

# FUNPROBE: Probing Functions from Binary Code through Probabilistic Analysis

Soomin Kim SoftSec Lab., KAIST Daejeon, Korea soomink@softsec.kaist.ac.kr Hyungseok Kim SoftSec Lab., KAIST Daejeon, Korea hskim@softsec.kaist.ac.kr Sang Kil Cha SoftSec Lab., KAIST Daejeon, Korea sangkilc@softsec.kaist.ac.kr

# ABSTRACT

Current function identification techniques have been mostly focused on a specific set of binaries compiled for a specific CPU architecture. While recent deep-learning-based approaches theoretically can handle binaries from different architectures, they require significant computation resources for training and inference, making their use less practical. Furthermore, due to the lack of interpretability of such models, it is fundamentally difficult to gain insight from them. Hence, in this paper, we propose FUNPROBE, an efficient system for identifying functions from binaries using probabilistic inference. In particular, we identify 16 architecture-neutral hints for function identification, and devise an effective method to combine them in a probabilistic framework. We evaluate our tool on a large dataset consisting of 19,872 real-world binaries compiled for six major CPU architectures. The results are promising. FUNPROBE shows the best accuracy compared to five state-of-the-art tools we tested, while it takes only 6 seconds on average to analyze a single binary. Notably, FUNPROBE is 6× faster on average in identifying functions than XDA, a state-of-the-art deep-learning tool that leverages GPU in its inference phase.

## **CCS CONCEPTS**

 $\bullet$  Software and its engineering  $\rightarrow$  Software testing and debugging.

# **KEYWORDS**

binary code analysis, function identification, probabilistic analysis

#### **ACM Reference Format:**

Soomin Kim, Hyungseok Kim, and Sang Kil Cha. 2023. FUNPROBE: Probing Functions from Binary Code through Probabilistic Analysis. In Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE '23), December 3–9, 2023, San Francisco, CA, USA. ACM, New York, NY, USA, 12 pages. https: //doi.org/10.1145/3611643.3616366

# **1 INTRODUCTION**

Function identification is a pivotal task in binary analysis. Major decompilation techniques, such as variable recovery [5], type recovery [35], and high-level control-flow restructuring [15, 52, 53, 62],

ESEC/FSE '23, December 3–9, 2023, San Francisco, CA, USA © 2023 Copyright held by the owner/author(s). ACM ISBN 979-8-4007-0327-0/23/12.

https://doi.org/10.1145/3611643.3616366

operate at a function level, assuming Control-Flow Graph (CFG) of a function is given. Binary-level control-flow integrity approaches, such as [65], adopt a function-level protection mechanism, too.

Conventional approaches to binary-level function identification employ various target-specific *heuristics*, which are often specifically devised to handle specific types of binaries, to overcome the intrinsic difficulty of binary analysis. For example, Ghidra [42] and FETCH [45] identify functions by leveraging C++ exception handling information, which does not necessarily exist in all binaries. Similarly, DDisasm [13] and FunSeeker [27] leverage target-specific syntactic patterns, and Jima [2] harvests function pointers from analyzing specific data sections. While such approaches are effective in identifying functions in certain binaries, they are not generally applicable to all binaries.

Learning-based approaches [6, 32, 60, 63] have been proposed to overcome the limitations of conventional approaches. Particularly, deep-learning-based approaches [48, 54, 64] recently demonstrate high accuracy in identifying functions. While these approaches do not require target-specific heuristics, they still suffer from three critical limitations. First, their detection performance is highly dependent on the training dataset used. Second, they require significant computation resources for both training and inference. Thus, applying them to large-scale binary analyses is not trivial in practice. Furthermore, those models are *not* interpretable, making it difficult to gain useful insight into the learned models. As such, one cannot easily understand why a particular function is misclassified, hence further improving the model is challenging.

Therefore, in this paper, we seek to develop a function identification algorithm that is *general*, *efficient*, and *interpretable*. Our approach should work for a wide range of binaries compiled for various CPU architectures; should run fast on a regular desktop machine; and should not rely on a deep-learning model that is difficult to understand.

To this end, we identify 16 architecture-neutral *hints*, and combine them to form a probabilistic model for function identification. The key intuition of our approach is that every function identification heuristic has a certain level of uncertainty, which can naturally be represented as a probability. For example, many tools heuristically regard a target of a call instruction as a function entry point, but call targets are not always a function entry point in reality. Such a heuristic is not always correct, but does provide a probabilistic *hint* for function identification. Hence, we combine those hints to form a Bayesian Network (BN) representing causal relationships between them. Our BN-based approach naturally allows us to reason about why a particular address is identified as a function entry point because the posterior probabilities let us know the influence of each hint on the address.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

There are mainly two challenges in designing our algorithm. First, our approach should be generally applicable to any type of compiler-generated binaries built with varying compilers, compiler options, and target architectures. Second, our model, i.e., a BN, naturally includes cyclic dependencies, which is expensive to handle, if not impossible [30, 31, 51, 56]. Hence, we should devise an efficient way to perform probabilistic inference on our BN.

We address the first challenge by carefully selecting architectureand compiler-independent heuristics, such as those found in control flows between functions and basic blocks [25]. FUNPROBE employs 16 intuitive function identification hints obtained either from existing work or from our own observations. We also tackle the second challenge by exploiting the fact that our BN includes many unnecessary edges. Specifically, we devise a novel approach, named bogus dependency pruning, which heuristically removes edges as well as loops from a BN. Our technique enables 12.6× faster function identification with negligible accuracy drop as shown in our experiments.

To realize these ideas, we implement FUNPROBE, a novel function identification tool that runs on raw (i.e., stripped) binaries. It leverages our bogus dependency pruning technique to quickly decide whether every instruction in the target binary should be considered to be a function entry point or not. We evaluate FUNPROBE on a large-scale benchmark that includes 19,872 real-world binaries compiled for six major CPU architectures. Our evaluation confirms that FUNPROBE shows the best accuracy compared to the five stateof-the-art binary analysis tools, including IDA Pro, Ghidra, and XDA, while it takes only 6 seconds on average to analyze a single binary. Notably, FUNPROBE is 6× faster than XDA, a state-of-the-art deep-learning tool, without the need for GPU resources. We also demonstrate that FUNPROBE is complementary to XDA. By simply feeding the output of XDA into FUNPROBE, we were able to achieve 99.9% of F1-score on our benchmark.

Overall, this paper makes the following contributions.

- We present a novel function identification algorithm that leverages a BN.
- We propose bogus dependency pruning to efficiently derive solutions from a BN.
- We design and implement FUNPROBE that incorporates our novel function identification algorithm.
- We publicize our tool to support open science: https://gith ub.com/B2R2-org/FunProbe.

## 2 BACKGROUND AND MOTIVATION

This section describes the basic concept of Bayesian Network (BN) and belief propagation, and motivates our research.

*Notation.* In this paper, P(X) denotes the probability distribution of a Boolean random variable *X*. For simplicity, we let P(x) be the probability of *X* being true, i.e., P(x) = P(X = 1).

#### 2.1 Bayesian Network & Belief Propagation

Bayesian Network (BN) is a set of Directed Acyclic Graphs (DAGs), each of which represents causal dependencies between random variables [46]. Each node in a BN is a random variable, and each directed edge represents a dependency between two random variables. For example, an edge  $X \rightarrow Y$  means that *X* causes *Y* (or *Y* is dependent on *X*).

We use a BN to model the relationships between probabilistic statements, such as "the probability of an address being a function entry point". Consider a basic block located at address  $\alpha$ . Let  $F_{\alpha}$  be a Boolean random variable indicating whether or not  $\alpha$  is a function entry point, and let  $C_{\alpha}$  be a Boolean random variable denoting whether or not  $\alpha$  is a target of a call instruction. We can then represent a causal relationship between  $F_{\alpha}$  and  $C_{\alpha}$  with an edge  $C_{\alpha} \rightarrow F_{\alpha}$  in a BN: If  $C_{\alpha}$  is true (there is a call instruction whose target is  $\alpha$ ), then  $F_{\alpha}$  is likely to be true ( $\alpha$  is a function entry point). Such a relationship can be denoted by a conditional probability  $P(f_{\alpha} | c_{\alpha})$ , and we call it a **hint** throughout this paper because it provides a *hint* for identifying functions.

Notice not every random variable in a BN is observable. In the above example,  $F_{\alpha}$  is a *hidden* random variable, whose value cannot be directly observed from analyzing a binary. On the other hand,  $C_{\alpha}$  can be easily observed: We can disassemble a binary and see if there is a call instruction whose target is  $\alpha$ . Therefore, our goal in this paper is to compute the marginal probabilities of hidden random variables, e.g.,  $P(f_{\alpha})$ , based on the probability distributions of the observed variables. This process is often referred to as *probabilistic inference* [10], which does *not* scale well with the number of random variables—the number of terms to consider for marginalization grows exponentially to the number of hidden variables.

Belief propagation [46] is a technique that addresses the scalability challenge with a dynamic programming method. When a BN contains a loop, however, belief propagation cannot compute the exact solution [40]. Thus, *loopy belief propagation* is used as an alternative [22, 30, 40, 51], which iteratively runs belief propagation until it converges or reaches a fixed time limit. However, when a graph is large and highly connected, (loopy) belief propagation is still computationally expensive [12, 22]. In this paper, we propose an efficient way to perform belief propagation on a large-size cyclic BN by reducing the graph size (see §4.3).

#### 2.2 Motivation

While *hints* are useful to identify functions, they do *not* provide a definite answer. In this section, we motivate the need for a probabilistic method by illustrating how one can combine various hints collected from a binary to make a precise decision. To support our claim, we ran Nucleus on dir, a binary taken from GNU Coreutils, to identify functions from it. Figure 1a shows the disassembled code snippets of the binary. To ease the explanation, we explicitly mark function entry points in Figure 1a with symbols, although we used a stripped binary when we ran Nucleus.

In this example, Nucleus identifies functions using two hints: (1) the target of a call instruction is likely to be a function entry point, and (2) a no-op-like instruction is unlikely to be a function entry point (as it is often a padding sequence). Unfortunately, both hints suffer from a certain level of uncertainty. The first hint seems legitimate as call instructions are designed to call a subroutine in a program. However, there are cases where a call instruction is used to retrieve the current Program Counter (PC). For example, a code pattern "call +5; pop ebx", where the target of call +5 is FUNPROBE: Probing Functions from Binary Code through Probabilistic Analysis

1	0x3bb0 <mai< th=""><th><b>n</b>&gt;:</th><th></th><th></th></mai<>	<b>n</b> >:		
2	0x3bb0:	push	r15	
3	0x3bb2:	push	r14	
4	0x3bb4:	mov	r14d,	edi
5	0x3bb7:	push	r13	
6	0x3bb9:	push	r12	
7				
8	0x4944:	mov	edi,	0xd
9	0x4949:	call	58b0	<is_colored></is_colored>
10				
11	0x5830 <dev< th=""><th>_ino_f</th><th>ree&gt;:</th><th></th></dev<>	_ino_f	ree>:	
12	0x5830:	jmp	1b770	<rpl_free></rpl_free>
13				
14	0x58b0 <is_< th=""><th>colore</th><th>d&gt;:</th><th></th></is_<>	colore	d>:	
15	0x58b0:	mov	edi,	edi
16	0x58b2:	lea	rax,	[rip+0x2217a7]
17	0x58b9:	shl	rdi,	0x4
18	0x58bd:	add	rdi,	rax
19	0x58c0:	xor	eax,	eax
20	0x58c2:	mov	rdx,	QWORD PTR [rdi]
21	0x58c5:	test	rdx,	rdx
22	0x58c8:	je	58df	
23				

(a) Disassembled x86-64 binary code.

1 0000f4 000010 0000c8 FDE cie=000030 pc=000058b0..00005909 2 DW\_CFA\_nop

3 ...

(b) Frame Description Entry (FDE) for is\_colored.

Figure 1: Our motivating example taken from dir, a GNU Coreutils binary compiled with GCC.

the pop instruction, is often used to obtain the current PC. Similarly, the latter hint is not always correct: Compilers sometimes emit a no-op-like instruction at a function entry point.

The primary challenge here is that there are cases where two hints are in conflict. Conventional approaches, such as Nucleus, would simply follow one of them. In our example, the function at 0x58b0 is a target of the call instruction at 0x4949. Thus, Nucleus can correctly identify the function with the first hint. However, the function begins with a no-op instruction, which is deemed as an invalid instruction by Nucleus based upon the second hint. Note that mov edi, edi does not change the CPU state except for the program counter, and is often used as a padding byte. The second hint makes Nucleus disregard the function, thereby causing a false negative error.

FUNPROBE handles such an intrinsic challenge with a probabilistic framework. Intuitively, each different *hint* provides a different *clue* about identifying function entry points, and we can represent their relationships with a BN. If we collectively consider the observed hints in the BN, we can make a decision in a more *systematic* and *holistic* way.

For example, let us consider an additional heuristic developed by FETCH [45], which leverages exception handling information stored in the .eh\_frame section of ELF binaries. The section stores a sequence of Frame Description Entries (FDEs), each of which corresponds to a consecutive code chunk in the binary. Typically, each code chunk represents a function (although there are exceptional cases). Therefore, in our example, the FDE shown in Figure 1b indicates that there is a code chunk located at 0x58b0, and it is likely that the address indicates a function entry point. With such an additional hint, we can say that 0x58b0 is more likely to be a ESEC/FSE '23, December 3-9, 2023, San Francisco, CA, USA



Figure 2: FUNPROBE architecture.



Figure 3: Illustration of the BN generated from our example.

function entry point (two positives vs. one negative). As we will discuss, FUNPROBE provides a way to systematically make an informed decision by gathering all the observed hints.

## **3 OVERVIEW**

In this section, we first present the overall architecture of FUN-PROBE, and describe its workflow with a running example shown in Figure 1. We then discuss several technical challenges in designing FUNPROBE.

## 3.1 **FUNPROBE** Architecture

Figure 2 illustrates the overall architecture of FUNPROBE, which takes in a binary as input, and produces a set of function entry points as output. FUNPROBE consists of three components: (1) Hint Collector, (2) Model Builder, and (3) Inferencer. Hint Collector harvests hints by analyzing the given binary. Model Builder constructs a BN using the collected hints. Inferencer runs belief propagation on the BN to infer the marginal probabilities for each address in the binary can be a function entry point.

We now demonstrate the component-wise workflow of FUN-PROBE using the example binary shown in Figure 1. Note our system runs on a stripped version of the binary.

#### 3.2 Hint Collector

First, Hint Collector takes in a binary as input, and outputs a set of hints taken from the binary. Formally, a *hint* is a conditional probability for an address being a function entry point. Let  $F_{\alpha}$  be a hidden Boolean random variable indicating whether or not  $\alpha$  is a function entry point. We can then define a *hint* as a conditional probability of the form  $P(f_{\alpha} | x)$  where X is a Boolean random variable and x is a shortcut for X = 1 as we defined in §2. Hint Collector analyzes the given binary to construct an inter-procedural CFG (§4.1), and, it collects 16 different types of hints from the CFG as well as the metadata stored in the binary (§4.2).

From our example binary shown in Figure 1, Hint Collector outputs a variety of hints including the following four:  $P(f_{58b0} | c_{58b0})$ ,  $P(f_{58b0} | e_{58b0})$ ,  $P(f_{58b0} | n_{58b0})$ , and  $P(f_{1b770} | f_{58b0})$ ,

where  $C_{58b0}$  indicates that  $0 \times 58b0$  is a call target,  $E_{58b0}$  means that there is an FDE for the instruction located at  $0 \times 58b0$ , and  $N_{58b0}$ shows that the instruction located at  $0 \times 58b0$  is semantically a no-op instruction; It does not change any CPU state except for the PC.

Note that two *hidden* random variables, e.g.,  $F_{1b770}$  and  $F_{58b0}$ , can also have a causal relationship. Let us consider the tail-call heuristic used by many tools [11, 42, 45, 49], which regards a jump as a tail-call when it crosses over a function entry point. This strategy is based on the observation that compilers usually emit a function body in a consecutive region. Therefore, if a jump instruction passes over another function entry point, it means the jump target is likely to belong to another function. In our example binary, there is a jump instruction located at 0x5830, and there is another function entry point at 0x58b0 in between the jump instruction and its target (0x1b770). Therefore, if there is a function entry point between 0x5830 and 0x1b770, then 0x1b770 is also likely to be a function entry point.

This observation provides at most 89,919 distinct hints ( $P(f_{1b770} | f_{5831})$ ,  $P(f_{1b770} | f_{5832})$ ,  $\cdots$ ,  $P(f_{1b770} | f_{1b76f})$ , in case the addresses are filled with single-byte instructions), which can easily bloat up the BN, rendering our probabilistic inference inefficient. Thus, we reduce the number with a novel technique, named bogus dependency pruning (§4.3).

## 3.3 Model Builder

Model Builder first builds a BN from the observed hints. Specifically, each hint forms an edge in the resulting BN. Figure 3 illustrates the BN obtained from our example binary. The BN only shows four hints (i.e., edges) for simplicity.

Next, Model Builder assigns probabilities for each edge based on the characteristic of each corresponding hint. Specifically, we use two user-configurable probability values: one for positive hints and another for negative hints. We say a hint is positive if it increases the likelihood of the corresponding address being a function entry point, and negative if otherwise. For example,  $P(f_{58b0} | c_{58b0})$  increases the likelihood of 0x58b0 being a function entry point because it is used as a call target, thus, it is a positive hint. On the contrary,  $P(f_{58b0} | n_{58b0})$  decreases the likelihood of 0x58b0 being a function entry point because the likelihood of 0x58b0 being a function entry point because the likelihood of 0x58b0 being a function entry point because there is a no-op instruction at 0x58b0, thus, it is a negative hint.

Obviously, positive hints, such as  $P(f_{58b0} | c_{58b0})$ , should have a high value (close to 1), and negative hints, such as  $P(f_{58b0} | n_{58b0})$ , should have a low value (close to 0) in order for them to respectively give a positive and negative influence to the probability of  $0 \times 58b0$ being a function entry point. We introduce two user-configurable parameters  $\mathcal{P}^+$  and  $\mathcal{P}^-$  in order to assign the probability values for positive and negative hints. In our current implementation, we use  $\mathcal{P}^+ = 0.65$  and  $\mathcal{P}^- = 0.4$  by default, which are empirically chosen best parameter values (see §5.2). Thus, in the previous example, we assign 0.4 to  $P(f_{58b0} | n_{58b0})$ , and 0.65 to the other edges. The marginal probabilities, i.e.,  $P(f_{58b0})$  and  $P(f_{1b770})$ , are then computed in the next step based on the assigned probabilities.

## 3.4 Inferencer

Inferencer takes in as input the initialized BN, and performs belief propagation on the BN to compute the marginal probabilities. That is, our goal here is to obtain  $P(F_{58b0})$  and  $P(F_{1b770})$  shown in Figure 3. Once we have the marginal probabilities, we can decide whether each address is a function entry point or not. In our current implementation, we say an address  $\alpha$  is a function entry point if  $P(f_{\alpha}) > 0.5$ . In our example, using belief propagation, Inferencer returns  $P(f_{58b0}) = 0.9216$  and  $P(f_{1b770}) = 0.9999$ . Since these probabilities are higher than our threshold 0.5, we regard both of them as a function entry point.

Recall from §3.2 that a tail-call detection strategy can easily produce thousands of hints, which makes the resulting BN complex and highly connected with loops. Therefore, Inferencer does not scale well with large binaries with many functions. We overcome this challenge by introducing a novel technique, named *bogus dependency pruning*, as we detail in §4.3.

## **4 FUNPROBE DESIGN**

This section details the design of FUNPROBE. We first show how FUNPROBE prepares an inter-procedural CFG to extract hints from the binary (§4.1). Next, we present all the function identification hints that FUNPROBE utilizes and discuss how we make them general (§4.2). Finally, we present bogus dependency pruning, a technique to simplify BNs to make probabilistic inference efficient (§4.3).

## 4.1 Control Flow Graph Construction

To collect hints from binaries in architecture-independent and compiler-independent manner, we mainly focus on structural features that can be obtained from a CFG instead of observing instruction patterns. One can leverage existing binary analysis frameworks, such as Ghidra and angr, to obtain CFGs from a binary, but those tools typically involve heavy-cost analyses. Thus, we perform our own lightweight analysis to quickly recover CFGs so that we can apply our function identification strategies to them. Note that our CFG recovery is a preprocessing step for function identification, and our goal is far from constructing precise CFGs. Our analysis runs in the following four steps.

- (1) Linearly disassemble the given binary (§4.1.1).
- (2) Detect non-returning function calls (§4.1.2).
- (3) Resolve indirect branch targets (§4.1.3).
- (4) Build inter-procedural CFG (§4.1.4).

4.1.1 *Linear-Sweep Disassembly.* Hint Collector first linearly disassembles instructions from a given binary. One could use superset disassembly [7] to completely recover all possible instructions in the binary, but linear-sweep disassembly is widely known to be efficient in achieving high instruction coverage without many false positives [3]. Therefore, we chose linear disassembly, a simpler technique for our implementation.

One challenge here is to avoid disassembling embedded data in code sections, which are commonly found in ARMv7 binaries. To distinguish code and data, we examine every memory load from the disassembled instructions, and filter out instructions that are referenced from another instruction.

4.1.2 Non-Returning Function Call Detection. Not every function call returns. Thus, we cannot simply connect a CFG edge from a call instruction to its fall-through instruction. Having a bogus edge

#### Table 1: Function Identifcation Hints that FUNPROBE used.

	Function Identification Hints	R. V.	Influence	Used By
Data	<ol> <li>Pointers in known pointer-array sections (e.g., .init_array) refer to a function entry point.</li> <li>Non-strippable function symbols point to a function entry point.</li> <li>Relocation entries can specify a function entry point.</li> <li>An FDE (Frame Description Entry) can point to a function entry point.</li> <li>Pointer-like values in data sections (e.g., .data) can refer to a function entry point.</li> </ol>	$P_{\alpha} \\ D_{\alpha} \\ R_{\alpha} \\ E_{\alpha} \\ V_{\alpha}$	+++++	[13] [42, 55] [9, 13] [2, 13, 27, 42, 45, 49] [13, 59]
Code	<ul> <li>6 Call targets can be a function entry point.</li> <li>7 When a jump edge crosses a function entry point, then the jump is likely to be a tail-call.</li> <li>(8) If a call target is a jump instruction, then the jump target is likely to be a function entry point.</li> <li>(9) Unreachable basic blocks are likely to be a function.</li> <li>(10) An unreachable basic block surrounded by two connected basic blocks is unlikely to be a function.</li> <li>(11) Padding sequences, such as no-op, are unlikely to be a function entry point.</li> <li>(12) If there are only no-op instructions between a jump instruction and its target then the jump is likely to be a no-op.</li> <li>(13) Inlined PC-getters are unlikely to be a function.</li> <li>(14) Conditional jump targets are unlikely to be a function entry point.</li> <li>(15) Basic block before the main entry point of a binary are unlikely to be a function.</li> <li>(16) Basic block leaders are unlikely to be a function entry point.</li> </ul>	$\begin{array}{c} C_{\alpha} \\ F_{\alpha} \\ U_{\alpha} \\ U_{\alpha} \\ S_{\alpha} \\ I_{\alpha} \\ I_{\alpha} \\ G_{\alpha} \\ J_{\alpha} \\ A_{\alpha} \\ B_{\alpha} \end{array}$	++++	All [11, 42, 45, 49] New [2, 4, 13, 49, 55] New [4, 13, 49] [42, 55] New All

R. V. stands for Random Variable.

can be problematic especially when a new function starts immediately after a non-returning function call. For example, a function may start right after a call to exit, because exit will never return. Thus, it is imperative to identify non-returning function calls to obtain precise CFGs.

Hint Collector employs a widely used mechanism for detecting non-returning functions, which simply finds calls to a known non-returning function name [36]. This is possible because even a stripped binary maintains the symbols of imported functions. In our implementation, we use the same list of non-returning function names used by Ghidra [42].

4.1.3 Indirect Branch Resolution. It is crucial to resolve indirect branch targets to improve the coverage of a CFG. For example, when an edge from an indirect branch to a basic block is missing, one may falsely identify the missing block as a function. To recover the targets of indirect jumps, we leverage a pattern-based heuristic used by Ghidra [42]. Specifically, we find the corresponding jump table based on the instruction patterns near the indirect jump instruction, and parse the jump table to recover the indirect jump targets.

4.1.4 Inter-procedural CFG Building. Given a list of disassembled instructions (§4.1.1), a set of non-returning function call sites (§4.1.2), and a set of jump targets for each indirect branch (§4.1.3), Hint Collector finally builds an inter-procedural CFG. Additionally, we parse exception handling information and connect edges from a try block to its corresponding catch block(s).

## 4.2 Function Identification Hints

FUNPROBE gathers 16 kinds of hints listed in Table 1. The first column shows the information source. The second column describes each hint. The third column shows which random variable is used to represent each hint. For example, a random variable  $X_{\alpha}$  for an address  $\alpha$  can represent a hint  $P(f_{\alpha} \mid x_{\alpha})$ . The fourth column indicates how each hint influences function identification: whether it is a positive hint ( $\clubsuit$ ) or a negative hint ( $\frown$ ). And the last column specifies which tool is using the corresponding heuristic. We mark with "All" when it is used by every tool that we studied, and "New" when we are unaware of a tool that employs a similar heuristic.

We note that all these hints rely on architecture-neutral metadata and CFG-structural features. Such a design choice makes FUNPROBE perform well across various binaries obtained from different architectures and compilers as we will show through our experiments.

4.2.1 *Hints from Data.* Hints (1) and (2) are derived from ELF metadata. First, ELF binaries have special sections that contain function pointers, such as .init\_array and .fini\_array. (1) states that those values are always function addresses, i.e.,  $P(f_{\alpha} | p_{\alpha}) = 1.0$ . Second, there are function symbols that remain intact even after stripping, e.g., GOT-based indirect jump targets in MIPS. (2) suggests collecting such function addresses by analyzing symbols. Both (1) and (2) provide definite evidences for a function located at  $\alpha$ , so we assign the probability 1.0, i.e.,  $P(f_{\alpha} | \cdot) = 1.0$ . For other hints, we assign probability values based on the parameters  $\mathcal{P}^+$  and  $\mathcal{P}^-$ .

Hint ③ provides a positive prediction for a code address  $\alpha$  if it is pointed to by a relocation entry:  $P(f_{\alpha} | r_{\alpha}) = \mathcal{P}^+$ . Relocation entries often contain function pointers. For example, when a PIE (Position Independent Executable) has a global function pointer, it should be relocated at runtime by the loader. Hence, a relocation section, e.g., .rela.dyn, of the binary should store a relocation entry for the function pointer.

Hint ④ states that addresses found in .eh\_frame are likely to be a function address. Modern compilers have recently started to provide exception handling information for every C function in order to support C++ interoperability [20, 21]. Particularly, the .eh\_frame section contains a list of Frame Description Entries (FDEs) to support stack unwinding when an exception occurs, and such information allows us to infer function addresses [45].

Hint (5) suggests that a data value is likely to be a function address if it is within a valid code address range. This heuristic is employed by several reassemblers, such as Ramblr [59] and DDisasm [13]. The intuition is that a constant value in a data section is likely to be a function pointer if it is within a valid address range.

4.2.2 *Hints from Code.* Hint (6) says that a call target is likely to be a function entry point. Hint (7) states that a jump target is likely to be a tail-call, i.e., the target is a function entry point, if the jump crosses over a function entry point. Both hints are discussed in §3.2.

Hint (8) handles a special case where a function body solely consists of a single (unconditional) jump instruction (see Line 11– 12 of Figure 1a). Such functions usually act as a trampoline, passing arguments to another function.

Hint (9) helps discover functions from unreachable parts of our CFG. Since our CFG reconstruction (§4.1) is incomplete, it may leave several basic blocks within the range of a function unreachable. Such a code chunk is often referred to as a gap [33], and this hint says that every basic block found in a gap is likely to be a function entry point.

However, not every unreachable basic block is a function entry point, and (9) alone can produce many false positives. Thus, we have to employ several negative hints to filter out false cases. Hint (10) says that gaps surrounded by connected basic blocks in our CFG are unlikely to be a function entry point. This is because a function body is usually a consecutive chunk of code. Thus, such a gap is unlikely to be another function. Compiler-generated padding bytes may form a gap, too. Thus, hint (11) helps disregard padding bytes from being considered a function entry point. Compilers can also use jump instructions to fill a gap. For example, consider a code pattern "ret; jmp LBL; nop; LBL: push ebp". In this example, the ret instruction is the end of a function, and push ebp is the start of the next function. The jmp instruction between these two functions is unreachable, and it merely acts as a dummy padding. (12) helps detect such a pattern to disregard the gap from being considered as a function entry point.

Recall from §2.2, there are cases where a call instruction is used to retrieve the current PC. Hint (3) helps prevent a call target of a PC-getter from being considered as a function entry point.

Hint 14 states that jump targets of a conditional jump are unlikely to be a function entry point because a conditional jump is rarely a tail-call [44].

Hint (15) suggests that basic blocks before the main entry point, i.e., \_start, are unlikely to be a function entry point. This is because GCC often puts rarely executed parts of a function, named with the suffix .cold in the separate text sections. And the sections are often relocated before the main entry point.

Hint (16) says that while a leader of a basic block is a potential function entry point, most leaders are not. This is the nature of a CFG; only the root node is a function entry, and the rest are not.

4.2.3 Generality of Function Identification Hints. We claim that our function identification hints in Table 1 are general enough to analyze various kinds of binaries. Hint (1-(4) use metadata commonly found in ELF binaries. Although our current implementation handles only ELF binaries, other file formats provide similar information, and supporting them should be straightforward. Hint (5) exploits a general characteristic of function pointers found in binaries. Hint (6)-(6), except (1) and (12), are based on structural properties found in CFGs. (1) and (12) are based on semantic properties of instructions, which can be captured by lifted Intermediate Represent (IR) [28].

It is worth noting that modern compilers tend to put regular functions after the entry point, i.e., \_start. We found that GCC often puts a part of a function (such as .part and .cold snippets) before the entry, but the main body of the function is mostly after the entry. We confirmed that this is also the case for Clang, too.

Algorith	m 1: Bogus	Dependenc	y Pruning A	lgorithm.
----------	------------	-----------	-------------	-----------

1	function	BogusDenpendencyPruning( $\mathcal{G}$ )	

```
2 \quad \mathcal{G}', \mathcal{V}, \mathcal{E} \leftarrow \mathsf{RemoveObservedNodes}(\mathcal{G})
```

4  $\mathcal{G}'_t \leftarrow \mathsf{RestoreObservedNodes}(\mathcal{G}_t, \mathcal{V}, \mathcal{E})$ 

# 4.3 Bogus Dependency Pruning

Although FUNPROBE employs only 16 hints, they can produce an extremely large BN with many loops, which makes traditional belief propagation significantly slow, if not impossible. In our dataset, FUNPROBE builds a BN of 46K nodes and 255K edges per binary on average. Therefore, we devise a novel approach to reduce the size of a BN, while not sacrificing much the accuracy.

We note that hint  $\bigcirc$  is the major source of complexity. First, it produces too many bogus dependencies as we discussed in §3.2. Formally, we say a hint  $P(f_{\alpha} | f_{\beta})$  is *bogus* if  $\alpha$  or  $\beta$  is not a function entry point. From our dataset, we found that about 92% of dependencies introduced by  $\bigcirc$  were indeed bogus dependencies. Second,  $\bigcirc$  can create cyclic dependencies, i.e., loops, in the resulting BN. For example, three hints  $F_{\alpha} \rightarrow F_{\beta}$ ,  $F_{\beta} \rightarrow F_{\gamma}$ , and  $F_{\gamma} \rightarrow F_{\alpha}$  form a loop, which prevents us from using the traditional belief propagation.

Therefore, we propose a novel method, named bogus dependency pruning, to find out and exclude such bogus dependencies as well as loops, so that the resulting graph becomes simple and loop-free, allowing us to use the traditional belief propagation. Algorithm 1 describes the overall process of bogus dependency pruning. It takes in as input a BN ( $\mathcal{G}$ ), which is a set of DAGs, and returns a modified BN ( $\mathcal{G}'_t$ ), which is a set of polytrees, as output. In Lines 2 and 4, observed nodes in the given BN are temporarily removed ( $\mathcal{G}'_t$ ) and restored ( $\mathcal{G}'_t$ ) in order to make the polytree calculation efficient (§4.3.1). In Line 3, we reduce each DAG in the modified BN to a polytree (§4.3.2).

4.3.1 Removing and Restoring Nodes. RemoveObservedNodes and RestoreObservedNodes are respectively preprocessing and postprocessing steps of our dependency pruning process (§4.3.2). When RemoveObservedNodes temporarily removes all the observed nodes and their outgoing edges, removed nodes ( $\mathcal{V}$ ) and edges ( $\mathcal{E}$ ) are returned and they are restored in RestoreObservedNodes. Note that our polytree computation (§4.3.2) is not affected by removing observed nodes because each of those nodes has a degree of one. Therefore, the removed nodes and edges will always be included in the resulting polytree anyways. By temporarily removing all the nodes of degree one, we can make our polytree computation fast.

4.3.2 Computing Polytree. Given the modified BN G', a function ComputePolytree transforms each DAG in G' into a polytree. The transformation runs in three steps. First, we convert a DAG into an undirected graph. Second, we run Kruskal's minimum spanning tree algorithm [34] using our custom weight that prefers edges with more positive hints. Formally, we define a weight  $w(\alpha, \beta)$  for an edge  $F_{\alpha} \rightarrow F_{\beta}$  as

$$w(\alpha,\beta) = -\left(\sum_{i=1}^{m} \mathsf{W}(X_{\alpha}^{i}) + \sum_{i=1}^{n} \mathsf{W}(X_{\beta}^{i})\right),$$

 $<sup>\</sup>mathcal{G}_t \leftarrow \mathsf{ComputePolytree}(\mathcal{G}')$ 

<sup>5</sup> **return**  $\mathcal{G}'_t$ 

where  $X_{\alpha}^{i}$  means the *i*-th observed Boolean random variable with regard to the address  $\alpha$ , and *m* and *n* are the numbers of observed Boolean random variables for  $\alpha$  and  $\beta$ , respectively. The function W outputs 1 when the given random variable represents a positive hint ( $\clubsuit$ ) and -1 for a negative hint ( $\frown$ ), respectively.

For example, consider the edge  $F_{58b0} \rightarrow F_{1b770}$  shown in Figure 3. We can compute the weight of the edge as below.

$$w(58b0, 1b770) = - (W(C_{58b0}) + W(E_{58b0}) + W(N_{58b0}))$$
  
= - (1 + 1 - 1)  
= -1.

This way, we can prefer edges that provide more positive influence in detecting functions. Notice that we negatize the sum as Kruskal's algorithm prefers a smaller weight.

## 4.4 Implementation

FUNPROBE is written in 5.5K SLoC of F#. To parse ELF file headers and disassemble instructions for various architectures, we used B2R2 [23], which provides an efficient binary analysis front-end.

# **5 EVALUATION**

In this section, we evaluate FUNPROBE to answer the following research questions.

- **RQ1.** How do the probability parameters affect the effectiveness of function identification? (§5.2)
- **RQ2.** How much performance gain does bogus dependency pruning provide? (§5.3)
- **RQ3.** How well does FUNPROBE perform against the conventional function identification tools? (§5.4)
- **RQ4.** How well does FUNPROBE perform against the learning-based function identification tools? (§5.5)
- **RQ5.** Can FUNPROBE and learning-based approaches be complementary to each other? (§5.6)

## 5.1 Evaluation Setup

*5.1.1 Benchmark.* We build our benchmark by compiling three popular packages, GNU Coreutils (v9.0, 107 programs), GNU Binutils (v2.37, 15 programs), and SPEC CPU2017 (16 programs) written in C and C++ languages. We consider the following compiler configurations to build our benchmark: (1) target architectures, (2) position independence, and (3) code optimization levels. These configurations can largely affect the shape of resulting binaries, thereby impacting the function identification results.

Specifically, we chose 6 popular CPU architectures (x86, x86–64, ARMv7, AArch64, MIPS, and MIPS64), 2 major compilers (GCC v8.4.0 and Clang v13.0.1), both PIE and non-PIE options, and six different compiler optimization levels (O0, O1, O2, O3, Os, and Ofast). This gives us a total of 144 (=  $6 \times 2 \times 2 \times 6$ ) different configurations. As a result, our benchmark consists of **19,872** binaries. To the best of our knowledge, this is the *largest benchmark* used to evaluate function identification algorithms [32].

5.1.2 *Ground Truth.* To evaluate function identification tools, we need to obtain the ground truth of our benchmark, i.e., a set of function entry points for each binary. Specifically, our ground truth data deal with functions located in the .text section. We mainly leveraged debugging information to gather ground truth data. Also,

Table 2: F1-scores achieved with different parameter values.

						$\mathcal{P}^+$				
		0.55	0.6	0.65	0.7	0.75	0.8	0.85	0.9	0.95
	0.05	23.99	23.98	23.97	23.97	44.23	45.02	78.00	78.43	78.56
	0.1	23.99	23.99	24.00	44.50	48.01	78.23	78.47	96.10	96.11
	0.15	23.99	24.00	44.06	45.08	78.25	78.50	78.58	96.10	96.19
	0.2	24.04	24.02	44.89	78.22	78.50	96.13	96.12	96.17	82.19
<i>р</i> -	0.25	24.05	44.46	78.04	78.45	78.62	96.14	96.19	82.22	82.21
	0.3	24.06	45.06	78.44	78.63	96.17	96.20	82.23	82.21	79.45
	0.35	24.08	78.31	78.63	96.18	96.22	82.23	82.22	79.47	79.19
	0.4	45.10	78.65	96.23	82.23	82.27	79.51	79.37	79.06	78.80
	0.45	96.17	82.27	79.54	79.12	78.91	78.73	78.60	78.40	78.18

we manually filtered out several function symbols with .cold or .part suffixes from GCC-compiled binaries [27]. Additionally, we manually added the addresses of compiler intrinsic functions to our ground truth data as they do not have debugging information.

5.1.3 Comparison Target. We selected five state-of-the-art tools for comparison. Four of them use conventional binary analyses: IDA Pro (v7.7.220118), Binary Ninja (v3.0.3233), Ghidra (v10.1.5) [42], and Nucleus (commit e3ab49d) [4]. One of them uses a deep-learning model to identify functions: XDA (commit 068007c) [48]. We fine-tuned XDA on a subset of our benchmark following the recommendation of the authors [47]. Specifically, we randomly selected 20% (1296 binaries) of x86 and x86-64 binaries in our benchmark to create about 640k training byte sequences. XDA is a representative ML-based solution that is publicly available. To our knowledge, it had demonstrated the best function identification accuracy so far in the literature. While there is another noteworthy tool named DeepDi [64], it is closed-source, hence, we were not able to train a model using our benchmark.

5.1.4 Our Environment. We used a server machine with 88 Intel Xeon E5 cores, 512 GB of RAM, and 8 TITAN Xp GPUs to run our experiments except for fine-tuning XDA (as discussed in §5.5). We used Docker 20.10.3 for running comparison targets. We used Ubuntu 20.04 containers for FUNPROBE, Binary Ninja, and Ghidra, and a Ubuntu 16.04 container for Nucleus as the authors suggested. Lastly, we used a Windows 10 VM, but not a Docker container, for IDA Pro due to the license issue.

#### 5.2 Probability Parameter Selection

Recall from §3.3, FUNPROBE provides two user-configurable parameters  $\mathcal{P}^+$  and  $\mathcal{P}^-$  to assign probabilities for positive and negative hints, respectively. How do these parameters affect the accuracy of FUNPROBE? Which values should we use to maximize the accuracy? To answer these questions, we ran FUNPROBE on a subset of our benchmark with varying parameter values, and measured the accuracy (F1-score) for each setting.

We picked 10 random binaries from GNU Coreutils for each build configuration (out of 144 build configurations with varying compilers, architectures, and compiler options). This gives us a total ESEC/FSE '23, December 3-9, 2023, San Francisco, CA, USA



(b) Average F1-score per subset of benchmarks.

Figure 4: Comparison between FunProbe and FunProbe-lbp in terms of both execution time and accuracy.

of 1,440 binaries corresponding to 7% of the entire benchmark. We then ran FUNPROBE on these binaries with 81 different combinations of  $\mathcal{P}^+$  and  $\mathcal{P}^-$  values. Note that we chose to use only a small number of binaries to show that our parameter selection was not biased to any specific configuration.

Table 2 presents the measured F1-scores for each parameter combination. The X-axis shows  $\mathcal{P}^+$  values, and the Y-axis shows the  $\mathcal{P}^-$  values. Each cell is filled with a grey-scale color, where the darker color indicates a higher F1-score.

The results show that  $\mathcal{P}^+$  and  $\mathcal{P}^-$  are negatively correlated. The bigger  $\mathcal{P}^+$  value is chosen, the smaller  $\mathcal{P}^-$  value needs to be selected to make FUNPROBE perform well. Therefore, the diagonal part of the table shows higher accuracy results. In addition, F1-scores shown in the table indicate that the performance of FUNPROBE is not too sensitive to the choice of parameters.

Since the highest value was achieved by ( $\mathcal{P}^+ = 0.65$ ,  $\mathcal{P}^- = 0.4$ ), we chose them as our default parameter values. Our experimental results on the entire benchmark (see §5.4) also show that our choice of default parameters is adequate. We leave it as future work to find a better combination of parameters with more fine-grained values.

## 5.3 Impact of Bogus Dependency Pruning

Recall from §4.3, bogus dependency pruning simplifies the given BN to be *loop-free*, which enables us to use traditional belief propagation. To measure the impact of bogus dependency pruning, we modified FUNPROBE to run *loopy* belief propagation without bogus dependency pruning, which is dubbed "FUNPROBE-lbp", and compared its performance against it of the original FUNPROBE.

Although loopy belief propagation is designed to handle BNs with loops, the solution may not converge in some cases. Therefore, it typically runs the algorithm only for a fixed number of iterations. In our implementation of FUNPROBE-lbp, we use the following convergence criteria for loopy belief propagation for maximum 10 iterations: For all basic block address  $\alpha$  in the given binary, if  $|P_{i-1}(f_{\alpha}) - P_i(f_{\alpha})| < 0.0001$  holds, then we consider the algorithm

is converged, where  $P_i(f_\alpha)$  is a marginal probability obtained after *i* iterations of the algorithm. The algorithm stops when *i* > 10.

We ran FUNPROBE and FUNPROBE-lbp on the same benchmark shown in §5.1.1 using Docker containers assigned with one CPU core. Each binary in the benchmark was analyzed for maximum 12 hours. Figure 4 reports both their running time as well as F1-score for 12 different architecture-compiler combinations. Each bar in Figure 4a represents the average running time in seconds, and each bar in Figure 4b represents the average F1-score in percentage.

Overall, FUNPROBE was 12.6× faster than FUNPROBE-lbp. The graph shows the computation cost of FUNPROBE is almost negligible. Moreover, FUNPROBE-lbp failed to run 32 binaries within 12 hours of timeout. On the contrary, FUNPROBE was able to analyze all the binaries in our benchmark without any timeout. The average F1-scores were nearly the same for both. The total average difference was only 0.07%: FUNPROBE and FUNPROBE-lbp recorded 98.99% and 99.06% of F1-score, respectively. These results confirm the impact of bogus dependency pruning: it makes FUNPROBE scalable while not sacrificing much of the accuracy.

## 5.4 Comparison with Conventional Tools

To see how well FUNPROBE performs compared to conventional approaches, we selected four state-of-the-art tools: Binary Ninja, IDA Pro, Ghidra, and Nucleus. We ran each tool on our entire benchmark (§5.1.1). We set a timeout of one hour for each binary. When reporting the final accuracy results, we excluded those binaries that did not meet the timeout requirement.

Table 3 shows the precision (P), recall (R), F1-score (F1), execution time (ET), and the number of binaries failed to analyze due to the timeout (TO). Each row summarizes the results for each architecture and compiler combination. A dash (-) mark indicates that the tool does not support the architecture.

Overall, FUNPROBE achieves the *best performance* for all the criteria. On average, FUNPROBE showed 99.37% precision, 98.70% recall, and 99.03% F1-score. Notably, FUNPROBE outperformed all our comparison targets for every architecture and compiler combination in terms of F1-score.

It is also noteworthy that FUNPROBE shows consistent performance on every architecture and compiler combination we tested. This result aligns well with our design principle that FUNPROBE should be *architecture-* and *compiler-agnostic*. For example, on x86, every other tool except FUNPROBE shows better performance on GCC-compiled binaries than on Clang-compiled binaries. We can also note how existing tools are fine-tuned to handle GCC-compiled binaries, which are more common in practice.

Notably, Ghidra showed the highest precision on MIPS binaries. By investigating those function addresses falsely identified by FUN-PROBE, we found that most of them are due to the imprecise CFG recovery of FUNPROBE. Since MIPS binaries extensively use the GOT base address to compute relative addresses, which can be temporarily stored on the stack or registers, precisely recovering CFG requires a sound data-flow analysis. Since FUNPROBE relies only on a lightweight analysis, it could miss out on many true edges.

In terms of execution time (ET), FUNPROBE spent only 6.30 seconds on average for analyzing a binary in our benchmark. Moreover, FUNPROBE was able to analyze all the binaries without hitting

Table 3: Function identification performance in terms of Precision (P), Recall (R), F1-score (F1), Execution Time (ET), and # of binaries failed due to TimeOut (TO).

Arch. Compiler		FunProbe			Binary Ninja				IDA Pro					Ghidra					Nucleus							
		P (%)	R (%)	F1 (%)	ET (s)	то	P (%)	R (%)	F1 (%)	ET (s)	ТО	P (%)	R (%)	F1 (%)	ET (s)	ТО	P (%)	R (%)	F1 (%)	ET (s)	ТО	P (%)	R (%)	F1 (%)	ET (s)	то
x86	GCC Clang	99.93 99.95	99.82 98.69	99.88 99.32	6.24 5.42	0 0	89.23 86.90	97.44 96.09	93.16 91.26	92.42 96.68	0 1	90.67 91.83	90.09 77.45	90.38 84.03	13.04 11.62	0 0	97.66 98.84	98.89 67.79	98.27 80.42	153.97 129.08	1 0	95.78 67.36	93.34 96.95	94.54 79.49	1.60 1.69	0 0
x86-64	GCC Clang	99.84 99.97	99.69 99.75	99.77 99.86	5.55 6.00	0 0	92.76 89.48	96.79 94.93	94.73 92.12	92.90 82.64	0 0	93.12 93.05	85.93 83.93	89.38 88.26	10.33 10.17	0 0	97.72 99.82	98.33 98.63	98.02 99.22	151.54 143.19	0 9	95.52 92.07	91.71 94.52	93.58 93.28	1.60 1.63	0 0
ARMv7	GCC Clang	<b>99.65</b> 97.15	96.76 <b>96.69</b>	98.18 96.92	11.41 11.67	0 0	93.69 91.82	<b>98.34</b> 95.85	95.96 93.80	56.34 62.95	0 0	92.34 89.56	96.97 92.99	94.60 91.24	30.87 22.01	0 0	99.41 <b>99.76</b>	61.44 66.99	75.94 80.15	138.54 94.27	2 5	-	-	-	-	-
AArch64	GCC Clang	99.91 99.87	99.67 99.77	99.79 99.82	5.30 5.09	0 0	92.00 88.33	93.50 95.21	92.74 91.64	79.87 108.33	0 3	92.78 92.63	98.44 98.39	95.52 95.42	24.29 20.69	0 0	89.26 99.78	98.15 98.36	93.50 99.07	95.50 129.27	0 7	88.92 91.46	76.56 80.55	82.28 85.66	1.26 1.20	0 0
MIPS	GCC Clang	97.39 99.68	98.24 98.53	97.81 99.10	5.23 4.16	0 0	84.28 84.37	96.42 95.82	89.94 89.73	112.58 103.79	10 0	86.70 89.04	81.79 76.18	84.18 82.11	38.93 48.89	0 1	99.31 99.75	79.67 75.09	88.41 85.68	176.35 159.06	6 4	-	-	-	-	-
MIPS64	GCC Clang	99.63 99.48	98.06 98.74	98.84 99.11	4.42 5.09	0 0	-	-	-	-	-	86.22 86.09	83.94 76.80	85.06 81.18	34.84 38.64	8 0	99.81 99.83	64.69 63.80	78.50 77.84	181.57 165.53	4 3	-	-	-	-	-
	Total	99.37	98.70	99.03	6.30	0	89.26	96.04	92.53	88.83	14	90.04	86.94	88.64	25.19	9	97.85	81.10	88.69	143.16	58	86.81	89.19	87.98	1.50	0

The numbers in bold represent the best result per row.

Table 4: Performance of FunProbe and XDA in terms of Precision (P), Recall (R), F1-score (F1), and Execution Time (ET).

Arch	Compiler		Fun	PROBE		XDA						
		P (%)	R (%)	F1 (%)	ET (s)	P (%)	R (%)	F1 (%)	ET (s)			
x86	GCC Clang	99.93 99.95	<b>99.82</b> 98.69	<b>99.88</b> 99.32	6.24 5.42	99.77 99.84	99.62 <b>99.89</b>	99.70 <b>99.8</b> 7	35.29 37.02			
x86-64	GCC Clang	99.84 99.97	99.69 99.75	99.77 99.86	5.55 6.00	99.48 99.82	99.35 99.73	99.41 99.78	35.83 36.58			
	Total	99.92	99.49	99.71	5.80	99.73	99.65	99.69	36.18			

The numbers in bold represent the best result per row.

the timeout. Although Nucleus records the best performance (1.5 seconds on average), the difference is less than 5 seconds, while FUNPROBE significantly outperforms Nucleus in terms of function identification accuracy.

While Binary Ninja, IDA Pro, and Ghidra show significantly slow running time overall, it is important to note that those tools perform not only function identification, but also other complex analyses, making the comparison not entirely fair. However, we argue that our technique can practically be used to improve binary analysis results as a preprocessor of other binary analyzers because FUNPROBE can correctly find more function entry points in a reasonable amount of time.

#### 5.5 Comparison with Learning-based Tools

How does FUNPROBE compare to an existing learning-based tool? We now compare the performance of FUNPROBE against XDA to answer this question. The comparison is made by running the tools on x86 and x86-64 architecture binaries as XDA only supports these architectures. In addition, we excluded the binaries that we used to fine-tune XDA for a fair comparison. In total, we used 5,412 binaries for the comparison.

We downloaded the pre-trained model of XDA from the official repository, and then used it to fine-tune our own model using our training dataset (see §5.1.3). We used a separate machine with a powerful GPU (GeForce 3090 Ti), but the fine-tuning process took more than 2.5 days. On the other hand, FUNPROBE does not require

any training process, and this makes FUNPROBE more practical to use on any benchmarks. For the inference process, we used our server machine described in §5.1.4. Specifically, we assigned a single CPU of our server machine to run both FUNPROBE and XDA, but we had to additionally assign a single GPU to run XDA as it also uses GPU in the inference phase.

Table 4 summarizes the comparison between FUNPROBE and XDA on x86 and x86-64 binaries. Overall, both tools show comparable performance, but FUNPROBE had a slightly better F1-score than XDA. It records 99.7% F1-score on average whereas XDA shows slightly a low (99.69%), F1-score.

In terms of execution time, FUNPROBE was considerably faster than XDA even though XDA utilizes additional GPU resources. To analyze a single binary, FUNPROBE spent 5.8 seconds on average whereas XDA spent 36.18 seconds on average. That is, FUNPROBE was overall 6.2× faster than XDA without regard to the training cost. Therefore, we conclude that FUNPROBE is more scalable than XDA while being as precise as XDA.

#### 5.6 Combining FUNPROBE and XDA

Can FUNPROBE benefit deep-learning-based approach or vice versa? To answer this question, we conduct an additional experiment by simply using FUNPROBE as a post-processor of XDA. Specifically, we first ran XDA on the binaries used in §5.5. We then created a positive hint (with  $\mathcal{P}^+$ ) for each function entry point identified by XDA, and fed in those hints to FUNPROBE. This means FUNPROBE will consider those XDA-generated hints as positive evidence to find function entry points within the same probabilistic framework we used.

As a result, the combined tool achieved 99.85%, 99.90%, and 99.88% in precision, recall, and F1-score, respectively. Given that both FUNPROBE and XDA already accomplished high accuracy, it is surprising that the combined tool achieves phenomenal performance. XDA supplies function entry points that are not covered by FUNPROBE. At the same time, FUNPROBE provides precision based on our probabilistic hints. The synergy increases F1-score by 0.17%. Additionally, running FUNPROBE as a post-processor of XDA is a matter of a few seconds for each binary. This result confirms that FUNPROBE and the state-of-the-art deep-learning approaches are indeed complementary to each other.

We further analyzed what kind of functions that the combined tool was able to identify while the original XDA failed. The most common error cases corrected by the combined tool were where a function contains only a few instructions. For example, a tiny function containing two instructions (mov followed by ret) was not identified by XDA. Although we cannot directly understand the reason for failure due to the lack of interpretability of the model, we conjecture that this is because such a function is considered to be a function epilogue than a prologue by XDA.

## 6 DISCUSSION AND FUTURE WORK

Supporting other file formats. Currently, FUNPROBE only supports ELF binaries. However, extending our support for more binary types, such as Mach-O and PE, is straightforward because all the hints we leveraged (summarized in Table 1) are derived from common information that can be found in any binary format. For example, we can apply ④ for PE binaries by analyzing RUNTIME\_FUNCTION [38] entries instead of FDE entries, because exception handling information should be stored in any binary format.

Parameter tuning. Although FUNPROBE currently uses empirically chosen probability parameters (0.65 for positive hints and 0.4 for negative hints), one may use different probabilities for each different hint. This is because each hint may have a different impact on function identification. We believe fine-tuning probability parameters for each different hint will further improve the performance of FUNPROBE, and it is a promising direction for future research.

Better pruning strategy. While our bogus dependency pruning significantly improves the performance of FUNPROBE, it can potentially remove critical (i.e., non-bogus) dependencies in our model. Although it is beyond the scope of this paper to find a better pruning strategy, one may consider using a more sophisticated pruning strategy that can identify true bogus dependencies.

*Obfuscation*. FUNPROBE does not consider obfuscated binaries. We leave it as future work to combine FUNPROBE with existing deobfuscation techniques such as [61].

Integration with other binary analysis tools. FUNPROBE can enhance other binary analysis tools by serving as a preprocessor. For instance, modern binary analysis frameworks, such as IDA Pro [19] or Ghidra [41], provide APIs for creating a function at a specific address. With such APIs, one may feed in the results from FUNPROBE to those frameworks to accurately construct CFGs for each function because more accurate function identification will enable more precise CFG reconstruction. Furthermore, having precise CFGs is crucial for complex binary analysis tasks such as reassembly [26] and type recovery [35].

#### 7 RELATED WORK

#### 7.1 Binary Function Identification

There has been significant research effort on binary-level function identification for more than two decades. Early research focuses on systematically recovering CFGs [9], thereby naturally identifying functions by following call edges from recovered CFGs. For example, Jakstab [29] performs data flow analyses to determine indirect call edges. There are several recent papers following this direction [2, 4, 11, 49, 63]. For example. Nucleus [4] constructs interprocedural CFGs to find call targets as well as unreachable targets. rev.ng [11] converts a binary to an LLVM IR, and performs a static analysis on top of IR code to recover CFGs. All these approaches suffer from the accuracy of the analyses, and thus, employ various heuristics to improve the effectiveness of their analyses. However, such heuristics are inherently architecture and compiler dependent.

Pattern-based approaches are widely adopted in mainstream binary analysis tools [1, 8, 16, 36, 42, 57, 58]. For example, BAP [8] and Dyninst [36] utilize pre-trained decision trees to identify functions. IDA Pro [17, 18] and Ghidra [43] provide their own pattern databases to match well-known function patterns. Recently Fun-Seeker [27] shows that one can precisely detect functions from Intel CET binaries by leveraging the usage patterns of endbr instructions. However, all these techniques rely on previously-known patterns, and thus, inherently suffer from handling binaries with unknown patterns. On the other hand, the hints that we use are *not* specific to an architecture nor a compiler.

There is a recent metadata-based function identification technique, named FETCH [45], which leverages exception handling information to identify functions. Their technique significantly outperforms existing techniques, except for binaries without exception information, such as C binaries generated from Clang.

There also have been various deep-learning-based approaches [48, 54, 60, 64]. Shin et al. [54] use a bi-directional and multi-layer RNN to predict function boundaries. FID [60] leverages symbolic execution to extract feature vectors from binaries, to make its model robust against highly optimized binaries. XDA [48] first pre-trains a transformer model using Masked Language Modeling and selfattention layers, and fine-tunes the model for specific disassembly tasks, such as function identification. DeepDi [64] uses the R-GCN model to learn instruction embeddings on Instruction Flow Graph, and identifies function entry points with trained instruction patterns near the function. Although those approaches usually show good accuracy, they suffer from the generalization problem, that is, their performance relies on their training dataset [4, 32]. Moreover, they consume a lot of computational resources during both training and inference stages. They often require the use of GPUs for efficiency, and even more resources when the model size is large. On the other hand, FUNPROBE does not have any dependency on training data, or need for computation resources.

#### 7.2 Probabilistic Binary Analysis

There are also diverse probabilistic-model-based approaches on binary code analysis [6, 39, 50, 66]. Rosenblum *et al.* [50] find function entry point idioms, which are instruction patterns of function prologues. While they use a probabilistic model to learn a probability distribution of such patterns, FUNPROBE directly analyzes various hints in order to compute the probability of each instruction being a function entry point. ByteWeight [6] creates a weighted prefix tree based on the patterns of function start instructions (or bytes). Miller *et al.* [39] suggest a probabilistic disassembly algorithm. They collect probabilistic hints for valid instructions to disassemble truepositive instructions. OSPREY [66] utilizes a probabilistic graph FUNPROBE: Probing Functions from Binary Code through Probabilistic Analysis

model to recover variables and their type information from stripped binaries.

## 7.3 Probabilistic Inference

Various approaches to derive the desired distribution from a given probabilistic graphical model are suggested. Belief propagation [46] can solve tree-like graphs to get an exact solution. However, it may fail to converge or converge to the wrong solution if the graph contains loops. To handle loopy graphs, loopy belief propagation [40] iteratively runs belief propagation algorithm until it converges or a certain number of iterations is reached. To approximate the solution from loopy graphs, Junction Tree algorithm [24] transforms the graphs into junction trees, and solves the graph. However, it does not scale to handle large graphs. Variational Bayesian inference [14] finds an alternative distribution to approximate the complicated distribution. It is known to converge fast, but it requires heavy computation. Monte Carlo method [37] is another approach to approximate the exact distribution based on random sampling. Despite its simplicity, it is hard to apply to our system because it does not work well if the space is high-dimensional.

## 8 CONCLUSION

In this paper, we introduced a novel way to identify function entry points from a stripped binary. The key idea was to regard a function identification heuristic as a producer of a probabilistic hint. We can then combine those hints in a graphical model, i.e., a Bayesian Network, and compute the probability of each address in the binary being a function entry point. To boost the speed of our probabilistic inference, we proposed a novel approach, named bogus dependency pruning, and empirically proved its effectiveness. We implemented FUNPROBE to realize these ideas, and showed its practical impact by comparing its performance against five state-of-the-art tools in terms of both speed and accuracy. Our experiments were performed based on the benchmark consisting of 19,872 binaries compiled for 6 major CPU architectures, which is, to our knowledge, the largest benchmark used in the field. As a result, FUNPROBE outperformed all of the state-of-the-art tools we tested.

## DATA AVAILABILITY

Our tool is available at https://github.com/B2R2-org/FunProbe.

# ACKNOWLEDGEMENTS

We thank anonymous reviewers for their invaluable feedback. This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2021-0-01332, Developing Next-Generation Binary Decompiler).

## REFERENCES

- [1] [n.d.]. Radare2. https://github.com/radare/radare2.
- [2] Jim Alves-Foss and Jia Song. 2019. Function Boundary Detection in Stripped Binaries. In Proceedings of the Annual Computer Security Applications Conference. 84-96.
- [3] Dennis Andriesse, 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 Proceedings of the USENIX Security Symposium 583-600
- In Proceedings of the USENIX Security Symposium. 583–600.
  [4] Dennis Andriesse, Asia Slowinska, and Herbert Bos. 2017. Compiler-Agnostic Function Detection in Binaries. In Proceedings of IEEE European Symposium on Security and Privacy. 177–189.

- [5] Gogul Balakrishnan, Radu Gruian, Thomas Reps, and Tim Teitelbaum. 2005. CodeSurfer/x86—A Platform for Analyzing x86 Executables. In Proceedings of the International Conference on Compiler Construction. 250–254.
- [6] Tiffany Bao, Jonathan Burket, Maverick Woo, Rafael Turner, and David Brumley. 2014. ByteWeight: Learning to recognize functions in binary code. In Proceedings of the USENIX Security Symposium. 845–860.
- [7] Erick Bauman, Zhiqiang Lin, and Kevin Hamlen. 2018. Superset Disassembly: Statically Rewriting x86 Binaries Without Heuristics. In Proceedings of the Network and Distributed System Security Symposium.
- [8] David Brumley, Ivan Jager, Thanassis Avgerinos, and Edward J. Schwartz. 2011. BAP: A Binary Analysis Platform. In Proceedings of the International Conference on Computer Aided Verification. 463–469.
- [9] Cristina Cifuentes and Mike Van Emmerik. 2000. UQBT: Adaptable Binary Translation at Low Cost. Computer 33, 3 (2000), 60-66.
- [10] Gregory F Cooper. 1990. The computational complexity of probabilistic inference using Bayesian belief networks. Artificial intelligence 42, 2-3 (1990), 393–405.
- [11] Alessandro Di Federico, Mathias Payer, and Giovanni Agosta. 2017. Rev.Ng: A Unified Binary Analysis Framework to Recover CFGs and Function Boundaries. In Proceedings of the International Conference on Compiler Construction. 131-141.
- [12] Frank DiMaio and Jude Shavlik. 2006. Improving the efficiency of belief propagation in large, highly connected graphs. Working Paper 06-1, UWML Research Group (2006).
- [13] Antonio Flores-Montoya and Eric Schulte. 2020. Datalog Disassembly. In Proceedings of the USENIX Security Symposium. 1075–1092.
- [14] Charles W Fox and Stephen J Roberts. 2012. A tutorial on variational Bayesian inference. Artificial intelligence review 38, 2 (2012), 85–95.
- [15] Andrea Gussoni, Alessandro Di Federico, Pietro Fezzardi, and Giovanni Agosta. 2020. A Comb for Decompiled C Code. In Proceedings of the ACM Symposium on Information, Computer and Communications Security. 637–651.
- [16] Hex-Rays SA. [n.d.]. IDA Pro. https://www.hex-rays.com/products/ida/.
- [17] Hex-Rays SA. [n. d.]. IDA Pro FLIRT. https://hex-rays.com/products/ida/tech/flirt/.
- [18] Hex-Rays SA. [n. d.]. IDA Pro Lumina Server. https://hex-rays.com/products/id a/lumina/.
- [19] Hex-Rays SA. [n. d.]. Module ida\_funcs. https://www.hex-rays.com/products/id a/support/idapython\_docs/ida\_funcs.html#ida\_funcs.add\_func.
- [20] H.J. Lu. [n. d.]. gcc/ChangeLog-2010. https://github.com/gcc-mirror/gcc/blob/ master/gcc/ChangeLog-2010.
- [21] H.J. Lu. [n. d.]. Turn on -fomit-frame-pointer by default for 32bit Linux/x86. https://gcc.gnu.org/legacy-ml/gcc-patches/2010-08/msg00922.html.
- [22] Saehan Jo, Jaemin Yoo, and U Kang. 2018. Fast and scalable distributed loopy belief propagation on real-world graphs. In Proceedings of the ACM International Conference on Web Search and Data Mining. 297–305.
- [23] Minkyu Jung, Soomin Kim, HyungSeok Han, Jaeseung Choi, and Sang Kil Cha. 2019. B2R2: Building an Efficient Front-End for Binary Analysis. In Proceedings of the NDSS Workshop on Binary Analysis Research.
- [24] David Kahle, Terrance Savitsky, Stephen Schnelle, and Volkan Cevher. 2008. Junction tree algorithm. *Stat* 631 (2008).
- [25] Dongkwan Kim, Eunsoo Kim, Sang Kil Cha, Sooel Son, and Yongdae Kim. 2022. Revisiting Binary Code Similarity Analysis using Interpretable Feature Engineering and Lessons Learned. *IEEE Transactions on Software Engineering* (2022), 1–23. https://doi.org/10.1109/TSE.2022.3187689
- [26] Hyungseok Kim, Soomin Kim, Junoh Lee, Kangkook Jee, and Sang Kil Cha. 2023. Reassembly is Hard: A Reflection on Challenges and Strategies. In Proceedings of the USENIX Security Symposium.
- [27] Hyungseok Kim, Junoh Lee, Soomin Kim, Seungll Jung, and Sang Kil Cha. 2022. How'd Security Benefit Reverse Engineers? The Implication of Intel CET on Function Identification. In Proceedings of the International Conference on Dependable Systems and Networks. 559–566.
- [28] Soomin Kim, Markus Faerevaag, Minkyu Jung, Seungil Jung, DongYeop Oh, JongHyup Lee, and Sang Kil Cha. 2017. Testing Intermediate Representations for Binary Analysis. In Proceedings of the International Conference on Automated Software Engineering. 353–364.
- [29] Johannes Kinder and Helmut Veith. 2008. Jakstab: A Static Analysis Platform for Binaries. In Proceedings of the International Conference on Computer Aided Verification. 423–427.
- [30] Alec Kirkley, George T Cantwell, and MEJ Newman. 2021. Belief propagation for networks with loops. Science Advances 7, 17 (2021), eabf1211.
- [31] Mieczysław A Kłopotek. 2006. Cyclic Bayesian network: Markov process approach. Studia Informatica: systems and information technology 1, 7) (2006), 47-55.
- [32] Hyungjoon Koo, Soyeon Park, and Taesoo Kim. 2021. A Look Back on a Function Identification Problem. In Proceedings of the Annual Computer Security Applications Conference. 158–168.
- [33] Christopher Kruegel, William Robertson, Fredrik Valeur, and Giovanni Vigna. 2004. Static Disassembly of Obfuscated Binaries. In Proceedings of the USENIX Security Symposium. 340–353.

ESEC/FSE '23, December 3-9, 2023, San Francisco, CA, USA

- [34] Joseph B Kruskal. 1956. On the shortest spanning subtree of a graph and the traveling salesman problem. Proceedings of the American Mathematical society 7, 1 (1956), 48–50.
- [35] JongHyup Lee, Thanassis Avgerinos, and David Brumley. 2011. TIE: Principled Reverse Engineering of Types in Binary Programs. In Proceedings of the Network and Distributed System Security Symposium. 251–268.
- [36] Xiaozhu Meng and Barton P. Miller. 2016. Binary Code is Not Easy. In Proceedings of the International Symposium on Software Testing and Analysis. 24–35.
- [37] Nicholas Metropolis and Stanislaw Ulam. 1949. The monte carlo method. Journal of the American statistical association 44, 247 (1949), 335–341.
- [38] Microsoft. [n. d.]. x64 exception handling. https://learn.microsoft.com/enus/cpp/build/exception-handling-x64?view=msvc-170.
- [39] Kenneth Miller, Yonghwi Kwon, Yi Sun, Zhuo Zhang, Xiangyu Zhang, and Zhiqiang Lin. 2019. Probabilistic Disassembly. In Proceedings of the International Conference on Software Engineering. 1187–1198.
- [40] Kevin Murphy, Yair Weiss, and Michael I. Jordan. 1999. Loopy Belief Propagation for Approximate Inference: An Empirical Study. In Proceedings of the Converence on Uncertainty in Artificial Inteligence. 467–476.
- [41] National Security Agency. [n. d.]. Class CreateFunctionCmd. https://ghidra.re/ ghidra\_docs/api/ghidra/app/cmd/function/CreateFunctionCmd.html.
- [42] National Security Agency. [n. d.]. Ghidra. https://ghidra-sre.org.
- [43] National Security Agency. [n. d.]. Ghidra x86 patterns. https://github.com/Natio nalSecurityAgency/ghidra/tree/master/Ghidra/Processors/x86/data/patterns.
- [44] 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 Proceedings of the IEEE Symposium on Security and Privacy. 833–851.
- [45] Chengbin Pang, Ruotong Yu, Dongpeng Xu, Eric Koskinen, Georgios Portokalidis, and Jun Xu. 2021. Towards Optimal Use of Exception Handling Information for Function Detection. In Proceedings of the International Conference on Dependable Systems and Networks. 338–349.
- [46] Judea Pearl. 1988. Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference. Morgan Kaufmann Publishers Inc.
- [47] Kexin Pei. [n. d.]. xda/issue-4. https://github.com/CUMLSec/XDA/issues/4.
- [48] Kexin Pei, Jonas Guan, David Williams King, Junfeng Yang, and Suman Jana. 2021. XDA: Accurate, Robust Disassembly with Transfer Learning. In Proceedings of the Network and Distributed System Security Symposium.
- [49] Rui Qiao and R. Sekar. 2017. Function interface analysis: A principled approach for function recognition in COTS binaries. In Proceedings of the International Conference on Dependable Systems and Networks. 201–212.
- [50] Nathan Rosenblum, Xiaojin Zhu, Barton Miller, and Karen Hunt. 2008. Learning to Analyze Binary Computer Code. In Proceedings of the AAAI Conference on Artificial Intelligence. 798–804.
- [51] Victor Garcia Satorras and Max Welling. 2021. Neural enhanced belief propagation on factor graphs. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 685–693.
- [52] Edward J. Schwartz, Cory F. Cohen, Michael Duggan, Jeffrey Gennari, Jeffrey S. Havrilla, and Charles Hines. 2018. Using Logic Programming to Recover C++

Classes and Methods from Compiled Executables. In Proceedings of the ACM Conference on Computer and Communications Security. 426–441.

- [53] Edward J. Schwartz, JongHyup Lee, Maverick Woo, and David Brumley. 2013. Native x86 Decompilation Using Semantics-preserving Structural Analysis and Iterative Control-flow Structuring. In *Proceedings of the USENIX Security Sympo*sium. 353–368.
- [54] Eui Chul Richard Shin, Dawn Song, and Reza Moazzezi. 2015. Recognizing Functions in Binaries with Neural Networks. In Proceedings of the USENIX Security Symposium. 611–624.
- [55] Yan Shoshitaishvili, Ruoyu Wang, Christopher Salls, Nick Stephens, Mario Polino, Andrew Dutcher, John Grosen, Siji Feng, Christophe Hauser, Christopher Kruegel, and Giovanni Vigna. 2016. (State of) The Art of War: Offensive Techniques in Binary Analysis. In *Proceedings of the IEEE Symposium on Security and Privacy*. 138–157.
- [56] Alexander L Tulupyev and Sergey I Nikolenko. 2005. Directed cycles in Bayesian belief networks: probabilistic semantics and consistency checking complexity. In Mexican International Conference on Artificial Intelligence. Springer, 214–223.
- 57] UCSB SecLab. [n. d.]. Angr. https://github.com/angr/angr.
- [58] Vector 35. [n. d.]. Binary Ninja. https://binary.ninja/.
- [59] Ruoyu Wang, Yan Shoshitaishvili, Antonio Bianchi, Aravind Machiry, John Grosen, Paul Grosen, Christopher Kruegel, and Giovanni Vigna. 2017. Ramblr: Making Reassembly Great Again. In Proceedings of the Network and Distributed System Security Symposium.
- [60] Shuai Wang, Pei Wang, and Dinghao Wu. 2017. Semantics-Aware Machine Learning for Function Recognition in Binary Code. In Proceedings of IEEE International Conference on Software Maintenance and Evolution. 388–398.
- [61] Babak Yadegari, Brian Johannesmeyer, Ben Whitely, and Saumya Debray. 2015. A Generic Approach to Automatic Deobfuscation of Executable Code. In Proceedings of the IEEE Symposium on Security and Privacy. 674–691.
   [62] Khaled Yakdan, Sebastian Eschweiler, Elmar Gerhards-Padilla, and Matthew
- [62] Khaled Yakdan, Sebastian Eschweiler, Elmar Gerhards-Padilla, and Matthew Smith. 2015. No More Gotos: Decompilation Using Pattern-Independent Control-Flow Structuring and Semantics-Preserving Transformations. In Proceedings of the Network and Distributed System Security Symposium.
- [63] Xiaokang Yin, Shengli Liu, Long Liu, and Da Xiao. 2018. Function Recognition in Stripped Binary of Embedded Devices. *IEEE Access* 6 (2018), 75682–75694.
- [64] Sheng Yu, Yu Qu, Xunchao Hu, and Heng Yin. 2022. DeepDi: Learning a Relational Graph Convolutional Network Model on Instructions for Fast and Accurate Disassembly. In Proceedings of the USENIX Security Symposium.
- [65] Mingwei Zhang and R. Sekar. 2013. Control Flow Integrity for COTS Binaries. In Proceedings of the USENIX Security Symposium. 337–352.
- [66] Zhuo Zhang, Yapeng Ye, Wei You, Guanhong Tao, Wen chuan Lee, Yonghwi Kwon, Yousra Aafer, and Xiangyu Zhang. 2021. OSPREY: Recovery of Variable and Data Structure via Probabilistic Analysis for Stripped Binary. In Proceedings of the IEEE Symposium on Security and Privacy. 813–832.

Received 2023-02-02; accepted 2023-07-27