

1-to-1 or 1-to-n?

Investigating the effect of function inlining on binary similarity analysis

ANG JIA, MING FAN, WUXIA JIN, XI XU, and ZHAOHUI ZHOU, Xi'an Jiaotong University, China

QIYI TANG, SEN NIE, and SHI WU, Tencent Security Keen Lab, China

TING LIU, Xi'an Jiaotong University, China

Binary similarity analysis is critical to many code-reuse-related issues, where function matching is its fundamental task. “1-to-1” mechanism has been applied in most binary similarity analysis works, in which one function in a binary file is matched against one function in a source file or binary file. However, we discover that the function mapping is a more complex problem of “1-to-n” (one binary function matches multiple source functions or binary functions) or even “n-to-n” (multiple binary functions match multiple binary functions) due to the existence of **function inlining**, different from traditional understanding.

In this paper, we investigate the effect of function inlining on binary similarity analysis. To support our study, a scalable and lightweight identification method is designed to recover function inlining in binaries. 88 projects (compiled in 288 versions and resulting in 32,460,156 binary functions) are collected and analyzed to construct 4 inlining-oriented datasets for various similarity analysis tasks, including code search, OSS (Open Source Software) reuse detection, vulnerability detection, and patch presence test. Then, we further study the extent of function inlining, the performance of existing works under function inlining, and the effectiveness of existing inlining-simulation strategies. Results show that the proportion of function inlining can reach nearly 70%, while most existing works neglect it and use “1-to-1” mechanism. The mismatches cause 30% loss in performance during code search and 40% loss during vulnerability detection. Moreover, most inlined functions would be ignored during OSS reuse detection and patch presence test, thus leaving these functions risky. We further analyze 2 inlining-simulation strategies on our dataset. It is shown that they miss nearly 40% of the inlined functions, and there is still a large space for promotion. By precisely recovering when function inlining happens, we discover that inlining is usually cumulative when optimization increases. Thus, conditional inlining and incremental inlining are suggested to design a low-cost and high-coverage inlining-simulation strategy.¹

CCS Concepts: • **Software and its engineering** → **Search-based software engineering; Maintaining software**; • **Security and privacy** → *Software and application security*.

Additional Key Words and Phrases: binary similarity analysis, function inlining, 1-to-1, 1-to-n

1 INTRODUCTION

According to a report published by Gartner [53], over 90% of the development organizations stated that they rely on open source components. Although proper use of open-source code helps reduce the development cost, improper use of open-source code may bring legal and security risks. Developers may unintentionally violate open-source licenses, consequently causing great financial loss to software companies. For example, Cisco and VMware were exposed to serious legal issues because they did not adhere to the licensing terms of the Linux kernel [1, 3]. Moreover, vulnerabilities are easy to exploit in open-source software since its code is accessible. Software systems that use the vulnerable code and do not patch in time will suffer from security risks [82]. Binary code similarity analysis works are proposed to resolve these code-reuse-related issues.

The fundamental task of binary code similarity analysis is to match the query binary function and target source/binary functions. Here, the query function is the function in a binary absence of its origin, and the target function represents functions that may be related to a license, vulnerability, or patch. Given the query binary function, finding its similar binary functions is called binary2binary matching. Binary2binary matching [39, 51, 63, 68, 84] has various applications

¹New Paper

including binary code search [43, 44, 64, 76], vulnerability detection [40, 52] and patch presence test [78]. Given the query binary function, finding its similar source functions is called binary2source matching. Binary2source matching [35, 45, 49, 50, 54, 58, 59, 65, 71, 81] is critical to many code-reuse-related tasks, including binary2source code search [81] and OSS reuse detection [35, 50]. The matching mechanism of both binary2binary matching and binary2source matching is always considered as “1-to-1” in most existing works, i.e., one function in a binary file is matched against one function in a source file or binary file. However, we discover that such mapping is usually a more complex problem of “1-to-n” or “n-to-n” due to the existence of **function inlining** [14], which differs from traditional assumptions.

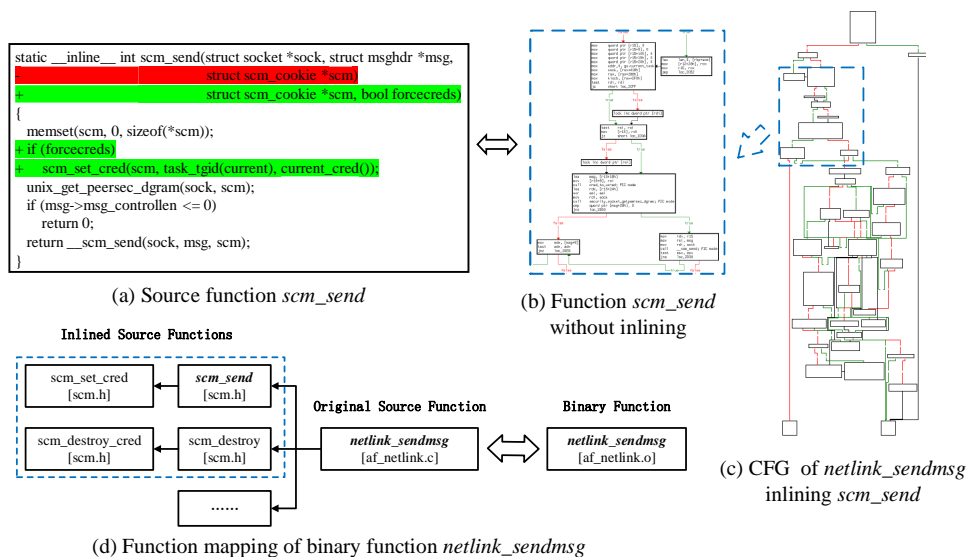


Fig. 1. “1-to-n” matching cases brought by function inlining

Figure 1 shows a motivating example when detecting vulnerable functions under inlining. Figure 1 (a) shows the source code of function *scm_send* in vulnerable version and patched version in the project Linux [26] for CVE-2020-16593 [2]. As code reuses of this vulnerable function may exist, the vulnerability would be propagated to disparate binaries. Detecting the influenced binaries requires searching whether there is a function in the binary similar to the vulnerable function in the source code. To support this process, Figure 1 (c) depicts the complete CFG (Control Flow Graph) of a binary function *netlink_sendmsg* that will be checked, with a partial CFG shown in Figure 1 (b). Each node represents a basic block, and the array represents the control flow between them.

Following the “1-to-1” mechanism, we cannot find an individual binary function that is matching to the function *scm_send*. Analyzing the debug information, we found the *scm_send*’s code is compiled into another function named *netlink_sendmsg*. However, the existing works calculate a low similarity score between the function *scm_send* and the target binary function *netlink_sendmsg*, indicating they are unmatched and thus leading to a failure that the vulnerability in *netlink_sendmsg* is missed.

Function inlining causes the failure of the “1-to-1” matching mechanism. Figure 1 (d) shows inlining happen to function *netlink_sendmsg*. Function *scm_send*, called by *netlink_sendmsg*, is inlined into *netlink_sendmsg*. When a function inlining occurs, it replaces the call-site statement inside a caller with the function body of the callee. As a result, the produced binary function will contain multiple source functions. Even more, the function inlining can be **nested**, meaning that an inlined function may further inline another function. As shown in Figure 1 (d), the *scm_send*, inlined by

`netlink_sendmsg`, inlines `scm_set_cred`. We can observe that the mapping between binary function `netlink_sendmsg` and its original source functions is no longer a “1-to-1” mapping but a “1-to-n” issue — *one binary function maps multiple source functions, sometimes even in a nested inlining manner*. Thus, the existing “1-to-1” matching mechanism, which tries to match two functions with unequal semantics, suffered the expected failure.

“1-to-n” issues bring a great challenge to existing binary2source matching works, as we have shown. Caused by the function inlining, the function mapping in binary2binary matching has become “1-to-n” and even “n-to-n”. However, to the best of our knowledge, few existing works have systematically studied the effect of function inlining on binary2source and binary2binary matching tasks (together, we call binary code similarity analysis). It is still unclear the extent of function inlining, and its impact on existing binary similarity analysis works. This paper will address this issue, and concretely we will investigate the three research questions.

RQ1: *To what extent will inlining happen during compilation?* To answer this RQ, we will evaluate the frequency and degree of function inlining, which can demonstrate the wild presence of function inlining in binary code similarity analysis tasks. This RQ will be answered in Section 4.

RQ2: *What is the effect of function inlining on the performance of existing binary code similarity analysis works?* To answer this RQ, we will evaluate the performance of existing works that do not consider function inlining or do not fully handle function inlining (which account for the majority). The decrease in performance indicates the importance of considering function inlining. We will answer this RQ in Section 5.

RQ3: *Can existing inlining-simulation strategies solve them?* Existing works have proposed strategies to simulate the results of function inlining (here, we call them inlining-simulation strategies). To answer this RQ, we will evaluate the effectiveness of existing inlining-simulation strategies from the inlining cost and coverage of inlined functions. The shortcomings of existing inlining-simulation strategies indicate directions for improvement. This RQ will be answered in Section 6.

We first construct four inlining-oriented datasets for four binary code analysis tasks, including code search, OSS reuse detection, vulnerability detection, and patch presence test to support our study. Furthermore, we propose a scalable and lightweight identification method that can automatically recover function inlining in binaries. Based on our labeled dataset, we then investigate the extent of function inlining, the impact of function inlining on existing binary code similarity works, and the effectiveness of existing inlining-simulation strategies. Finally, we discuss the interesting findings and give our suggestions for designing more effective inlining-simulation strategies.

Our main contributions are listed below:

- To the best of our knowledge, we are the first to comprehensively evaluate the effect of function inlining on binary code similarity analysis tasks.
- We create the first function inlining dataset by designing a scalable and lightweight identification method that can automatically recover function inlining in binaries (Section 3).
- Our study recovers the high frequency and high degree of function inlining (Section 4). Averagely the proportion of function inlining can range from 40% to 70% and a binary function can inline 2-4 source functions.
- We conduct five experiments to evaluate the performance of existing binary code similarity works under function inlining (Section 5). Code search suffers a 30% decline in detecting functions with inlining while OSS reuse detection and patch presence test miss most inlined functions.
- And we take a deep analysis of existing inlining-simulation strategies (Section 6) and point out the directions for designing more effective strategies (Section 7).

To facilitate further research, we have made the source code and dataset publicly available [32].

2 PRELIMINARY

This section will introduce existing static binary code similarity analysis works, commonly-used inlining strategies, and the research scope of our work.

2.1 Binary Code Similarity Analysis

2.1.1 Binary2binary code matching. According to the granularity to conduct matching, existing binary2binary matching can be classified into instruction level, basic block level, and function level method.

Bindiff[8], Kam1n0[43] and Asm2Vec[44] are instruction level methods. To conduct matching, Bindiff and Kam1n0 are working on the raw assembly instructions and conducting sub-graph matching to find similar functions. Instead, Asm2Vec trains the embedding for the assembly code and matching functions using function-level embedding by lifting instruction-level embedding.

Deepbindiff[46] is a basic block level method to conduct binary diffing. To conduct matching, Deepbindiff first constructs a connected control flow graph of two binaries by merging their common-used library calls and strings. Then it applies TADW[80] to learn embeddings of basic blocks. Finally, they obtain the matching by calculating the similarities of embeddings.

Function level matching methods [37, 40, 52, 64, 76, 78, 79] are the most widely-used methods when conducting binary2binary code matching. Tracy [40] extract the data flow strands from functions and matching functions by the hash of strands. Gemini[76] uses the CFG (control flow graph) augmented with block attributes as the representation of a function and trains a Siamese network for function matching. SAFE [64] use a self-attentive network [61] directly on instruction embeddings to obtain the function embeddings. Vulseeker [52] works on labeled semantic flow graphs (LSFG) and uses its embedding to find vulnerable functions. BinXray [78] uses diffs between vulnerable version and patch version to conduct path presence test.

Among these binary2binary matching works, Bingo [37, 79] and Asm2Vec [44] especially propose selective inlining-simulation strategies to tackle the challenges that function inlining brings. However, our work is different from theirs. We aim to provide a systematical and general understanding of function inlining and its impacts. Moreover, in our work, we also evaluate the effectiveness of their strategies leveraging our unique insights of function inlining.

2.1.2 Binary2source code matching. Existing binary2source code matching works can be classified into feature-based methods and learning-based methods.

Feature-based methods [45, 49, 50, 54, 59, 65, 71] are commonly used by early works. BAT [54] considers textual strings as features to match source and binary code. BinPro [65] extracts similar features but utilizes machine learning algorithms to compute optimal code features for the matching. B2SFinder [50], a state-of-the-art feature-based method, considers three types of features, including string, integer, and control-flow features. As these features are sometimes rare at the function level, most feature-based methods conduct matching at the file level.

CodeCMR [81] is a learning-based method for function-level binary2source matching. CodeCMR models binary2source matching as a cross-modal retrieval task, with the source function code and binary function code as two neural network inputs. The model output is the similarity between them. CodeCMR also extracts string and integer features to supply the inputs. It has achieved fast and precise matching by developing a deep learning method.

2.2 Function Inlining

Function inlining [56] conducts function expansion, replacing a function call with the callee body. Besides that function inlining benefits decreasing the cost of method calls, its major purpose is to enable additional compilation optimizations

[70]. Though there are passes that can be used for inter-procedural optimizations, inlining the callee function can convert inter-procedural optimizations into intra-procedural optimizations, facilitating more optimizations.

On the other hand, function inlining can cause an increase in binary code size and compilation time. It is a trade-off to determine when to inline by a balance between these benefits and costs. Prior researches have proposed lots of inlining strategies [33, 38, 41, 42, 47, 55, 57, 83] from different aspects. We will introduce inlining strategies adopted in GCC [15] and LLVM [25], two of the most popular compilers.

GCC and LLVM follow a similar workflow for function inlining, including inlining decision and inlining conduction. To decide which call-sites should be inlined, they consider user-forced inlining, user-suggested inlining, and normal inlining candidates. User-forced inlining, such as “`__attribute__((always_inline))`”, will first be considered to be inlined, as long as the call-site of this function does not violate fundamental requirements. User-suggested inlining is specified by users using the keyword “inline” to decorate function definitions. Such decorated functions will be possible candidates to be inlined, but whether to inline still depends on the strategy of the compiler. Normal inlining candidates will be identified and determined according to the strategies used in GCC and LLVM themselves.

GCC uses a priority-guided inlining algorithm to decide which call-sites should be inlined. They first define the inlining priority of a function call by pre-estimating the benefit of this inlining. Then they conduct inlining following the priorities. At the same time, the cost of inlining also increases. Once the total cost of inlined methods reaches a maximal threshold as configured, the inlining process will stop.

Instead, LLVM uses a prediction-driven inlining method to determine function inlining. It first classifies functions into hot and cold ones according to the times they are called. Then each call-site will be assigned with measurements of inlining benefit and inlining cost [4]. If the cost value is lower than the benefit value, this call-site will be inlined [18].

Besides, compilers also implement specific settings for small-size functions and functions that are called only once. For example, when compiling a program with the option of “-O1”, GCC-10 applies “-finline-functions-called-once” to conduct inlining. This option will inline all static functions called only once no matter whether they are marked with “inline” or not. Moreover, if a function called once is inlined, then the function will not be compiled as an individual function in the binary[19], such as the function `scm_send` in Figure 1.

2.3 Research Scope

Before introducing our work, we first clarify our research scope. Our work investigates the impact of inlining strategies used in common compilers on static function-level binary code similarity analysis works. Other function inlining strategies or instruction level, basic block level and file level method is beyond our scope.

We focus on C/C++ programs compiled by GCC and Clang with O0-O3. The main architecture of binaries is in X86-64, but we also analyze datasets compiled in 8 architectures. Moreover, we only consider the inlining of user-defined functions in this paper. The inlining of library functions is similar to user-defined functions, but library functions may differ when the compilation environment changes.

3 DATASET CONSTRUCTION

As few studies have researched the impact of function inlining on binary code similarity analysis work, an inlining-oriented dataset is needed. To fill the gap, we will first collect 88 projects for different applications and propose a scalable and lightweight identification method to recover function inlining in binaries automatically. As illustrated in Figure 2, we first collect source projects and compile them into binaries under diverse configurations (Section 3.1).

Based on the source code and binaries, we propose an automatic identification method to identify function inlining in binaries (Section 3.2).



Fig. 2. Workflow of dataset construction

3.1 Dataset Collection

In this work, we focus on C/C++ programs and two popular compilers, GCC [15] and Clang [10]. We have collected four datasets by employing and extending Binkit [22, 60], a binary code similarity analysis benchmark. The compositions and compilation settings of these dataset are listed in Table 1.

Table 1. Composition of four datasets

Dataset	# of Projects	Representative Projects	Compilation Setting	Versions	Binary Functions
I	51	Coreutils[11], Findutils[13]	Table 2	288	18,783,986
II	24	bzip2[29], minizip[30]	gcc-8.2.0, X86-64, O0-O3	4	190,611
III	12	Openssl[31], Binutils[9]	gcc-8.2.0, X86-64, O0-O3, inline, noinline	8	8,635,362
IV	1	Chromium[23]	clang, debug, release	2	4,850,197

Table 2. Compilers, optimizations and architectures used in dataset I

Compilers	gcc-4.9.4, gcc-5.5.0, gcc-6.4.0, gcc-7.3.0, gcc-8.2.0, clang-4.0, clang-5.0, clang-6.0, clang-7.0
Optimizations	O0, O1, O2, O3
Architectures	X86-32, X86-64, ARM-32, ARM-64, MIPS-32, MIPS-64, MIPS64-32, MIPS64-64

Dataset I - for code search. Dataset I was comprised by the 51 packages from GNU projects[24], compiled by 9 compilers, 4 optimizations, to 8 architectures (shown in Table 2), resulting in 288 compilation versions containing 67,680 binaries and 18,783,986 functions. Dataset I is collected from Binkit[60]. This dataset contains many widely-used packages such as *Coreutils* [11] which have been extensively used in binary similarity detection works [37, 48, 62, 75, 77, 81].

Dataset II - for OSS reuse detection. Dataset II was extended from a partial reuse dataset used in ISRD[77] containing 24 programs with 74 partial reuses. It was originally compiled by GCC with O2 in X86-64. We extend it to 4 optimizations by using gcc-8.2.0. Dataset II is not extended to other architectures because existing OSS reuse detection works do not include them. Dataset II was comprised of 96,085 source functions and 190,611 binary functions.

Dataset III - for vulnerability/patch detection. Dataset III was constructed from the dataset used in BinXray[78], containing 479 CVEs and 6238 vulnerable functions. We generate two versions of Dataset III, including Dataset III (without additional compilation flags) and Dataset III-NI (with flag “*-fnoinline*” to forbid inlining). Dataset III is comprised of 8,635,362 binary functions.

Dataset IV - for application binaries. Apart from these three datasets, we also construct a dataset by compiling chromium[23] with Clang in its debug and release optimization to X86-64 to form the dataset IV. Dataset IV is comprised of 4,850,197 binary functions. This dataset will provide a vision of function inlining in large application binaries.

3.2 Function Inlining Identification

In this section, we will first illustrate terminologies that we will use throughout our paper and then introduce our function inlining identification method.

Algorithm 1: Method to generate function-level mappings

Input: Source project S , Binary with debug info B
Output: Function mapping FM_{b2s}

```

1 Function Preprocessing( $B, S$ ):
2    $LM_{b2s}$  = Extract_Address_to_Line_Mapping( $B$ ) // return a list of (address, line) mapping
3    $SM_{l2f}$  = Extract_Line_to_SF_Mapping( $S$ ) // return a dict of line to function mapping
4    $BM_{a2f}$  = Extract_Address_to_BF_mapping( $B$ ) // return a dict of address to function mapping
5   return  $LM_{b2s}, SM_{l2f}, BM_{a2f}$ ;
6  $LM_{b2s}, SM_{l2f}, BM_{a2f}$  = Preprocessing( $B, S$ );
7  $FM_{b2s}$  = {};
8 for ( $address, line$ ) in  $LM_{b2s}$  do // traverse all address-to-line mappings
9    $SF$  =  $SM_{l2f}[line]$ ;
10   $BF$  =  $BM_{a2f}[address]$ ;
11  if  $BF$  not in  $FM_{b2s}$  then
12    |  $FM_{b2s}[BF]$  = [];
13  end
14  if  $SF$  not in  $FM_{b2s}[BF]$  then
15    |  $FM_{b2s}[BF].append(SF)$  // add mapped source functions to the binary function dict
16  end
17 end
18 return  $FM_{b2s}$  ;

```

3.2.1 Terminology. To make the representation more clear, we first illustrate our notions of the functions from both source and binary. Usually, as shown in Figure 1 (d), when function inlining happens, a binary function (*netlink_sendmsg*) is compiled from the original source function (*netlink_sendmsg*) and some inlined source functions (*scm_send*, *scm_set_cred* and etc.). Here, we named the produced binary function (*netlink_sendmsg*) as BFI (binary function with inlining), the original source function (*netlink_sendmsg*) as OSF, and the inlined source functions (*scm_send*, *scm_set_cred* and etc.) as ISFs. Similarly, binary function without inlining is named as NBF (normal binary function), NBF mapped source function as NSF (normal source function). Moreover, binary function as BF and source function as SF. These notions will be used throughout the paper.

3.2.2 Function inlining identification method. To identify function inlining, we first construct function mappings from source code to binaries, which can also serve as the ground truth for binary code similarity analysis. When compiling a program in debug mode, compiler will record the mapping from binary instructions to source instructions. Following that mapping, we further find the binary function that binary instructions belong to and the source function that source instructions belong to. Thus, function-to-function mapping can be generated.

Algorithm 1 formalizes our method for generating function-level mappings for datasets. The method leverages the line mapping implied by the relation between binary address and source code line, provided in the debug mode. By extending line mapping results, the method will generate a function-level mapping between binary functions and source functions. We will introduce the algorithm with an example in Figure 3.

Our method first compiles source code with the “-g” option and leverages Dwarf [12] to produce the `.debug_line` section in binary. To extract the line-level mapping, we uses Readelf [20] to parse binaries (line 2). As shown in Figure 3, the line mapping contains the mapping between source file lines and binary addresses.

Next, the method extracts line-to-function mapping in the source code by employing Understand [21] and extracts address-to-function mapping in the binary code by employing IDA Pro [17]. Understand and IDA Pro are two state-of-art

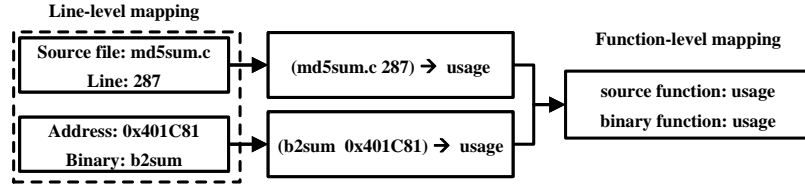


Fig. 3. Example of extending line-level mapping to function-level mapping

commercial tools and have been widely used in academia [46, 67, 69] and industry [5, 6, 28]. Understand is a static analysis tool to identify source entities (e.g., files, functions) and dependencies. By exploring source entities, we can infer which entity the code line belongs to (line 3). IDA Pro recovers the function boundary [34] from binaries, outputting binary address to binary function (line 4). For example in Figure 3, source line 287 belongs to the source function *usage*, and binary address 0x401C81 belongs to binary function *usage*.

Finally, our method extends the above line-level mapping results to function-level mapping by combining these three mappings (lines 8-17). As shown in Figure 3, now we can get the function mapping between source function *usage* and binary function *usage*.

Our datasets collected in Section 3.1 are labeled by our automatic method. According to the mapping results between BFs and SFs, we can easily identify BFs produced by function inlining (i.e., BFI). Simply, we identify BFs mapped to more than one SFs as BFIs and those mapped to one SF as NBFs. By identifying the OSF or NSF from which BFs are compiled, we can also easily establish binary2binary function-level mappings. Our datasets with function-level mappings can support the study in our work.

Besides, our method does not rely on any modification of the compilers and can be applied to any binaries compiled in debug mode. Note that Understand and IDA Pro can be substituted with other tools with similar functionalities. We also implement a tool using tree-sitter[27] and Ghidra[16], facilitating free usages.

4 INVESTIGATION ON THE EXTENT OF FUNCTION INLINING

Before analyzing the extent of inlining, we first discover that inlining is not the same in all circumstances. Firstly, inlining is defined by the compiler. Thus, inlining is influenced by the compilation settings. Secondly, the compiler decides whether to inline by considering the features of a call site, which will differ in different projects. Thus inlining also depends on the projects. Finally, we noticed that different applications also care about different parts of inlining. Thus, in this section, we propose three research questions (RQs) to help understand the extent of function inlining.

4.1 RQ1.1: What Is the Extent of Function inlining under different compilations?

Figure 4 shows the statistics of function inlining in dataset I compiled by 8 architectures, 9 compilers, and 4 optimizations. Due to page limit, we only present results in clang-7.0 and gcc-8.2.0. In Figure 4, “Inlining ratio” shows the percent of BFIs in BFs, “ISFs/BFI” represents the number of ISFs that a BFI inlines, “BFIs/ISF” represents the number of BFIs that an ISF is inlined into. Here we calculate these values by average. These metrics measure three aspects associated with function inlining: happening frequency in binary functions (Inlining Ratio), inlining degree in binary functions (ISFs/BFI), and inlining degree in source functions (BFIs/ISF).

4.1.1 Statistics in default O_x optimizations. For statistics of all 64 compilation options, we can see that the inline ratios range from 0.94% compiling with clang-7.0 in O1 to ARM64 to 41.67% compiling with gcc-8.2.0 in O3 to MIPS64. And when a binary function is generated with function inlining, at least 1.39 ISFs are inlined averagely. Moreover,

for Inlining Ratio, B2S for ISFs/BFI, and S2B for BFIs/ISF. Compared with binaries without LTO, binaries enabled with LTO have 30% more inlining ratios. And when compiling in O3 enabled with LTO, more than 70% of the BF are BFIs. Apart from the increasing frequency of inlining, 2 more ISFs are inlined into the BFIs when using LTO averagely. The increasing frequency and extent of inlining bring more challenges to existing binary code similarity works.




Fig. 5. Function inlining statistics with LTO

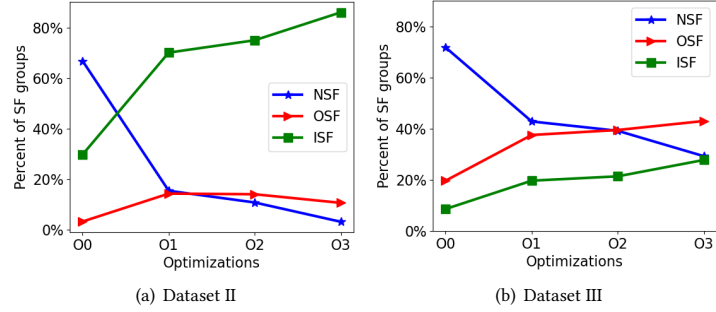


Fig. 6. Distribution of critical SF in BFIs

Answering RQ1.1: Inlining happens in high optimizations, and its ratio usually ranges from 30% to 40% when applying O3. However, inlining under different compilers may differ at the same optimizations. Besides, when combining O3 with LTO, inlining ratios can raise to 70%, and averagely a BFI can inline 4 ISFs.

4.2 RQ1.2: What Is the Extent of Function Inlining in Different Projects?

Different projects may have different inlining statistics due to their coding style and design rule. To avoid getting a biased understanding, we first analyze inlining statistics in four kinds of projects, including system code from dataset I, zip/unzip, timing, and data transformation code from dataset II, cryptography, database, and image processing code from dataset III, and browsers from dataset IV.

Table 3. Function inlining statistics in different kinds of projects

D	I				II				III				IV	
	Opt	O0	O1	O2	O3	O0	O1	O2	O3	O0	O1	O2	O3	Og
IR	2.7%	18.4%	26.4%	39.3%	6.5%	35.6%	40.2%	48.1%	22.2%	34.8%	39.0%	46.5%	6.4%	57.9%
B2S_avg	2.35	2.25	2.27	2.67	1.86	2.08	2.16	2.45	2.27	3.20	3.23	3.39	1.89	3.53
B2S_max	62	76	89	111	63	70	145	121	2756	1727	1781	2073	36	139
S2B_avg	1.41	1.35	1.53	1.86	2.99	1.90	1.94	2.17	2.72	3.02	2.89	2.80	1.67	2.58
S2B_max	25	74	82	76	46	106	106	139	525	892	872	900	73	229

Table 3 shows the inlining statistics of these four datasets. Among the first three datasets, dataset II has the highest inlining frequency, and dataset III has the largest inlining degree. In particular, we notice that sometimes a BFI in dataset III can inline more than 2000 ISFs, and an ISF can be inlined into more than 800 BFIs. This is caused by some extremely-long call chains in the cryptography module of *Openssl* [31], and we noticed that long call chains are also common in the dataset I and II. And in dataset IV, we observe higher inlining ratios in the released binary compared with dataset I-III which are compiled with default O_x optimizations, indicating large applications may use additional custom optimizations to further ensure the efficiency of the released executable.

Answering RQ1.2: Generally speaking, inlining widely exists in different kinds of projects. And in some large application binaries, inlining may happen more frequently.

4.3 RQ1.3: What Parts of Inlining Are Cared for in Different Applications?

Different applications care about different parts of inlining due to their objectives. For example, code search aims to find target functions with similar functionality of the query function; thus, it needs to search BFs originated from the same OSF when function inlining happens. OSS reuse detection and vulnerability/patch detection aim to search functions containing the code of the reused or vulnerable functions; thus, they need to search BFs containing the same ISFs. Here, we focus on the inlining under OSS reuse detection and vulnerability/patch detection, and analyze the existence and proportion of target SFs in the query BF.

As shown in Figure 6(a), we divided the relation between query BF and target SF into three classes: NSF, OSF, and ISF. For example, OSF means that the target SF is the OSF of the query BF, while ISF means that the target SF is an ISF of the query BF. Similar to Figure 6(a), in Figure 6(b), OSF means that query BF and target BF share the same OSF while ISF means the NSF of target BF (which is an NBF) is an ISF of the query BF.

For OSS reuse detection, we noticed that when optimization goes from O0 to O3, ISF is gradually counting for the largest part, and NSF decreases from more than 60% to nearly 2%, which indicates the majority reuse between source projects will be inlined when applying high optimization, making it a new challenge for existing works.

For vulnerability/patch detection, when applying high optimization, OSF becomes the majority where NBFs and BFIs need to be matched. Besides, 30% of vulnerable functions are also inlined into BFIs, and finding these BFIs is also a task of existing works.

Answering RQ1.3: Code search works cares about the BFIs with the same OSF, while OSS reuse detection and vulnerability/patch detection also care about the BFIs with reused or vulnerable ISFs. The proportion of reused or vulnerable ISFs is not rare. Thus, they are also needed to be considered when conducting matching.

4.4 Answer to RQ1

In summary, function inlining widely exists in high optimizations and different projects. For released binaries, the inlining ratios range from 35% to more than 70%, and a BFI inlines 2-4 ISFs on average. Code search works should consider the unequal semantics between functions that inlining has brought. Besides, 80% of reused SFs and 30% of vulnerable SFs are inlined to BFIs when applying high optimization. OSS reuse detection and vulnerability/patch detection should also care about the hidden risky functions inlined into binary functions.

5 PERFORMANCE OF EXISTING WORKS UNDER INLINING

This section will carry out several evaluations on the performance of existing works under inlining. Directly, we first split existing binary code similarity analysis works into binary2source matching and binary2binary matching works. Then we set different objectives for different applications. As we illustrated in Section 4.3, code search only needs to match BFs with the same functionalities, while OSS reuse detection and vulnerability/patch detection need to match every reused function. Thus, in our evaluation, apart from BFIs with the same OSF that code search works are required to accomplish, existing OSS reuse detection works and vulnerability/patch detection works are also evaluated on finding BFIs containing reused or vulnerable ISFs.

As inlining always comes with high optimizations, excluding the influence from other added optimization options is necessary. To evaluate the effect merely from inlining, we first select SFs that can produce BFs under all compilation settings. Then we split these SFs into two classes: OSFs whose BFs compiled in high optimization are BFIs and NSF

whose BFs are all NBFs. The only difference for BFs generated from these two classes is whether inlining has happened. Thus evaluations conducted on these two split datasets only reflect the influence of inlining.

For all works selected for evaluation, we use its default setting in its code or its paper and present its direct output.

5.1 Evaluation of Binary2source Matching Works

For binary2source matching works, we select CodeCMR[81] targeting two application areas: code search and OSS (open-source-software) reuse detection. To suit the applications, we use dataset I to evaluate the task in code search and dataset II for OSS reuse detection, and here are the results:

5.1.1 Binary2source code search . We select CodeCMR because it is the state-of-the-art function level binary2source code searching work. And its application binaryai[7] has been released for public usage. We analyze the impact of inlining on the performance of binaryai at a downstream task of code search: retrieving similar SFs for the query BF.

Cross-optimization evaluation for CodeCMR. Figure 7(a) shows our evaluation of CodeCMR in BFIs and NBFs. We use the same metrics (recall@1 and recall@10) used in CodeCMR to evaluate its performance. As shown in Figure 7(a), blue lines present the performance of matching OSFs with their BFIs, and red lines present the performance of matching NSF with their NBFs. When matching functions without inlining, we can see a smooth curve for both recall@1 and recall@10, though they suffered a subtle loss when matching NSF with NBFs from O0 to O3. But when matching functions with inlining, we can see a drastic loss from O0 to O3, about 30% for recall@1 and 25% for recall@10, which indicates that function inlining has brought a great challenge to CodeCMR.




Fig. 7. Evaluation of CodeCMR under inlining

Inlining-degree evaluation for CodeCMR. To further analyze the influence of inlining on CodeCMR, we rearrange the dataset into classes by their inlining influence degrees. In detail, we calculate the ratio between the total length of ISFs and the length of the OSF to define the influence degree that inlining has on matching a BFI with an OSF. And for matching pairs with different inlining degrees, we calculate their average recall@1 and recall@10 to form Figure 7(b).

In Figure 7(b), blue lines show the change of recall@1 and red lines for the recall@10. Inlining influence degree as “1” represents matching pairs that ISFs of BFI have a 1x to 2x length equal with the OSF of BFI. As shown in Figure 7(b), both recall@1 and recall@10 decrease dramatically when inlining influence degree increase. In detail, recall@1 reaches nearly 80% when the ISFs have a smaller length than the OSF but falls directly to 0 when the total length of ISFs is larger than the length of OSF. Recall@10 decreases when the inlining influence degree ranges from 0 to 10 and experiences a fluctuation between 10 and 30. But in general, when the inlining influence degree is larger than 1, it has become particularly difficult for CodeCMR to find the correct OSFs.

5.1.2 **OSS reuse detection.** As we illustrated in Section 2, existing feature-based binary2source matching works cannot precisely detect software reuse at the function level. Thus, we extend CodeCMR to detect partial reuses.

Table 4. Function similarity for OSS reuse detection

similarity	O0	O1	O2	O3
NSF	77.37%	73.31%	68.17%	58.12%
OSF	66.27%	71.42%	67.60%	58.08%
ISF	10.54%	9.01%	8.42%	5.93%

Table 5. Function similarity for vulnerability detection

similarity	O0	O1	O2	O3
NSF	89.36%	93.43%	92.68%	94.11%
OSF	82.69%	84.02%	85.92%	80.57%
ISF	46.16%	63.39%	64.34%	55.80%

Table 4 shows the results of CodeCMR in detecting three types of reuse corresponding to three classes in Section 4.3. When matching OSFs with BFIs, CodeCMR returns a similarity a little lower than matching NSF with NBFs. However, when searching ISFs for BFIs, CodeCMR can only return a similarity of less than 10% on average. As a result, CodeCMR can detect most NSF and OSF reused in NBFs and BFIs, but missed most ISFs that are inlined into BFIs.

5.2 Evaluation of Binary2binary Matching Works

For binary2binary matching works, we select SAFE [64] for code search, Vulseeker [52] for vulnerability detection, and BinXray [78] for patch presence test. SAFE is evaluated on dataset I, while Vulseeker and BinXray are evaluated on dataset III.

5.2.1 **Binary2binary code search.** We select SAFE for code search as it can represent the common workflow of most recent learning-based methods: training a model to produce embedding, input functions to obtain embeddings, and comparing embeddings for similarities.




Fig. 8. Cross-optimization function similarity for SAFE

Figure 8 shows the function similarity that SAFE obtained for cross-optimization code search. To reveal the impact of function inlining, we divide the matching cases into normal cases, which only compares two NBF and inlining cases where at least one function is a BFI. In Figure 8, we use blue lines to represent the similarity in normal matching and red lines for inlining matching. Besides, we use ‘G’ for GCC and ‘C’ for Clang. And x-coordinate shows the detailed optimization setting of the experiments. For example, the first O1 with “-O0-” below represents searching query BFIs generated by O0 in target BFIs generated by O1.

In general, SAFE suffers a 20% loss in similarity when detecting NBFs at the normal cross-optimization task (without inlining), but the loss increase to more than 40% when detecting BFIs (with inlining). Interestingly, we noticed that inlining in clang has a greater impact on SAFE than gcc, and this difference becomes more obvious when crossing O0, O1 to O2, O3.




Fig. 9. Cross-inlining results of BinXray

5.2.2 Vulnerability detection. We select Vulseeker for vulnerability detection. To see the effect of inlining, we recompile dataset III with “-fno-inline” to create a dataset without inlining (dataset III-NI). Then we use vulnerable functions in dataset III-NI to generate fingerprints and functions in dataset III as the target. Results are shown in three types corresponding to three classes in Section 4.3.

As shown in Table 5, Vulseeker obtained a similarity of more than 90% for normal matching cases, a similarity of 80% when searching BFI, but a similarity of less than 65% when searching inlined functions. As lower similarities indicate bigger difficulties in obtaining the target function, BFIs, which inline vulnerable ISFs, are very hard to find.

5.2.3 Patch presence test. We select BinXray for patch presence test. As BinXray targets detecting patches at the same optimization, we conduct two experiments: default and cross-inlining. The default experiment generates fingerprints from binary functions compiled in O2 with inlining and search in binaries compiled in the same setting. The cross-inlining experiment generates fingerprints from binary functions compiled in O2 without inlining but searches binary functions compiled in O2 with inlining. Here are the results.

Figure 9 shows the results of default experiment and cross-inlining experiment. The number of rightly identified functions drops when detecting NBFs in the cross-inlining experiment compared with the default experiment, and this gap grows much larger for BFIs. Only about 2/5 of the vulnerable BFIs are rightly classified in the cross-inlining experiment, while 3/5 fall to “cannot analyze”.

Interestingly, although the rightly identified cases decrease dramatically, wrongly identified only increase a little. Instead, more cases become hard to be analyzed. That is because the analysis of BinXray relies on its successful identification of the patch boundary. For BFIs with vulnerable OSF that we mention in section 4.3, BinXray cannot find the similar patch boundary derived from dataset III-NI.

Besides, another problem is that BinXray regards the functions having the same name with vulnerable function as the suspicious function to conduct detecting. Thus BFIs which has inlined vulnerable ISFs will all be missed. Note that these BFIs with vulnerable ISFs account for 30% when compiling in O3 as shown in Section 4.3.

5.3 Answering to RQ2

In summary, code search works suffer a 20%-30% performance loss when detecting functions with inlining. Ignorance of semantics from inlined functions reveals their inherent defect in detecting function similarities. Besides, OSS reuse detection and vulnerability/patch detection works cannot detect inlined functions. Thus, attackers or plagiarists can inline these risky functions into binary functions to bypass the detection.

6 EVALUATION OF EXISTING INLINING-SIMULATION STRATEGIES

In section 5, we have shown that existing works suffer from the challenges that inlining brings. To resolve these challenges, existing works propose inlining-simulation strategies to simulate the inlining results of compilers. In this section, we will introduce existing inlining-simulation strategies and evaluate their effectiveness against function inlining in two fundamental tasks.

6.1 Introduction of Existing Inlining-simulation Strategies

We first introduce the strategies used in two existing works: Bingo[37] and ASM2Vec[44].

Bingo. Bingo proposes a selective inlining-simulation strategy to recursively expand callee functions for more precise similarity detection. To decide whether a callee should be inlined, Bingo listed 6 patterns for inlining conduction, and we summarize them into the following 5 cases where 2-5 are towards UD (User-Defined) functions.

1. if the callee function is a library function, inline it.
2. If the UD callee function and its caller recursively call each other, inline it.
3. If the UD callee function only calls library functions, and more than half of library functions are not termination functions (e.g., exit and abort), inline it.
4. If the UD callee function only calls library functions, and more than half of library functions are termination functions, do not inline it.
5. If the UD callee function calls other UD functions, the inlining is decided by the following equation 1.

$$\alpha(f_c) = \text{outdegree}(f_c) / (\text{outdegree}(f_c) + \text{indegree}(f_c)) \quad (1)$$

In equation 1, in-degree and out-degree refer to the call relation between the callee f_c and other UD functions in the FCG (Function call graph). Bingo considers that the lower the value of $\alpha(f_c)$, the more likely the callee should be inlined. Thus when $\alpha(f_c)$ is lower than a threshold, this callee will be inlined. In Bingo's original setting, this threshold is set to 0.01.

ASM2Vec. The inlining-simulation strategy of ASM2Vec employs the strategy of Bingo, but with some modifications and additional filters. On the one hand, ASM2Vec change the recursively inlining-simulation strategy into a one-layer inlining-simulation strategy, where they only expand the first-order callees. On the other hand, they defined a metric to remove lengthy callees, as shown in equation 2.

$$\delta(f_s, f_c) = \text{length}(f_c) / \text{length}(f_s) \quad (2)$$

In equation 2, f_c represents for the callee function and f_s for the caller. The length of the function is measured by the number of lines of instructions. When δ is less than 0.6 or f_s is comprised of less than 10 lines of instructions (considered as a wrapper function), inlining will happen.

6.2 Evaluating Strategies on Finding Similar BFIs

As finding BFIs with the same OSF is the main objective of existing binary code similarity analysis works, we first evaluate their effectiveness in finding similar BFIs. According to their settings, the inlining-simulation strategies will be conducted for both the query BF and target BF, and the BFI' 's produced by the inlining-simulation strategies help find similar BFIs.

To evaluate their effectiveness, we implement the inlining-simulation strategies of Bingo and ASM2Vec and apply them to dataset I. For each BF, Bingo and ASM2Vec will generate the BFI' using their inlining-simulation strategies. In detail, we use IDA[17] to disassemble binaries and extract features used for inlining selection. And we provide two metrics to evaluate the effectiveness of their strategies.

$$\text{cost} = BFI'_1 - BF_1 + BFI'_2 - BF_2 \quad (3)$$

$$\text{similarity} = \frac{BFI'_1 \cap BFI'_2}{BFI'_1 \cup BFI'_2} \quad (4)$$

In Equation 3 and 4, BF_1 and BF_2 are two BFIs that are need to be compared, and BFI'_1 and BFI'_2 are the BFI' 's generated by the inlining-simulation strategies proposed by Bingo and ASM2Vec. Equation 3 calculates the number of inlined BFIs

for both BF_1 and BF_2 , and Equation 4 calculates the similarity of produced BFI'_1 and BFI'_2 . And we measure the BFI' in equation 3 by the BFs it contains and BFI' in equation 4 by the SFs it contains.

Figure 10 shows the evaluation results of Bingo and ASM2Vec. To facilitate comparison, we also present the cost and similarity without applying inlining strategies (without inlining in Figure 10). The dotted line represents the similarity and the solid line for the cost. In general, compared with not applying strategies, Bingo averagely selects 4-5 BFs for inlining to obtain a 5% increase in its obtained similarity. And ASM2Vec can reduce 0.25 BFs for inlining compared to Bingo while losing 1-2% in the similarity.




Fig. 10. Effectiveness of existing inlining-simulation strategies on binary2binary matching

From the results, we can notice that Bingo and ASM2Vec are selective towards callees as recursively inlining all callees will cause a cost to inline nearly 55 callees while they only need to inline 4-5 callees. And their effectiveness is quite visible — a 5% increase in similarity. However, there is still a large space to improve, as when detecting functions cross O0 to O3, existing inlining-simulation strategies can only help guarantee that 80% of SFs in BFI'_1 and BFI'_2 are matched. And we noticed that some BFs, which both are the callees of BF_1 and BF_2 , are also inlined by BFI'_1 and BFI'_2 , helping increase the function similarity.

In general, Bingo and ASM2Vec can help existing works improve their accuracy in finding similar BFIs, but there is still a large space for improvement.

6.3 Evaluating Strategies on Finding Inlined Functions

In our evaluation of existing binary similarity analysis works, we noticed new requirements in some specific areas, including OSS reuse detection, vulnerability detection, and patch presence test, where inlined reused and inlined vulnerable functions also need to be detected. In this section, we will evaluate the ability of Bingo and ASM2Vec in finding inlined functions.

Bingo. To facilitate evaluation, we first apply Bingo to dataset I and analyze its inlining decisions towards BFs. As Bingo has 3 cases (2, 3, 5) that callees (UD functions) will be inlined, to evaluate the effectiveness of its inlining-simulation strategies, we summarize the inlining decisions under the above 3 cases in O0-O3, which is shown in Table 6.

In table 6, R, L, and α respectively represents inlining for **R**ecursive call in case 2, inlining for callees with only **L**ibrary call in case 3 and inlining determined by α in case 5. BFNN represents BFs that do not need to be inlined, while BFN represents BFs needed. To facilitate analysis, we only present results of comparing BFs generated by low optimization to BFs generated by high optimizations, and we analyze its inlining conduction for the BFs in low optimizations. For example, the first value 20.56% indicates that, when matching BFs generated by O0 to BFs generated by O1, among all inlined functions, 20.56% are BFs which are not needed to be inlined but inlined, as it falls to case 5 and its α satisfies Bingo's requirements.

Table 6. Inlining decision of Bingo for BFs

Kind	Case	O0-O1	O0-O2	O0-O3	O1-O2	O1-O3	O2-O3
BFNN(BF Not Needed)	α	20.56%	21.57%	21.79%	20.26%	20.14%	25.07%
	L	68.56%	45.48%	42.53%	50.53%	46.68%	57.47%
	R	4.22%	4.57%	4.59%	5.56%	5.50%	3.21%
BFN(BF Needed)	α	0.14%	0.32%	0.32%	0.15%	0.15%	0.00%
	L	6.41%	27.85%	30.49%	22.83%	26.71%	8.84%
	R	0.12%	0.22%	0.28%	0.69%	0.81%	5.42%
coverage		78.51%	67.65%	60.68%	72.42%	65.01%	74.33%

For an effective inlining-simulation strategy, the percentage of BFNN should be low as they do not need to be inlined, and the percentage of BFN should be high as they need to be inlined. But from table 6, we can see that nearly 3/4 of inlined functions are BFNNs. To be more specific, among the three metrics, we noticed that only L had selected a large amount of BFNs while α and R rarely select BFNs. Bingo’s strategies waste a lot of cost in inlining unnecessary callees.

Besides, table 6 also shows the coverage of needed inlined functions in each matching. When the compilation setting is not so diverse, Bingo can cover more than 70% inlined functions. But when comparing O0 to O3, 40% of inlined functions remain uncovered, making the results less complete.

ASM2Vec. ASM2Vec proposed two measures to reduce the functions Bingo chooses to inline. They have a effect in reducing the inlining candidates (see Section 6.2). However, we argue that their restrictions are still so strict that they would suffer additional misses of ISFs. To provide a comprehensive understanding of boundaries for inlining, we conduct a study to analyze inlining from SFs that comprise the final BFs.

Figure 11(a) shows the distribution of depth for BFs. We take the example in Figure 1 (d) to illustrate the meaning of depth. In Figure 1 (d), BFI *netlink_sendmsg* has inlined four source functions, including *scm_set_cred* and *scm_send*. And *scm_send* is called by OSF *netlink_sendmsg* and *scm_set_cred* is called by ISF *scm_send*, where we recognize the depth of *scm_send* as 1 and the depth of *scm_set_cred* as 2. Finally, we use the maximum of depths (i.e.,2) as the depth of this BFI. As shown in Figure 11(a), hundreds of thousands of BFs inline ISFs with a calling depth of 2-3, and hundreds of BFs inline ISFs even through 8 layers of the call graph. 30% of the BFs have a depth of more than 1 layer.




Fig. 11. Distribution of metrics for inlined source functions

Figure 11(b) shows the distributions of ISFs lengths. Though lots of ISFs have fewer than 10 lines, we noticed that tens of thousands of ISFs are above 100 lines. Even thousands of ISFs with a length above 1000 are still inlined. Interestingly, the largest ISF is “*aarch64_opcode_lookup_1*” with 19083 lines of code which are inlined by the OSF “*aarch64_opcode_lookup*” with less than 10 lines of code. And this inlining happens during the compilation of Binutils compiling with GCC in O1. This is actually caused by the “-finline-functions-called-once” option as the ISF “*aarch64_opcode_lookup_1*” is only called by the OSF “*aarch64_opcode_lookup*”.

Figure 11(c) shows the distribution of ratios between ISF and OSF. For all ISFs, we noticed that ratios of ISF/OSF range from 0 to more than 1000. Only less than 1/3 of ISFs have a less than 0.6 of δ that will be inlined according to the ASM2Vec strategies, while 50% of the ISFs have a larger length than their OSF.

When using one-layer inlining and removing candidates whose $\delta \geq 0.6$, ASM2Vec additionally misses ISFs. The missing of these ISFs not only influences the performance in finding the OSF of BFI but also its ability to detect BFIs that have inlined reused or vulnerable ISFs.

6.4 Answer to RQ3

Bingo and ASM2Vec can help find similar BFIs, but there is still a large space for promotion. When applying to detecting inlined functions, they miss nearly 40% of the inlined functions. More effective inlining-simulation strategies are needed to resolve the effect of function inlining on binary code analysis works.

7 TOWARDS MORE EFFECTIVE STRATEGIES

This section will discuss how to design a more effective inlining-simulation strategy to help further reduce the inlining cost and find more ISFs. In detail, we first recover the matching patterns that inlining brings, and based on the findings, we propose several suitable measures to help develop a more effective strategy.

7.1 Investigation of Matching Patterns Under Inlining

The matching pattern inlining brings to binary2source matching is “1-to-n” compared with the existing “1-to-1” matching mechanism. However, this becomes much more complex for binary2binary matching. Here, we summarize the patterns of function-level matching under function inlining, as shown in Figure 12.

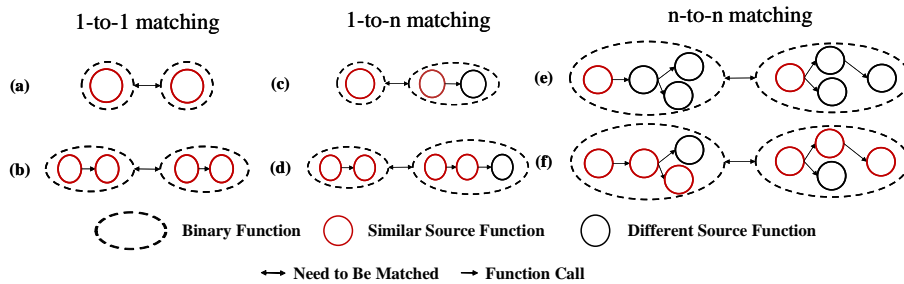


Fig. 12. Binary2binary matching patterns under inlining

1-to-1 matching. “1-to-1” matching is usually the pattern that most existing binary2binary work applies. Figure 12 (a) just shows the matching without function inlining. And another matching with function inlining is also “1-to-1” matching, where two BFs have inlined the same ISFs and thus have equal semantics.

1-to-n matching. “1-to-n” matching is similar to the pattern in binary2source matching. Figure 12 (c) shows the similar matching with binary2source matching. Apart from that, as shown in Figure 12 (d), matching the BFI A with the BFI B, which has inlined all ISFs in A and other ISFs, is also an “1-to-n” matching.

n-to-n matching. “n-to-n” matching reveals the most complex matching patterns when conducting binary2binary matching. Figure 12 (e) and (f) shows two classes of “n-to-n” matching. In detail, (e) shows two BFIs which have inlined different ISFs and the common semantics of these two BFIs only come from their shared OSF. (f) describe the BFI pairs which share the OSF and other ISFs.

Existing inlining-simulation strategies are conducted both at the query and target functions, so they can resolve these three matching patterns. However, we recovered that not all these patterns are common in all matching cases. Thus, we could focus only on one or two particular patterns to resolve the challenges that inlining brings.

Table 7. Distribution of binary2binary matching patterns

Pattern	a	b	c	d	e	f
Percent	63.80%	6.33%	25.70%	3.11%	0.82%	0.21%

Table 7 shows the proportion of binary2binary matching patterns by comparing BFs originated from the same OSF but compiled with different architectures, compilers, and optimizations in dataset I. Interestingly, pattern e and f account for only 1% even with distant compilation settings. Instead, “1-to-1” matching accounts for the most common part as 70%, and “1-to-n” matching accounts for the rest 29%. To be more specific, comparing BFIs compiled in low optimizations with BFIs compiled in high optimizations are mostly cases falling to the pattern (c) and (d), which indicates the inlining is usually cumulative when optimization increases.

7.2 Suggestions for Designing More Effective Strategies

Considering the evaluation results of Bingo and ASM2Vec and proportion of matching patterns in binary2binary matching, we find that existing strategies have three main shortcomings:

First, existing strategies conduct inlining for all cases on two sides. However, 70% of matching cases do not need inlining, while 29% only need inlining at one side. Only 1% of matching cases need inlining at two sides. Second, inlining unnecessary functions will bring difficulties in identifying the true match. For example, Bingo inlines some functions both in two functions, where different BFs with the same callees after inlining will also become similar. Third, many ISFs are missed as the side-effect of reducing the inlining candidates. For example, the restriction of Bingo makes it miss nearly 40% of inlined functions, and ASM2Vec applies two restrictions for candidates resulting in more missed ISFs. As a result, it cannot help detect BFIs with vulnerable ISFs inlined. To overcome the shortcomings listed above, we propose several suggestions to reduce the inlining cost and help find the inlined functions.

Necessary preprocessing to reduce candidates. In binary2source matching, user-forced inlining candidates should be first identified. For example, functions with “`__attribute__((always_inline))`” will be inlined. Functions with “`__attribute__((noinline))`” will not be inlined. Thus, these functions can be determined whether to inline before applying strategies. Further, callees of not inlined functions can be filtered, and callees of inlined ones can be added to candidates.

Inspection of compilation settings to decide whether to apply strategies. Leveraging some works designed for toolchain provenance recovery [36, 66, 72–74], the compilation setting of binaries can be inferred. Then we could decide whether to apply strategies according their compilation settings. For example, in binary2source matching, strategies are not needed if the binary is generated without inlining (such as “-O0”). In binary2binary matching, strategies are not needed if binaries are generated in the same compilation setting.

Relaxing strategies of Bingo and ASM2Vec to recover more ISFs. Section 6 reveals the effectiveness of Bingo and ASM2Vec in reducing candidates but at the cost of missing inlined functions. We find that the strategy of Bingo and ASM2Vec is too strict. We can relax their restriction to obtain more inlined functions. For example, changing the restriction from $\alpha < 0.01$ to $\alpha < 0.5$ helps increase the similarity by 10% while only inlining 5 more BFs and increasing the depth of BFIs from 1 to 2 can improve the coverage of ISFs from 70% to nearly 93%.

Incremental inlining-simulation strategy to help locate ISFs. Considering that the most matching pattern is “1-to-n” when inlining happens, we do not need to expend all suitable callees at both sides. Interestingly, we find that

we can conduct an inlining trail to filter inlining candidates if the method can distinguish between more similar and less similar function pairs. For function pairs such as the example in Figure 12 (c), we named that NBF as A and BFI as B. Initially, we discover many inlining candidates for both A and B. For A, we take a trail to inline one of the candidates into A to form A with inlining. Then we compare the similarity between A and B to the similarity between A with inlining and B. If the similarity increase with inlining, we regard it as a suitable candidate and continue. Otherwise, the candidate, along with its callees, can be filtered. By conducting the inlining trail incrementally, we can quickly reduce the unnecessary branches and focus on some branches with deeper depth.

8 THREATS TO VALIDITY

Dataset labeling method. As our dataset labeling method is based on existing tools, the accuracy of dataset labeling depends on the accuracy of existing tools. In fact, we notice that there are difficulties for *understand* in parsing source entities, so we use a conservative approach to label the source code. In detail, we run *understand* in its “strict” mode where it will accurately identify the entity or add it to the unresolved list. And we manually label the unresolved entities to complement them.

Inspecting inlining from the generated binaries. Some readers may wonder why we inspect inlining from the generated binaries instead of the compiler designs. There are three reasons. First, understanding the inlining strategies from compilers’ source code seems feasible, but this requires great manual effort. Moreover, as the inlining strategies of different compilers in different versions will differ, we regard directly extracting inlining statistics from binary as a more general way. Second, inlining results rely on not only the inlining strategies but also the features of the source project. For example, the inlining decision of GCC depends on the total cost of all inlined call-sites; thus, without the exact project, the inlining decision will not be clear. Third, we also combine some inlining strategies with the inlining statistics in binary. We think using strategies to assist in understanding function inlining is a good complement for inspecting function inlining in real binaries.

Metrics for evaluating existing works. In Section 5, we follow the following principles to determine the metrics for evaluating existing works. If the work has used a metric to evaluate its effectiveness, we directly use it. Otherwise, we use its direct output, such as the similarities between functions, to present its results. We do not choose to implement a new metric because the function similarities are enough to present their effectiveness under inlining, as the metrics used to evaluate downstream tasks are all based on function similarity.

9 CONCLUSION

For the first time, our work investigates the effect of function inlining on binary similarity analysis research. Four datasets are constructed, and an automatic identification method is proposed to facilitate the analysis of function inlining. Our analysis finds that 36%-70% binary functions have functions inlined in high optimizations, while most binary code similarity still regards function mapping as “1-to-1”. This mismatch causes 30% loss in performance during code search and 40% loss during vulnerability detection. Moreover, inlined functions are nearly all ignored during OSS reuse detection and patch presence test. Furthermore, existing inlining-simulation strategies usually produce less complete inlining results, leaving the binary similarity analysis under function inlining still challenged.

We propose several suggestions to design a more effective inlining-simulation strategy, but we still regard it as unfinished work. We hope that more researchers can pay attention to function inlining and conduct more research to resolve function-inlining-related issues.

REFERENCES

- [1] 2009. Cisco settles FSF GPL lawsuit. <http://arstechnica.com/information-technology/2009/05/cisco-settles-fsf-gpl-lawsuit-appoints-compliance-officer> [Online; accessed April 7, 2021].
- [2] 2012. af_netlink: force credentials passing [CVE-2012-3520]. <https://github.com/torvalds/linux/commit/e0e3cea46d31d23d40df0a49a7a2c04fe8edfea>. [Online; accessed 23-February-2022].
- [3] 2015. VMware sued for failure to comply with Linux license. <https://www.zdnet.com/article/vmware-sued-for-failure-to-comply-with-linux-license/> [Online; accessed April 7, 2021].
- [4] 2018. llvm/InlineCost.cpp. <https://github.com/llvm-mirror/llvm/blob/88ab6705571782fa664ecfa71b2f959a0daf2d78/lib/Analysis/InlineCost.cpp>. [Online; accessed 23-April-2021].
- [5] 2020. Reveiws 1 - SciTools. <https://news.sophos.com/en-us/2020/04/26/asnarok/>. [Online; accessed 3-September-2021].
- [6] 2020. What's up, Emotet? CERT Polska. <https://cert.pl/en/posts/2020/02/whats-up-emotet/>. [Online; accessed 3-September-2021].
- [7] 2021. BinaryAI. <https://binaryai.tencent.com/>. [Online; accessed 27-November-2021].
- [8] 2021. BinDiff. <https://www.zynamics.com/bindiff.html>. [Online; accessed 10-December-2021].
- [9] 2021. Binutils - GNU Project - Free Software Foundation. <https://www.gnu.org/software/binutils/>. [Online; accessed 4-August-2021].
- [10] 2021. Clang: a C language family frontend for LLVM. <https://clang.llvm.org/>. [Online; accessed 20-April-2021].
- [11] 2021. Coreutils - GNU core utilities. <https://www.gnu.org/software/coreutils/>. [Online; accessed 23-April-2021].
- [12] 2021. DWARF Standard. <http://dwarfstd.org/Dwarf5Std.php>. [Online; accessed 20-April-2021].
- [13] 2021. findutils - GNU Project - Free Software Foundation (FSF). <https://www.gnu.org/software/findutils/>. [Online; accessed 4-August-2021].
- [14] 2021. Function Inlining. https://compileroptimizations.com/category/function_inlining.htm. [Online; accessed 3-September-2021].
- [15] 2021. GCC, the GNU Compiler Collection. <https://gcc.gnu.org/>. [Online; accessed 16-April-2021].
- [16] 2021. Ghidra Software Reverse Engineering Framework - GitHub. <https://github.com/NationalSecurityAgency/ghidra>. [Online; accessed 27-August-2021].
- [17] 2021. IDA Pro Disassembler and Debugger - Hex Rays. <https://www.hex-rays.com/ida-pro/>. [Online; accessed 20-April-2021].
- [18] 2021. LLVM's Analysis and Transform Passes. <https://llvm.org/docs/Passes.html>. [Online; accessed 16-April-2021].
- [19] 2021. Optimize Options (Using the GNU Compiler Collection (GCC)). <https://gcc.gnu.org/onlinedocs/gcc/Optimize-Options.html>. [Online; accessed 16-April-2021].
- [20] 2021. readelf(1) — Linux manual page. <https://man7.org/linux/man-pages/man1/readelf.1.html>. [Online; accessed 20-April-2021].
- [21] 2021. SciTools. <https://www.scitools.com/>. [Online; accessed 20-April-2021].
- [22] 2021. SoftSec-KAIST/BinKit:Binary Code Similarity Analysis (BCSA) Benchmark. <https://github.com/SoftSec-KAIST/BinKit>. [Online; accessed 4-August-2021].
- [23] 2021. The Chromium Projects. <https://www.chromium.org/>. [Online; accessed 28-December-2021].
- [24] 2021. The GNU Operating System and the Free Software Movement. <https://www.gnu.org>. [Online; accessed 4-August-2021].
- [25] 2021. The LLVM Compiler Infrastructure. <https://llvm.org/>. [Online; accessed 16-April-2021].
- [26] 2021. torvalds/linux: Linux kernel source tree. <https://github.com/torvalds/linux>. [Online; accessed 24-August-2021].
- [27] 2021. Tree-sitter - GitHub. <https://github.com/tree-sitter/tree-sitter>. [Online; accessed 28-December-2021].
- [28] 2021. "Asnarök" Trojan targets firewalls. <https://www.scitools.com/us-air-force-case-study>. [Online; accessed 3-September-2021].
- [29] 2022. bzip2. <http://www.bzip.org/>
- [30] 2022. minizip. <https://github.com/nmoinvaz/minizip>
- [31] 2022. OpenSSL. <https://www.openssl.org/>.
- [32] 2022. Source code and dataset. <https://github.com/island255/TOSEM2022>.
- [33] Pär Andersson. 2009. Evaluation of inlining heuristics in industrial strength compilers for embedded systems.
- [34] Dennis Andriess, Xi Chen, Victor van der Veen, Asia Slowinska, and Herbert Bos. 2016. An In-Depth Analysis of Disassembly on Full-Scale x86/x64 Binaries. In *25th USENIX Security Symposium, USENIX Security 16, Austin, TX, USA, August 10-12, 2016*. USENIX Association, 583–600. <https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/andriess>
- [35] Gu Ban, Lili Xu, Yang Xiao, Xinhua Li, Zimu Yuan, and Wei Huo. 2021. B2SMatcher: fine-Grained version identification of open-Source software in binary files. *Cybersecurity* 4, 1 (2021), 1–21.
- [36] Sébastien Bardin, Tristan Benoit, and Jean-Yves Marion. 2021. Compiler and optimization level recognition using graph neural networks. In *MLPA 2020-Machine Learning for Program Analysis*.
- [37] Mahinthan Chandramohan, Yinxiang Xue, Zhengzi Xu, Yang Liu, Chia Yuan Cho, and Hee Beng Kuan Tan. 2016. Bingo: Cross-architecture cross-os binary search. In *Proceedings of the 2016 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering*. 678–689.
- [38] Keith D. Cooper, Timothy J. Harvey, and Todd Waterman. 2008. An Adaptive Strategy for Inline Substitution. In *Compiler Construction, 17th International Conference, CC 2008, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2008, Budapest, Hungary, March 29 - April 6, 2008. Proceedings (Lecture Notes in Computer Science, Vol. 4959)*. Springer, 69–84. https://doi.org/10.1007/978-3-540-78791-4_5
- [39] Yaniv David, Nimrod Partush, and Eran Yahav. 2016. Statistical similarity of binaries. In *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2016, Santa Barbara, CA, USA, June 13-17, 2016*. ACM, 266–280. <https://doi.org/10.1145/>

- 2908080.2908126
- [40] Yaniv David and Eran Yahav. 2014. Tracelet-based code search in executables. *Acm Sigplan Notices* 49, 6 (2014), 349–360.
- [41] Jack W. Davidson and Anne M. Holler. 1992. Subprogram Inlining: A Study of its Effects on Program Execution Time. *IEEE Trans. Software Eng.* 18, 2 (1992), 89–102. <https://doi.org/10.1109/32.121752>
- [42] Jeffrey Dean and Craig Chambers. 1994. Towards Better Inlining Decisions Using Inlining Trials. In *Proceedings of the 1994 ACM Conference on LISP and Functional Programming, Orlando, Florida, USA, 27-29 June 1994*. ACM, 273–282. <https://doi.org/10.1145/182409.182489>
- [43] Steven HH Ding, Benjamin CM Fung, and Philippe Charland. 2016. Kam1n0: Mapreduce-based assembly clone search for reverse engineering. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. 461–470.
- [44] Steven HH Ding, Benjamin CM Fung, and Philippe Charland. 2019. Asm2vec: Boosting static representation robustness for binary clone search against code obfuscation and compiler optimization. In *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 472–489.
- [45] Ruian Duan, Ashish Bijlani, Meng Xu, Taesoo Kim, and Wenke Lee. 2017. Identifying Open-Source License Violation and 1-day Security Risk at Large Scale. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS 2017, Dallas, TX, USA, October 30 - November 03, 2017*. ACM, 2169–2185. <https://doi.org/10.1145/3133956.3134048>
- [46] Yue Duan, Xuezixiang Li, Jinghan Wang, and Heng Yin. 2020. Deepbindiff: Learning program-wide code representations for binary diffing. In *Network and Distributed System Security Symposium*.
- [47] Irène A Durand and Robert Strandh. 2018. Partial Inlining Using Local Graph Rewriting. In *11th European Lisp Symposium*.
- [48] Manuel Egele, Maverick Woo, Peter Chapman, and David Brumley. 2014. Blanket execution: Dynamic similarity testing for program binaries and components. In *23rd {USENIX} Security Symposium ({USENIX} Security 14)*. 303–317.
- [49] Muyue Feng, Weixuan Mao, Zimu Yuan, Yang Xiao, Gu Ban, Wei Wang, Shiyang Wang, Qian Tang, Jiahuan Xu, He Su, Binghong Liu, and Wei Huo. 2019. Open-Source License Violations of Binary Software at Large Scale. In *26th IEEE International Conference on Software Analysis, Evolution and Reengineering, SANER 2019, Hangzhou, China, February 24-27, 2019*. IEEE, 564–568. <https://doi.org/10.1109/SANER.2019.8667977>
- [50] Muyue Feng, Zimu Yuan, Feng Li, Gu Ban, Shiyang Wang, Qian Tang, He Su, Chendong Yu, Jiahuan Xu, Aihua Piao, Jingling Xue, and Wei Huo. 2019. B2SFinder: Detecting Open-Source Software Reuse in COTS Software. In *34th IEEE/ACM International Conference on Automated Software Engineering, ASE 2019, San Diego, CA, USA, November 11-15, 2019*. IEEE, 1038–1049. <https://doi.org/10.1109/ASE.2019.00100>
- [51] Debin Gao, Michael K. Reiter, and Dawn Xiaodong Song. 2008. BinHunt: Automatically Finding Semantic Differences in Binary Programs. In *Information and Communications Security, 10th International Conference, ICICS 2008, Birmingham, UK, October 20-22, 2008, Proceedings (Lecture Notes in Computer Science, Vol. 5308)*. Springer, 238–255. https://doi.org/10.1007/978-3-540-88625-9_16
- [52] Jian Gao, Xin Yang, Ying Fu, Yu Jiang, and Jianguang Sun. 2018. VulSeeker: a semantic learning based vulnerability seeker for cross-platform binary. In *2018 33rd IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 896–899.
- [53] Dale Gardner. 2019. Technology Insight for Software Composition Analysis. <https://www.gartner.com/en/documents/3971011/technology-insight-for-software-composition-analysis>. [Online; accessed 27-November-2021].
- [54] Armijn Hemel, Karl Trygve Kalleberg, Rob Vermaas, and Eelco Dolstra. 2011. Finding software license violations through binary code clone detection. In *Proceedings of the 8th International Working Conference on Mining Software Repositories, MSR 2011 (Co-located with ICSE), Waikiki, Honolulu, HI, USA, May 21-28, 2011, Proceedings*. ACM, 63–72. <https://doi.org/10.1145/1985441.1985453>
- [55] Jan Hubicka. 2004. The GCC call graph module: a framework for inter-procedural optimization.
- [56] Wen-mei W. Hwu and Pohua P. Chang. 1989. Inline Function Expansion for Compiling C Programs. In *Proceedings of the ACM SIGPLAN'89 Conference on Programming Language Design and Implementation (PLDI), Portland, Oregon, USA, June 21-23, 1989*. ACM, 246–257. <https://doi.org/10.1145/73141.74840>
- [57] Wen-mei W. Hwu and Pohua P. Chang. 1989. Inline Function Expansion for Compiling C Programs. In *Proceedings of the ACM SIGPLAN'89 Conference on Programming Language Design and Implementation (PLDI), Portland, Oregon, USA, June 21-23, 1989*. ACM, 246–257. <https://doi.org/10.1145/73141.74840>
- [58] Yuede Ji, Lei Cui, and H Howie Huang. 2021. BugGraph: Differentiating Source-Binary Code Similarity with Graph Triplet-Loss Network. In *16th ACM ASIA Conference on Computer and Communications Security (ASIACCS)*.
- [59] Dongjin Kim, Seong-je Cho, Sangchul Han, Minkyu Park, and Ilsoon You. 2014. Open Source Software Detection using Function-level Static Software Birthmark. *J. Internet Serv. Inf. Secur.* 4, 4 (2014), 25–37. <https://doi.org/10.22667/JISIS.2014.11.31.025>
- [60] Dongkwan Kim, Eunsoo Kim, Sang Kil Cha, Soeul Son, and Yongdae Kim. 2020. Revisiting Binary Code Similarity Analysis using Interpretable Feature Engineering and Lessons Learned. (2020). arXiv:2011.10749 [cs.SE]
- [61] Zhouhan Lin, Minwei Feng, Cicero Nogueira dos Santos, Mo Yu, Bing Xiang, Bowen Zhou, and Yoshua Bengio. 2017. A structured self-attentive sentence embedding. *arXiv preprint arXiv:1703.03130* (2017).
- [62] Bingchang Liu, Wei Huo, Chao Zhang, Wenchao Li, Feng Li, Aihua Piao, and Wei Zou. 2018. α diff: cross-version binary code similarity detection with dnn. In *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering*. 667–678.
- [63] Lannan Luo, Jiang Ming, Dinghao Wu, Peng Liu, and Sencun Zhu. 2017. Semantics-Based Obfuscation-Resilient Binary Code Similarity Comparison with Applications to Software and Algorithm Plagiarism Detection. *IEEE Trans. Software Eng.* 43, 12 (2017), 1157–1177. <https://doi.org/10.1109/TSE.2017.2655046>
- [64] Luca Massarelli, Giuseppe Antonio Di Luna, Fabio Petroni, Roberto Baldoni, and Leonardo Querzoni. 2019. Safe: Self-attentive function embeddings for binary similarity. In *International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment*. Springer, 309–329.

- [65] Dhaval Miyani, Zhen Huang, and David Lie. 2017. BinPro: A Tool for Binary Source Code Provenance. *CoRR* abs/1711.00830 (2017). arXiv:1711.00830 <http://arxiv.org/abs/1711.00830>
- [66] Yuhei Otsubo, Akira Otsuka, Mamoru Mimura, Takeshi Sakaki, and Hiroshi Ukegawa. 2020. o-glassesX: compiler provenance recovery with attention mechanism from a short code fragment. In *Proceedings of the 3rd Workshop on Binary Analysis Research*.
- [67] Chengbin Pang, Ruotong Yu, Yaohui Chen, Eric Koskinen, Georgios Portokalidis, Bing Mao, and Jun Xu. 2021. Sok: All you ever wanted to know about x86/x64 binary disassembly but were afraid to ask. In *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, 833–851.
- [68] Jannik Pewny, Behrad Garmany, Robert Gawlik, Christian Rossow, and Thorsten Holz. 2017. Cross-architecture bug search in binary executables. *it Inf. Technol.* 59, 2 (2017), 83. <https://doi.org/10.1515/itit-2016-0040>
- [69] Andrea Pinna, Simona Ibba, Gavina Baralla, Roberto Tonelli, and Michele Marchesi. 2019. A massive analysis of ethereum smart contracts empirical study and code metrics. *IEEE Access* 7 (2019), 78194–78213.
- [70] Aleksandar Prokopec, Gilles Duboscq, David Leopoldseeder, and Thomas Würthinger. 2019. An Optimization-Driven Incremental Inline Substitution Algorithm for Just-in-Time Compilers. In *IEEE/ACM International Symposium on Code Generation and Optimization, CGO 2019, Washington, DC, USA, February 16-20, 2019*. IEEE, 164–179. <https://doi.org/10.1109/CGO.2019.8661171>
- [71] Ashkan Rahimian, Philippe Charland, Stere Preda, and Mourad Debbabi. 2012. RESource: A Framework for Online Matching of Assembly with Open Source Code. In *Foundations and Practice of Security - 5th International Symposium, FPS 2012, Montreal, QC, Canada, October 25-26, 2012, Revised Selected Papers (Lecture Notes in Computer Science, Vol. 7743)*. Springer, 211–226. https://doi.org/10.1007/978-3-642-37119-6_14
- [72] Ashkan Rahimian, Paria Shirani, Saed Alrbaee, Lingyu Wang, and Mourad Debbabi. 2015. Bincomp: A stratified approach to compiler provenance attribution. *Digital Investigation* 14 (2015), S146–S155.
- [73] Nathan Rosenblum, Barton P Miller, and Xiaojin Zhu. 2011. Recovering the toolchain provenance of binary code. In *Proceedings of the 2011 International Symposium on Software Testing and Analysis*. 100–110.
- [74] Zhenzhou Tian, Yaqian Huang, Borun Xie, Yanping Chen, Lingwei Chen, and Dinghao Wu. 2021. Fine-Grained Compiler Identification With Sequence-Oriented Neural Modeling. *IEEE Access* 9 (2021), 49160–49175.
- [75] Shuai Wang and Dinghao Wu. 2017. In-memory fuzzing for binary code similarity analysis. In *2017 32nd IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 319–330.
- [76] Xiaojun Xu, Chang Liu, Qian Feng, Heng Yin, Le Song, and Dawn Song. 2017. Neural network-based graph embedding for cross-platform binary code similarity detection. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. 363–376.
- [77] Xi Xu, Qinghua Zheng, Zheng Yan, Ming Fan, Ang Jia, and Ting Liu. 2021. Interpretation-enabled Software Reuse Detection Based on a Multi-Level Birthmark Model. In *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*. IEEE, 873–884.
- [78] Yifei Xu, Zhengzi Xu, Bihuan Chen, Fu Song, Yang Liu, and Ting Liu. 2020. Patch based vulnerability matching for binary programs. In *Proceedings of the 29th ACM SIGSOFT International Symposium on Software Testing and Analysis*. 376–387.
- [79] Yinxing Xue, Zhengzi Xu, Mahinthan Chandramohan, and Yang Liu. 2018. Accurate and scalable cross-architecture cross-os binary code search with emulation. *IEEE Transactions on Software Engineering* 45, 11 (2018), 1125–1149.
- [80] Cheng Yang, Zhiyuan Liu, Deli Zhao, Maosong Sun, and Edward Chang. 2015. Network representation learning with rich text information. In *Twenty-fourth international joint conference on artificial intelligence*.
- [81] Zeping Yu, Wenxin Zheng, Jiaqi Wang, Qiyi Tang, Sen Nie, and Shi Wu. 2020. CodeCMR: Cross-Modal Retrieval For Function-Level Binary Source Code Matching. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*. <https://proceedings.neurips.cc/paper/2020/hash/285f89b802bcb2651801455c86d78f2a-Abstract.html>
- [82] Hang Zhang and Zhiyun Qian. 2018. Precise and Accurate Patch Presence Test for Binaries. In *27th USENIX Security Symposium, USENIX Security 2018, Baltimore, MD, USA, August 15-17, 2018*. USENIX Association, 887–902. <https://www.usenix.org/conference/usenixsecurity18/presentation/zhang-hang>
- [83] Peng Zhao and José Nelson Amaral. 2003. To Inline or Not to Inline? Enhanced Inlining Decisions. In *Languages and Compilers for Parallel Computing, 16th International Workshop, LCPC 2003, College Station, TX, USA, October 2-4, 2003, Revised Papers (Lecture Notes in Computer Science, Vol. 2958)*. Springer, 405–419. https://doi.org/10.1007/978-3-540-24644-2_26
- [84] Fei Zuo, Xiaopeng Li, Patrick Young, Lannan Luo, Qiang Zeng, and Zhexin Zhang. 2019. Neural Machine Translation Inspired Binary Code Similarity Comparison beyond Function Pairs. In *26th Annual Network and Distributed System Security Symposium, NDSS 2019, San Diego, California, USA, February 24-27, 2019*. The Internet Society.