

GPT Conjecture: Understanding the Trade-offs between Granularity, Performance and Timeliness in Control-Flow Integrity

Zhilong Wang, Peng Liu
 College of Information Sciences and Technology
 The Pennsylvania State University, USA
 zzw169@psu.edu, pliu@ist.psu.edu

Abstract—Performance/security trade-off is widely noticed in CFI research, however, we observe that not every CFI scheme is subject to the trade-off. Motivated by the key observation, we ask three questions. Although the three questions probably cannot be directly answered, they are inspiring. We find that a deeper understanding of the nature of the trade-off will help answer the three questions. Accordingly, we proposed the GPT conjecture to pinpoint the trade-off in designing CFI schemes, which says that at most two out of three properties (fine granularity, acceptable performance, and preventive protection) could be achieved.

Keywords—Conjecture, Control-flow integrity, Trade-off.

I. INTRODUCTION

Along with the increased complexity of software, it becomes harder for the developers to ensure execution correctness in their software products, especially in those developed by the low-level programming languages, such as C/C++. A substantial amount of execution in-correctness is caused by the exploitation of software vulnerabilities in the real world. Softwares inevitably contain a wide variety of vulnerabilities, opening a window for attacks to compromise the system. Attackers have developed a series of attack methods, such as shellcode injection[1], return-to-libc[2], ROP[3] and so on, to exploit all kinds of vulnerabilities, e.g., buffer overflow, format string, use-after-free, and so on [4]. Among all kinds of attacks, the control-flow hijacking attack is the most dangerous one, because it allows the attacker to control the program’s execution, execute arbitrary malicious code and attain Turing-complete operation[3]. To mitigate the threats, many defense mechanisms, such as stack smashing protector (SSP)[5], address space layout randomization (ASLR)[6], data execution prevention (DEP)[7] and so on, have been put forward by researchers and applied in the real world software products.

Among all the defense techniques, security schemes based on the concept of control-flow integrity (CFI) have attracted many researchers’ attention because of its simplicity to implement, effectiveness to cope with the full spectrum of control-flow hijacking attacks, and flexibility to trade between security and efficiency. CFI schemes guarantee the correctness of the program by dynamically checking the control-flow transfer and confining the target address to a legal set.

Since CFI was introduced by Abadi et al. in 2005 [8], many researchers afterward were dedicated to enhance its runtime performance, security, scalability, compatibility and

so on. According to mainstream taxonomy, most CFI schemes can be clarified into two categories: fine-grained CFI schemes that provide more security guarantee, and coarse-grained CFI schemes that attain higher runtime performance. However, both fine-grained and coarse-grained CFI schemes have noticeable limitations that have not been addressed yet. As shown in previous survey papers [9], lightweight CFI schemes can not fully prevent sophisticated code reuse attack. The adversary’s attacking strategy is to search large gadgets chain whose starting addresses are allowed in a rough control-flow graph that coarse-grained CFI schemes adopted [10], [11]. Precise CFI schemes usually suffer from unacceptable runtime overhead. Hence, it is widely believed “performance/security trade-off” exists between runtime overhead and security in different CFI schemes [9], [12].

However, we observe that not every CFI scheme is subject to the trade-off between performance and security. In fact, several CFI schemes are “immunized” from doing such a trade-off. For instance, π CFI designed by Niu et al. achieves fine-grained security with a runtime overhead of 3.2% on average, which is fairly low and acceptable [13]. Victor et al. proposed a context-sensitive CFI scheme that achieves stronger security than conventional fine-grained ones with an overhead of less than some of the coarse-grained ones [14].

Key Observation. The trade-off between performance and security does *not* universally exist in meaningful CFI schemes. This intriguing observation motivates us to ask three questions: ❶ does trade-off really exist in different CFI schemes? ❷ if trade-off do exist, How do previous works comply with it? ❸ how can it inspire future research?

Although the questions probably cannot be directly answered, they are inspiring. On the other hand, we find that a deeper understanding of the nature of the trade-off will help answer these questions. Accordingly, we propose the GPT conjecture to pinpoint general trade-offs in CFI schemes: the impossibility of guaranteeing both fine granularity and acceptable performance in a Just-In-Time CFI scheme. We analyze its rationality through empirical study—surveying a series of representative CFI schemes and showing how existing CFI schemes comply with our conjecture. Finally, we give some recommendations for future researchers. We believe that our conjecture will help researchers have a more clear understanding of internal relations among properties of CFI schemes, thereby, motivating future research in this area.

II. BACKGROUND

When compiling source code written by low-level language (such as C or C++) into machine code, the compiler emits control data [15] (data that are loaded to processor program counter at some point in program execution, e.g., return addresses and function pointers) into the binary file without any protection. The security of control data depends on checks inserted by the programmer to enforce memory safety [16]. Along with program execution, attacker’s malicious tampering with control data through software vulnerabilities, such as buffer overflow, can transfer the program’s control-flow to any executable address in process space.

Based on this observation, researchers invented CFI to protect programs against control-flow hijacking attacks by checking programs’ control data before loading them into the program counter (EIP/RIP register in x86/x64 architecture). CFI’s strategy is to restrict the control-flow of a program to a pre-calculated CFG by checking indirect control-flow transfers at runtime [9]. Generally, most of CFI schemes follow a mainstream that consists of two phases.

In phase one, an analyzer statically computes the program’s control-flow graph (CFG). CFG is a representation in graph form of all legitimate control-flow transfers (also being called branch) in program space. It consists of sets of nodes and directed edges. Each node and edge denotes a basic block and a valid branch in the program respectively. For a comprehensive understanding, we refer the reader to the formal definition of CFG in work by Allen, et al. [17].

In phase two, a runtime control-flow checking (validation) component validates just fetched control data before each indirect-branching according to the legitimate CFG generated in phase one¹. An indirect-branch can pass checking only if it can be matched to a corresponding edge in the CFG. A failed validation will result in the process to terminate its execution and report an error. In such a fashion, control-flow attacks which usually introduce out-of-range branch are extremely prohibited. Researchers need to design efficient data structures to represent the CFG and enable runtime checking.

Despite its straightforward main idea, it is pretty challenging to design a CFI scheme with strong security, acceptable performance, high compatibility and so on [12], [9]. Researchers have designed hundreds of CFI schemes to explore its potential in different perspectives. The dominant difference of these various CFI schemes can be summarized into three aspects: 1) the precision of a CFG they employed. 2) the algorithm they designed to check indirect-branches. 3) the time point checking algorithm was activated.

1) *Precision of CFG Analyzer*: CFG can be obtained by analyzing the program’s source code or binary code. Like pointer analysis [18], perfect CFG generation is can not be fully achieved yet in many situations [10]. By now researchers have adopted several types of methods (insensitive analysis, context-sensitive analysis, and path-sensitive analysis) in their CFG analyzer and achieve different precisions. It is widely agreed that path-sensitive analysis is more precise than context-

sensitive analysis, and context-sensitive analysis is more precise than insensitive analysis [19].

2) *Algorithm to Enforce Checking*: The efficiency of different CFI schemes is largely dependent on their algorithms to enforce validation, which is tightly combined with their data structure that represents the CFG and enables runtime checking. Researchers have designed different types of algorithms and data structures in different CFI schemes. For example, the original CFI scheme proposed by the Abadi, et al. groups branch targets into different sets, assigns each set with a label, and inlines labels into each jump targets, i.e., the basic block’s in code. Based on this data structure, “guard instructions” are emitted before each indirect-branch instruction to compare its label with the one in target basic block [8]. A mismatch indicates that the control data is corrupted, then the program’s execution will be redirect to the error handling code accordingly.

π CFI [13] and MCFI [20] by Niu, et al. adopts two ID tables, namely Bary and Tary, to store target program’s CFG. In essence, Bary table and Tary table are hashmaps which can efficiently map indirect-branch points and target basic blocks to their corresponding IDs. Specifically, the Tary table is an array of IDs indexed by code addresses, mapping target basic block to their corresponding IDs. The Bary table uses a similar design, mapping indirect-branch points to their corresponding IDs. Two tables enable efficient ID look-ups and a indirect-branch is checked by comparing the IDs of branch point and target.

3) *Just-In-Time Checking vs. Lazy Checking*: Another difference among CFI schemes is how they schedule their checking operations. Most CFI schemes check the target address before indirect-branch occurs (we define it as *Just-In-Time checking*). While, to achieve better performance, some works log each indirect-branches at runtime and check them by employing another accompanying thread [21], [22], [23], [14], [22] (we define it as *Lazy checking*). For example, PITYPAT [22] enforces path-sensitive CFI by maintaining a “shadow” execution/analyzer, running concurrently with the protected process and checks its finished indirect-branches. Such a non-intrusive checking does not disturb the normal execution of the monitored process, hence achieves path-sensitive CFI with practical runtime overhead.

III. CONJECTURE

This section aims to answer Question❶ and Question❷. We observe that some terms, such as coarse-grained/fine-grained, have not been clearly defined. Before introducing the GPT conjecture, let us give a more precise definition of the terms and concepts that will be used throughout the paper. Then we propose the GPT conjecture which helps to answer Question❶. At last, some evidence is collected from an empirical study to answer the Question❷.

A. Terminology

Property 1. (Granularity)

Suppose a program has n indirect branch instructions. Let \mathbb{Z}_i ² denote the set of valid successors (basic blocks) of the

¹Direct control-flow transfers do not load any control data, their target addresses/offsets are hard-coded in their instructions.

²It is computed through mainstream insensitive control flow analysis. We admit the inaccuracy due to the difficulty of the pointer analysis.

i -th indirect branch instruction, and \mathcal{S} denote the set of all successor sets, namely,

$$\mathcal{S} = \{\mathbb{Z}_i : 1 \leq i \leq n\} \quad (1)$$

For a CFI scheme, let \mathbb{C}_i denote the checking set which is defined by the scheme and assigned to the i -th indirect branch instruction, then used to check the branch’s target at runtime. Only the elements in \mathbb{C}_i are valid successors authorized by the CFI schemes that the i -th branch instruction could jump to.

Definition 1. For arbitrary two sets $\mathbb{Z}_i, \mathbb{Z}_j$ from \mathcal{S} , satisfying $\mathbb{Z}_i \cap \mathbb{Z}_j \neq \emptyset \vee \mathbb{Z}_i \neq \mathbb{Z}_j$, as long as the CFI scheme merges $\mathbb{Z}_i, \mathbb{Z}_j$ when define its \mathbb{C}_i or \mathbb{C}_j , namely,

$$\mathbb{Z}_i \cup \mathbb{Z}_j \in \mathbb{C}_i \vee \mathbb{Z}_i \cup \mathbb{Z}_j \in \mathbb{C}_j \quad (2)$$

we define this scheme as a coarse-grained CFI scheme. Otherwise, we define it as a fine-grained CFI scheme. This definition enables us to determinate the granularity property of CFI schemes.

REMARK 1. According to Definition 1, both context-sensitive and path-sensitive CFI schemes belong to fine-grained CFI scheme. In essence, they reduce the size of their checking set \mathbb{C}_i for $i \in [1, n]$ based on context-sensitive or path-sensitive pointer analysis. Their protection is generally considered to be more powerful than that of insensitive fine-grained CFI scheme.

REMARK 2. Note that CFI schemes [24], [25] which adopt pointer encryption approach should be classified as coarse-grained CFI scheme. They cannot fully prevent code reuse attack because of two noticeable drawbacks. As discussed in Cryptographically Enforced Control Flow Integrity (CCFI) [24], it is still possible to replace the current encrypted pointer with another one from the program space and potentially disrupt control flow. The other drawback is that these schemes suffer from key leakage issues: the key can be inferred by brute-force attack or known-plaintext attack [26], especially for schemes which adopt a linear encryption/decryption method (XOR) [25].

REMARK 3. We remark that schemes that only provide partial protection—protecting subset of indirect branches in program space—belong to coarse-grained CFI scheme. For instance, `vfGuard` [27], `VTV` [28], and `SAFEDIS-PATCH` [29] only achieve strict protection for virtual function calls in COTS binaries;

Property 2. (Performance)

Evidence 1. As discussed in many papers [4], [9], [30], runtime performance is one of the most important determinants of whether a defense technique will be adopted by industry. Generally, to get adopted by industry, a defense technique should introduce less than 5% average overhead, such as `StackGuard`, `ASLR`, and `DEP`. Techniques incurring an overhead larger than 10% do not tend to gain wide adoption in production environments. Accordingly, the threshold should lie between 5%-10%.

Evidence 2. Other than runtime performance, space performance is another important index to measure a scheme. Program’s runtime memory consumption consists of four aspects,

i.e., code, global data, heap, and stack. Different programs have different ratios in four aspects, and a defense technique commonly increases memory consumption in one or more aspects. We observe that shadow based protections like shadow stack [31], shadow memory [32] and shadow processing [33], that double memory consumption in one or more aspects are unlikely to be deployed in practice.

Definition 2. Conservatively, we define a runtime overhead of less than 10% and a space overhead of less than 100% (in any of aforementioned four aspects) as an acceptable performance. Otherwise, it is an unacceptable performance. This definition enables us to determinate the performance property of CFI schemes.

Property 3. (Timeliness)

Observation 1. Whereas the term “integrity” in the context of CFI implies that it can prevent the attacks [8], some of the CFI schemes do not hit the mark. To achieve higher efficiency, some CFI schemes as mentioned in Section II-3 adopted a lazy checking mechanism, which checks programs’ control-flow following the programs execution rather than before each indirect branching. Generally, they log the program’s runtime control-flow transfer along with its execution, then check the control-flow offline or through an accompanying thread. In these designs, a sliding window exists between the program’s control-flow transfer and checking. The attacker can compromise the system without being perceived in the sliding window, which means this kind of CFI cannot protect software against such attacks.

Definition 3. We regard that the aforementioned design of CFI schemes provides less protection than CFI schemes that perform Just-In-Time checking. We define protection capability powered by lazy checking schemes as detective protection, the others that powered by Just-In-Time checking as preventive protection. This definition enables us to determinate the property timeliness of CFI schemes.

B. The Proposed Conjecture

GPT Conjecture: A control-flow integrity scheme can have at most two out of three properties:

- P1.** Fine granularity
- P2.** Acceptable performance
- P3.** Preventive protection

C. Some Evidence of the GPT Conjecture

In this section, we will reflect on our conjecture through several pieces of evidence. To verify the rationality of our conjecture, we conduct an empirical study on 32 representative works, and show the results in Table I. Three columns (**P1**, **P2** and **P3**) in the table display three properties respectively as we define in Section III-A. **P1** column denotes the granularity—check-mark indicates a fine-grained scheme whileas cross-mark represents a coarse-grained

TABLE I: Reflection of GPT conjecture in 32 control-flow integrity schemes.

Schemes		P1 ¹	P2	P3
CFIXX [34]		X	4.98%	✓
REINS [35]		X	2.40%	✓
vfGuard [27]		X	18.30%	✓
VTV [28]		X	9.60%	✓
LLVM-CFI [36]		X	1.10%	✓
VTI [37]		X	0.50%	✓
CFGuard [38]		X	2.30%	✓
IFCC [28]		X	-0.30%	✓
ROPecker [39]		X	2.60%	✓
bin-CFI [40]		X	8.50%	✓
ROPGuard [41]		X	0.48%	✓
SafeDispatch [29]		X	2.00%	✓
CCFIR [42]		X	2.08%	✓
kBouncer [43]		X	4.00%	✓
OCFI [44]		X	4.70%	✓
CFIMon [45]		X	6.10%	✓
τ CFI [46]		X	2.89%	✓
HCIC [25]		X	0.95%	✓
RAGuard [47]		X	1.86%	✓
HyperSafe [48]		X	5.00%	✓
BinCC [48]		X	4.00%	✓
CCFI [24]		X	52.00%	✓
KCoFI [49]		X	13.00%	✓
Original CFI [50]		✓	16.00%	✓
Lockdown [51]		✓	20.00%	✓
MCFI [20]		✓	5.00% & 4GB	✓
π CFI [13]		✓	3.20% & 4GB	✓
GRIFFIN [23]	H ²	✓	11.90%	X
PITTYPAT [22]	P, H	✓	12.73%	X
ECFI [52]		✓	1.50%	X
μ CFI [21]	C, H	✓	10.00%	X
PathArmor [14]	C	✓	3.00%	X

¹ If a CFI scheme supports different security levels, e.g. having both coarse-grained and fine-grained versions, we focus on its most secure version.

² ‘H’, ‘P’ and ‘C’ denote hardware-assisted CFI scheme, path sensitive CFI scheme, and context sensitive CFI scheme, respectively.

scheme. **P2** column shows the performance overheads which are reported in corresponding papers. Note that we prefer evaluation results which are based on SPEC CPU[®]2006 benchmarks [53]. **P3** column labels whether a CFI scheme provides preventive protection. We label the data in each column with *red* color when it fails to meet the requirement defined in the conjecture.

Evidence i. It can be clearly seen in Table I that all CFI schemes we surveyed comply with our conjecture—**no CFI schemes can achieve all three properties**. Also, some of unsophisticated schemes, such as PITTYPAT [22] and GRIFFIN [23], only achieve one property, i.e., fine granularity.

Evidence ii. MCFI [20] and π CFI developed by Niu, et al. achieve fine granularity with acceptable runtime overheads, i.e., 3.2% and 5.0%, respectively. However, researchers did not realize that their better runtime overhead is achieved through sacrificing their space performance. Even though they

did not report their space overhead in their paper explicitly, we can infer it in a reasonable manner.

As discussed in Section II-2, both of two schemes adopt two tables, namely Bary and Tary, to support their runtime checking. Accordingly, 1GB/4GB memory space on x86-32 and x86-64 operating system, respectively, need to be reserved in each process for the tables. As stated by the author, “On x86-32, memory segmentation is used, as in NaCl [54]. A 1GB segment is reserved for running the application code and another 1GB segment is reserved for the table region. x86-64, however, does not support memory segmentation. Instead, memory writes are instrumented so that they are restricted to the [0, 4GB) memory region. Another 4GB memory region is reserved for tables.” In view of the size of memory consumption of typical programs (mostly less than 1GB [53]), their space overhead has already reached 100% except for code bloat caused by extra no-op instructions inserted to enforce four-byte alignment on indirect-branch targets.

Evidence iii. GRIFFIN [23] is a hardware-assisted CFI, which leverages Intel PT to record control-flow of a monitored program. It supports multiple types of CFI policies to enable flexible trade-offs between security and performance. The fine-grained scheme incurs an average of 11.9% overhead. It leverages idle cores on a multi-core system for security checking by having multiple worker threads to check runtime control-flow simultaneously. In most of the time, it performs non-blocking checking which analyzes trace buffer of Intel PT whenever it becomes full; In a few cases when security-sensitive system calls are invoked, it performs blocking checking which stops the target thread until all the control transfers in the buffer have been checked. It can only provide the detective protection for software according to Definition 3. This case indicates that GPT conjecture is applicable to hardware-assisted CFI schemes.

Evidence iv. PITTYPAT [22], μ CFI [21] and PathArmor [14] are path/context sensitive CFI schemes which adopt path-sensitive or context-sensitive analysis to generate their CFG. However, path-sensitive and context-sensitive analysis is generally considered to be more time-consuming and space-consuming than insensitive analysis [19]. We find that all three CFI schemes adopt two common features: hardware-assisted branch recording and lazy checking. Specifically, PITTYPAT and μ CFI employ Intel PT—a brand new hardware feature in Intel CPUs—to efficiently record conditional and indirect branches taken by a program at runtime while PathArmor adopts Last Branch Record (LBR) registers available in Intel processors to monitor recently exercised control-flow transfers in an efficient way. Their control-flow checking is achieved through accompanying threads. This case indicates that both path-sensitive and context-sensitive CFI schemes conform to the claim of GPT conjecture.

REMARK 4. Our observations indicate that the GPT conjecture is universally applicable in all kinds of scenarios. Further, four pieces of evidence are not meant to be exhaustive and more evidence are easy to find.

IV. IMPLICATIONS OF THE GPT CONJECTURE

In this section, we will focus on answering Question ③: how can GPT conjecture inspire future research?

First of all, GPT conjecture illustrates the inherent trade-offs of three important properties (fine granularity, acceptable performance, and preventive protection) in CFI schemes. It helps researchers to have a deeper understanding of the nature of CFI based protection. Accordingly, future researchers should make a necessary sacrifice before designing new CFI schemes. In the broader context, GPT conjecture provides insights into the feasible design space for CFI schemes, shedding some light on the manner in which algorithm designers and software engineers have circumvented the conjecture.

Second, for decades, security researchers have been focused on CFI scheme's runtime performance and made their best effort to improve it. Evidence ii shows that in some cases, better runtime performance is achieved by sacrificing its space performance. Just as Gerhard states, "For some problems, we can reach an improved time complexity, but it seems that we have to pay for this with an exponential space complexity" [55]. Therefore, performance evaluation in future research should not merely be limited to runtime performance and researchers should have a more comprehensive evaluation of their schemes.

Third, Evidence iii that even powerful hardware support cannot eliminate the runtime overhead of Just-In-Time CFI schemes to an acceptable level, which implies that the challenge in the implementation of CFI cannot be solved only through engineering efforts, instead, it may relate to computational complexity theory [56]. In a broader sense, we observe that indirect branching poses not only challenge in the security field, but also challenges to many others: precise pointer analysis is an NP-complete problem [57]; indirect branch prediction is a performance-limiting factor for current computer systems [58]. Hence, GPT conjecture implies the complexity of the CFI problem, which deserves to be investigated through theoretical methods.

At last, despite the inspiring implications that Gpt conjecture gives to us, we admit that we still cannot prove the conjecture at this time.

V. CONCLUSION

Control-flow integrity is a popular defence technique for detecting and defeating control-flow hijacking attacks. Since its inception in the decade, researchers have put great efforts to explore its potential regarding security, performance, compatibility and so on. Even though performance/security trade-off is widely noticed in CFI research, we observe that not every CFI scheme is subject to it. In this paper, we propose the GPT conjecture to illustrate the general trade-offs in CFI schemes. The conjecture points out the impossibility of guaranteeing both fine granularity and acceptable performance in a Just-In-Time CFI schemes. We have verified the rationality of our conjecture based on an empirical study on existing works. Even though we cannot prove the conjecture at this time, we believe that GPT conjecture will help researcher to have a deeper understanding of the nature of CFI problem and it will direct future research in this area.

REFERENCES

- [1] J. Erickson, *Hacking: The Art of Exploitation*. No Starch Press, 2008.
- [2] R. Wojtczuk, "The Advanced Return-into-Libc Exploits: PaX Case Study." *Phrack Magazine*, 2001.
- [3] H. Shacham *et al.*, "The Geometry of Innocent Flesh on the Bone: Return-into-libc without Function Calls (on the x86)." in *ACM conference on Computer and communications security*. New York., 2007, pp. 552–561.
- [4] L. Szekeres, M. Payer, T. Wei, and D. Song, "Sok: Eternal War in Memory." in *2013 IEEE Symposium on Security and Privacy (S&P)*. IEEE, 2013, pp. 48–62.
- [5] C. Cowan, C. Pu, D. Maier, J. Walpole, P. Bakke, S. Beattie, A. Grier, P. Wagle, Q. Zhang, and H. Hinton, "Stackguard: Automatic Adaptive Detection and Prevention of Buffer-overflow Attacks." in *USENIX Security Symposium*. San Antonio, TX, 1998.
- [6] H. Shacham, P. Matthew, P. Ben, G. Eu-Jin, M. Nagendra, and B. Dan, "On the Effectiveness of Address-Space Randomization." in *Proceedings of the 11th ACM conference on Computer and Communications Security*. ACM, 2004.
- [7] S. Andersen and V. Abella, "Data Execution Prevention. Changes to Functionality in Microsoft Windows XP Service Pack 2, Part 3: Memory Protection Technologies." 2004.
- [8] M. Abadi, M. Budiu, U. Erlingsson, and J. Ligatti, "Control-flow Integrity." in *Proceedings of the 12th ACM Conference on Computer and Communications Security*, ser. CCS '05. New York, NY, USA: ACM, 2005, pp. 340–353.
- [9] N. Burow, S. A. Carr, J. Nash, P. Larsen, M. Franz, S. Brunthaler, and M. Payer, "Control-Flow Integrity: Precision, Security, and Performance." *ACM Computing Surveys (CSUR)*, vol. 50, no. 1, p. 16, 2017.
- [10] E. Göktaş, E. Athanasopoulos, H. Bos, and G. Portokalidis, "Out of Control: Overcoming Control-Flow Integrity." in *2014 IEEE Symposium on Security and Privacy (S&P)*. IEEE, 2014.
- [11] D. Lucas, A.-R. Sadeghi, D. Lehmann, and F. Monrose, "Stitching the Gadgets: On the Ineffectiveness of Coarse-Grained Control-Flow Integrity Protection." in *23rd USENIX Security Symposium (USENIX Security 14)*. USENIX, 2014.
- [12] X. Xu, M. Ghaffarinia, W. Wang, K. W. Hamlen, and Z. Lin, "CONFIRM: Evaluating Compatibility and Relevance of Control-flow Integrity Protections for Modern Software." in *28th USENIX Security Symposium (USENIX Security 19)*. Santa Clara, CA: USENIX Association, Aug. 2019, pp. 1805–1821.
- [13] B. Niu and G. Tan, "Per-Input Control-Flow Integrity." in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '15. New York, NY, USA: ACM, 2015, pp. 914–926.
- [14] V. Van der Veen, D. Andriess, E. Göktaş, B. Gras, L. Sambuc, A. Slowinska, H. Bos, and C. Giuffrida, "Practical Context-Sensitive CFI." in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2015, pp. 927–940.
- [15] S. Chen, J. Xu, E. C. Sezer, P. Gauriar, and R. K. Iyer, "Non-Control-Data Attacks Are Realistic Threats." in *USENIX Security Symposium*, vol. 5, 2005.
- [16] S. G. Nagarakatte, *Practical Low-Overhead Enforcement of Memory Safety for C Programs*. University of Pennsylvania, 2012.
- [17] F. E. Allen, "Control Flow Analysis." in *Proceedings of a Symposium on Compiler Optimization*. New York, NY, USA: ACM, 1970, pp. 1–19. [Online]. Available: <http://doi.acm.org/10.1145/800028.808479>
- [18] M. Hind, "Pointer Analysis: Haven't We Solved This Problem Yet?" in *Proceedings of the 2001 ACM SIGPLAN-SIGSOFT Workshop on Program Analysis for Software Tools and Engineering*, ser. PASTE '01. New York, NY, USA: ACM, 2001, pp. 54–61.
- [19] U. Khedker, A. Sanyal, and B. Sathe, *Data Flow Analysis: Theory and Practice*. CRC Press, 2017.
- [20] B. Niu and G. Tan, "Modular Control-flow Integrity." in *Proceedings of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 14)*. New York, NY, USA: ACM, 2014, pp. 577–587. [Online]. Available: <http://doi.acm.org/10.1145/2594291.2594295>

- [21] H. Hu, C. Qian, C. Yagemann, S. P. H. Chung, W. R. Harris, T. Kim, and W. Lee, "Enforcing Unique Code Target Property for Control-Flow Integrity." in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2018, pp. 1470–1486.
- [22] R. Ding, C. Qian, C. Song, B. Harris, T. Kim, and W. Lee, "Efficient Protection of Path-Sensitive Control Security." in *26th USENIX Security Symposium (USENIX Security 17)*. Vancouver, BC: USENIX Association, Aug. 2017, pp. 131–148.
- [23] X. Ge, W. Cui, and T. Jaeger, "GRIFFIN: Guarding Control Flows Using Intel Processor Trace." in *Proceedings of the Twenty-Second International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS 17)*. New York, NY, USA: ACM, 2017, pp. 585–598.
- [24] A. J. Mashtizadeh, A. Bittau, D. Boneh, and D. Mazières, "CCFI: Cryptographically Enforced Control Flow Integrity." in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2015, pp. 941–951.
- [25] J. Zhang, B. Qi, Z. Qin, and G. Qu, "HCIC: Hardware-Assisted Control-Flow Integrity Checking." *IEEE Internet of Things Journal*, vol. 6, no. 1, pp. 458–471, Feb 2019.
- [26] X. Peng, P. Zhang, H. Wei, and B. Yu, "Known-Plaintext Attack on Optical Encryption Based on Double Random Phase Keys." *Optics Letters*, vol. 31, no. 8, pp. 1044–1046, 2006.
- [27] A. Prakash, X. Hu, and H. Yin, "vGuard: Strict Protection for Virtual Function Calls in COTS C++ Binaries." in *Symposium on Network and Distributed System Security (NDSS)*, 2015.
- [28] C. Tice, T. Roeder, P. Collingbourne, S. Checkoway, Ú. Erlingsson, L. Lozano, and G. Pike, "Enforcing Forward-Edge Control-Flow Integrity in GCC & LLVM." in *23rd USENIX Security Symposium (USENIX Security 14)*, 2014, pp. 941–955.
- [29] D. Jang, Z. Tatlock, and S. Lerner, "SafeDispatch: Securing C++ Virtual Calls from Memory Corruption Attacks." in *the Network and Distributed System Security Symposium (NDSS)*, 2014.
- [30] S. Andersen and V. Abella, "The BlueHat prize contest official rules." 2012. [Online]. Available: <http://www.microsoft.com/security/bluehatprize/rules.aspx>.
- [31] T. H. Dang, P. Maniatis, and D. Wagner, "The Performance Cost of Shadow Stacks and Stack Canaries." in *Proceedings of the 10th ACM Symposium on Information, Computer and Communications Security (ASIACCS 15)*. New York, NY, USA: ACM, 2015, pp. 555–566.
- [32] J. Newsome and D. X. Song, "Dynamic Taint Analysis for Automatic Detection, Analysis, and Signature Generation of Exploits on Commodity Software." in *the Network and Distributed System Security Symposium (NDSS)*, vol. 5. Citeseer, 2005, pp. 3–4.
- [33] H. Patil and C. N. Fischer, "Efficient Run-Time Monitoring Using Shadow Processing." in *Proceeding of Automated and Algorithmic Debugging (AADEBUG)*, vol. 95, 1995, pp. 1–14.
- [34] N. Burow, D. McKee, S. A. Carr, and M. Payer, "CfIXX: Object Type Integrity for C++ Virtual Dispatch." in *Proceedings of Network and Distributed System Security Symposium (NDSS)*. <https://hexhive.epfl.ch/publications/files/18NDSS.pdf>, 2018.
- [35] R. Wartell, V. Mohan, K. W. Hamlen, and Z. Lin, "Securing Untrusted Code Via Compiler-Agnostic Binary Rewriting." in *Proceedings of the 28th Annual Computer Security Applications Conference*. ACM, 2012, pp. 299–308.
- [36] "LLVMControl Flow Integrity." 2015. [Online]. Available: <http://clang.llvm.org/docs/ControlFlowIntegrity.html>
- [37] D. Bounov, R. G. Kici, and S. Lerner, "Protecting C++ Dynamic Dispatch Through VTable Interleaving." in *the Network and Distributed System Security Symposium (NDSS)*, 2016.
- [38] Microsoft, "Visual Studio 2015 compiler options enable control flow guard." 2015. [Online]. Available: <https://msdn.microsoft.com/en-us/library/dn919635.aspx>
- [39] Y. Cheng, Z. Zhou, Y. Miao, X. Ding, and R. H. Deng, "ROPecker: A Generic and Practical Approach for Defending Against ROP Attack." in *Symposium on Network and Distributed System Security (NDSS)*. Internet Society, 2014.
- [40] M. Zhang and R. Sekar, "Control Flow Integrity for COTS Binaries." in *Proceeding of the 22nd USENIX Security Symposium (USENIX Security 13)*, 2013, pp. 337–352.
- [41] I. Fratrić, "ROPGuard: Runtime Prevention of Return-Oriented Programming Attacks." Technical report, Tech. Rep., 2012.
- [42] C. Zhang, T. Wei, Z. Chen, L. Duan, L. Szekeres, S. McCamant, D. Song, and W. Zou, "Practical Control Flow Obfuscation and Randomization for Binary Executables." in *2013 IEEE Symposium on Security and Privacy (S&P)*. IEEE, 2013, pp. 559–573.
- [43] V. Pappas, M. Polychronakis, and A. D. Keromytis, "Transparent ROP Exploit Mitigation Using Indirect Branch Tracing." in *Proceeding of the 22nd USENIX Security Symposium (USENIX Security 13)*, 2013, pp. 447–462.
- [44] V. Mohan, P. Larsen, S. Brunthaler, K. W. Hamlen, and M. Franz, "Opaque Control-Flow Integrity." in *the Network and Distributed System Security Symposium (NDSS)*, vol. 26, 2015, pp. 27–30.
- [45] Y. Xia, Y. Liu, H. Chen, and B. Zang, "CFIMon: Detecting Violation of Control Flow Integrity Using Performance Counters." in *IEEE/IFIP International Conference on Dependable Systems and Networks (DSN 2012)*. IEEE, 2012, pp. 1–12.
- [46] P. Muntean, M. Fischer, G. Tan, Z. Lin, J. Grossklags, and C. Eckert, "τ CFI: Type-Assisted Control Flow Integrity for x86-64 Binaries." in *International Symposium on Research in Attacks, Intrusions, and Defenses*. Springer, 2018, pp. 423–444.
- [47] J. Zhang, R. Hou, J. Fan, K. Liu, L. Zhang, and S. A. McKee, "RAGuard: A Hardware Based Mechanism for Backward-Edge Control-Flow Integrity." in *Proceedings of the Computing Frontiers Conference*, ser. CF'17. New York, NY, USA: ACM, 2017, pp. 27–34.
- [48] W. Zhi and J. Xuxian, "HyperSafe: A Lightweight Approach to Provide Lifetime Hypervisor Control-Flow Integrity." in *2010 IEEE Symposium on Security and Privacy*, May 2010, pp. 380–395.
- [49] J. Criswell, N. Dautenhahn, and V. Adve, "KCoFI: Complete Control-Flow Integrity for Commodity Operating System Kernels," in *2014 IEEE Symposium on Security and Privacy*, May 2014, pp. 292–307.
- [50] M. Abadi, M. Budiu, Ú. Erlingsson, and J. Ligatti, "Control-Flow Integrity Principles, Implementations, and Applications." *ACM Transactions on Information and System Security (TISSEC)*, vol. 13, no. 1, p. 4, 2009.
- [51] M. Payer, A. Barresi, and T. R. Gross, "Fine-grained Control-Flow Integrity Through Binary Hardening." in *International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment*. Springer, 2015, pp. 144–164.
- [52] A. Abbasi, T. Holz, E. Zambon, and S. Etalle, "ECFI: Asynchronous Control Flow Integrity for Programmable Logic Controllers," in *Proceedings of the 33rd Annual Computer Security Applications Conference*, ser. ACSAC 2017. New York, NY, USA: ACM, 2017, pp. 437–448. [Online]. Available: <http://doi.acm.org/10.1145/3134600.3134618>
- [53] "SPEC CPU 2006 system requirements," Available: <https://www.spec.org/cpu2006/>, Standard Performance Evaluation Corporation.
- [54] B. Yee, D. Sehr, G. Dardyk, J. B. Chen, R. Muth, T. Ormandy, S. Okasaka, N. Narula, and N. Fullagar, "Native Client: A Sandbox for Portable, Untrusted x86 Native Code." in *2009 30th IEEE Symposium on Security and Privacy (S&P)*, May 2009, pp. 79–93.
- [55] G. J. Woeginger, "Space and Time Complexity of Exact Algorithms: Some Open Problems." in *Parameterized and Exact Computation*, R. Downey, M. Fellows, and F. Dehne, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 281–290.
- [56] O. Goldreich, "Computational Complexity: A Conceptual Perspective." *SIGACT News*, vol. 39, no. 3, pp. 35–39, Sep. 2008. [Online]. Available: <http://doi.acm.org/10.1145/1412700.1412710>
- [57] S. Horwitz, "Precise Flow-insensitive May-alias Analysis is NP-hard." *ACM Trans. Program. Lang. Syst.*, vol. 19, no. 1, pp. 1–6, Jan. 1997. [Online]. Available: <http://doi.acm.org/10.1145/239912.239913>
- [58] O. J. Santana, A. Falcón, E. Fernández, P. Medina, A. Ramírez, and M. Valero, "A Comprehensive Analysis of Indirect Branch Prediction." in *High Performance Computing*, H. P. Zima, K. Joe, M. Sato, Y. Seo, and M. Shimasaki, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2002, pp. 133–145.