

Live Path CFI Against Control Flow Hijacking Attacks

Mohamad Barbar^{*}, Yulei Sui^{*}, Hongyu Zhang[†], Shiping Chen[‡], Jingling Xue[‡]

^{*}University of Technology Sydney, Australia

[†]University of Newcastle, Australia

[‡]CSIRO/Data61, Australia

[‡]University of New South Wales, Australia

Abstract. Through memory vulnerabilities, control flow hijacking allows an attacker to force a running program to execute other than what the programmer has intended. Control Flow Integrity (CFI) aims to prevent the adversarial effects of these attacks. CFI attempts to enforce the programmer's intent by ensuring that a program only runs according to a control flow graph (CFG) of the program. The enforced CFG can be built statically or dynamically, and Per-Input Control Flow Integrity (PICFI) represents a recent advance in dynamic CFI techniques. PICFI begins execution with the empty CFG of a program and lazily adds edges to the CFG during execution according to concrete inputs. However, this CFG grows monotonically, i.e., edges are never removed when corresponding control flow transfers become illegal. This paper presents LPCFI, Live Path Control Flow Integrity, to more precisely enforce forward edge CFI using a dynamically computed CFG by both adding and removing edges for all indirect control flow transfers from indirect callsites, thereby raising the bar against control flow hijacking attacks.

Keywords: Control flow integrity

1 Introduction

Programs written in low-level languages, such as C and C++, make up the majority of performance-critical system software (e.g., web browsers and language runtimes) running on most computing platforms. In some domains, like embedded systems, these languages are almost ubiquitous. However, these unsafe languages are prone to memory corruption vulnerabilities (e.g., use-after-free and buffer overflows). An attacker may leverage these vulnerabilities to launch control flow hijacking attacks by changing the target of an indirect branch instruction to force a running program to execute at a location of the attacker's choice. In realistic scenarios, attackers may be able to perform Turing complete computation by abusing memory vulnerabilities and using techniques like return oriented programming [1] and counterfeit object-oriented programming [2].

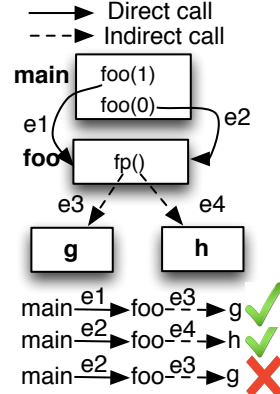
Control Flow Integrity (CFI) has been proposed to prevent control flow hijacking [3]. CFI typically works by enforcing a control flow graph (CFG), which represents the programmer's intent - or rather, what can be inferred as legal and illegal control flow from the program. Edges in the CFG represent control flow transfers, and CFI aims to protect indirect control flow edges from being taken illegally. The protection offered by CFI is more effective if a more precise CFG is used. The CFG can be computed statically and this does not consider the fact that the legal status of indirect control flow transfers constantly changes during

```

1: void (*fp)(void);
2: void foo(int n) {
3:   if (n) {
4:     lpcfi_assign_const(fp, &g);
5:     fp = &g;
6:   } else {
7:     lpcfi_assign_const(fp, &h);
8:     fp = &h;
9:   }
10://unsafe: modify the value of fp
11: lpcfi_check(fp);
12: fp();
13: }
14: void main(void) {
15:   foo(1)
16:   foo(0)
17: }

```

(a) Otherwise unsafe code protected by LPCFI.



(b) PICFI's CFG and feasible paths (green) and infeasible path (red) not protected by PICFI.

Fig. 1: A motivating example to demonstrate the limitation of PICFI.

runtime. For example, when a function pointer is reassigned to a new value, an indirect call via that function pointer will call a new function target and calling the previous target would be illegal.

Insights. Per-Input Control Flow Integrity (PICFI) [4] represents a recent dynamic approach to forward edge CFI. PICFI first pre-computes a static CFG as the upper bound for its dynamic one. PICFI starts with the empty CFG of a program, and during runtime, once a function address is taken (e.g., $p = \&func$), it will add an edge from each indirect callsite to $func$ if this edge is also found in the static CFG. Hence PICFI provides better security guarantees than CFI techniques which enforce a statically computed CFG. However, PICFI's dynamic CFG grows monotonically, i.e., edges added to the CFG are never removed. Hence, edges become permanently legal to take regardless of whether their legality changes over time. The conservatively constructed dynamic CFG used by PICFI leaves an attack surface: when an indirect transfer remains on the monotonically growing CFG but can never be legally executed again.

Motivating Example. Figure 1 illustrates this limitation of PICFI via a proof-of-concept attack. Note that the lines marked in blue are instrumentation calls from our LPCFI approach to protect against this attack, and will be explained below. PICFI begins execution with an empty CFG. Initially the indirect callsite $fp()$ at line 12 cannot invoke any function legally. After executing the *if* branch via $foo(1)$ at line 15, g becomes a legitimate target (e_3 is added to the CFG). After executing the *else* branch via $foo(0)$ at line 16, h becomes a legitimate target (e_4 is added to the CFG).

Figure 1b gives PICFI's CFG constructed immediately before the indirect callsite $fp()$ at line 12 when foo is invoked for a second time via $foo(0)$ at line 16. Unfortunately, the indirect call edge $fp() \xrightarrow{e_3} g$, which was added during the first execution of foo , has already become illegal to take since fp only points to h during the second execution at the time of calling $fp()$. However, this spurious

edge $\text{fp}() \xrightarrow{e^3} \mathbf{g}$ remains on PICFI’s CFG. This conservative CFG allows attackers to redirect $\text{fp}()$ to \mathbf{g} by modifying fp ’s value to be \mathbf{g} via a memory corruption error [5], despite `foo` not being allowed to call \mathbf{g} when `n`’s value is 0. Therefore, PICFI still provides an attacker opportunities to launch control flow hijacking attacks by treating ”out-of-date” control flow edges as legitimate. This paper presents LPCFI, Live Path Control Flow Integrity, which aims to overcome this limitation of PICFI by both adding and removing CFG edges, allowing at most one outgoing forward edge from every indirect callsite at any one program point.

Let us revisit the example in Figure 1 whilst taking into consideration LPCFI’s instrumentation (highlighted in blue). During the first call to `foo`, $\text{fp}() \xrightarrow{e^3} \mathbf{g}$ is added to the CFG via `lpcfi_assign_const`. A check is then performed to ensure that the indirect call transfer from $\text{fp}()$ will reach the only legitimate target \mathbf{g} . During the second call to `foo`, `lpcfi_assign_const` in the *else* branch updates the CFG by first removing invalid edge $\text{fp}() \xrightarrow{e^3} \mathbf{g}$ from the CFG, and then adding $\text{fp}() \xrightarrow{e^4} \mathbf{h}$. This removal is important since the second call to `foo` via `foo(0)` is not allowed to call \mathbf{g} , which PICFI ignores. LPCFI ensures only one legitimate (live) function target is allowed at any call path to an indirect callsite.

Challenges. Designing a CFI technique that overcomes the aforementioned limitation is challenging. Firstly, precise static analysis is required to find statements which may require instrumentation as the precision of static analysis directly correlates with the overhead reduction achieved. Only the statements which may modify or read the value of a function pointer should be identified by static pointer analysis for instrumentation. Secondly, function pointer values need to be correctly maintained in safe memory and the metadata data structure needs to be well designed to ensure efficient lookup and runtime checks.

Our Solution. LPCFI aims to ensure only edges which are currently ”live” - can be legally taken - exist within the CFG. We have designed and implemented a new instrumentation approach which tracks function pointers and the address-taken function which they point to at any program point. A function pointer may only ever point to a single function object, so our instrumentation correctly updates which pointers point to which function objects in an efficient data structure in safe memory. We apply pointer analysis [6] to identify all statements which may potentially access the value of a function pointer, and instrument only those statements to minimise runtime overhead. Any callsite from a function pointer is checked to ensure the runtime value matches the value stored in safe memory.

This paper makes the following key contributions:

- We present LPCFI, a new dynamic control flow integrity technique that can protect against attacks undetected by the conservative monotonically growing CFG used by PICFI.
- We propose a new instrumentation approach coupled with a data structure to allow only one function to be a legal target for any indirect callsite.
- We have developed a proof-of-concept attack and defence to demonstrate the effectiveness of LPCFI in mitigating control flow attacks which are not protected by PICFI. This is publicly available at <https://github.com/mbarbar/lpcfi>.

2 Related works

Often, CFI implementations determine policy (i.e. valid targets for an indirect control flow transfer at a particular time) according to only static information. This is limited in that some properties are impossible to determine statically, for example, the value of a function pointer reliant on user input.

Per-Input Control Flow Integrity (PICFI) is a CFI implementation which uses dynamic information to gradually build the CFG [4]. PICFI begins execution with an empty CFG; all indirect transfers of control are illegal. The CFG is gradually constructed by discovering valid targets for indirect control flow transfers during runtime according to program inputs. For example, when a function is called, that callsite becomes a valid target of return instructions, or when a function pointer is assigned a value, that value becomes a valid target for indirect callsites (constrained by the static CFG). However, these additions to the CFG are permanent; the CFG grows *monotonically*. This means that changes in target legality are not reflected in the CFG, and hence not enforced by PICFI. A target which is made legal by PICFI is regarded as legal for the rest of execution.

Offering improvements over PICFI, PittyPat [7], a very recent work, uses dynamic path-sensitive points-to analysis to further restrict the set of allowed function pointers at indirect callsites during runtime. Rather than considering just address activation, or the static points-to sets at a particular program point, PittyPat considers the points-to set of a function pointer at a particular program point only based on the *executed* program path. Hence PittyPat avoids PICFI's limitation of keeping previously legal targets which have become illegal. However, PittyPat has a strong dependency on specific hardware and a modified kernel. In contrast, LPCFI is a portable purely software-based approach without any hardware dependency.

A shadow stack is a *second* stack existing in memory used to ensure return instructions jump to the correct address [8, 9]. Shadow stacks work by mirroring return addresses pushed onto the execution stack. Upon returning, the value on top of the shadow stack is compared with that on the execution stack, and if the comparison fails, an error is detected. If the shadow stack is safe from manipulation, shadow stacks perfectly protect return transfers. However, shadow stacks only protect backward edges but not forward edges like virtual calls.

CFI techniques have recently been used to protect against virtual table hijacking attacks in low-level object-oriented languages like C++. VTV [10], VTrust [11], and SafeDispatch [12] apply Class Hierarchy Analysis (CHA) to analyse virtual calls to enforce CFI. ShrinkWrap [13] aims to improve CHA based CFI by considering multiple and diamond inheritance. VIP [14] is a recent CFI technique that enforces a more precise call graph than CHA based approaches by using pointer analysis and a fast index-based instrumentation.

This work builds on an earlier version of our work [15].

3 LPCFI Approach

This section details our Live Path Control Flow Integrity approach designed to reduce the attack surface left by PICFI. Section 3.1 describes the program

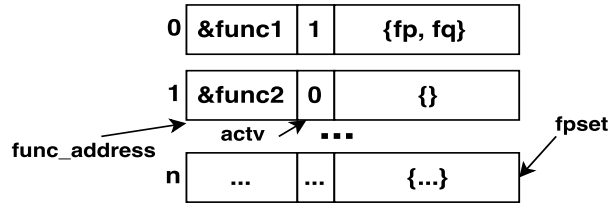


Fig. 2: Internal representation of the fp-table.

representation of a C/C++ program. Section 3.2 introduces the *fp-table*, the internal metadata design. Finally, Section 3.3 describes the instrumentation which operates on the *fp-table* to precisely update the dynamic CFG at runtime.

3.1 Program Representation

We represent programs in LLVM’s SSA form following [6, 16]. The set of all variables is separated into two subsets: top-level pointers (registers) whose addresses are not taken, and all potential targets, i.e., all address-taken objects of a pointer. In SSA, a program is represented by five statement types: `const` ($p = \&o$), `copy` ($p = q$), `store` ($*p = q$), `load` ($p = *q$), and `call` ($fp(\dots)$). Passing arguments into and returning results from functions are modeled by copies. A global variable initialisation is translated into one of the four types of assignments and analysed immediately at the beginning of the main function. For a `const` statement $p = \&o$ (allocation sites), o is a stack or global variable, or a dynamically created abstract heap object. We only analyse statements which access (modify or read) the value of a function pointer according to static pointer analysis [6].

3.2 Data structures

LPCFI needs to store metadata in the *fp-table* (Figure 2) to perform bookkeeping to update the dynamic CFG. The metadata is stored in a safe memory region which is accessed frequently for both reading and writing following [11].

LPCFI maintains the *fp-table* as shown in Figure 2, which is a fixed size array (size is the number of address-taken functions in the program) where each element holds: (1) the address of a function `func_address`, (2) an *activation* bit `actv`, and (3) a set `fpset` of function pointers which legally point to `func_address` at a particular program point during runtime. `pt(fp_table, fp)` returns the function that pointer `fp` points to. Overloaded `lookup(fp_table, &func)` returns the index of `&func` in the `fp_table`, and `lookup(fp_table, &fp)` returns the index of the function which `&fp` points to in the `fp_table`.

The *fp-table* is a simple yet efficient solution for fast lookup using a one dimensional array. The *fp-table* uses function addresses as keys for various reasons. Firstly, it can be a fixed size since functions which may have their addresses taken (`const` statements) at runtime are known statically. Secondly, the checking operation can perform lookups on the function that is about to be called (the runtime value) and retrieve its `fpset`. Finally, a data structure with function addresses as the key is required regardless to keep track of whether functions have been address-taken (activated) to guarantee a lower security bound of PICFI.

3.3 Instrumentation

LPCFI’s instrumentation is placed immediately before the five statement types. We insert instrumentation for an assignment **only if** it may read/write a function pointer value as determined by Andersen’s pointer analysis [6]. All instrumentations except the checking instrumentation write to the fp-table.

```

1: update(fp, &o) {
2:     // Check function object
3:     if(o not a function obj) return;
4:     // Search for index of object which fp points to
5:     oldInd = lookup(fp_table, &pt(fp_table,fp));
6:     // Remove fp from set of fp_table[oldInd]
7:     if (oldInd!=-1) remove(fp_table[oldInd].fpset, fp);
8:     // Search the index of function o in fp-table
9:     newInd = lookup(fp_table, &o);
10:    if (newInd==-1) error('not found');
11:    // Add fp to the new function pointer set
12:    add(fp_table[newInd].fpset, fp);
13: }

```

Fig. 3: Helper function `update` to remove and add pointers in the fp-table.

The four assignment instrumentations share helper function `update(fp, &o)` in Figure 3 which updates a function pointer `fp` to correctly point to function `o` by removing `fp` from the `fpset` of `fp`’s old points-to target (if it is a member) at line 7, and adding `fp` to `o`’s `fpset` at line 12. Note that pointer analysis is always an over-approximation. A pointer `q` resolved to point to a function statically, may not point to such at runtime. LPCFI will not perform any runtime update if the right hand side expression of an assignment (e.g., `... = q`) does not refer to a function object as shown at line 3 in Figure 3.

Handling constant assignments `fp = &func`: This case, as carried out by `lpcfi_assign_const` shown in Figure 4, is simple as it is a direct assignment of a function address `func` to a function pointer `fp`. Upon executing this statement, LPCFI requires that, (1) `func` be regarded as activated, and (2) `fp` exclusively points to `func` in the fp-table.

Assignments of this form may execute multiple times for the same RHS value. Hence, functions will be *activated* multiple times. This does not affect correctness and runtime overhead for the activation operation is negligible.

```

1: lpcfi_assign_const(fp, &func) {
2:     // Search the index of &func in fp_table
3:     ind = lookup(fp_table, &func);
4:     if(ind==-1) error('not found');
5:     // Mark func as activated
6:     fp_table[ind].actv_bit = 1;
7:     // Update fp to point to func
8:     update(fp, &func);
9: }
fp = &func;

```

Fig. 4: Handling const statements using `lpcfi_assign_const`.

Handling copy assignments `p = q`: Represented by `lpcfi_assign_copy`, the second case is also straightforward as shown in Figure 5. First, we obtain `pt(fp_table,q)`, the points-to target `o` of the RHS pointer `q` derived from the fp-table. Then, `p` is made to exclusively point to `q`’s pointee `o` if `o` is a function object, so both `p` and `q` are put into the `fpset` of object `o`.

```

1: lpcfi_assign_copy(p, q) {
2:     // Get the object that q points to in fp-table
3:     o = pt(fp_table, q);
4:     // Update p to point to o only if o is a function
5:     update(p, &o);
6: }
p = q;

```

Fig. 5: Handling copy statements using `lpcfi_assign_copy`.

Handling load assignments `p = *s`: `lpcfi_assign_load`'s implementation is shown in Figure 6. Similar to handling the copy case, we first retrieve points-to target `o` of `*s` from the `fp-table`. `o` is checked to ensure that it has been activated (lines 5-8) (for a lower bound protection of PICFI, further discussed in Section 4.3). Then, `p` is made to exclusively point to the object `o` (line 10).

```

1: lpcfi_assign_load(p, *s) {
2:     // Get the object that *s points to
3:     o = pt(fp_table, *s);
4:     // Search for the index of &o in fp-table
5:     ind = lookup(fp_table, &o);
6:     if(ind == -1) error('not found');
7:     // Ensure o has been activated
8:     assert(fp_table[ind].actv);
9:     // Update p to point to o
10:    update(p, &o);
11: }
p = *s;

```

Fig. 6: Handling load statements using `lpcfi_assign_load`.

Handling store assignments `*r = q`: `lpcfi_assign_store`'s implementation is shown in Figure 7. This case is similar to the copy case. The points-to target `o` of the RHS pointer `q` is retrieved via `pt(fp_table, q)`. Then, runtime value `*r` is made to exclusively point to the same as that which `q` does in the `fp-table`.

```

1: lpcfi_assign_store(*r, q) {
2:     // Get the object that *r points to
3:     o = pt(fp_table, q);
4:     // Update q to point to o
5:     update(*r, &o);
6: }
*r = q;

```

Fig. 7: Handling store statements using `lpcfi_assign_store`.

Handling calls `fp(...)`: As shown in Figure 8, whenever a call is made from a function pointer, the runtime value of the function pointer needs to be checked against its saved value in the `fp-table`. Furthermore, a check confirming that a callsite-to-target edge is within the static CFG must also be performed to guarantee a security lower bound of PICFI. If either check fails, LPCFI will report an error indicating an attempted illegal control flow transfer.

```

1: lpcfi_check(fp, callsite) {
2:     // Get the object that fp points to
3:     o = pt(fp_table, fp);
4:     // Enforce control flow integrity
5:     assert(runtimeVal(*fp) == &o
6:            && edge(callsite, &o) ∈ static CFG);
7: }
fp(...);

```

Fig. 8: Handling call statements using `lpcfi_check`.

4 Implementation

We have developed a prototype with a step-by-step live demo to illustrate examples (those in Figure 1 and Figure 9) that can be protected by LPCFI but not by PICFI. They are publicly available at <https://github.com/mbarbar/lpcfi>.

4.1 Instrumentation and data structure

In our open-source prototype, LPCFI’s data structure (Figure 2) and its instrumentation are implemented in an equivalent yet less efficient manner as a standalone library (i.e., `lpcfi.h`, `lpcfi.c`, `fp-table.h` and `fp-table.c`). In order to demonstrate the key idea and techniques easily, our prototype performs manual instrumentation for the motivating example (Figure 1) as available in `demo.c`.

At indirect callsites, a lookup operation through `lpcfi_check` is performed as discussed in Section 4.2. Assignment instrumentations are not idempotent so PICFI’s optimisation strategy of patching out instrumentation can not be achieved. `lpcfi_assign_const` performs both function activation (which is idempotent) and fp-table modification. Function activation results in a bit being set and is negligible to the total runtime overhead.

Andersen’s pointer analysis [6] is used to check whether pointer dereferences can read or write a function pointer value. This is conservative, so any statement determined to not access such a value is safe without runtime bookkeeping.

4.2 Lookup operation on the fp-table

The lookup operation is important to LPCFI’s metadata manipulation. This happens often, especially since checking function pointer callsites requires this search. During initialisation, the fp-table is sorted according to the `func_address` field for efficient searching. Then, a binary search can be performed on the fp-table with the `func_address` field as the key, an $O(\log n)$ operation.

Overhead mainly comes from the `update` helper function due to the search operation on the `fp-table` for assignment instrumentations. Optimisations can be implemented to improve the performance of the search, e.g., caching common searches with a hash map.

4.3 Security guarantee

LPCFI guarantees security at a lower bound of PICFI but reduces the attack surface by removing spurious CFG edges during runtime. Following PICFI, LPCFI only allows an indirect call to target a function whose address has been taken (activated) if such callsite-target edge exists in the static CFG. However, LPCFI places a further restriction: that function pointers hold their last assigned value.

Calling a function pointer after it has been modified outside the standard assignment statements results in a raised assertion because assignment instrumentations are the only way to write to the fp-table, which the check operation relies on. Like PICFI, LPCFI enforces control flow integrity, not data flow integrity [17, 18]. LPCFI does not ensure memory safety for code and data pointers (e.g., the pointers dereferenced in `load/store` statements are unprotected).

5 Proof-of-concept attack and defence

Figure 9 demonstrates LPCFI’s effectiveness over PICFI with a proof-of-concept example in the presence of loops. This is a permission access scenario that allows a user to access a privileged or non-privileged call depending upon the


```

1: #include "privileges.h"
2: /* The header file contains function pointers */
3: /* 'volatile (void)(*priv)(void)' and 'volatile (void)(*nopriv)(void)' */
4: /* for accessing privileged and non-privileged system methods. */
5: int main(void) {
6:     (void)(*op)(void);
7:     char password[7];
8:     while (true) {
9:         fgets(password, 7, stdin);
10:        if (strcmp(password, "secret") == 0) {
11:            lpcfi_assign_copy(op, priv);
12:            op = priv;
13:        } else {
14:            lpcfi_assign_copy(op, nopriv);
15:            op = nopriv;
16:        }
17:        // memory corruption vulnerability: modify the value of op
18:        lpcfi_check(op);
19:        op();
20:    }
21: }

```

Fig. 9: Password verification cope that is safe with LPCFI, but not with PICFI. password entered. This demo (including `extended-demo.c`, `privileges.c`, and `privileges.h`) is publicly available in the `extended-demo` folder in our release.

LPCFI's instrumentation is shown in blue (discussed below). In an infinite loop, a user is prompted for a password. If correct, function pointer `op` is set to function pointer `priv`, a privileged operation. If not, `op` is set to function pointer `nopriv`, a non-privileged operation. Finally, `op` is called and the loop begins anew. A memory vulnerability before the call allows an attackers to modify `op`.

If not instrumented, an attacker may change the value of `op` to any value, and the call will target that location. If the code was instrumented by PICFI, initially, the `op` call is deemed unable to target any location legally. The first time the password is entered incorrectly, the `op` call may reach the value pointed to by `nopriv`. Similarly, the first time the password is entered correctly, the `op` call may reach the value pointed to by `priv`. When *both* possible values have been activated, PICFI will see the `op` call as being able to legally take on either value until program exit. If a user enters the password incorrectly, they may modify the value of `op` to be that of the privileged function pointer, and PICFI will allow this call to be made. This is a problem when a malicious user uses the system after a privileged user.

When the code is instrumented with LPCFI (as shown in blue), this problem is remedied. When `op` is set to `priv` at line 12, the `op` call will only succeed if `op` retains the value it was assigned (`priv`). Similarly, when `op` is set to `nopriv` at line 15, for the `op` call to succeed, `op` must retain its value (`nopriv`). The fp-table is storing a **single** value - the most recently assigned value.

6 Conclusion

This paper presents LPCFI, a new dynamic control flow integrity technique that can protect against attacks undetected when using the monotonically growing CFG used by PICFI. LPCFI achieves a lower bound security guarantee of that

promised by PICFI but reduces the attack surface left by PICFI using a new instrumentation approach and, with a specially designed data structure, ensures that indirect callsites from function pointers can only target at most one function.

References

1. Hovav Shacham. The geometry of innocent flesh on the bone: return-into-libc without function calls (on the x86). In *CCS '07*, pages 552–561.
2. Felix Schuster, Thomas Tendyck, Christopher Liebchen, Lucas Davi, Ahmad-Reza Sadeghi, and Thorsten Holz. Counterfeit object-oriented programming: on the difficulty of preventing code reuse attacks in C++ applications. In *S&P '15*, pages 745–762.
3. Martín Abadi, Mihai Budiu, Úlfar Erlingsson, and Jay Ligatti. Control-flow integrity principles, implementations, and applications. *ACM Trans. Inf. Syst. Secur.*, 13(1):4:1–4:40, 2009.
4. Ben Niu and Gang Tan. Per-Input Control-Flow Integrity. In *CCS '15*.
5. Isaac Evans, Fan Long, Ulziibayar Otgonbaatar, Howard Shrobe, Martin Rinard, Hamed Okhravi, and Stelios Sidiroglou-Douskos. Control Jujutsu: on the weaknesses of fine-grained control flow integrity. In *CCS '15*, pages 901–913.
6. Yulei Sui and Jingling Xue. SVF: Interprocedural static value-flow analysis in LLVM. In *CC '16*, pages 265–266, 2016.
7. Ren Ding, Chenxiong Qian, Chengyu Song, Bill Harris, Taesoo Kim, and Wenke Lee. Efficient protection of path-sensitive control security. In *USENIX Security '17*, pages 131–148.
8. Saravanan Sinnadurai, Qin Zhao, and Weng-Fai Wong. Transparent runtime shadow stack: protection against malicious return address modifications. 2008.
9. Úlfar Erlingsson, Martín Abadi, Michael Vrable, Mihai Budiu, and George C. Necula. XFI: software guards for system address spaces. In *OSDI '06*, pages 75–88.
10. Caroline Tice, Tom Roeder, Peter Collingbourne, Stephen Checkoway, Úlfar Erlingsson, Luis Lozano, and Geoff Pike. Enforcing forward-edge control-flow integrity in GCC & LLVM. In *USENIX Security '14*, pages 941–955.
11. Chao Zhang, Scott A Carr, Tongxin Li, Yu Ding, Chengyu Song, Mathias Payer, and Dawn Song. VTrust: regaining trust on virtual calls. In *NDSS '16*.
12. Dongseok Jang, Zachary Tatlock, and Sorin Lerner. SafeDispatch: securing C++ virtual calls from memory corruption attacks. In *NDSS '14*.
13. Istvan Haller, Enes Göktaş, Elias Athanasopoulos, Georgios Portokalidis, and Herbert Bos. ShrinkWrap: VTable protection without loose ends. In *ACSAC '15*, pages 341–350.
14. Xiaokang Fan, Yulei Sui, Xiangke Liao, and Jingling Xue. Boosting the precision of virtual call integrity protection with partial pointer analysis for C++. In *ISSTA '17*, pages 329–340.
15. Mohamad Barbar, Yulei Sui, Hongyu Zhang, Shiping Chen, and Jingling Xue. Live Path Control Flow Integrity. In *ICSE '18*.
16. Yulei Sui and Jingling Xue. On-demand strong update analysis via value-flow refinement. In *FSE '16*, pages 460–473.
17. Miguel Castro, Manuel Costa, and Tim Harris. Securing software by enforcing data-flow integrity. In *OSDI '16*, pages 147–160.
18. Volodymyr Kuznetsov, László Szekeres, Mathias Payer, George Candea, R Sekar, and Dawn Song. Code-pointer integrity. In *OSDI '14*, pages 147–163.