three main phases: (i) identifying cryptographic primitives, (ii) scoping changes, and (iii) augmenting and rewriting changes.

```

1 void MD5(const unsigned char* input,
2         size_t inputlen, unsigned char*
3         output) {
4     MD5_CTX ctx;
5     MD5Init(&ctx);
6     MD5Update(&ctx, input, inputlen);
7     MD5Final(output, &ctx);
8 }
9
10 int main(void) {
11     // Initialize input and output buffers
12     char input[] = "Hello, world!";
13     size_t inputlen = strlen(input);
14     unsigned char digest[16];
15     char hexdigest[33] = {0};
16
17     // Compute: output = MD5(input)
18     MD5(input, inputlen, digest);
19
20     // Convert output digest to hex string
21     for(int i=0; i < 16; i++) {
22         sprintf(hexdigest+2*i, "%02x",
23             digest[i]);
24     }
25
26     // Print digest in hex format
27     for(int i=0; i < 16; i++) {
28         printf("%02x", digest[i]);
29     }
30
31     // Print hexdigest string
32     printf("\n%s\n", hexdigest);
33     return 0;
34 }

```

(a) Simple program utilizing a cryptographic primitive (the MD5 hash function)

```

<main_fn_prologue>:
400629: push  rbp
40062a: mov   rbp, rsp
40062d: sub   rsp, 0x60
...
<call_to_md5>:
400684: call  4004c0 <strlen@plt>
400689: mov   rcx, rax
40068c: lea  rdx, [rbp-0x40]
400690: lea  rax, [rbp-0x50]
400694: mov  rsi, rcx
400697: mov  rdi, rax
40069a: call  400616 <MD5>
...
<sprintf_loop_body>:
4006ad: movzx eax, BYTE PTR [rbp+rax*1-0x40]
4006b2: movzx eax, al
4006b5: mov  edx, DWORD PTR [rbp-0x54]
4006b8: add  edx, edx
4006ba: movsxd rdx, edx
4006bd: lea  rcx, [rbp-0x30]
4006c1: add  rcx, rdx
4006c4: mov  edx, eax
4006c6: mov  esi, 0x4007d4
4006cb: mov  rdi, rcx
4006ce: mov  eax, 0x0
4006d3: call 400500 <sprintf@plt>
<sprintf_loop_condition>:
4006d8: add  DWORD PTR [rbp-0x54], 0x1
4006dc: cmp  DWORD PTR [rbp-0x54], 0xf
4006e0: jle  4006a8 <main+0x7f>
...
<printf_loop_body>:
4006eb: mov  eax, DWORD PTR [rbp-0x54]
4006ee: cdqe
4006f0: movzx eax, BYTE PTR [rbp+rax*1-0x40]
4006f5: movzx eax, al
4006f8: mov  esi, eax
4006fa: mov  edi, 0x4007d4
4006ff: mov  eax, 0x0
400704: call 4004e0 <printf@plt>
<printf_loop_condition>:
400709: add  DWORD PTR [rbp-0x54], 0x1
40070d: cmp  DWORD PTR [rbp-0x54], 0xf
400711: jle  4006eb <main+0xc2>

```

(b) Corresponding disassembly of main function in (a), compiled with O0 flag

Fig. 1: An example of a simple program utilizing a cryptographic primitive (the MD5 hash function in this case)

4.1 Identifying Cryptographic Primitives (Hash Functions)

We designed ALICE to target non-malicious (i.e., unobfuscated) binary programs. The first phase of ALICE leverages this characteristic and identifies hash functions by first detecting static features that are known of the target hash function.

Observation 1 (Constants) *A common design approach for hash functions is to initialize a digest buffer using well-known constants. As an example, MD5 uses 32-bit constant words: $X_1|X_2|X_3|X_4$ with the following hex values:*

$$X_1 = 0x67452301, X_2 = 0xEFCDAB89, X_3 = 0x98BADCFE, X_4 = 0x10325476$$

If we locate such constants in a binary, there is a high chance that a routine enclosing those constants implements part(s) of the MD5 hash function. Our approach starts by scanning a binary program to find the addresses where known constants appear. We then mark a routine in which those constants are enclosed as a candidate implementation of the target hash function.

Observation 2 (Context Initialization) *The example in Figure 1a illustrates a typical usage of a hash function in practice. An application function – `main()` – calls the `MD5()` function, which in turn calls `MD5Init()` to initialize a digest buffer with known constants. Having a dedicated function to setup an initial context (e.g., `MD5Init()`) is common practice when implementing most cryptographic primitives and can be found in several open-source libraries such as the `OpenSSL` or `libcrypt`.*

This observation suggests that the identified candidate implementation will typically correspond to the initialization routine – `Init()`. However, it is not always the case as `Init()` could be inlined. For example, `MD5Init()` in Figure 1 will be inlined inside `MD5()` when the program is compiled with the optimization flag `O3`. In this scenario, the identified routine will instead correspond to the implementation of the target hash function – `MD5()`. One could use a simple heuristic based on the size of the routine to distinguish between the two scenarios. However, we found this approach to produce a lot of false-negatives in practice. Instead, we adopt a more conservative approach and consider routines produced in both scenarios as candidates. More specifically, ALICE analyzes the program’s callgraph and determines all routines that invoke the previously identified routine. It then includes those caller routines into the list of candidates.

Now, ALICE needs a mechanism to eliminate false-positives, which can arise due to two reasons. The first reason is that our approach so far focuses on ensuring no false-negatives by accepting all routines possibly implementing a target hash routine. The second reason is that static features such as a constant vector are not always unique to a single hash function. It is not uncommon for different hash functions to share the same constant vectors. Examples of a pair of hash functions that use the same constant vectors are `BLAKE2b` – `SHA-512` and `MD5` – `MD4`. In Figure 1, even if we successfully determine that `MD5Init()` is inlined, we still cannot easily distinguish whether the identified hash routine implements `MD4` or `MD5` hash function.

Observation 3 (Input/Output Uniqueness) *A cryptographic hash function is deterministic, i.e., for a given input string, it always generates the same digest as output. The input/output pair is usually (in practice) unique to the hash function that produces them.*

With this observation, the best way to test whether a candidate implements the target hash function is to execute the identified routine, and compare the resulting output with the expected output. Since we expect the identification phase to be an offline computation, naturally we would base this step of our approach on an offline dynamic execution technique, which allows us to execute a given routine with any concrete chosen input. This is in contrast with online dynamic execution, which requires running the entire binary program with test cases. To perform offline dynamic execution, it is necessary to setup a call stack with proper parameters that will be passed into that routine.

Observation 4 (Parameters) *An implementation of a cryptographic function generally takes a fixed number of function parameters. For example, in Figure 1a,*

MD5() include three parameters: an input string, the input length and an output digest. Some implementations do not mandate the use of input length as it can be inferred from the input string (e.g., via `strlen()`). Thus, a hash routine in such implementations will take only two parameters. It is also worth noting that, even though a number is fixed, the order in which these parameters appear may not be the same for all implementations.

Based on this observation, ALICE enumerates all possible combinations of a hash routine's parameters and prepares 8 ($= 3!+2!$) call stacks, each initialized with different combinations of parameters. It then executes a candidate routine on each call stack and observes the output buffer after the execution. The first phase of ALICE finishes by outputting candidates producing the expected output as well as the parameter information obtained from corresponding call stacks.

4.2 Scoping Changes

After locating target hash routines, ALICE must determine changes (throughout the binary) that are required for replacing such routines. We now describe three categories of such changes using the illustrative example in Figure 1, and later outline how to identify a subset of such changes (at the binary level).

- C1 Routine Replacement:** The first category is a change in the hash routine itself. Code/instructions implementing the target hash routine are replaced by code/instructions that implement a more secure hash function. For example, if our goal is to replace the MD5 function with SHA-256 in Figure 1, instructions corresponding to MD5() (i.e., the ones starting from address 0x400616) need to be replaced by SHA-256 instructions. We will discuss how ALICE augments this type of change into a binary in Section 4.3.
- C2 Changes in Buffers Sizes:** Depending on the digest size of both the replacement and the target hash functions, other related memory buffers may need to be enlarged to correctly accommodate the new replacement routine. For instance, replacing MD5 with SHA-256 in Figure 1 would also require enlarging the size of variables storing the output of the hash function (i.e., `digest` variable) from 16 bytes to 32 bytes. This change in buffer size affects other memory buffers that consume the output digest, e.g., `hexdigest` also needs to be expanded by 16 bytes. Such changes have to be scoped and propagated throughout the entire binary. We discuss how to identify this type of change in the remaining of this section and how ALICE performs augmentation on such changes in Section 4.3.
- C3 Changes in Logic:** This category refers to changes that need to be applied to the underlying binary logic in order to have a correct resulting binary function. For example, in Figure 1, simply replacing MD5() with SHA-256() and enlarging `hexdigest` and `digest` variable do not suffice to produce the desired binary. One would have to also edit a loop terminating condition from $i < 16$ to $i < 32$ in line 19 and 24 to reflect the replacement SHA-256. At the binary-level, this change corresponds to modifying the instructions at

addresses 0x4006dc and 0x40070d from `[cmp DWORD PTR [rbp-0x54], 0xf]` to `[cmp DWORD PTR [rbp-0x54], 0x1f]`. This requires knowing that the constant 0xf in those instructions is related to the digest size. However, it is hard, in some cases impossible, to locate and augment this type of changes without any prior knowledge of the correct behavior of the resulting binary. Therefore, we do not consider this category of changes in this work.

Of the three categories of changes, **C1** is identified in the previous phase of ALICE in Section 4.1 while **C3** is out of scope in this work. The remainder of this section will focus on how ALICE locates the changes from **C2**.

ALICE leverages dynamic taint analysis to determine the change in buffer size **C2**. Typically, dynamic taint analysis starts by marking any data that comes from an untrusted source as tainted. It then observes program execution to keep track of the flow of tainted data in registers and memory. We adapt this idea to identify all *memory buffers* that are affected (or tainted) by the output digest of the target hash routine. In particular, our dynamic taint analysis executes a binary on test inputs with the following taint policies:

Taint introduction. At the beginning of execution, ALICE initializes all memory locations to be non-tainted. During the execution, whenever entering the target hash routine, ALICE reads the value in the parameter registers to observe the base address of the output digest. Since the digest size is well-known and deterministic for any given hash function, ALICE can also identify the entire address range of the digest buffer. Upon exiting the routine, ALICE then assigns a taint label to all memory locations in the digest buffer. ALICE uses three taint labels to differentiate 3 types of memory allocations:

- **Static allocation.** In our target executables, static memory is allocated at compile time before the program is executed. Thus, the location of this type of memory is usually deterministic, stored in either `.data` or `.bss` segment of the associated binary. Detecting whether a given memory location is statically allocated is simply done by checking whether its address lies within the boundaries of those segments.
- **Heap-based allocation.** ALICE traces all heap-based memory allocations by intercepting a call to three well-known C routines: `malloc()`, `calloc()` and `realloc()`. Whenever each of these routines is called, ALICE learns the size of allocated memory by reading values of its parameter registers. Upon exiting the same routine, ALICE then learns the base address of allocated memory via the return value. With this information, ALICE later can determine whether memory at a given location is allocated on the heap.
- **Stack-based allocation.** ALICE maintains stack-related information of the execution via a *shadow* stack. Specifically, after executing any `call` instruction, ALICE pushes into the shadow stack: a pair of the current stack pointer and an address of the function being called. Upon returning from a routine (via a `ret` instruction), ALICE pops the shadow stack. This information allows ALICE to reconstruct stack frames at any point during the execution

of dynamic taint analysis. ALICE determines whether memory at a given address is on the stack by checking it against all stack frames.

Taint Propagation. ALICE’s taint propagation rules are enforced at the word-level granularity. While we could use a more precise granularity such as the bit-level [33], we did not find such approach to be cost-effective as test inputs may require a timely interaction with a remote server; having a significantly long delay in the dynamic taint analysis can cause the remote server to timeout and consequently the analysis may not be performed as expected.

In addition to the general taint propagation rules, ALICE also considers the *taint-through-pointer* scenario: if a register A is tainted and a register B is assigned with the referenced value of A, i.e., $B := *A$, then B is considered tainted. Such rule is necessary to accurately capture the data-flow in a common usage of a hash function, where the raw digest value is converted to human-readable format via a look-up table, e.g., the use of `sprintf()` in line 20 of Figure 1a.

Using these rules, ALICE’s dynamic taint analysis can determine, and assign taint labels to, all memory locations affected by the output digest. At the end of the analysis, ALICE aggregates individual tainted memory locations into unified memory buffers. Our aggregation rule is simple: ALICE considers contiguous memory locations to be a memory buffer if their address range is at least as long as the target hash function’s digest size. Lastly, in this phase, ALICE outputs the types (i.e., either stack-based, heap-based or static), locations (e.g., a stack offset or global address), and relevant instructions (e.g., an instruction address of a call to `malloc`) of memory buffers that are derived from the output digest.

4.3 Augmenting and Rewriting Changes

ALICE can incorporate several rewriting approaches. Since runtime of rewritten binaries is our primary concern, we mainly use static binary rewriting that has been previously shown to have minimal impact on the runtime [22].

To reduce the size of rewritten binaries, we rewrite at *routine level* rather than at section level³. If there is at least one instruction that needs to be edited in a particular routine, we rewrite the binary as follows:

1. Create a new empty section in the binary
2. Apply changes from `C1` and `C2` to the routine
3. Modify the routine with respect to the placement of the new section
4. Insert the entire rewritten routine into the new section
5. Insert a `jump` instruction to the new section at original routine’s entry point

Steps (1), (4) and (5) are explained in Section 2. We focus in this section on steps (2) and (3). We refer to the full version of this paper [12] for implementation details of all steps.

³ We intentionally avoid rewriting at the instruction level as this can potentially incur significant run-time overhead for the rewritten/output binaries.

For step (3), we only need to ensure that the rewritten routine maintains the correct control flow targets. Doing so requires editing all instruction operands in the routine that use `rip`-relative addressing. The displacement of such operands is recomputed based on the address of the new location:

$$new_disp = old_disp + old_inst_addr - new_inst_addr$$

In step (2), to apply changes from **C1**, ALICE generates a patch from a user-supplied C code that implements the replacement hash function, and adds them to the new empty section of the binary. To ensure correctness of the rewritten binaries, implementation of the user-supplied replacement hash function must have the same parameter order as that of the target hash function as well as be self-contained. We also ensure that a call to the target hash routine is redirected to this new code by simply rewriting the first instruction of the target hash routine to: `jmp [new_code_entry_point]`. For each memory buffer identified in **C2**, ALICE computes the new buffer size based on the ratio of the digest sizes of the target hash function and that of the replacement hash function, i.e., $new_size = \lceil old_size \times |digest_{secure}| / |digest_{target}| \rceil$. ALICE rewrites the binary to support the expanded buffers by employing different techniques for each type of buffers:

Static Buffer. As a static buffer is allocated at a fixed address, we expand such buffer by creating another buffer at a new location and modify all instruction operands that access the original buffer to this newly allocated buffer. Specifically, ALICE first allocates a new data segment in the binary, and creates a mapping of the address of the original buffer to the address in the new data segment. To ensure that the rewritten binary uses the new address instead of the original, ALICE scans through all instructions in the original binary and edits the ones that contain an access to the original address by using information obtained from the previously computed address mapping.

Heap-based buffer. Unlike a static buffer, this type of buffer is allocated dynamically through a call to `malloc()`, `alloc()` or `realloc()` routine. Fortunately, ALICE learns when this type of buffer is allocated through the dynamic taint analysis in Section 4.2. Thus, expanding a heap-based buffer only requires ALICE to trace back to the instruction allocating such buffer, i.e. a `call` instruction to `malloc()`, `alloc()` or `realloc()`, and update the parameter register value storing the allocation size information to the new buffer size.

Stack-based buffer. Figure 2 shows how ALICE modifies the `main` routine in the example from Figure 1b to support the expansion of stack-based buffers: `digest` and `hexdigest` by 16 and 32 bytes respectively. Intuitively, expanding a buffer allocated on the stack at the binary level requires: (i) locating the routine that uses the corresponding stack frame, (ii) enlarging the frame to be large enough to hold the new buffers, and (iii) adjusting every access to memory inside the frame accordingly. ALICE’s previous phase, in Section 4.2, provides necessary information to satisfy the first requirement (via the shadow stack). To achieve the second requirement, ALICE rewrites the instructions that are responsible for increasing and decreasing the stack pointer in the prologue/epilogue of

<pre> <main_fn_prologue>: 400629: push rbp 40062a: mov rbp,rsp 40062d: sub rsp,0x60 ... <call_to_md5>: 400684: call 4004c0 <strlen@plt> 400689: mov rcx,rax 40068c: lea rdx,[rbp-0x40] 400690: lea rax,[rbp-0x50] 400694: mov rsi,rcx 400697: mov rdi,rax 40069a: call 400616 <MD5> ... <sprintf_loop_body>: 4006ad: movzx eax,PTR [rbp+rax*1-0x40] 4006b2: movzx eax,al 4006b5: mov edx,DWORD PTR [rbp-0x54] 4006b8: add edx,edx 4006ba: movsxd rdx,edx 4006bd: lea rcx,[rbp-0x30] 4006c1: add rcx,rdx ... 4006d3: call 400500 <sprintf@plt> <sprintf_loop_condition>: 4006d8: add DWORD PTR [rbp-0x54],0x1 4006dc: cmp DWORD PTR [rbp-0x54],0xf 4006e0: jle 4006a8 <main+0x7f> ... <printf_loop_body>: 4006eb: mov eax,DWORD PTR [rbp-0x54] 4006ee: cdqe 4006f0: movzx eax,PTR [rbp+rax*1-0x40] 4006f5: movzx eax,al 4006f8: mov esi,eax 4006fa: mov edi,0x4007d4 4006ff: mov eax,0x0 400704: call 4004e0 <printf@plt> <printf_loop_condition>: 400709: add DWORD PTR [rbp-0x54],0x1 40070d: cmp DWORD PTR [rbp-0x54],0xf 400711: jle 4006eb <main+0xc2> </pre>	<pre> <main_fn_prologue>: 400629: push rbp 40062a: mov rbp,rsp 40062d: sub rsp,0x90 ... <call_to_md5>: 400684: call 4004c0 <strlen@plt> 400689: mov rcx,rax 40068c: lea rdx,[rbp-0x70] 400690: lea rax,[rbp-0x80] 400694: mov rsi,rcx 400697: mov rdi,rax 40069a: call 400616 <MD5> ... <sprintf_loop_body>: 4006ad: movzx eax,PTR [rbp+rax*1-0x70] 4006b2: movzx eax,al 4006b5: mov edx,DWORD PTR [rbp-0x84] 4006b8: add edx,edx 4006ba: movsxd rdx,edx 4006bd: lea rcx,[rbp-0x50] 4006c1: add rcx,rdx ... 4006d3: call 400500 <sprintf@plt> <sprintf_loop_condition>: 4006d8: add DWORD PTR [rbp-0x84],0x1 4006dc: cmp DWORD PTR [rbp-0x84],0xf 4006e0: jle 4006a8 <main+0x7f> ... <printf_loop_body>: 4006eb: mov eax,DWORD PTR [rbp-0x84] 4006ee: cdqe 4006f0: movzx eax,PTR [rbp+rax*1-0x70] 4006f5: movzx eax,al 4006f8: mov esi,eax 4006fa: mov edi,0x4007d4 4006ff: mov eax,0x0 400704: call 4004e0 <printf@plt> <printf_loop_condition>: 400709: add DWORD PTR [rbp-0x84],0x1 40070d: cmp DWORD PTR [rbp-0x84],0xf 400711: jle 4006eb <main+0xc2> </pre>
---	---

(a) Original Binary

(b) Rewritten Binary

Fig. 2: Disassembly of `main` before & after increasing `digest` and `hexdigest` buffers by 16 and 32 bytes, respectively. Lines containing rewritten instructions are highlighted in green and changes are in red.

the located routine, e.g., `[40062d: sub rsp,0x60]` in Figure 2a. For the third requirement, ALICE iterates through all instructions in the routine and inspects the ones that use the stack offset, i.e., via `rsp` or `rbp` registers. ALICE then recomputes the stack offset with respect to the increased frame size and rewrites those instructions if the newly computed offset differs from the original. In Figure 2, ALICE identifies all instructions that access a stack element and rewrites the ones highlighted in green.

5 Experimental Evaluation

5.1 Experimental Setup

Goals & Datasets. The goal of this evaluation is two-fold, first, to assess whether ALICE can accurately identify and replace different implementations of hash functions, and second, to measure ALICE’s effectiveness on real-world applications. Different implementations may include different hash function structures (e.g., with or without `Init()`), different parameter orders, or simply different

implementation details. We apply ALICE to a dataset that consists of four popular cryptographic libraries: `OpenSSL`, `libgcrypt`, `mbedtls`, and `FreeBL`. We compile each library with different optimization levels, including `00`, `01`, `02`, `03` and `0s`, into a static library. We then create a simple C application (similar to the one in Figure 1a) that calls exactly one hash function located in the static library. We compile this application without debugging/relocation symbols and link it with each individual static library. We also assess ALICE’s effectiveness on 6 real-world applications: `smd5_mkpass` and `ssha_mkpass` – github projects for creating LDAP passwords ⁴, `md5sum` and `sha1sum` – string/file checksum programs, `lighttpd` – a lightweight webserver program, and `curl` – a webclient command line tool. Similar to the first dataset, each program was compiled without debugging symbols and with various optimization levels, and statically linked to the cryptographic library used within the program.

Insecure Hash Functions. We consider MD2, MD4, MD5, SHA1, and RIPEMD-160 as insecure hash functions to be replaced in our experiments. Our objective is to identify implementations of such hash functions in binaries and replace them with stronger ones, i.e., SHA-256. A list of all (insecure) hash functions in each dataset is shown in Table 1.

Environment. Experiments are performed on a virtual machine with Ubuntu 16.04.5 OS, 4GB of RAM and 2 cores of 3.4GHz of CPU running on top of an Intel i7-3770 machine. The following versions of required tools are used in the experiments: `gcc-5.4.0`, `angr-7.8.2.21` [34], `Triton-0.6` [35] and `Pin-2.14.71313` [28].

Dataset	Version	MD2	MD4	MD5	SHA1	RIPEMD160
Crypto Libraries	OpenSSL-1.1.1	✗	ILO	ILO	ILO	ILO
	libgcrypt-1.8.4	✗	✗	✗	OIL	OIL
	mbedtls-2.16.0	ILO	ILO	ILO	ILO	✗
	FreeBL-3.42	OI	✗	OIL,OI	OIL,OI	✗
Real-world Programs	<code>smd5_mkpass</code>	✗	✗	OIL	✗	✗
	<code>ssha_mkpass</code>	✗	✗	✗	OIL	✗
	<code>md5sum-5.2.1</code>	✗	✗	ILO	✗	✗
	<code>sha1sum-5.2.1</code>	✗	✗	✗	ILO	✗
	<code>curl-7.56.0</code>	✗	✗	OI	✗	✗
	<code>lighttpd-1.4.49</code>	✗	✗	✗	ILO	✗

Table 1: Hash functions used in our test datasets. ✗ indicates no hash function while ILO, OIL and OI denote a function with the parameter order (input, inputlen, output), (output, input, inputlen) and (output, input).

5.2 Evaluation Results: Cryptographic Libraries

As described in Section 4.2, we only consider automated inference of required changes from categories `C1` and `C2` in this work. In order to properly evaluate ALICE, we perform manual analysis to identify all changes required in `C3` and supply them to ALICE. The manually supplied changes for this case (in this dataset) consist of only a couple of instructions that typically specify loop

⁴ <https://github.com/pellucida/ldap-passwords>

termination condition(s). For example, in the binary from Figure 1b, we instruct ALICE to modify two instructions at addresses 0x4006dc and 0x40070d from `[cmp DWORD PTR [rbp-0x54],0xf]` to `[cmp DWORD PTR [rbp-0x54],0x1f]`.

Correctness of Rewritten Binary. To simplify illustration, we first describe behaviors of the rewritten binaries that are considered *incorrect* in this dataset. First, if ALICE misidentifies any of necessary changes in the input binary, the resulting binary will not display the *correct* SHA-256 digest of the input variable. For instance, it may display nonsensical data, a digest produced by the original hash function or an incomplete version of the SHA-256 digest. Second, if ALICE’s rewriting phase does not function properly (e.g., it expands the buffer size by a different amount or adjusts memory access to the stack incorrectly), it likely results in a runtime error for the output binary. We consider the correctness of the rewritten binaries from this dataset to be the converse of the aforementioned behaviors, i.e., execution of the output binary must terminate without any errors and it must result in displaying the correct SHA-256 digest of the input variable. All binaries produced by ALICE in this dataset work as expected.

Binary Size and Execution Overhead. ALICE adds around 3-13KB to the output binaries. On further inspection, we found two main reasons for this overhead. First, ALICE statically adds a patch implementing the replacement SHA-256 hash function, which contributes around 3KB to the output binary. Second, our underlying binary rewriter, `patchkit` [36], expects code and data of this patch to be aligned to a page size (i.e., 4KB in our testing machine), which can add up to another 8KB to the output binary.

In terms of execution overhead, we implement a simple Pintool [28] to count the number of instructions executed by the output binaries. We then compare the result with the baseline, where manual editing is performed on the original binary’s source code in order to replace the insecure hash function and the modified source-code is properly optimized by the standard `gcc` compiler. The results in Figure 3a show that the binaries produced by ALICE have low execution overhead with an average of 300 added instructions, or only an increase of 0.3%, compared to the baseline. We also did not observe any noticeable increase in execution-time for the output binaries.

Toolchain Runtime. Figure 3b shows runtime of ALICE to produce the output binaries. The total runtime heavily depends on the size of input binaries, and is dominated by the runtime of the identification phase. This bottleneck happens because the identification phase involves heavy-weight analysis such as disassembling the entire binary and/or recovering the binary call graph. It is also worth noting that such analysis is performed only once and ALICE re-uses the analysis results in latter phases; this leads to lower runtimes in the latter phases.

5.3 Evaluation Results: Real-World Binaries

Similar to the previous dataset, we manually inspect the binaries in this dataset to identify changes required in `C3`, then supply them to ALICE. Such changes are mainly related to the digest size that is hard-coded in the source code.

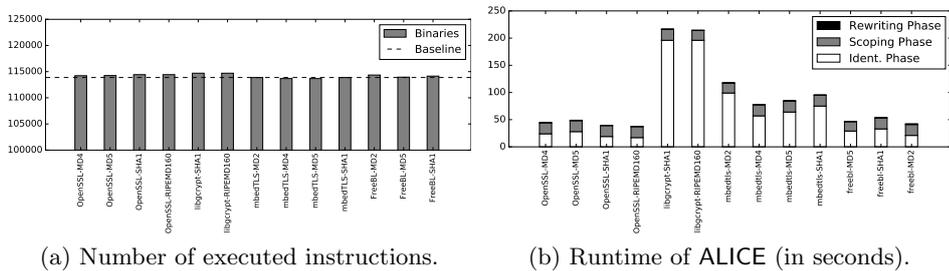


Fig. 3: Evaluation results of ALICE on cryptographic libraries. Original binaries are compiled with O2. Results for binaries with different optimization flags are similar and thus omitted.

Correctness of Rewritten Binary. We consider rewritten binaries to be correct if changes performed by ALICE: (1) correctly implement new functionalities with respect to the target SHA-256 hash function and (2) do not interfere with the remaining functionalities. For instance, the former enforces the rewritten binary of `md5sum` to be able to perform `sha256sum` of a given input string. We realize the latter requirement by executing the binaries produced by ALICE with all test cases (except the ones that use insecure hash functions) provided in their original respective project repository. Indeed, our definition of correctness is an approximation. Satisfying this definition does not guarantee that the rewritten/output binaries are truly correct. However, we believe this definition is a good approximation of what is required in practice, especially with access to a dataset with enough test cases that one can evaluate against results produced by the rewritten binaries.

Table 2 shows the correctness of output binaries produced by ALICE in this dataset. All output binaries pass all test cases in their original project repository while only one output binary fails to pass the expected functionality. We manually examined the failed binary and found that ALICE misidentified an insecure hash routine. Our further inspection reveals that the main culprit appears to be our underlying dynamic concrete executor, which fails to output the expected MD5 digest even if ALICE sets up a proper call stack. As a result, ALICE’s identification phase did not detect this insecure hash function in the input binary, and the output binary remained unchanged.

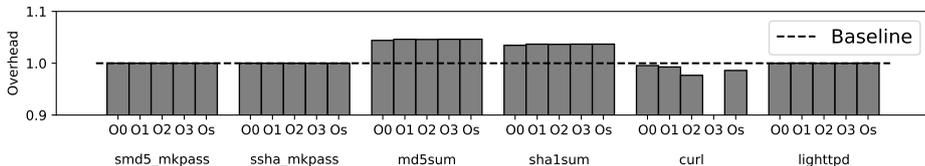


Fig. 4: Overhead in terms of executed instructions of real-world binaries.

Program	OFLAG	Binary Size			Correctness of Output Binaries
		O	R	Δ	
md5sum	-O0	43.6KB	51.9KB	8.3KB	✓
	-O1	35.4KB	45.3KB	9.9KB	✓
	-O2	35.4KB	45.0KB	9.6KB	✓
	-O3	39.5KB	49.1KB	9.6KB	✓
	-Os	31.3KB	40.4KB	9.1KB	✓
sha1sum	-O0	43.6KB	51.9KB	8.3KB	✓
	-O1	35.4KB	45.3KB	9.9KB	✓
	-O2	35.4KB	45.0KB	9.6KB	✓
	-O3	39.5KB	49.1KB	9.6KB	✓
	-Os	31.3KB	40.4KB	9.1KB	✓
smd5_mkpass	-O0	22.8KB	29.9KB	7.1KB	✓
	-O1	18.7KB	25.8KB	7.1KB	✓
	-O2	18.7KB	25.8KB	7.1KB	✓
	-O3	22.8KB	29.8KB	7.0KB	✓
	-Os	18.7KB	25.8KB	7.1KB	✓
ssha_mkpass	-O0	22.8KB	29.9KB	7.1KB	✓
	-O1	18.7KB	25.8KB	7.1KB	✓
	-O2	18.7KB	25.8KB	7.1KB	✓
	-O3	22.8KB	29.8KB	7.0KB	✓
	-Os	18.7KB	25.8KB	7.1KB	✓
curl	-O0	929.6KB	937.3KB	7.7KB	✓
	-O1	589.6KB	596.1KB	6.5KB	✓
	-O2	614.2KB	620.7KB	6.5KB	✓
	-O3	675.6KB	N/A	N/A	✗
	-Os	528.1KB	534.5KB	6.4KB	✓
lighttpd	-O0	720.2KB	724.0KB	3.8KB	✓
	-O1	522.1KB	529.1KB	7.0KB	✓
	-O2	534.7KB	545.8KB	11.2KB	✓
	-O3	584.7KB	592.0KB	7.4KB	✓
	-Os	466.4KB	473.1KB	6.6KB	✓

Table 2: Size of original (O) and rewritten (R) binaries of real-world applications.

Binary Size and Execution Overhead. Table 2 shows the increase in binary size in this dataset. ALICE adds around 4 to 11KB to the original binary. As mentioned in Section 5.2, up to 8KB of this overhead is caused by the underlying binary rewriter performing a patch alignment. The remaining overhead stems from rewritten functions that are appended at the end of the new binary. We note that we excluded the result for the `curl` binary compiled with `O3` in Table 2 (and subsequent figures) as ALICE could not produce the correct output binary in that case.

We did not observe any noticeable execution overhead in terms of execution-time when running the output binary against the provided test cases from the project’s repository. In addition, we measure the number of executed instructions for the output binary to perform the expected functionality with respect to the SHA-256 function and compare it to the baseline, where we manually edit the source code to replace the insecure hash function. The result, shown in Figure 4, also indicates negligible ($< 5\%$) increase in execution-time in this case. Note that even though execution of the rewritten `curl` binaries becomes faster (requires 2% fewer instructions), this improvement is still negligible. As such, we do not claim that ALICE helps producing a more efficient output binary.

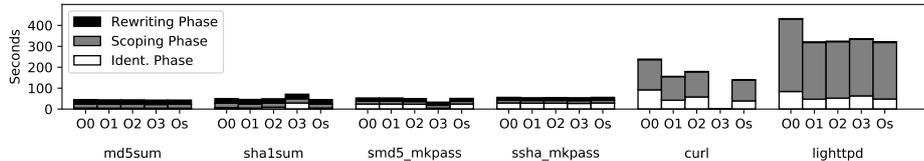


Fig. 5: Runtime of the ALICE toolchain on real-world binaries.

Toolchain Runtime. Figure 5 illustrates ALICE’s runtime to produce each output binary. In simpler programs (e.g., `md5sum` or `smd5_mkpass`), ALICE identifies and replaces weak primitives in less than a minute. For more complex programs (e.g., `lighttpd`), ALICE’s runtime can be a bit slower – up to 5 minutes. Most of the runtime overhead comes from the scoping phase because it needs to instrument a large number of instructions (e.g., $\approx 500k$ instructions for `lighttpd`) while execution of simpler programs contains significantly fewer instructions. We consider ALICE’s runtime to be acceptable since the entire process only needs to be performed once, making the toolchain runtime not a primary concern.

Program		md5sum	sha1sum	smd5_mkpass	ssha_mkpass	curl	lighttpd	Avg.
Changes:	C1	1024	1806	932	1676	424	1712	1262
	C2	1	1	3	9	8	2	4
	C3	1	1	2	3	1	2	1.7
$(C1 + C2) / \text{Total}$		99.9%	99.94%	99.79%	99.85%	99.78%	99.88%	99.87%
$C2 / (C2 + C3)$		50%	50%	60%	75%	87.5%	50%	70.17%

Table 3: Number of rewritten instructions required for each category of changes.

Reduction in Manual Efforts. While ALICE currently does not automatically identify changes from the **C3** class, it still saves considerable manual effort. We quantify such savings in Table 3 as the number of instructions required to be rewritten in order to implement changes for each category. On average, ALICE automatically identifies and rewrites 1,266 instructions, which translates into 99.87% reduction in manual efforts. However, we acknowledge that it may be possible to use existing cryptographic identification tools (with some modifications) to locate changes from **C1**. Even when such tools exist, our toolchain still significantly reduces manual work by 70.17%. It is worth emphasizing again that no existing tools are capable of identifying changes from **C2** and **C3**.

6 Limitations and Future Work

The current version of ALICE has limitations. First, ALICE relies on some underlying (open source) building-block tools and inherits their limitations. For example, we encountered instances where the underlying x86-64 assembler, *Keystone* [37], fails to translate uncommon instructions whose operands contain `fs` registers or a `rep` instruction. Whenever we encounter such an issue, we manually

fixed it by directly hard-coding the correct behavior into ALICE. Furthermore, the first phase of ALICE relies on the `angr` framework [34] for disassembly of stripped binaries; `angr` does not perform static disassembly with correctness guarantees. In fact, static disassembly of stripped binaries with correctness guarantees is still an open problem [38]. Thus, `angr` may produce incorrect results in ALICE’s first phase, which affects outcomes of output binaries.

ALICE also assumes that the routine implementing an insecure hash function and necessary changes are statically included in the main application. ALICE does not currently support identifying and replacing insecure cryptographic primitives located in a dynamic library. Expanding ALICE’s functionalities to dynamic libraries is possible since most of our underlying tools are capable of locating and analyzing dynamic libraries used by the main application.

ALICE does not automatically identify nor rewrite changes from `C3` and relies on the user to supply them to the toolchain. In practice, some manual effort is required to locate changes in binary logic for a large binary. Automating this process is a challenging problem and is an interesting avenue for future work.

Finally, we design ALICE to target non-malicious legacy binaries and assume that such binaries are not obfuscated. In practice, even legitimate software may make use of obfuscation techniques, e.g., to protect intellectual property. Extending ALICE to support obfuscated but non-malicious binaries is also an interesting future direction.

Acknowledgments

This work was sponsored by the U.S. Department of Homeland Security (DHS) Science and Technology (S&T) Directorate under Contract No. HSHQDC-16-C-00034. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of DHS and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DHS or the U.S. government. The authors thank the anonymous reviewers for their valuable comments.

7 Conclusion

We have developed ALICE, a toolchain for identifying weak or broken cryptographic primitives and replacing them with secure ones without relying on source code or debugging symbols. We have implemented a prototype of ALICE that can detect several cryptographic primitives while only requiring access to the binaries containing them. Our implementation of ALICE can also automatically replace weak and/or broken implementations of cryptographic hash functions in ELF-based x86-64 binaries. We have demonstrated ALICE’s effectiveness on various open-source cryptographic libraries and real-world applications utilizing cryptographic hash functions. Our experimental results show that ALICE can successfully locate and replace insecure hash functions while preserving existing functionalities in the original binaries.

References

1. N. Heninger, Z. Durumeric, E. Wustrow, and J. A. Halderman, “Mining your Ps and Qs: Detection of widespread weak keys in network devices,” *USENIX Security Symposium*, 2012.
2. M. Egele, D. Brumley, Y. Fratantonio, and C. Kruegel, “An empirical study of cryptographic misuse in android applications,” in *ACM Conference on Computer and Communications Security*, 2013.
3. N. P. Smart, “Physical side-channel attacks on cryptographic systems,” *Software Focus*, vol. 1, no. 2, pp. 6–13, 2000.
4. J. Li, Z. Lin, J. Caballero, Y. Zhang, and D. Gu, “K-hunt: Pinpointing insecure cryptographic keys from execution traces,” in *ACM Conference on Computer and Communications Security*, 2018.
5. US-CERT, “Openssl ‘heartbleed’ vulnerability (cve-2014-0160),” <https://www.us-cert.gov/ncas/alerts/TA14-098A>, 2014.
6. —, “Ssl 3.0 protocol vulnerability and padding oracle on downgraded legacy encryption(poodle) attack,” <https://www.us-cert.gov/ncas/alerts/TA14-290A>, 2014.
7. S. Calzavara, R. Focardi, M. Nemec, A. Rabitti, and M. Squarcina, “Postcards from the post-http world: Amplification of https vulnerabilities in the web ecosystem,” in *IEEE Symposium on Security and Privacy*, 2019.
8. D. Adrian, K. Bhargavan, Z. Durumeric, P. Gaudry, M. Green, J. A. Halderman, N. Heninger, D. Springall, E. Thomé, L. Valenta *et al.*, “Imperfect forward secrecy: How diffie-hellman fails in practice,” in *ACM Conference on Computer and Communications Security*, 2015.
9. M. Stevens, “Counter-cryptanalysis,” in *Annual Cryptology Conference*, 2013.
10. G. Leurent and T. Peyrin, “SHA-1 is a shambles - first chosen-prefix collision on SHA-1 and application to the pgp web of trust,” 2020, <https://eprint.iacr.org/2020/014>.
11. F. Li and V. Paxson, “A large-scale empirical study of security patches,” in *ACM Conference on Computer and Communications Security*, 2017.
12. K. Eldefrawy, M. Locasto, N. Rattanavipanon, and H. Saidi, “Towards automated augmentation and instrumentation of legacy cryptographic executables: Extended version.” [Online]. Available: <https://arxiv.org/abs/2004.09713>
13. aldeid, “Ida-pro/plugins/findcrypt2,” <https://www.aldeid.com/wiki/IDA-Pro/plugins/FindCrypt2>, 2019.
14. igNorAMUS, snaker, Maxx, and pusher, “Kanal - krypto analyzer for peid,” <http://www.dcs.fmph.uniba.sk/zri/6.prednaska/tools/PEiD/plugins/kanal.htm>, 2019.
15. apponic, “Hash & crypto detector,” <https://hash-crypto-detector.apponic.com/>, 2019.
16. P. Lestringant, F. Guihéry, and P.-A. Fouque, “Automated identification of cryptographic primitives in binary code with data flow graph isomorphism,” in *ACM ASIA Conference on Computer and Communications Security*, 2015.
17. J. Calvet, J. M. Fernandez, and J.-Y. Marion, “Aligot: Cryptographic function identification in obfuscated binary programs,” in *ACM Conference on Computer and Communications Security*, 2012.
18. N. Lutz, “Towards revealing attacker’s intent by automatically decrypting network traffic,” *Mémoire de maîtrise, ETH Zürich, Switzerland*, 2008.
19. F. Gröbert, C. Willems, and T. Holz, “Automated identification of cryptographic primitives in binary programs,” in *Recent Advances in Intrusion Detection*. Springer, 2011, pp. 41–60.

20. D. Xu, J. Ming, and D. Wu, "Cryptographic function detection in obfuscated binaries via bit-precise symbolic loop mapping," in *2017 IEEE Symposium on Security and Privacy*, May 2017.
21. G. Hunt and D. Brubacher, "Detours: Binary interception of Win32 functions," in *3rd USENIX Windows NT Symposium*, 1999.
22. E. Bauman, Z. Lin *et al.*, "Superset disassembly: Statically rewriting x86 binaries without heuristics," in *Network and Distributed System Security Symposium*, 2018.
23. K. Anand, M. Smithson *et al.*, "A compiler-level intermediate representation based binary analysis and rewriting system," in *ACM European Conference on Computer Systems*, 2013.
24. A. Edwards, H. Vo, A. Srivastava, and A. Srivastava, "Vulcan binary transformation in a distributed environment," Microsoft Research, Tech. Rep., 2001.
25. S. Wang, P. Wang, and D. Wu, "Reassembleable disassembling," in *USENIX Security Symposium*, 2015, pp. 627–642.
26. R. Wang, Y. Shoshitaishvili, A. Bianchi, A. Machiry, J. Grosen, P. Grosen, C. Kruegel, and G. Vigna, "Ramblr: Making reassembly great again," in *Network and Distributed System Security Symposium*, 2017.
27. N. Nethercote and J. Seward, "Valgrind: A program supervision framework," *Electronic Notes in Theoretical Computer Science*, vol. 89, no. 2, 2003.
28. C.-K. Luk, R. Cohn *et al.*, "Pin: building customized program analysis tools with dynamic instrumentation," in *ACM SIGPLAN Notices*, vol. 40. ACM, 2005, pp. 190–200.
29. <http://www.dynamorio.org/>, "Dynamic instrumentation tool platform," <http://www.dynamorio.org/>, 2017.
30. J. H. Perkins, S. Kim *et al.*, "Automatically patching errors in deployed software," in *ACM SIGOPS*, 2009.
31. V. Klima, "Tunnels in hash functions: Md5 collisions within a minute." *IACR Cryptology ePrint Archive*, vol. 2006, p. 105, 2006.
32. Y. Sasaki and K. Aoki, "Finding preimages in full md5 faster than exhaustive search," in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2009, pp. 134–152.
33. B. Yadegari and S. Debray, "Bit-level taint analysis," in *2014 IEEE 14th International Working Conference on Source Code Analysis and Manipulation (SCAM)*. IEEE, 2014, pp. 255–264.
34. Y. Shoshitaishvili, R. Wang, C. Salls, N. Stephens, M. Polino, A. Dutcher, J. Grosen, S. Feng, C. Hauser, C. Kruegel *et al.*, "Sok:(state of) the art of war: Offensive techniques in binary analysis," in *2016 IEEE Symposium on Security and Privacy*. IEEE, 2016, pp. 138–157.
35. F. Saudel and J. Salwan, "Triton: A dynamic symbolic execution framework," in *Symposium sur la sécurité des technologies de l'information et des communications, SSTIC, France, Rennes, June 3-5 2015*. SSTIC, 2015, pp. 31–54.
36. R. Hileman, "Binary patching from python," <https://github.com/linuxbochs/patchkit>, 2018.
37. N. A. Quynh, "Keystone – the ultimate assembler," <http://www.keystone-engine.org/>, 2019.
38. D. Andriessse, X. Chen *et al.*, "An in-depth analysis of disassembly on full-scale x86/x64 binaries," in *USENIX Security Symposium*, 2016.