

BinMatch: A Semantics-based Hybrid Approach on Binary Code Clone Analysis

Yikun Hu*, Yuanyuan Zhang*, Juanru Li*, Hui Wang*, Bodong Li*, Dawu Gu^{†*}

*Shanghai Jiao Tong University, Shanghai, China

[†]Shanghai Institute for Advanced Communication and Data Science, Shanghai, China

{yixiaoxian, yyjess, jarod, tony-wh, uchihal, dwgu}@sjtu.edu.cn

Abstract—Binary code clone analysis is an important technique which has a wide range of applications in software engineering (e.g., plagiarism detection, bug detection). The main challenge of the topic lies in the semantics-equivalent code transformation (e.g., optimization, obfuscation) which would alter representations of binary code tremendously. Another challenge is the trade-off between detection accuracy and coverage. Unfortunately, existing techniques still rely on semantics-less code features which are susceptible to the code transformation. Besides, they adopt merely either a static or a dynamic approach to detect binary code clones, which cannot achieve high accuracy and coverage simultaneously.

In this paper, we propose a semantics-based hybrid approach to detect binary clone functions. We execute a template binary function with its test cases, and emulate the execution of every target function for clone comparison with the runtime information migrated from that template function. The semantic signatures are extracted during the execution of the template function and emulation of the target function. Lastly, a similarity score is calculated from their signatures to measure their likeness. We implement the approach in a prototype system designated as BINMATCH which analyzes IA-32 binary code on the Linux platform. We evaluate BINMATCH with eight real-world projects compiled with different compilation configurations and commonly-used obfuscation methods, totally performing over 100 million pairs of function comparison. The experimental results show that BINMATCH is robust to the semantics-equivalent code transformation. Besides, it not only covers all target functions for clone analysis, but also improves the detection accuracy comparing to the state-of-the-art solutions.

I. INTRODUCTION

Binary code clone analysis is a fundamental technique in software engineering. It has important applications in fields of software maintenance and security, for example, plagiarism detection [20], [45], patch code analysis [5], code searching [6], program comprehension [18], malware lineage inference [32], [30], [1], known vulnerability detection [34], [12], [13], etc.

The main challenge that affects the accuracy of binary code clone analysis stems from the semantics-equivalent *code transformation* (C1), typically including link-time optimization of compilers and code obfuscation [11]. The transformation modifies the representations of binary code. Even though two pieces of code are compiled from the same code base, the resulting binaries after the transformation would differ significantly on the syntax or structure level (e.g., instructions, control flow graphs). Another challenge is the *trade-off* between detection accuracy and coverage (C2), which

corresponds to analyzing binary code in which manner, dynamic or static [39]. Dynamic methods procure rich semantics from code execution to ensure high accuracy, but they analyze only the executed code, leading to low coverage. In contrast, static methods are able to cover all program components, while they rely more on syntax and structure features which lack semantics. Additionally, static methods cannot decide the targets of indirect jumps and calls. Thus, the analysis accuracy of static methods is relatively low.

In the literature, binary code clone analysis has drawn much attention. However, existing techniques adopt either static method which depends on semantics-less features or dynamic method which merely cares about executed code. For example, static methods discovRE [12], Genius [13], and KamIn0 [10] extract features from control flow graphs, and measure similarity of binary functions basing on graph isomorphism. Multi-MH [34] and BinGo [6] capture behaviors of a binary function by sampling it with random values. Since the random inputs lack semantics and are usually illegal for the function, they could hardly trigger the real semantics of a function. For dynamic methods, Ming et al. [32], Jhi et al. [20], and Zhang et al. [45] perform analysis merely on executed code. BLEX [11] pursues high code coverage at the cost of breaking normal execution of a binary function, distorting the semantics inferred from its collected features. Therefore, it is necessary to propose a method which depends only on semantics and takes advantages of both static and dynamic techniques to detect binary code clones.

In this paper, we propose BINMATCH, a semantics-based hybrid approach, to detect binary clone functions. Given a template function, BINMATCH firstly instruments and executes it with test cases to record its runtime information (e.g., function argument values). It then migrates the information to *each* candidate target function and emulates the execution of the function. During the execution and emulation, semantic signatures of the template and target functions are recorded. Finally, BINMATCH compares signatures of the template function and each target function to measure their similarity. To overcome C1 of semantics-equivalent code transformation, BINMATCH only relies on semantic signatures extracted from the whole template or target function. To address C2 of the trade-off between accuracy and coverage, BINMATCH adopts the hybrid method which captures semantic signatures in both static and dynamic manners. By executing the template function,

BINMATCH captures its signature of rich semantics. Then, it emulates every candidate target function with the runtime information of the template function to extract their signatures, which takes all target functions into consideration.

BINMATCH is evaluated with eight real-world projects compiled with various compilation configurations and obfuscation settings, totally performing over 100 million pairs of function comparison. The experimental results indicate that BINMATCH not only is robust to semantics-equivalent code transformation, but also outperforms the state-of-the-art solutions.

In summary, the contributions of this paper are as followed:

- We propose a semantics-based hybrid approach to analyze binary code clones. The approach captures the semantic signature of a binary function in either dynamic (execution) or static (emulation) manner. Thus, it could not only detect clone functions accurately with signatures of rich semantics, but also cover all target functions under analysis.
- To smooth the migration of runtime information and the emulation of a function, we propose novel strategies to handle global variable reading, indirect calling/jumping, and library function invocation.
- We implement the approach in a prototype system BINMATCH which supports IA-32 binary code clone analysis on the Linux platform. BINMATCH is evaluated with eight real-world projects which are compiled with different compilation configurations and obfuscation settings. The experimental results show that BINMATCH is robust to the semantics-equivalent code transformation. Besides, it covers all candidate target functions for clone analysis, and outperforms the state-of-the-art solutions from the perspective of accuracy.

II. MOTIVATION AND OVERVIEW

In this section, we firstly present an example to illustrate the limitations of previous work on binary code clone analysis, which motivate our research. Then, we explain the basic idea of our approach and show the system overview.

A. Motivating Example

It is a typical application of binary code clone detection to locate known vulnerable code in binary programs [34], [12], [13]. Given a piece of code which contains a known bug, it is possible to locate the corresponding clone (or similar) code in other programs to check whether those programs are also vulnerable.

NConvert [42] is a closed-source image processor which supports multiple formats. It statically links the open-source library `libpng` [29] to handle files of the PNG format. Function `png_set_unknown_chunks` of `libpng` is found to contain an integer overflow vulnerability before the version of 1.5.14 (CVE-2013-7353). It is necessary to locate the statically-linked function in NConvert to verify whether the function is vulnerable and ensure the program security. Since the source code of `libpng` is available, it is reasonable

to fulfill the target with the clone code detection technique. However, only the executable of NConvert is accessible that its compilation configuration is unknown. Even though executables are compiled from the same code base, different compilation configurations would lead to semantics-equivalent binary code of equal semantics. Hence, methods relying on syntax or structural features (e.g., control flow graph isomorphism) become ineffective. Besides, it is challenging to not only locate `png_set_unknown_chunks` accurately, but also achieve high code coverage of NConvert (C2). The target function is statically-linked, mixing with the user-defined functions of NConvert. Static methods of binary code clone detection could cover all functions in NConvert to find `png_set_unknown_chunks`. However, they leverage semantics-less features, generating inaccurate results. In contrast, dynamic methods depend on semantic features which are extracted via code execution, while they merely focus on the executed code. It even requires huge extra work for dynamic methods to generate test cases in order to cover the target function. Unfortunately, code coverage is still an issue for dynamic analysis of binaries [24].

B. System Overview of BINMATCH

We propose BINMATCH to perform binary function clone analysis. Given a binary function (the *template*), BINMATCH finds its clone match in the target binary program, returning a list of functions (the *targets*) from the program, which is ranked basing on the semantic similarity.

Figure 1 presents the work flow of BINMATCH. Given the template function which has been well analyzed or understood (`png_set_unknown_chunks`), BINMATCH instruments and executes it with test cases, capturing its semantic signature (§III-A). Meanwhile, runtime information is recorded during the execution as well (§III-B). Then, BINMATCH migrates the runtime information to each target function of the target binary program (NConvert). It emulates the execution of the target function to extract the semantic signature (§III-C). Afterward, BINMATCH compares the signature of the template function to that of each target function and computes their similarity score (§III-D). Lastly, a list of target functions is generated, which is ranked by the similarity scores in descending order.

In summary, to overcome C1, BINMATCH completely depends on semantic signatures to detect binary function clones. Additionally, the signatures are captured in a hybrid manner, which addresses C2. BINMATCH firstly extracts the signature of the template function via executing its test cases. We assume that the template function has been well studied that its test cases are available. In above example, the vulnerability of `png_set_unknown_chunks` has been known, and its test cases could be found in the `libpng` project as well as from the vulnerability database. Then, with the runtime information of the template function, BINMATCH generates the signature of each target function of the binary program under analysis (NConvert) via emulation. Therefore, BINMATCH is able




Fig. 1: System Architecture of BINMATCH

to cover all target functions to detect their clone matches with signatures of rich semantics.

III. METHODOLOGY

In this section, we firstly introduce the semantic signatures adopted by BINMATCH, then discuss how it captures the signatures of binary functions and measures their similarity.

A. Semantic Signatures

For each binary function, BINMATCH captures behaviors during the execution or emulation as its signature. Given a specific input, the signature indicates how the function processes the input and generates the output, reflecting the semantics of that function. The signature consists of following features:

- *Read and Written Values*: The feature consists of global (or static) variable values read from or written to the memory during an (emulated) execution. It contains the input and output values of the function when provided with a specific input, indicating the semantics of the function.
- *Comparison Operand Values*: The feature is composed of values for comparison operations whose results decide the following control flow of an (emulated) execution. It indicates the path of the function followed by an input to generate the output. Thus, it is semantics related as well.
- *Invoked Standard Library Functions*: Standard library functions provide fundamental operations for implementing user-defined functions (e.g., `malloc`, `memcpy`). The feature has been shown to be semantics-related and effective for code clone analysis [41], [40]. Therefore, it is adopted as complement to the semantic signature of BINMATCH.

During the execution or emulation, BINMATCH captures the sequence of above features, and considers the sequence as the signature of a binary function for latter similarity comparison.

B. Instrumentation and Execution

In this step, BINMATCH instruments a binary function F to generate its signature by running test cases. Meanwhile, runtime information for Emulation (§III-C) is recorded as well.

Algorithm 1 presents the pseudo-code of instrumentation. BINMATCH traverses each instruction (\mathcal{I}) of F . If \mathcal{I} accesses global variables, performs comparison operations, or calls a standard library function, BINMATCH injects code before \mathcal{I}

Algorithm 1: Algorithm of Instrumentation

Input: Instruction under Analysis \mathcal{I}
Output: Instruction after Instrumentation \mathcal{I}_r

```

1 Algorithm Instrumentation ( $\mathcal{I}$ )
2    $\mathcal{I}_r \leftarrow \mathcal{I}$ 
3   // capture features for the signature
4   if  $\mathcal{I}$  accesses global/static data then
5      $\mathcal{I}_r \leftarrow \text{record\_data\_val}(\mathcal{I}_r)$ 
6   if  $\mathcal{I}$  performs comparison then
7      $\mathcal{I}_r \leftarrow \text{record\_opr\_val}(\mathcal{I}_r)$ 
8   if  $\mathcal{I}$  calls a standard library function then
9      $\mathcal{I}_r \leftarrow \text{record\_libc\_name}(\mathcal{I}_r)$ 
10  // record runtime information
11  if  $\mathcal{I}$  reads an argument of the function then
12     $\mathcal{I}_r \leftarrow \text{record\_arg\_val}(\mathcal{I}_r)$ 
13  else if  $\mathcal{I}$  calls a function indirectly then
14     $\mathcal{I}_r \leftarrow \text{record\_func\_addr}(\mathcal{I}_r)$ 
15  else if a function returns then
16     $\mathcal{I}_r \leftarrow \text{record\_ret\_val}(\mathcal{I}_r)$ 
17  return  $\mathcal{I}_r$ 

```

to capture corresponding features and generate the signature of F (Line 4-9).

Line 11-16 present code for recording runtime values of F 's execution. According to *cdecl*, the default calling convention of IA-32 binaries, function arguments are prepared by callers and passed through the stack. Therefore, if \mathcal{I} reads a variable which is pushed onto the stack before the invocation of F , BINMATCH considers the variable as a function argument and records its value (Line 11-12). Besides, BINMATCH records the addresses of subroutines invoked by F indirectly (Line 13-14). The return values of all subroutines, including both user-defined functions and standard library functions, are recorded as well (Line 15-16).

C. Emulation

For every target function T to be compared with the template function F , BINMATCH emulates its execution with the runtime information extracted from the last step. The semantic signature of T is captured simultaneously. Clone functions should behave similarly if they are executed with the same input [11]. Namely, if T is the clone match of F , their signatures should be similar. Algorithm 2 presents the pseudo-code of emulation. BINMATCH provides T with the arguments of F (Line 2), and emulates it with the runtime

Algorithm 2: Algorithm of Emulation

Input: Emulated Memory Space of the Target Function \mathcal{M}
Input: Runtime Value Set of the Template Function \mathcal{S}

```
1 Algorithm Emulation ( $\mathcal{M}, \mathcal{S}$ )
2   assign_func_arg ( $\mathcal{M}, \mathcal{S}$ )
3   foreach instruction  $I$  to be emulated do
4     if  $I$  reads global variables then
5        $addr \leftarrow$  get_var_addr ( $I$ )
6       if  $addr$  is accessed for the first time then
7         migrate_var_val ( $\mathcal{M}, addr, \mathcal{S}$ )
8     if  $I$  calls a function indirectly then
9        $addr \leftarrow$  get_tar_addr ( $I$ )
10      if  $addr \in \mathcal{S}$  then
11        migrate_ret_val ( $\mathcal{M}, addr, \mathcal{S}$ )
12      else exit_emulation()
13    if  $I$  invokes a standard library function then
14       $libc \leftarrow$  get_func_name ( $I$ )
15      if  $libc$  needs system supports then
16        migrate_ret_val ( $\mathcal{M}, libc, \mathcal{S}$ )
17    // capture features for the signature
18    if  $I$  captures features then
19      record_feat_val ( $\mathcal{M}, I$ )
20    emulate_inst ( $\mathcal{M}, I, \mathcal{S}$ )
```

information of F (Line 20). Besides, BINMATCH records the features of T to generate its signature (Line 18-19). Next, we discuss the algorithm for emulation in more details.

1) *Function Argument Assignment:* In our scenario, binary functions for comparison are compiled from the same code base, i.e., clone functions have the same number of arguments. According to the calling convention, BINMATCH recognizes the arguments of the target function T . If the argument number of T equals to that of F , BINMATCH assigns argument values of F to those of T in order. Otherwise, BINMATCH skips the emulation of T which cannot be the match of F . For example, F and T have the following argument lists:

```
F(farg_0, farg_1, farg_2)
T(targ_0, targ_1, targ_2)
```

If BINMATCH has the values of $farg_0$ and $farg_2$ that F only accesses the tow arguments in the execution, BINMATCH assigns their values to $targ_0, targ_2$ separately. To make the emulation smoothly, arguments without corresponding values ($targ_1$) are assigned with a predefined value (e.g., 0xDEADBEEF).

2) *Global Variable Reading:* In the execution of the template function F , it might read global (or static) variables whose values have been modified by former executed code. To emulate the target function T in the memory space of F , BINMATCH migrates global variable values of F to corresponding addresses which T reads from. BINMATCH needs to consider two points: i) getting the global variable addresses which T reads from (Line 5 of Algorithm 2), and ii) migrating the corresponding global variable value from the memory of F to that of T (Line 7 of Algorithm 2).

```
1  mov    ecx, gvar1          1  mov    ecx, gvar1'
2  test   ecx, ecx           2  mov    ebp, gvar2'
3  mov    eax, gvar2         3  test   ebp, ebp
4  add    ecx, eax           4  add    ebp, ecx
```

(a) Template Function (F)

(b) Target Function (T)

Fig. 2: Global Variable Value Migration

```
1  mov    edx, [ebp-0E4h] ; load a local variable
2  lea   eax, [edx-0Ah] ; get the index
3  cmp   eax, 2Ah
4  ja    loc_8052880 ; the default case
5  jmp   dword ptr [eax*4+808F630h]; indirect jump
```

Fig. 3: Indirect Jump of a Switch

Global variables are stored in specific sections of a binary program (e.g., `.data`). The size of each variable is decided by the source code. The location of the variable, including the base address of a global data structure (e.g., array), is determined in the binary code after compilation and not changed afterward. Thus, global variables are accessed with hard-coding addresses. Each member of a global data structure is accessed by adding its corresponding offsets to the constant base address, and the offset is generated from the input (function arguments). Hence, BINMATCH is able to obtain global variable addresses of T easily during the emulation.

BINMATCH migrates global variable values according to their usage order. Figure 2 shows an example of two functions for global variable value migration. During the execution of F , two global variables `gvar1` and `gvar2` are read at Line 1 and Line 3 separately in Figure 2a. `gvar1` is used to test its value at Line 2, and `gvar2` is used for the addition operation at Line 4. So the usage order of the two variable is `[gvar1, gvar2]`. When emulating T in Figure 2b, BINMATCH identifies `ecx` and `ebp` are loaded with global variables `gvar1'` and `gvar2'` at Line 1 and Line 2. Then, it finds `ebp` is used for testing at Line 3, and `ecx` is used for the addition at Line 4 afterward. The usage order of the global variables in Figure 2b is `[gvar2', gvar1']`. Therefore, BINMATCH assigns the value of `gvar1` to `gvar2'`, and `gvar2` to `gvar1'` accordingly. If there are not enough global values to assign (e.g., T reads two global variables but F reads only one), BINMATCH provides the surplus variables of T with predefined values (e.g., 0xDEADBEEF).

3) *Indirect Calling/Jumping:* Targets of indirect calls are decided by the input at runtime. Since the target function T is emulated in the memory space of the template function F , if T is the clone match of F , the indirect call targets of T should be those invoked during the execution of F . BINMATCH then migrates the return values of F to corresponding indirect calls of T (Line 10-11 in Algorithm 2). Otherwise, the target function under emulation cannot be the match of F . BINMATCH stops the process and exits (Line 12 in Algorithm 2).

An indirect jump (or branch) is implemented with a jump table which contains an ordered list of target addresses. Jump tables are stored in `.rodata`, the read-only data section of an executable. Therefore, similar to the reading of a global

data structure, a jump table entry is accessed by adding the offset to the base address of the jump table. The base address is a constant value, and the offset is computed from the input.

Figure 3 shows an indirect jump of a switch structure. At Line 2, the index value is computed with `edx`, a value of an input-related local variable, and stored in `eax`. If the index value is not above `0x2A`, which represents the default case, an indirect jump is performed according to the jump table whose base address is `0x808F630` (Line 5). As entries of a jump table are sorted, with identical input, clone code would have equal offset and jumps to the path of the same semantics. BINMATCH just follows the emulation and has no need to do extra work for indirect jumps.

4) *Standard Library Function Invocation*: If the target function T calls a standard library function which requests the system support (e.g., `malloc`), BINMATCH skips its emulation and assigns it with the result of the corresponding one invoked by the template function F (Line 15-16). For example, F and T calls following library functions in sequence:

```
F: malloc_0, memcpy, malloc_1
T: malloc_0', memset, malloc_1'
```

BINMATCH assigns return values of `malloc_0`, `malloc_1` to `malloc_0'`, `malloc_1'` separately, and skips the emulation. `memset` is emulated normally, because it has no need for the system support.

D. Similarity Comparison

BINMATCH has captured the semantic signature (feature sequence) of the template function via execution, and those of target functions via emulation. In this step, it computes the similarity score of the template function signature and that of each target function in pairs. We utilize the Longest Common Subsequence (LCS) algorithm [4] to the similarity measurement. On one hand, a signature is captured from the (emulated) execution of a function. The appearance order of each entry in the signature is a feature as well. On the other hand, a signature is captured from optimized or obfuscated binary programs that it contains diverse or noisy entries in the sequence. LCS not only considers the element order of two sequences for comparison, but also allows skipping non-matching elements, which tolerates code optimization and obfuscation. Hence, the LCS algorithm is suitable for signature similarity comparison of BINMATCH.

The similarity score is measured by the Jaccard Index [15]. Given two semantic signatures S_f and S_t , the Jaccard Index is calculated as followed:

$$J(S_f, S_t) = \frac{|S_f \cap S_t|}{|S_f \cup S_t|} = \frac{|S_f \cap S_t|}{|S_f| + |S_t| - |S_f \cap S_t|} \quad (1)$$

Here, $|S_f|$ and $|S_t|$ are the lengths of sequence S_f and S_t . $|S_f \cap S_t|$ is the LCS length of the two sequences. $J(S_f, S_t)$ ranges from 0 to 1, which is closer to 1 when S_f and S_t are considered more similar.

After this step, BINMATCH generates a target function list which is ranked by the similarity scores in descending order.

IV. IMPLEMENTATION

Currently, BINMATCH supports IA-32 binary function clone analysis of ELF (Executable and Linkable Format) files. Next, we discuss the key aspects of the implementation.

A. Binary Function Boundary Identification

BINMATCH requires addresses and lengths of binary functions to perform clone analysis. Given an ELF file, We leverage `IDA Pro v6.6` [7], an industrial strength reverse engineering tool, to disassemble it, identifying the boundaries of each binary function. The plugin of `IDA Pro`, `IDAPython`, provides interfaces to obtain addresses of functions, i.e., `Functions(start, end)` which returns a list of function first addresses between `start` and `end`. Therefore, we develop a script with `IDAPython` to acquire function addresses of binary files automatically. Although the resulting disassembly of `IDA Pro` is not perfect [2], it is sufficient for BINMATCH.

B. Instrumentation and Emulation

We implement the instrumentation module of BINMATCH with `Valgrind` [33], a dynamic instrumentation framework. `Valgrind` unifies binary code under analysis into `VEX-IR`, a RISC-like intermediate representation (IR), and injects instrumentation code into the IR code. Then, it translates the instrumented IR code into binaries for execution. IR translation unifies the operations of binary code and facilitates the process of signature extraction. For example, memory reading and writing operations are all unified with `Load` and `Store`, the opcodes defined by `VEX-IR`. Hence, we just concentrate on the specific operations of IR and ignore the complex instruction set of IA-32.

The step of emulation is implemented basing on `angr` [36], a static binary analysis framework. `angr` borrows `VEX-IR` from `Valgrind`, translating binary code to be analyzed into IR statically. Given a user-defined initial state, it provides a module named `SimProcedure` to emulate the execution of IR code. `SimProcedure` allows injecting extra code to monitor the emulation of the IR code. It actually emulates the process of instrumentation. Besides, `angr` maintains a database of standard library functions to ease the emulation of those functions (§III-C4). Thus, we develop a script of monitoring code, which is similar to the instrumentation code developed with `Valgrind`, to capture semantic signatures during the emulation with `angr`.

C. Similarity Comparison

As the length of a signature sequence might have the scale over 10^4 , the possibility is high for the traditional LCS algorithm, whose memory complexity is $O(mn)$, to run out the memory. We implement BINMATCH to compute LCS with the Hirschberg's Algorithm [17] which needs only $O(\min(m, n))$ memory space.

TABLE I: Object Projects of Evaluation

Program	Version	Description
convert	6.9.2	Command-line interface to the ImageMagick image editor/converter
curl	7.39	Command-line tool for transferring data using various protocols
ffmpeg	2.7.2	Program for transcoding multimedia files
gzip	1.6	Program for file compression and decompression with the DEFLATE algorithm
lua	5.2.3	Scripting parser for Lua, a lightweight, multi-paradigm programming language
mutt	1.5.24	Text-based email client for Unix-like systems
openssl	1.0.1p	Toolkit implementing the TLS/SSL protocols and a cryptography library
wget	1.15	Program retrieving content from web servers via multiple protocols

V. EVALUATION

We conduct empirical experiments to evaluate the effectiveness and capacity of BINMATCH. Firstly, BINMATCH is evaluated with binaries compiled with different compilation configurations, including variant optimization options and compilers. The results are then compared to those of existing solutions (§V-B). Secondly, we evaluate the effectiveness of BINMATCH in handling obfuscation by comparing binary programs with their obfuscated versions (§V-C). Lastly, with the motivating example of NConvert described in §II-A, we show how BINMATCH locates the statically-linked defective function of NConvert (§V-D).

A. Experiment Setup

The evaluation is performed in the system Ubuntu 16.04 which is running on an Intel Core i5-2320 @ 3GHz CPU with 8G DDR3-RAM.

1) *Dataset*: We adopt programs of eight real-world projects as objects of the evaluation, as listed in Table I. The object programs have various functionalities, such as data compression (gzip), code parsing (lua), email posting (mutt), etc. With those object programs, the effectiveness of BINMATCH is shown to be not limited by the type of programs and functions under analysis.

In the first group of experiments (§V-B), the object programs are compiled with different compilers, i.e., gcc v4.7 and clang v3.8.0, and variant optimization options, i.e., -O3 and -O0. In the second group of experiments (§V-C), we adopt Obfuscator-LLVM v4.0.1 (OLLVM) [22] to obfuscate the object programs for comparison. OLLVM provides three widely used techniques for obfuscation. We leverage the three techniques to obfuscate the object programs which are optimized with -O3 and -O0 respectively. Therefore, we compile 10 (= 2 · 2 + 3 · 2) unique binary executables for each object program, overall 80 (= 10 · 8) for the evaluation.

For each experiment, we select two from the 10 executables of an object program, i.e., E_{tem} (the template executable) and E_{tar} (the target executable). BINMATCH executes E_{tem} with test cases obtained from its project, considering each

executed function as a template function. Then it compares every template function to all target functions of E_{tar} in pairs to find the clone match. On average, 291 functions are triggered in an execution of E_{tem} , and E_{tar} contains 4,353 functions. As a result, BINMATCH totally performs over 100 million pairs of function comparison in all the experiments.

2) *Ground Truth*: All the 80 executables are stripped that their debug and symbol information is discarded for the evaluation. To verify the correctness of the experimental results, we compile *extra unstripped copies* of the 80 executables, and establish the ground truth with their debug and symbol information.

For each template function, BINMATCH generates a list of target functions ranked by the similarity scores in descending order (as described in §III-D). According to the ground truth, if the symbol name of the *Top 1* function (with the highest similarity score) in the resulting list is the same as that of the template function, the match is considered to be *correct*. Besides, we manually verify cases of function inline. For example, function A invokes B in E_{tem} , while the corresponding function B' is inlined into A' becoming $A'B'$ in E_{tar} , and B' disappears. If BINMATCH matches B with $A'B'$, we consider it *correct* as well.

3) *Evaluation Metrics*: Similar to previous work [11], [18], we measure the performance of BINMATCH with accuracy, the percentage of executed template functions whose correct matches are found. The formula is as followed:



$$Accuracy = \frac{|Correct\ Matches|}{|Template\ Functions|} \quad (2)$$

B. Accuracy across Compilation Configurations

1) *Cross-optimization Analysis*: In this section, we leverage BINMATCH to match clone functions of different optimizations. For a compiler, higher optimization options contain all strategies specified by lower ones. Taking gcc as an example, the option -O3 enables all 88 optimizations of -O2, and turns on another 14 optimization flags in addition. Thus, we only discuss the case of -O3 (E_{tem}) versus -O0 (E_{tar}), which has larger differences than any other pair of cross-optimization analysis.

Figure 4 shows the results of cross-optimization analysis for each object program compiled by gcc (Figure 4a) and clang (Figure 4b) separately. In Figure 4a, BINMATCH achieves the accuracy over 82.0% for each object program, and the average accuracy is 91.5%. For every executable compiled by clang in Figure 4b, BINMATCH correctly detects over 80.0% functions of each object as well, and the average accuracy is 92.0%.


We observe that *function inline* is a reason leading to the incorrect matches. For example, template A calls B , while the corresponding target function B' is inlined into A' becoming $A'B'$. Because the semantic signature of $A'B'$ contains those of both A' and B' , signature length of A is shorter than that of $A'B'$. Hence, the similarity score of function pair ($A, A'B'$) might be relative small, and BINMATCH reports an incorrect match.




(a) Gcc -O3 vs -O0

(b) Clang -O3 vs -O0

Fig. 4: Accuracy of Cross-optimization Analysis



(a) Gcc -O3 vs Clang -O0



(b) Clang -O3 vs Gcc -O0

Fig. 5: Accuracy of Cross-compiler Analysis

2) *Cross-compiler Analysis*: In this section, BINMATCH is evaluated with binaries compiled by different compilers. Similar to the cross-optimization analysis, only the case of `-O3` versus `-O0` is considered. The results are presented in Figure 5. For comparisons between `gcc -O3` (E_{tem}) and `clang -O0` (E_{tar}), BINMATCH gives the accuracy all over 84.0%, and the average accuracy is 90.3%, as shown in Figure 5a. Additionally, in Figure 5b, BINMATCH achieves an average accuracy of 91.4% for the setting of `clang -O3` (E_{tem}) versus `gcc -O0` (E_{tar}). The accuracy of each object program exceeds 85.0%.

In addition to function inline introduced by different optimizations, we find *floating-point number* is another reason resulting in incorrect matches. `gcc` leverages `x87` floating-point instructions to implement corresponding operations, while `clang` uses the SSE (Streaming SIMD Extensions) instruction set. `x87` adopts the FPU (floating point unit) stack to assist in processing floating-point numbers. The operations deciding whether the stack is full or empty add redundant entries to the semantic signature of comparison

TABLE II: Accuracy compared with the state-of-the-art solution Kam1n0 and the industrial standard tool BinDiff

Setting	BINMATCH	Kam1n0	BinDiff
<code>gcc -O3 vs. gcc -O0</code>	0.915	0.294	0.331
<code>clang -O3 vs. clang -O0</code>	0.920	0.252	0.506
<code>gcc -O3 vs. clang -O0</code>	0.903	0.216	0.385
<code>clang -O3 vs. gcc -O0</code>	0.914	0.273	0.501
Average Accuracy	0.916	0.265	0.409

operand values (§III-A). In contrast, SSE directly operates on a specific register set (e.g., XMM registers) and has no extra operations. Besides, `x87` could handle single precision, double precision, and even 80-bit double-extended precision floating-point calculation, while SSE mainly processes single-precision data. Due to the different precision of representations, even though the floating-point numbers are the same, their values generated by the two compilers are not equal, which therefore results in the incorrect matches.

3) *Comparison with Existing Work*: In this section, we compare BINMATCH to the state-of-the-art solution Kam1n0 [10] and the industrial standard tool BinDiff v4.2.0 supported by Google [14], [47] from the perspective of detection accuracy. Because Kam1n0 and BinDiff are both made available to the public, we could use the two solutions to detect binary clone functions with the same settings as BINMATCH. BINMATCH is evaluated by the detection accuracy of executed template functions. To perform fair comparison, we measure the accuracy of Kam1n0 and BinDiff by detecting clones of those template functions as well. The results are presented in Table II. Obviously, BINMATCH outperforms Kam1n0 and BinDiff in detecting binary clone functions across compilation configurations.

Kam1n0 and BinDiff are typical solutions which rely on syntax and structure features to detecting binary clone functions. Kam1n0 captures features of a function from its control flow graph (CFG), and encode the features as a vector for indexing. Thus, essentially, it detects clone functions by analyzing graph isomorphism of CFG. The relatively low accuracy of Kam1n0 indicates that compilation configurations indeed affect representations of binaries, even though two pieces of code are compiled from the same code base. In addition to measuring the similarity of CFG, BinDiff considers other features to detect clone functions, such as function hashing which compares the hash of raw function bytes, call graph edges which match functions basing on the dependencies in the call graphs, etc. By carefully choosing suitable features to measure the similarity of functions, BinDiff becomes resilient towards variant compilers as well as optimization options to an extent. Therefore, it performs better than Kam1n0, but is still at an apparent disadvantage compared to BINMATCH.

4) *Processing Time*: BINMATCH analyzes binary function clones in three steps: instrumentation & execution, emulation, and similarity comparison (as shown in Figure 1). According to Algorithm 1, BINMATCH injects code only to record runtime information and semantic signatures, and does not do online analysis. Thus, the overhead of instrumentation is low.

As described in Emulation (§III-C), BINMATCH merely emulates the user code of a target function, and borrows runtime values directly from the template function, skipping the emulation of specific system operations, such as allocating memory, reading or writing a file, etc. As a result, it would not take much time to emulate a function. In above experiments, BINMATCH spends 4.3 CPU seconds emulating a function on average. Since BINMATCH adopts LCS to compute similarity scores whose time complexity is relative high, the step of similarity comparison occupies the most processing time. In the experiments, it costs 573.9 CPU seconds on average to complete a pair of function comparison.

C. Accuracy of Matching Obfuscated Code

In this section, we conduct experiments to compare normal binary programs with their corresponding obfuscated code. We compile the object programs with the setting of `gcc -O3` as the normal code (E_{tem}). We use all the three obfuscation meth-

TABLE III: Accuracy of analyzing obfuscated code. The template binaries are compiled with `gcc -O3`. OLLVM adopts `clang` as its compiler.

Target	Obfuscation	BINMATCH	BinDiff
ollvm -O3 ollvm -O0	Instructions Substitution	0.891	0.676
		0.887	0.302
ollvm -O3 ollvm -O0	Bogus Control Flow	0.843	0.295
		0.796	0.281
ollvm -O3 ollvm -O0	Control Flow Flattening	0.874	0.464
		0.791	0.323
Average Accuracy	/	0.847	0.411

ods provided by OLLVM to obfuscate the object programs generated with `clang -O3` and `clang -O0` separately (E_{tar} , OLLVM adopts `clang` as its compiler).

The experimental results are shown in Table III. Results of BinDiff are also presented as references. *Instruction substitution* replaces standard operators (e.g., addition operators) with sequences of functionality-equivalent, but more complex instructions. It obfuscates code on the *syntax* level, affecting the detection accuracy of BinDiff, but posing few threats to BINMATCH which is semantics-based.

Bogus control flow (BCF) adds opaque predicates to a basic block, which breaks the original basic block into two. *Control flow flattening* (FLA) generally breaks a function up into basic blocks, then encapsulates the blocks with a selective structure (e.g., the switch structure) [27]. It creates a state variable for the selective structure to decide which block to execute next at runtime via conditional comparisons. BCF and FLA both changes the *structure* of the original function, i.e., modifying the control flow. They insert extra code which is irrelevant to the functionality of the original function, generating redundant semantic features which are indistinguishable from normal ones (e.g., comparison operand values of opaque predicates). Thus, they affect the detection accuracy of BINMATCH. When analyzing binaries optimized with `-O0`, it correctly detects 79.6% of functions obfuscated by BCF, and 79.1% by FLA, while the average accuracy of `gcc -O3` vs. `clang -O0` is 90.3%. However, BINMATCH still achieves twice the average accuracy of BinDiff, i.e., 84.7% of BINMATCH and 41.1% of BinDiff.

D. Case Study: *libpng* vs. *NConvert*

As described in §II-A, before the version of 1.5.14, *libpng* contains an integer overflow vulnerability in function *png_set_unknown_chunks*. *NConvert*, a closed-source image processor, statically links the library to handle files of the PNG format. In this section, we download the source code of *libpng* v1.5.12 and the executable of *NConvert* v6.17 from their home pages, aiming to locate *png_set_unknown_chunks* in *NConvert* with BINMATCH.

We compile *libpng* v1.5.12 with the default configurations, i.e., `gcc -O2`. Then the clone analysis is performed in following steps:

- 1) We run a test case of *libpng* from its project to cover the template function *png_set_unknown_chunks*,

and recored its semantic signature as well as runtime information with BINMATCH. As a result, the signature of the template function contains 133 entries (features).

- 2) `png_set_unknown_chunks` has 4 arguments. Thus, BINMATCH only emulates 337 target functions of NConvert whose identified argument number is 4 as well. Overall, the process of emulation takes 5987.0 CPU seconds.
- 3) BINMATCH totally spends 57.0 CPU hours to compare the signature of the template function and those of 337 target functions. It reports `func_81ad770` (the function at 0x81AD770 in NConvert) achieves the highest similarity score $0.378 (= \frac{99}{133+228-99})$, the signature length is 228 and LCS length is 99).

By manual verification, we find the result is correct that `func_81ad770` is the clone match of `png_set_unknown_chunks`. After locating the target function, analysts could do further analysis on it, e.g., checking whether the function is vulnerable, which is out of the scope of this paper.

E. Threats to Validity

BINMATCH is implemented with `Valgrind` and `angr` which both adopt VEX-IR as the intermediate representation (§IV-B). However, VEX-IR is not perfect that 16% x86 instructions could not be lifted, although only a small subset of instructions is used in executables in practice and VEX-IR could handle most cases [26]. The incompleteness of VEX-IR might affect the accuracy of semantics signature extraction, while BINMATCH still produces promising results in above experiments.

VI. DISCUSSION AND FUTURE WORK

A. Scope of Application

In the evaluation, BINMATCH is shown to be effective in analyzing obfuscated binary clone code which is generated by OLLVM (§V-C). The robust of BINMATCH is due to the nature of dynamic analysis and the adoption of semantic signatures. However, that does not mean BINMATCH could handle all kinds of obfuscations. Besides, the OLLVM code actually affects the accuracy of BINMATCH in the experiments. When analyzing benign code, BINMATCH achieves higher average accuracy which is 91.6%, while the ratio of obfuscated code is 84.7%. In the literature, deobfuscation has been well studied [38], [43], [44]. Therefore, if BINMATCH fails to detect an obfuscated function, it is a better choice to deobfuscate it firstly, then perform further analysis.

In this paper, we present BINMATCH to analysis binary programs of ELF (Executable and Linkable Format) on the IA-32 architecture. Because the method is semantics-based, BINMATCH could be ported to other platforms, for example, PE (Portable Executable) files on Windows. Besides, BINMATCH is implemented basing on `Valgrind` and `angr` which both support cross architecture analysis. Hence, BINMATCH is applicable for multiple architectures, such as x86-64, ARM, MIPS, etc. We leave it as future work.

B. Inline Function Detection

As discussed in the section of Evaluation Metrics (§V-B1), function inline poses a threat to the accuracy of BINMATCH. Empirically, a compiler inlines a function because the function is short and invoked for numerous times. Namely, size and invocation times might be features of an inline function. Thus, it is possible to detect the inline functions with machine learning techniques. If a function is considered as the potential inline function, we could combine it to its callers and capture the signature. That is left as future work.

C. Scalability

The step of similarity comparison (§III-D) is the performance bottleneck of BINMATCH. It calculates the similarity score of two signatures with the LCS algorithm whose time complexity is high. However, the comparisons of function pairs are unrelated to each other. The step could be implemented in parallel to reduce the total processing time. Besides, MinHash [37] is a possible solution for the similarity comparison. It calculates the Jaccard Index directly without computing the LCS of the two signatures. However, MinHash treats the signature sequence as a set, discarding the order information of elements in the sequence, which is a potential semantic feature (as discussed in §III-D). Therefore, MinHash is a trade-off between accuracy and efficiency.

VII. RELATED WORK

Code clone (or similarity) analysis is a classic topic of software engineering. Due to the code reuse of software development, automatically identifying clone code becomes a common requirement of software maintenance (e.g., bug detection). The technique is also applied in other fields, e.g., malware analysis of security. In the last twenty years, researchers have made much effort into source code clone analysis, typically including CCFinder [23], DECKARD [21], CloneDR [3], CP-Miner [28], etc. As the focus of this paper is clone analysis on binary code, which has its own challenges and scenarios, we would not talk about work on source code in more details. Next, we mainly discuss the related work on binary code clone analysis.

Syntax and structural features are widely adopted to detect binary clone code. Sæbjørnsen et al. [35] detect binary clone code basing on opcode and operand types of instructions. Hemel et al. [16] treat binary code as text strings and measure similarity by data compression. The higher the compression rate is, the more similar the two pieces of binary code are. Khoo et al. [25] leverage n-gram to compare the control flow graph (CFG) of binary code. David et al. [9] measure the similarity of binaries with the edit distances of their CFGs. BinDiff [14] and Kam1n0 [10] extract features from the CFG and call graphs to search binary clone functions.

As discussed earlier in this paper, the main challenge of binary code clone analysis is semantics-equivalent code transformation, such as link-time optimization, obfuscation, etc. Because of the transformation, representations of binary code are altered tremendously, even though the code is compiled

from the same code base. Therefore, syntax- and structure-based methods become ineffective, and semantics-based methods prevail. Jhi et al. [20] and Zhang et al. [45] leverage runtime invariants of binaries to detect software and algorithm plagiarism. Ming et al. [32] infer the lineage of malware by code clone analysis with the system call traces as the semantic signature. However, those solutions require the execution of binary programs and cannot cover all target functions. Egele et al. [11] propose blanket execution to match binary functions with full code coverage which is achieved at the cost of detection accuracy. Luo et al. [31] and Zhang et al. [46] detect software plagiarism by symbolic execution. Although their work is resilient to code transformation, symbolic execution is trapped in the performance of SMT/SAT solvers which cannot handle all cases, e.g., indirect calls. David et al. [8] decompose the CFG of a binary function into small blocks, and measure the similarity of the small blocks basing on a statistical model. However, the boundaries of CFG blocks would be changed by code transformation, affecting the accuracy of the method.

More recently, with the prevalence of IoT devices, binary code clone analysis is proposed to perform across architectures. Multi-MH [34], discovRE [12], Genius [13] are proposed to detect known vulnerabilities and bugs in multi-architecture binaries via code clone analysis. BinGo [6] and CACompare [19] are proposed to analyze the similarity of binary code across architectures as well. However, discovRE and Genius still depend heavily on the CFG of a binary function. Multi-MH, BinGo and CACompare sample a binary function with random values to capture corresponding I/O values as the signature, while the random values are meaningless that they merely trigger limited behaviors of the function. Therefore, it is difficult for them to cover the core semantics of binary code.

To sum up, the topic of binary code clone analysis mainly focuses on two points: i) what signature to adopt, such as opcodes and operand types (syntax), CFG (structure) and system calls (semantics); ii) how to capture the signatures, such as statically disassembling, dynamically running and sampling, etc. BINMATCH leverage the combination of read and written values, comparison operand values, and invoked library functions as the signature which is able to better reveal the semantics of binary code. Besides, it captures the signature via both execution and emulation, which not only ensures the richness of semantics, but also covers all target functions to be analyzed.

VIII. CONCLUSION

In this paper, we propose BINMATCH, a hybrid approach, to detect binary clone functions. BINMATCH executes the template function with its test cases, and migrates the runtime information to target functions in order to emulate their executions. During the execution and emulation, BINMATCH captures semantic signatures of the functions for clone analysis. The experimental results show that BINMATCH is robust to semantics-equivalent code transformation, including different compilation configurations and commonly-used obfuscations.

Besides, we show that BINMATCH performs better than the state-of-the-art solutions to binary code clone analysis, such as BinDiff and KamIn0.

IX. ACKNOWLEDGMENTS

We would like to thank the anonymous reviewers for their insightful comments which greatly help to improve the manuscript. This work is partially supported by the Key Program of National Natural Science Foundation of China (Grant No.U1636217), the National Key Research and Development Program of China (Grant No.2016YFB0801200), and a research grant from the Ant Financial Services Group.

REFERENCES

- [1] W. Andrew and L. Arun. The software similarity problem in malware analysis. In *Duplication, Redundancy, and Similarity in Software*. Internationales Begegnungs- und Forschungszentrum für Informatik (IBFI), Schloss Dagstuhl, Germany, 2007.
- [2] D. Andriessie, X. Chen, V. van der Veen, A. Slowinska, and H. Bos. An in-depth analysis of disassembly on full-scale x86/x64 binaries. In *25th USENIX Security Symposium (USENIX)*. USENIX Association, 2016.
- [3] I. D. Baxter, A. Yahin, L. Moura, M. Sant'Anna, and L. Bier. Clone detection using abstract syntax trees. In *Proceedings of the 6th International Conference on Software Maintenance (ICSM)*, 1998.
- [4] L. Bergroth, H. Hakonen, and T. Raita. A survey of longest common subsequence algorithms. In *String Processing and Information Retrieval, 2000. SPIRE 2000. Proceedings. Seventh International Symposium on*, pages 39–48. IEEE, 2000.
- [5] D. Brumley, P. Poosankam, D. Song, and J. Zheng. Automatic patch-based exploit generation is possible: Techniques and implications. In *Security and Privacy, 2008. SP 2008. IEEE Symposium on*, pages 143–157. IEEE, 2008.
- [6] M. Chandramohan, Y. Xue, Z. Xu, Y. Liu, C. Y. Cho, and H. B. K. Tan. Bingo: Cross-architecture cross-os binary search. In *Proceedings of the 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering (FSE)*, 2016.
- [7] R. Data. Ida pro disassembler. <https://www.datarescue.com/ibase/>.
- [8] Y. David, N. Partush, and E. Yahav. Statistical similarity of binaries. In *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI'16*, pages 266–280, New York, NY, USA, 2016. ACM.
- [9] Y. David and E. Yahav. Tracelet-based code search in executables. In *Proceedings of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI'14*, 2014.
- [10] S. H. Ding, B. Fung, and P. Charland. KamIn0: Mapreduce-based assembly clone search for reverse engineering. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 461–470. ACM, 2016.
- [11] M. Egele, M. Woo, P. Chapman, and D. Brumley. Blanket execution: dynamic similarity testing for program binaries and components. In *Proceedings of the 23rd USENIX Security Symposium (USENIX Security)*, 2014.
- [12] S. Eschweiler, K. Yakdan, and E. Gerhards-Padilla. discovRE: Efficient cross-architecture identification of bugs in binary code. 2016.
- [13] Q. Feng, R. Zhou, C. Xu, Y. Cheng, B. Testa, and H. Yin. Scalable graph-based bug search for firmware images. In *23rd ACM Conference on Computer and Communications Security (CCS)*, 2016.
- [14] H. Flake. Structural comparison of executable objects. *Proceedings of the First International Conference on Detection of Intrusions and Malware and Vulnerability Assessment (DIMVA)*, 2004.
- [15] L. Hamers et al. Similarity measures in scientometric research: The jaccard index versus salton's cosine formula. *Information Processing and Management*, 25(3):315–18, 1989.
- [16] A. Hemel, K. T. Kalleberg, R. Vermaas, and E. Dolstra. Finding software license violations through binary code clone detection. In *Proceedings of the 8th Working Conference on Mining Software Repositories (MSR)*, 2011.

- [17] D. S. Hirschberg. A linear space algorithm for computing maximal common subsequences. *Communications of the ACM*, 18(6):341–343, 1975.
- [18] Y. Hu, Y. Zhang, J. Li, and D. Gu. Cross-architecture binary semantics understanding via similar code comparison. In *2016 IEEE 23rd International Conference on Software Analysis, Evolution, and Reengineering (SANER)*, volume 1, pages 57–67. IEEE, 2016.
- [19] Y. Hu, Y. Zhang, J. Li, and D. Gu. Binary code clone detection across architectures and compiling configurations. In *Program Comprehension (ICPC), 2017 IEEE/ACM 25th International Conference on*, pages 88–98. IEEE, 2017.
- [20] Y.-C. Jhi, X. Wang, X. Jia, S. Zhu, P. Liu, and D. Wu. Value-based program characterization and its application to software plagiarism detection. In *Proceedings of the 33rd International Conference on Software Engineering (ICSE)*, 2011.
- [21] L. Jiang, G. Misherghi, Z. Su, and S. Glondu. Deckard: Scalable and accurate tree-based detection of code clones. In *Proceedings of the 29th International Conference on Software Engineering (ICSE)*, 2007.
- [22] P. Junod, J. Rinaldini, J. Wehrli, and J. Michielin. Obfuscator-LLVM – software protection for the masses. In B. Wyseur, editor, *Proceedings of the IEEE/ACM 1st International Workshop on Software Protection, SPRO’15, Firenze, Italy, May 19th, 2015*, pages 3–9. IEEE, 2015.
- [23] T. Kamiya, S. Kusumoto, and K. Inoue. Cefinder: a multilinguistic token-based code clone detection system for large scale source code. *IEEE Transactions on Software Engineering*, 2002.
- [24] U. Kargén and N. Shahmehri. Turning programs against each other: high coverage fuzz-testing using binary-code mutation and dynamic slicing. In *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering*, pages 782–792. ACM, 2015.
- [25] W. M. Khoo, A. Mycroft, and R. Anderson. Rendezvous: A search engine for binary code. In *Proceedings of the 10th Working Conference on Mining Software Repositories (MSR)*, 2013.
- [26] S. Kim, M. Faerevaag, M. Jung, S. Jung, D. Oh, J. Lee, and S. K. Cha. Testing intermediate representations for binary analysis. In *Proceedings of the 32nd IEEE/ACM International Conference on Automated Software Engineering*, pages 353–364. IEEE Press, 2017.
- [27] T. László and Á. Kiss. Obfuscating c++ programs via control flow flattening. *Annales Universitatis Scientiarum Budapestinensis de Rolando Eötvös Nominatae, Sectio Computatorica*, 30:3–19, 2009.
- [28] Z. Li, S. Lu, S. Myagmar, and Y. Zhou. Cp-miner: Finding copy-paste and related bugs in large-scale software code. *IEEE Transactions on software Engineering*, 32(3):176–192, 2006.
- [29] libpng.org. libpng home page. <http://www.libpng.org/pub/png/libpng.html/>, January 2018.
- [30] M. Lindorfer, A. Di Federico, F. Maggi, P. M. Comparetti, and S. Zanero. Lines of malicious code: Insights into the malicious software industry. In *Proceedings of the 28th Annual Computer Security Applications Conference (ACSAC)*. ACM, 2012.
- [31] L. Luo, J. Ming, D. Wu, P. Liu, and S. Zhu. Semantics-based obfuscation-resilient binary code similarity comparison with applications to software plagiarism detection. In *Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering (FSE)*, 2014.
- [32] J. Ming, D. Xu, Y. Jiang, and D. Wu. Binsim: Trace-based semantic binary diffing via system call sliced segment equivalence checking. In *Proceedings of the 26th USENIX Security Symposium. USENIX Association*, pages 253–270, 2017.
- [33] N. Nethercote and J. Seward. Valgrind: a framework for heavyweight dynamic binary instrumentation. In *Proceedings of the 28th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, 2007.
- [34] J. Powny, B. Garmany, R. Gawlik, C. Rossow, and T. Holz. Cross-architecture bug search in binary executables. In *2015 IEEE Symposium on Security and Privacy*, pages 709–724. IEEE, 2015.
- [35] A. Sæbjørnsen, J. Willcock, T. Panas, D. Quinlan, and Z. Su. Detecting code clones in binary executables. In *Proceedings of the 18th International Symposium on Software Testing and Analysis (ISSTA)*, 2009.
- [36] 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 *Security and Privacy (SP), 2016 IEEE Symposium on*, pages 138–157. IEEE, 2016.
- [37] A. Shrivastava and P. Li. In defense of minhash over simhash. In *Artificial Intelligence and Statistics*, pages 886–894, 2014.
- [38] S. K. Udupa, S. K. Debray, and M. Madou. Deobfuscation: Reverse engineering obfuscated code. In *Proceedings of the 12th Working Conference on Reverse Engineering (WCORE)*, 2005.
- [39] S. Wang and D. Wu. In-memory fuzzing for binary code similarity analysis. In *Proceedings of the 32nd IEEE/ACM International Conference on Automated Software Engineering*, pages 319–330. IEEE Press, 2017.
- [40] X. Wang, Y. Dang, L. Zhang, D. Zhang, E. Lan, and H. Mei. Can i clone this piece of code here? In *Proceedings of the 27th IEEE/ACM International Conference on Automated Software Engineering*, pages 170–179. ACM, 2012.
- [41] X. Wang, Y.-C. Jhi, S. Zhu, and P. Liu. Behavior based software theft detection. In *Proceedings of the 16th ACM Conference on Computer and Communications Security (CCS)*. ACM, 2009.
- [42] XnSoft. Nconvert. <https://www.xnview.com/en/nconvert/>, January 2018.
- [43] B. Yadegari and S. Debray. Symbolic execution of obfuscated code. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pages 732–744. ACM, 2015.
- [44] B. Yadegari, B. Johannesmeyer, B. Whitely, and S. Debray. A generic approach to automatic deobfuscation of executable code. In *Security and Privacy (SP), 2015 IEEE Symposium on*, pages 674–691. IEEE, 2015.
- [45] F. Zhang, Y.-C. Jhi, D. Wu, P. Liu, and S. Zhu. A first step towards algorithm plagiarism detection. In *Proceedings of the 21th International Symposium on Software Testing and Analysis (ISSTA)*, 2012.
- [46] F. Zhang, D. Wu, P. Liu, and S. Zhu. Program logic based software plagiarism detection. In *Proceedings of the 25th IEEE International Symposium on Software Reliability Engineering (ISSRE)*, 2014.
- [47] zynamics. Bindiff. <https://www.zynamics.com/index.html>.