

# EC-CFI: Control-Flow Integrity via Code Encryption Counteracting Fault Attacks

Pascal Nasahl<sup>\*†</sup>, Salmin Sultana<sup>\*</sup>, Hans Liljestrand<sup>\*</sup>, Karanvir Grewal<sup>\*</sup>,  
Michael LeMay<sup>\*</sup>, David M. Durham<sup>\*</sup>, David Schrammel<sup>†</sup>, Stefan Mangard<sup>†</sup>

<sup>\*</sup>Intel Labs

<sup>†</sup>Graz University of Technology, {firstname.lastname}@iaik.tugraz.at

Fault attacks enable adversaries to manipulate the control-flow of security-critical applications. By inducing targeted faults into the CPU, the software’s call graph can be escaped and the control-flow can be redirected to arbitrary functions inside the program. To protect the control-flow from these attacks, dedicated fault control-flow integrity (CFI) countermeasures are commonly deployed. However, these schemes either have high detection latencies or require intrusive hardware changes.

In this paper, we present EC-CFI, a software-based cryptographically enforced CFI scheme with no detection latency utilizing hardware features of recent Intel® platforms. Our EC-CFI prototype is designed to prevent an adversary from escaping the program’s call graph using faults by encrypting each function with a different key before execution. At runtime, the instrumented program dynamically derives the decryption key, ensuring that the code only can be successfully decrypted when the program follows the intended call graph. To enable this level of protection on Intel® commodity systems, we combine Intel®’s TME-MK with the virtualization technology to achieve function-granular encryption. We open-source our custom LLVM-based toolchain automatically protecting arbitrary programs with EC-CFI. Furthermore, we evaluate EPT aliasing with the SPEC CPU2017 and Embench-IoT benchmarks and discuss and evaluate potential TME-MK hardware changes minimizing runtime overheads.

**Index Terms**—fault attacks, control-flow integrity, encryption

## I. INTRODUCTION

Fault attacks are active, physical attacks where an adversary injects one or multiple faults into a chip. While these attacks originally required physical access to the device under attack, new attack methodologies, such as Plundervolt [27], CLKSCREW [47], or VoltJockey [36], [37], have demonstrated that faults can also be injected remotely in software. The effects of injected faults, which comprise transient bit-flips and permanent stuck-at effects, can be exploited by an adversary to manipulate the control-flow of software [7], [32], [50]. In this scenario, the adversary arbitrarily redirects the control-flow of a program by injecting bit errors into the CPU.

Control-flow integrity [3] is a well-established countermeasure to protect the control-flow from software vulnerabilities, e.g., memory safety vulnerabilities. The goal of this mitigation concept is to detect control-flow deviations from the legitimate

control-flow graph (CFG) of the program. As CFI assumes a software adversary in its threat model, these schemes [21], [24], [26] only protect control-flow edges, such as indirect branches and returns, from control-flow manipulations. However, as the fault attack threat model comprises a broader attack surface, i.e., any control-flow edge including direct branches, these attacks can bypass state-of-the-art CFI countermeasures. In addition to hijacking control-flow edges, faults also enable the attacker to redirect the control-flow at any execution point, e.g., by manipulating the instruction pointer [31], [48], [49].

Dedicated CFI schemes offering protection against faults [29], [33], [38], [40], [55] cover all control-flow edges in their protection. These signature-based approaches maintain a global signature during runtime and compare this signature with the compile time precalculated signature value. On control-flow deviations, the signature check fails and a control-flow attack is detected. This mechanism allows these CFI schemes to verify that the control-flow of a program follows the intended control-flow. However, as the signature checks are only conducted at certain points in the program, control-flow violations are detected with some latency. For example, a fault into the instruction pointer redirecting the control-flow of the program could enable the adversary to still execute security-sensitive code before the signature check detects the violation. Hence, this detection latency can have severe security implications, limiting the practicability of these schemes.

To overcome this detection latency, [8], [52] implicitly conduct the signature check on each executed instruction. Here, the code is encrypted in memory and can only be decrypted when the signature matches the precalculated signature. On a signature mismatch, the instructions are decrypted with a wrong key, yielding garbled instructions which, with a high probability, trigger an exception. However, these schemes currently require intrusive hardware changes in the processor’s pipeline, which makes it hard to deploy them on a larger scale.

Hence, to protect the control-flow of software against fault attacks, new countermeasures achieving minimal detection latencies without intrusive hardware changes on commodity systems are required.

## Contribution

This paper introduces EC-CFI, a cryptographically enforced control-flow integrity scheme designed to counteract fault

attacks aiming to redirect the control-flow of programs outside of their call graph. In EC-CFI, each function is encrypted with a different encryption key before the program’s execution. At runtime, EC-CFI-instrumented programs dynamically derive the active decryption key before each control-flow edge, i.e., direct or indirect function calls. This derivation produces the correct decryption key only if the control-flow matches the statically determined call graph that was used to derive the encryption keys at load-time. When a fault redirects a control-flow edge to another function outside of the call graph, the code is decrypted with the wrong key, which can be immediately detected with a high probability. Moreover, the protection of EC-CFI comprises not only control-flow edges, any redirection to other functions, e.g., by instruction pointer manipulations, can be mitigated.

To enable this level of protection on recent commodity Intel® platforms without hardware modifications, we utilize the *total memory encryption - multi key* (TME-MK) feature for the function encryption. However, as Intel®’s TME-MK so far is only used for page-granular memory encryption, which is too coarse-grain for function encryption, we introduce a new concept based on extended page table (EPT) aliasing. This mechanism allows us to leverage TME-MK for fine-granular, in the case of EC-CFI, function-granular, memory encryption. Moreover, our approach based on EPT aliasing, which is a combination of Intel®’s virtualization technology (VT) and TME-MK, enables us to frequently switch the key used for encryption and decryption.

We showcase how to implement EC-CFI using the generic EPT aliasing approach and introduce a prototype implementation. We open-source our custom LLVM toolchain, which is responsible for automatically instrumenting programs with the key derivation mechanism without any user interaction. Furthermore, we measure the performance impact of EPT aliasing on a recent Intel® CPU using the SPEC CPU2017 and Embench-IoT benchmarks. Finally, we discuss potential minimal-invasive hardware changes decreasing the runtime overhead of EC-CFI.

In summary, our contributions are:

- We present a CFI scheme that is designed to hinder a fault adversary from escaping the call graph of a protected program by encrypting each function with a different encryption key. By dynamically deriving the decryption key at runtime, the code of a function can only be successfully decrypted if it is reached by following the static call graph used to encrypt the function.
- We introduce a fine-granular encryption approach for recent Intel® platforms based on EPT aliasing consisting of a novel combination of TME-MK and VT. This approach enables us to achieve function-granular encryption and to use different encryption keys for different functions without hardware changes.
- We showcase how to implement EC-CFI with EPT aliasing on recent Intel® platforms. Here, we open-source our LLVM-based toolchain capable of automatically protecting programs with EC-CFI.

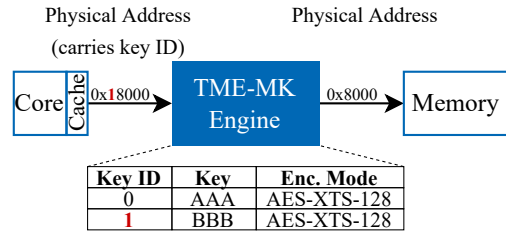


Fig. 1: Overview of the TME-MK engine.

- We evaluate the performance impact of our EPT aliasing approach and analyze security benefits of EC-CFI.
- Finally, we discuss minimal TME-MK hardware changes and showcase that these changes minimize the runtime overhead of EC-CFI.

## II. BACKGROUND

### A. Signature-based Control-Flow Integrity

Signature-based control-flow integrity schemes aim to detect fault-induced control-flow manipulations at a certain granularity. The main idea of these schemes is to check whether the control-flow of the executed program follows the control-flow statically extracted at compile time. On a mismatch, an attack manipulating the control-flow is detected.

To implement this concept, these CFI schemes [11], [15], [22], [33], [38], [40], [51], [54], [55] assign each code block at the protection granularity, e.g., function or basic-block granularity, a unique identifier at compile time. Moreover, the executed program is instrumented with routines responsible for updating a global signature  $S$  on each control-flow transfer at this granularity. As this update function is accumulative, the entire execution history of the program is stored in a compressed form in this signature. By comparing the signature derived at runtime with the signature defined at compile time, control-flow hijacks are detected.

However, as these control-flow checks are costly, they are only placed at certain locations. For example, FIPAC [40] performs these checks at the end of each basic-block, function, or before exiting the protected program. Hence, depending on the chosen checking policy, the attacker can still execute code before the control-flow manipulation is detected.

To overcome this detection latency, schemes such as SOFIA [8] and SCFP [52] implicitly perform this check using code encryption. However, these approaches require intrusive hardware changes in the CPU pipeline.

### B. Intel TME-MK

Total memory encryption (TME) [17] is a feature provided on recent Intel® CPUs allowing the system to transparently encrypt all data passed from the CPU to the external memory. By using a single secret encryption key, TME is capable of preserving data confidentiality in different threat scenarios, e.g., cold-boot attacks [14].

Intel® TME-MK [20] is an extension allowing the system to use multiple keys to encrypt data. As shown in Figure 1, the encryption engine for TME-MK resides between the caches

and the memory controller and uses AES-XTS with either 128- or 256-bit keys. Internally, the engine consists of a table containing a mapping from key identifier to encryption key. On each memory request, i.e., read or write, the key identifier is embedded into the upper bits of the physical address, which are usually not used. By using different key identifiers, different pages can be encrypted or decrypted with different encryption keys. To define which key is used, software can set the key identifier in the page table entry (PTE) of a page. On address translation, the key identifier is then automatically set in the physical address. Using this approach, TME-MK can provide page-granular encryption. One use case of TME-MK is the cryptographic isolation of different virtual machines on a host system.

### C. Intel Virtualization Technology

Intel® virtualization technology (VT) [18] is a set of features allowing the processor to efficiently and securely share computing resources among different workloads. One key feature is the hardware-based second level address translation mechanism allowing each guest to have its own virtual address space. Here, the guest system is responsible for the first level address translation, i.e., guest linear addresses (GLA) to guest physical addresses (GPA), by using page tables. For each guest, the host then provides a mapping from guest physical addresses to host physical addresses (HPA) using extended page tables (EPTs). The `vmfunc` instruction allows the guest to set the current active EPT from an extended page table pointer (EPTP) list stored in the virtual machine control structure (VMCS).

## III. THREAT MODEL

In our threat model, we consider an adversary capable of injecting a targeted fault into the processor or the external memory. We assume that this fault is either injected remotely, e.g., by using Plundervolt [27] or CLKSCREW [47], or locally, e.g., by using laser fault injection. The goal of the adversary is to redirect the control-flow of a program outside of the call graph of the corresponding program. Figure 2 depicts the presumed attack scenario. In the illustrated call graph, function *A* can call function *B* and function *B* can either call function *A* or *C*. During the execution of function *A*, the attacker injects a fault redirecting the control-flow from *A* to *C*.

A fault attacker can redirect the control-flow outside of the call graph by either targeting the **control-flow edges** between functions, flipping bits in any other **instruction**, or manipulating the **instruction pointer** of the CPU. For the **control-flow edges** between functions, the attacker can target direct or indirect branches. To manipulate the execution of indirect calls, a fault attacker can flip bits in addresses stored in registers used by these calls. Furthermore, the adversary also can manipulate the address used by direct calls by injecting a fault into the address generation unit (AGU) of the CPU. Moreover, by flipping bits in the program memory of the application, addresses of direct calls or the registers used by indirect calls [12] can be manipulated.




Fig. 2: Call graph with manipulated control-flow.

In addition, the attacker can also flip bits in any **instruction** of the program in such a way that the control-flow is redirected, e.g., the opcode is changed to a branch [31]. Finally, a redirection of the control-flow also can be performed by injecting faults directly into the **instruction pointer** of the CPU [13], [49]. In summary, this attacker model is stronger than threat models used by traditional CFI targeting a software-only adversary, where only indirect branches and returns are considered to be vulnerable.

For our work, we exclude side-channel and microarchitectural attacks and assume that the operating system and the hypervisor are trusted by the system.

## IV. DESIGN

EC-CFI aims to hinder an adversary from redirecting the control-flow to arbitrary points in the program by encrypting each function with a different encryption key at load-time. At runtime, EC-CFI restricts the set of callable functions for the current execution context to the set of call targets defined in the call graph by dynamically deriving the decryption key. When the attacker redirects the control-flow to a function outside of the call graph, the encrypted code is decrypted with a wrong key. As this decryption yields garbled code, the instruction decoding fails with a high probability. Although it could be possible that decrypting an instruction with an invalid key could produce a valid instruction, the likelihood of decrypting multiple instructions correctly is low [4]. Hence, EC-CFI is capable of detecting control-flow manipulations with no or minimal detection latency. EC-CFI achieves this level of protection on recent Intel® commodity hardware by combining a signature-based control-flow integrity scheme with fine-granular memory encryption.

### A. Fine-Granular Memory Encryption

EC-CFI encrypts each function *F* with a different encryption key  $K_F$  using Intel®’s TME-MK memory encryption engine. However, as highlighted in Section II-B, in the intended usage mode, TME-MK only provides the possibility to encrypt entire memory pages (e.g., 4 kB pages) with different encryption keys. Although increasing the code sizes of functions to page sizes would enable the processor to encrypt each function with a different key, this approach would also significantly increase the memory overhead.

To overcome this limitation, we introduce a novel fine-granular memory encryption approach based on a combination of TME-MK with the extended page table (EPT) feature

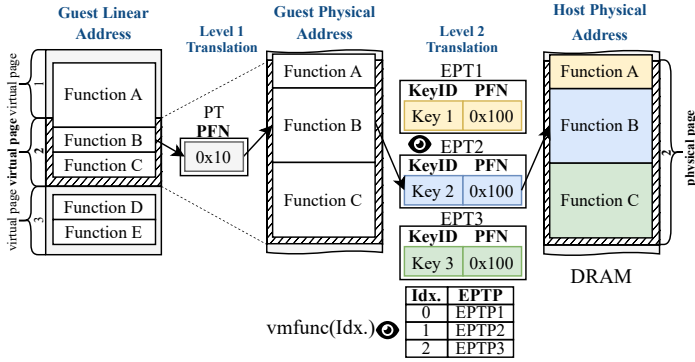


Fig. 3: EPT aliasing combined with memory encryption for fine-grained memory encryption.

of Intel® VT. Hereby, EC-CFI achieves sub-page granular memory encryption by combing **EPT aliasing** with memory encryption. With this approach, the encryption granularity is only limited by the encryption primitive, e.g., 128 bit block size for AES. Such a small encryption granularity was previously only possible using custom CPU designs [30], [45].

Figure 3 illustrates the core idea of EPT aliasing combined with memory encryption based on an example with three functions A, B, and C located inside the 4 kB virtual page 2. The first level address translation mechanism translates the guest linear addresses (GLA) of functions A, B, and C to the guest physical addresses (GPA) using the page frame number (PFN) of the page table (PT). In our example, a PFN of  $0x10$  is used to translate the addresses. Now, our approach based on EPT aliasing establishes separate extended page tables ( $EPT1$ ,  $EPT2$ , and  $EPT3$ ) for each encryption domain using a different key, i.e., key 1 for function A, key 2 for function B, and key 3 for function C. In the EPT entries of these EPTs, the guest physical to host physical address (HPA) mapping is identical, i.e.,  $EPT1$ ,  $EPT2$ , and  $EPT3$  use the PFN  $0x100$  for functions A, B, and C. However, the key identifier fields in the EPT entries are different, i.e., key 1 for  $EPT1$ , key 2 for  $EPT2$ , and key 3 for  $EPT3$ .

This approach allows us to have different views (👁️) on the memory by switching the current, active extended page table. For example, when  $EPT2$  is active (👁️), the GPA of function B is translated by the second level address translation mechanism to the HPA with the address translation information stored in the entries of  $EPT2$ . As the key identifier key 2 is embedded into the upper bits of the HPA during the address translation, TME-MK now encrypts or decrypts function B with the key assigned to this key identifier. Note that for the actual physical memory access, the key identifier bits are stripped from the physical address. When accessing function C with  $EPT2$ , which was encrypted with key identifier key 3 in the  $EPT3$  memory view, only garbled code is retrieved as the wrong decryption key 2 is used for the access.

To switch between these EPTs, the extended page table pointer (EPTP) that specifies the active EPT can be changed.




Fig. 4: Signature init, update, and key switch for a direct call.

Such an EPTP switch is initialized with the `vmfunc` instruction. By passing the EPTP index, e.g., 0, 1, or 2, as shown in Figure 3, to this instruction, the CPU switches the EPTP to the corresponding EPTP in the configured EPTP list.

As shown in Figure 3, EC-CFI does not restrict the locations of functions in memory, i.e., multiple functions inside of a page or functions occupying multiple pages are supported.

### B. Signature-Based Control-Flow Integrity Scheme

EC-CFI uses a signature-based control-flow integrity approach to automatically derive the decryption keys for each encrypted function at runtime. In our scheme, a random signature  $S_F$  is assigned to each function  $F$  and the current, active signature is stored in the global signature register  $S$ .

EC-CFI uses the approach based on **EPT aliasing** (cf. Section IV-A) for fine-grained encryption of code blocks. For each encryption domain, EC-CFI initiates a separate EPT with a different encryption key embedded into the extended page table entry. As each encryption key only is used in one EPT, we have a bijective mapping  $EPTP \mapsto K$ . EC-CFI now passes the signature  $S$  to the `vmfunc` instruction to select the active EPT and, therefore, the current encryption key, i.e.,  $S \mapsto EPTP \mapsto K$ . Note that the signature  $S$  is not a signature in the cryptographic sense, instead, it is the index (cf. Figure 3) to the extended page table pointer (EPTP), which points to an extend page table.

EC-CFI consists of three major runtime primitives: (i) signature init, (ii) switch key, and (iii) signature update. At the start of the program, the signature register is initialized (i) with the signature of the entry function. Then, EC-CFI activates the key for decrypting the entry function by switching (ii) the EPTP to the EPT containing the corresponding key. Due to the bijective mapping from the signature to the key over the EPTP, the signature  $S$  automatically selects the correct key and the function can be decrypted. During the program’s execution, the current signature is updated (iii) before each control-flow transfer to a different function, i.e., on direct or indirect calls.

$$S = S \oplus C \quad (1)$$

Equation 1 shows the used accumulative update function. The compiler selects the position-dependent constant  $C$  for the call in such a way that the resulting signature matches the signature of the called function. After updating the signature, the key for the called function is activated by switching (ii) the EPTP




Fig. 5: Security implications when the call sites A and C derive an identical key for the multi-call target B.

using the current  $S$ . When the derived signature matches the signature of the call target, the function can be successfully decrypted. After returning from the callee, the signature is again updated and the key is switched such that the code of the caller can be decrypted.

Figure 4 depicts the signature init (Line 1), switch key (Lines 3 and 7), and signature update (Lines 2 and 6) required by EC-CFI to correctly derive and switch the key for both functions. Hereby, the color highlights the corresponding code encrypted with the different keys  $K_{Main}$  and  $K_A$ . Due to the mapping  $S \mapsto K$ , correctly deriving  $S$  and calling the intended function allows the CPU to successfully decrypt the code with key  $K$ . Note that the key immediately becomes active after executing the key switching routine. Therefore, the instruction calling function A (Line 4) already needs to be encrypted with the key for this function.

For indirect calls, accurately determining the caller target at compile time is not possible. Hence, EC-CFI determines the possible set of call targets, which then share an encryption key. To ensure that the same key is derived, EC-CFI induces signature collisions, i.e.,  $C$  is accordingly chosen to derive the same signature  $S$  for different indirect calls.

1) *Multi-Call Targets*: Assigning multi-call targets, i.e., functions that can be called from multiple other functions, an identical encryption key enables the adversary to escape the call graph. Figure 5 describes the security implications of deploying a shared key for the multi-call target B. By inducing a fault during the execution of this function, the adversary can redirect the control-flow either to A or C, independently from the original call site. EC-CFI mitigates this security weakness by adapting the concept of call headers introduced in [40].

Figure 6 shows our approach of securely handling multi-call targets using call headers. Each function is assigned the corresponding signature, i.e.,  $S_A$ ,  $S_B$ , and  $S_C$  and these functions are encrypted with the corresponding key. Furthermore, a call header encrypted with a distinct key, i.e.,  $S_{AH} \mapsto K_{AH}$  and  $S_{CH} \mapsto K_{CH}$ , is added to the multi-call target function B. Before calling function B, the key is switched to this call header key. Then, the function is called and the execution flow is redirected to the corresponding call header. Inside this header, the key is updated to the key of the called function, i.e.,  $S_B \mapsto K_B$ . Additionally, a return constant  $C_{Ret}$  for each header is set. When returning from the function, this constant is used to switch the key back to the call header key. This ensures that the program only can return to the original call site. After the call instruction, the key is switched back to the signature  $S_A$  or  $S_C$  of the corresponding function.

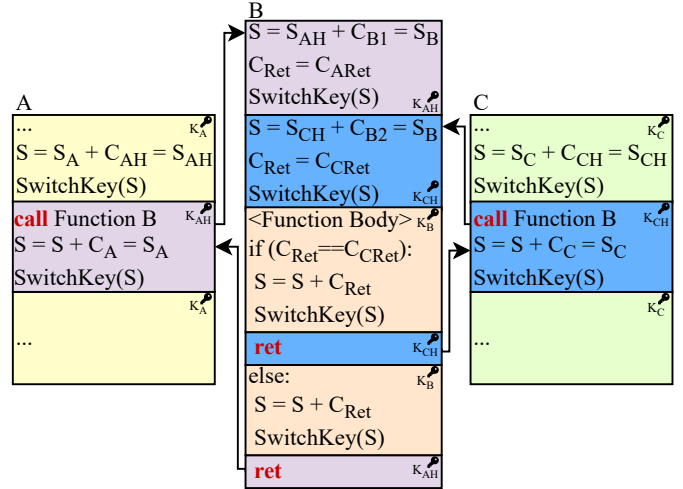


Fig. 6: Secure handling of multi-call targets.

For indirect calls, the headers of the possible set of call targets share a common signature, i.e., they are encrypted with the same key.

## V. IMPLEMENTATION

The prototype implementation of EC-CFI consists of three major building blocks (cf. Figure 7). The (i) compiler is responsible for instrumenting binaries, the (ii) hypervisor provides multiple EPTs, and the (iii) loader uses the hypervisor and metadata provided in the instrumented binary to encrypt each code block with a different key before execution.




Fig. 7: Overview of our EC-CFI prototype implementation.

### A. Compiler

To automatically protect programs without user interaction, we integrate EC-CFI into a custom LLVM-based toolchain [23]. Our backend pass of the custom toolchain is responsible for 1) assigning signatures to functions, 2) instrumenting calls, 3) inserting call headers, and 4) aligning the code blocks to the cache line size.

1) *Signature Assignment*: The first step the compiler conducts is the assignment of the signatures  $S_F$  to all functions in the program. Here, the compiler chooses a random ID between  $S_{LOW}$  and  $S_{HIGH}$  for each function and stores this information into a compiler-internal structure. The signature range is configured by the user compiling a program and needs to reflect the number of available TME-MK key identifiers and EPTs of the targeted processor.

```

1  push %rax          # Save rax & rcx to stack.
2  push %rcx
3  xor  %rax, %rax    # Set rax to 0.
4  mov  %r13, %rcx   # Move signature to rcx.
5  vmfunc            # Switch EPTP.
6  pop  %rcx         # Restore rax & rcx from
7  pop  %rax         # stack.

```

Listing 1: Key switch instruction sequence.

Afterward, the compiler defines a constant used to initialize the signature register with the chosen  $S_F$  in the program’s entry point. Hereby, the signature is moved into the signature register and the `key_switch` routine instructions, which are shown in Listing 1, are inserted. This routine first preserves the content of registers `rax` and `rcx` by pushing them on the stack, sets up the arguments and invokes the `vmfunc` instruction, and restores `rax` and `rcx` from the stack. The argument `rax = 0` for `vmfunc` instructs the CPU to switch the EPTP to the EPTP specified in `rcx = S`. Note that the compiler reserves the callee-saved register `r13` exclusively for the EC-CFI signature.

2) *Call Instrumentation*: As functions are encrypted with different encryption keys, the correct decryption key needs to be in place when calling functions. Our toolchain finds all direct and indirect calls and calculates the constant  $C = S_{Current} \oplus S_{Target}$ . For direct calls, the target signature  $S_{Target}$  is the signature of the call header of the corresponding function. As an indirect call can have multiple possible call targets, a points-to analysis is needed to reveal these targets. For external function calls into unprotected programs, e.g., shared libraries, a default target signature using the TME-MK default encryption key is used.

```

1  xor  $C, %r13      # Update signature S=S⊕C.
2  key_switchRoutine # Switch the key.

```

Listing 2: Call prologue and epilogue.

Then, right before the call instruction, the compiler inserts the call prologue. As shown in Listing 2, this prologue consists of the signature update, i.e., XORing the constant  $C$  to the current signature  $S$ , and the `key_switch` routine (cf. Listing 1). After the call instruction, the identical instruction sequence, i.e., the call epilogue, is inserted to switch back to the key of the caller function.

3) *Call Headers and Footers*: To handle multi-call targets (cf. Section IV-B1), EC-CFI inserts call headers in front of each function.

```

1  call_header_1:
2  xor  $C, %r13      # Update signature S=S⊕C.
3  mov  $Cret, %r14   # ret_c = Cret.
4  key_switchRoutine # Switch the key.
5  jmp  $function_body # Jump to function begin.
6  call_header_2:
7  ...
8  function_body:
9  ...

```

Listing 3: Call header.

As illustrated in Listing 3, the call header first updates the signature with a constant to match the signature of the function body. Then, the return constant  $C_{ret}$  is loaded into the reserved `r14` register and the key for the function body is activated. A jump to the function body jumps over the call headers of other callees. In the callee, the compiler rewrites the addresses of calls to point to the corresponding call header.

```

1  xor  %r14, %r13    # Update signature S=S⊕ret_c.
2  key_switchRoutine # Switch the key.

```

Listing 4: Call footer.

Before each return in the function body, the modified compiler adds, for each call header, the call footer instructions shown in Listing 4. These instructions update the signature with the return constant such that the signature is identical to the signature in the call header.

4) *Code Block Alignment*: The `key_switch` routine (cf. Listing 1) switches the EPTP and, therefore, the current, active decryption key. Hence, as the key is immediately switched after the `vmfunc` instruction, the next fetched instruction is already decrypted with this key. To avoid that a cache line contains data encrypted with different encryption keys, which would trigger a cache miss and require a costly additional memory fetch, our toolchain ensures that `vmfunc` instructions are aligned to the end of a cache line. Note that this alignment needs to be done in the call prologues and epilogues as well as in the call headers and footers.

## B. Hypervisor

The hypervisor is responsible for setting up the EPT aliasing functionality and providing an interface for the binary loader to run protected programs.

1) *System Setup*: When booting the system, the hypervisor puts the operating system into the guest mode and creates a virtual machine control structure (VMCS). Furthermore, the hypervisor creates three default and  $NUM\_PROT\_EPTS$  EPTs and stores the pointer to them into the EPTP list of the VMCS. Note that  $NUM\_PROT\_EPTS$  is limited by the number of available TME-MK key identifiers and the number of EPTPs which can be stored in the EPTP list. In our prototype implementation, the hypervisor exclusively uses  $EPT0$ , the kernel  $EPT1$ , and the user mode  $EPT2$ . The remaining  $NUM\_PROT\_EPTS$  EPTs are utilized by protected programs. In the initialization phase, i.e., before starting a protected program, all of these EPTs are identical and use the default 0 TME-MK key identifier in the EPT entries.

2) *Setup of Protected Programs*: By using the `vmcall` instruction, the binary loader communicates with the hypervisor to configure EPT aliasing before starting a program. Here, the loader uses this interface to register the program and the used code pages to the hypervisor. The hypervisor uses this information, i.e., page address and size, to set the TME-MK key identifiers in the entries of the  $NUM\_PROT\_EPTS$  EPTs. To enable data sharing between functions, only encryption keys for code pages are set. For data pages, the key identifier

field in the EPT entries for all EPTs contains the default 0 key. When calling external functions, the compiler ensures that the EPTP is switched to the default user mode *EPT2*.

3) *Termination of Protected Programs*: After the execution of a protected program, a call to the hypervisor is used to deregister the program. Hereby, the hypervisor resets the TME-MK key identifier field in the EPT entries to the default 0 key.

4) *User and Kernel Mode Switches*: When switching from user mode to the kernel, the hypervisor needs to save the current, active EPTP, i.e., *EPT2* for unprotected programs and *EPT2* to *EPT2 + NUM\_PROT\_EPTS* for protected programs, and switch to the kernel *EPT1*. This saved EPTP is restored when switching back from kernel to user mode. To provide the hypervisor with an opportunity to perform these EPTP switches, the current prototype implementation triggers an EPT violation on each switch between user and kernel using Mode-Based Execution Control (MBEC) by marking pages in user views as only user-executable and pages in the kernel view as only supervisor-executable.

### C. Binary Loader

The modified binary loader is responsible for loading code blocks of an instrumented binary into memory and starting the program. An instrumented binary generated with our custom compiler contains metadata, i.e., address, size, and key identifier for each code block. The loader first allocates a page for the code using this metadata and registers this code page to the hypervisor. Now, the EPT entries of all EPTs for this code page contain the key identifiers specified in the program metadata. To encrypt code blocks with their corresponding key identifier, the loader first switches the EPTP with the `vmfunc` instruction to the EPT tagged with the key identifier. Then, the code is copied from the binary to the memory encrypting the code block with the key identifier of the current, active EPT. This procedure is repeated for each code block but with a different key. Finally, the loader activates the EPT of the program's entry point and passes execution to the application.

## VI. SECURITY DISCUSSION

This section discusses security benefits of EC-CFI in respect of the threat model introduced in Section III.

### A. Flipping Address Bits

To redirect the control-flow of a program outside of the call graph, the fault attacker can induce bit-flips into control-flow related addresses. These addresses comprise the instruction pointer `rip` and addresses stored in memory or registers and used by indirect calls. For direct calls, the adversary can induce a fault into the relative address encoded into the instruction, which is then translated to an address by the address generation unit. Here, in EC-CFI, the fault can affect guest linear, guest physical, and host physical addresses. The attacker could aim to redirect the control-flow to any point in the program by injecting faults into guest linear or guest physical addresses. However, when the current active decryption key does not match the encryption key for this

point in the program, the execution of these instructions fails. By manipulating both the address and the key identifier in the HPA, the attacker could redirect the control-flow. Nevertheless, this attack vector is hard to exploit, i.e., precisely manipulating both fields is challenging, and the effect is limited. More specifically, executing a single instruction could be possible when redirecting the control-flow by manipulating the address and the key identifier in the HPA. However, as the bit-flip in the key identifier field of the HPA is not permanent for transient faults, the key identifier of the subsequent instruction again is determined by the current EPT. Hence, the decryption of this instruction then fails.

### B. Manipulating EPT Entries

The attacker could try to permanently change the key identifier for an address region by manipulating these bits in the corresponding EPT entries. However, as TME-MK always encrypts the entire external memory using the default key identifier, also the EPTs stored in memory are encrypted. Hence, deterministically flipping key identifier bits in EPT entries without knowing the secret key is not possible. By targeting the translation lookaside buffer (TLB), the attacker could forge the key identifier used for addresses as long as the TLB entry is valid. Here, additional countermeasures, e.g., error detection or correction checks [41], could be added.

### C. Leaking Key Identifiers

When the attacker is capable of leaking key identifiers, control-flow manipulations with two precise faults can be possible. Here, the attacker would need to manipulate the key identifier in the EPT entry to the leaked identifier of the target function and redirect the control-flow to this function. However, as controlling a fault, i.e., timing and location, is extremely challenging on a complex Intel® CPU, the probability of successfully inducing two subsequent faults is low. Moreover, as the control-flow signature was not changed, it no longer matches the predefined signature. Therefore, the wrong decryption key is used at the next call instruction.

### D. Key Space

Ideally, each function in EC-CFI is encrypted with its own encryption key. Then, redirecting the control-flow to any other function outside of the call graph deterministically fails. However, the encryption key space is limited by the available key identifiers as well as the number of available extended page tables. According to the Intel® manual [20], in total, TME-MK supports up to  $2^{15}$  different key identifiers. However, as our EPT aliasing approach requires us to have multiple extended page tables, the encryption key space is also determined by the number of available EPTs. Currently, the `vmfunc` instruction allows the system to switch between 512 different EPTPs [18]. Hence, when there are more functions in a program than available EPTPs, TME-MK key identifier collisions can occur.

Note that the actual TME-MK key identifier space implemented by the platform could be smaller than the technical

upper limit of  $2^{15}$  different identifiers in the TME-MK engine. When the key identifier space is smaller than the limit of available entries in the EPTP list, i.e., 512, the following key assignment strategy could be used: The hypervisor assigns each EPT a random key identifier. As some EPTs share the same key identifier, the attacker could redirect the control-flow to other functions and successfully execute code. However, as the signature is accumulatively updated independently from the EPT key identifier, the next derived signature does not match the signature of the next called function. Hence, with a high probability, at this point, the control-flow manipulation can be detected by EC-CFI.

#### E. Decrypting Instructions with an Invalid Key

The instruction length in x86-64 is between 1 and 15 B and the opcode can utilize 1 to 3 B in an instruction. To form a valid instruction, both the opcode as well as the other bytes in the instruction need to be valid. Depending on the density of the x86-64 instruction set, which is hard to determine [9], it is possible to retrieve a valid instruction when using an invalid decryption key. Nevertheless, the security impact of decrypting an instruction with a wrong key is minimal due to two reasons. First, the attacker’s goal is to execute a specific instruction and not just a random one. Although some instructions could have multiple opcodes, e.g., `0x00` and `0x01` for an `add`, the remaining decrypted bytes of the instruction are either invalid, causing an instruction fetch failure, or change the behavior of the program. Second, while it might be possible that a single instruction was correctly decrypted, the probability that the subsequent instruction also is valid, is very low. Note that an encryption engine also providing integrity, such as used by Intel® TDX [19], could immediately detect decryption attempts with the wrong key.

#### F. Intra-Function Control-Flow Attacks

Control-flow hijacks within a function, e.g., skipping instructions, cannot be mitigated with the current protection granularity used by EC-CFI. This is in line with our threat model defined in Section III. However, as EC-CFI is a generic concept and not bound to the function-level protection granularity, future work could aim to encrypt code blocks at a finer granularity.

#### G. Control-Flow Attacks within the Call Graph

Similar to related work [8], [33], [38], [40], [53], EC-CFI aims to prevent control-flow manipulations outside of the call graph and not within the borders of the call graph. This is an inherent characteristic of CFI schemes as the compiler cannot exactly determine the targets of indirect calls [2]. When targeting conditional branches or data used by these instructions, EC-CFI, prevents control-flow redirections outside of the call graph. To mitigate redirections within the call graph, i.e., from one branch target to the other, orthogonal countermeasures [40], [42] are needed.

#### H. Shared Libraries

As shared libraries need to be accessible for unprotected programs, they are encrypted with the systemwide default 0 key identifier. To avoid that a fault attacker manipulates calls to external functions in libraries, programs can be statically linked, i.e., the libraries are then part of the binary and are, therefore, also protected. This protection behavior is in-line with related CFI schemes [33], [38], [40], [53].

### VII. PERFORMANCE EVALUATION

In this section, we first evaluate the code size overhead of protecting the Embench-IoT and SPEC CPU2017 benchmarks against fault-based control-flow manipulations using EC-CFI. Then, we analyze the runtime overheads of these benchmarks when using our extended page table aliasing approach. Here, our focus is on evaluating the impact of switching the extended page tables on the translation lookaside buffers (TLBs). We conduct our experiments without enabled TME-MK as the expected performance impact of the memory encryption is small and as we currently do not have access to a system supporting TME-MK for the performance evaluation. According to Intel® [6], TME-MK induces a performance impact of less than or equal to 2.2% for certain workloads.

#### A. Code Size Overhead

To measure the code size overhead of EC-CFI, we compiled the C-based SPEC CPU2017 [44] benchmarks without OpenMP support using our custom LLVM-based toolchain. We compiled all benchmarks twice, i.e., in the protected and unprotected mode, with identical compilation flags and enabled the `-O3` optimizations. Similarly, we compiled the Embench-IoT [35] benchmark with our custom toolchain. Then, we compared the code sizes of the protected binaries to the unprotected binaries with the GNU `size` utility.

The measured code size overhead for SPEC CPU2017 is between 61.36% and 143.89% with a geometric mean of 82.87%. For Embench-IoT, we measured a geometric mean of 22.93% for the code size overhead.

In general, the code size overhead consists of three parts: The (i) call headers and footers increase the code size for each function in the program. Similarly, EC-CFI adds (ii) a call epilogue and prologue responsible for switching the key before and after each direct and indirect call. Finally, the (iii) alignment of the code blocks and `vmfunc` instructions to cache lines increases the code size of protected programs as EC-CFI performs this alignment with `nop` instructions.

#### B. Runtime Overhead

To measure the performance impact of switching the extended page table pointers, and therefore the view on memory with the extended page tables, with `vmfunc`, we use the instrumented and uninstrumented binaries generated for the code size evaluation in Section VII-A. Here, we executed both versions of the binaries on an Intel® CPU supporting the VT building-block of EC-CFI without TME-MK.






Fig. 8: Runtime overhead for SPEC CPU2017.




Fig. 9: Runtime overhead for Embench-IoT.

Figure 8 illustrates the percentual runtime overhead of the instrumented binaries relative to the uninstrumented baseline. Here, we measured a performance impact between  $\times 1.18$  and  $\times 27.05$ , and a geometric mean of  $\times 6.63$  for SPEC CPU2017.

Furthermore, we compiled the Embench-IoT [35] benchmark with our custom toolchain and measured the number of cycles with `rdtsc` [34]. Figure 9 highlights the runtime overheads for the Embench-IoT benchmarks. When averaging the cycle count over 10000 runs and comparing it to the baseline without instrumentation, we measured a geometric mean for the runtime overhead of  $\times 5.97$ .

### C. TLB Misses

The EPT aliasing approach requires us to frequently switch the view on memory by switching the EPTP using the `vmfunc` instruction. This switching negatively affects the hit rate of the translation lookaside buffers (TLBs) as the address translation information stored inside these buffers is tagged with the EPTP [18]. Moreover, according to the Intel® manual [18], an EPT violation also invalidates the TLB entries associated with the current EPTP. Hence, EPT aliasing increases the pressure on the TLBs. Figure 10 depicts the TLB misses for the Embench-IoT benchmarks. Here, we measured with the `perf` tool a geometric mean of  $\times 36.46$  for the data TLB load misses,  $\times 7.08$  for the data TLB store misses, and  $\times 24.02$  for the instruction TLB load misses.

## VIII. TME-MK HARDWARE CHANGE

A minimal-invasive hardware change altering the TME-MK mode of operation can minimize the performance impact of our encryption-based control-flow integrity scheme. Currently, as described in Section II, the TME-MK engine leverages the upper bits of the physical address as the key identifier bits. Hence, to encrypt data with different keys, the identifier needs to be set in the page table or extended page table entries, which requires techniques such as EPT aliasing used in this paper.

In our proposed hardware change, the TME-MK engine retrieves the key identifier from a user-accessible key identifier register instead from the upper bits of a physical address. The key associated with the key identifier can be rapidly switched by writing to that register. Hence, EC-CFI could be implemented without the EPT aliasing approach, significantly

reducing the runtime overhead. With this proposed hardware change, EC-CFI does not induce any additional TLB pressure. Moreover, the key identifier space is no longer limited by the available bits in the physical address, i.e., up to 15 bit, and therefore could be increased to 32 bit.

To measure the runtime overhead of EC-CFI with this hardware modification, we emulated the key switch routine by replacing all `vmfunc` instructions in a protected program with a write to a register. As shown in Figure 11, the runtime overhead of EC-CFI for the SPEC CPU2017 benchmark is significantly lower than with the EPT aliasing approach. More specifically, we measured a runtime overhead between  $\times 1.02$  and  $\times 1.51$  and a geometric mean of  $\times 1.15$ . Similarly, as shown in Figure 12, the proposed hardware change minimizes the runtime overhead of the Embench-IoT benchmarks protected with EC-CFI to a geometric mean of  $\times 1.21$ . We measured a code size overhead between 63.99% and 151.73% and a geometric mean of 86.45% for SPEC CPU2017 and a geometric mean of 24.27% for Embench-IoT.

## IX. RELATED WORK

Control-flow integrity [3] is a generic countermeasure that also can be used to protect programs against software attacks. Here, these schemes assume that the adversary manipulates the control-flow by overwriting control-flow related addresses, such as function pointers or returns, by exploiting a memory safety vulnerability [46]. To mitigate this threat, CFI schemes [3], [21], [24], [26] aim to maintain the integrity of these addresses. However, as the underlying threat model of these countermeasures is weaker (cf. Section III), they can only provide limited protection against fault-induced control-flow hijacks. More specifically, contrary to a software adversary, a fault attacker can also manipulate direct calls and flip bits in the instruction pointer. Hence, even in the presence of a CFI scheme mitigating software attacks, fault attackers can still manipulate the control-flow.

Therefore, dedicated CFI schemes aiming to protect against faults are commonly used. These schemes [29], [33], [38], [40], [55] derive a signature and, in contrast to EC-CFI, explicitly compare this signature to the signature defined at compile-time. Hence, depending on the location of these checks, an adversary capable of redirecting the control-flow still can execute some instructions before the control-flow manipulation is detected. Although within a protection domain, i.e., intra-function, EC-CFI provides similar protection, the protection across function boundaries is stronger. More specifically, when the attacker redirects the control-flow to a function encrypted with a different key, the execution immediately fails. In other schemes, such as FIPAC [40], the attacker still can execute instructions until the signature is checked, e.g., at the end of a function. Moreover, EC-CFI, with the minimal hardware change, performs similar in terms of runtime overhead than FIPAC, i.e., 15% for EC-CFI and 22% for FIPAC (function end checking policy) for the SPEC CPU2017 geometric mean.

Similar to EC-CFI, SCFP [52] also implicitly conducts the signature checks by using code encryption. However, SCFP




Fig. 10: TLB misses for Embench-IoT.




Fig. 11: Emulated runtime overhead for SPEC CPU2017.




Fig. 12: Emulated runtime overhead for Embench-IoT.

adds a dedicated pipeline stage between the instruction fetch and decode stage for the protection. Moreover, dedicated instructions are added to the instruction set to interact with this pipeline stage. We argue that integrating such intrusive hardware changes into the pipeline of a complex Intel® CPU are not feasible as such changes negatively affect area and power consumption as well as they add complexity to the overall functionality of the processor. In comparison, EC-CFI requires no or only minimal-intrusive hardware changes, i.e., changing the behavior or TME-MK, not affecting the general structure of the CPU pipeline.

## X. FUTURE WORK AND LIMITATIONS

**Future Work.** The hardware change described in Section VIII could be implemented in a system emulator. However, as neither QEMU [5] nor gem5 [25] currently support TME-MK, this hardware extension first needs to be integrated. Moreover, as described in Section V-B4, we currently trigger an EPT violation when switching between kernel and user mode to save and restore the active EPT. As an EPT violation causes a costly `vmexit`, future work could investigate how to avoid these exits. One possibility would be to extend the hypervisor and the kernel. The hypervisor could store the current active EPTP into a kernel-accessible memory region. When entering the kernel from a protected program, the extended kernel then switches to the default kernel EPTP. When leaving the syscall, the kernel could fetch the last active EPTP from the memory region and restore the EPTP. Another option would be to map the kernel address space for each

EPTP. Then, when switching from kernel to the user and back, a EPTP switch would not be necessary.

**Limitations.** In our current prototype implementation, we do not perform a points-to analysis to identify all potential call targets of indirect calls. Instead, we use a default signature for these calls. Although this is a security limitation of the current prototype, this simplification accurately models the runtime and code size overhead. Our current implementation does not support the protection of multiple programs executed on a CPU. To overcome this limitation, the hypervisor needs to be extended to manage the key identifiers and EPTs for each process. Finally, as there are no system emulators available supporting TME-MK and the HDL description or netlist of a Intel® CPU is not publicly available, we focused on providing a security discussion instead of performing fault experiments with fault injection frameworks [10], [16], [28], [39], [43].

## XI. CONCLUSION

In this paper, we presented EC-CFI, a cryptographically enforced control-flow integrity scheme utilizing recent hardware features of Intel® platforms and effective against a fault adversary. EC-CFI prevents that an adversary escapes the call graph of a program by encrypting each function with a different encryption key before executing the application. Only when the execution history is identical to the statically determined control-flow and the call target is within the bounds of the call graph, the decryption key for the called function is correctly derived. On control-flow manipulations outside of the call graph, code is decrypted with the wrong key, which can be detected with no or minimal detection latency. To achieve function-granular encryption on Intel® commodity platforms, we introduced a novel combination of TME-MK and the virtualization technology. In our paper, we showcased how to utilize our approach based on EPT aliasing to implement EC-CFI and open-source our custom toolchain. Moreover, we analyzed the EPT switching mechanism using the Embench-IoT and SPEC CPU2017 benchmarks. Finally, we described and evaluated a TME-MK hardware modification that could significantly reduce the performance impact of EC-CFI.

## ACKNOWLEDGMENT

We would like to thank the anonymous reviewers for their review and feedback.

## REFERENCES

- [1] *IEEE International Symposium on Hardware Oriented Security and Trust, HOST 2021, Tysons Corner, VA, USA, December 12-15, 2021.*
- [2] Martín Abadi, Mihai Badiu, Úlfar Erlingsson, and Jay Ligatti. Control-flow integrity. In *CCS*, pages 340–353, 2005.
- [3] Martín Abadi, Mihai Badiu, Úlfar Erlingsson, and Jay Ligatti. Control-flow integrity principles, implementations, and applications. *ACM Trans. Inf. Syst. Secur.*, 13:4:1–4:40, 2009.
- [4] AMD. Amd memory encryption, 2016.
- [5] Fabrice Bellard. QEMU, a Fast and Portable Dynamic Translator. In *USENIX ATC*, pages 41–46, 2005.
- [6] Intel Corporation. Runtime encryption of memory with Intel® total memory encryption-multi-key. <https://www.intel.com/content/www/us/en/developer/articles/news/runtime-encryption-of-memory-with-intel-tme-mk.html>, 2012. [accessed 2023-01-13].
- [7] Ang Cui and Rick Housley. BADFET: Defeating Modern Secure Boot Using Second-Order Pulsed Electromagnetic Fault Injection. In *WOOT*, 2017.
- [8] Ruan de Clercq, Ronald De Keulenaer, Bart Coppens, Bohan Yang, Pieter Maene, Koen De Bosschere, Bart Preneel, Bjorn De Sutter, and Ingrid Verbauwhede. SOFIA: Software and control flow integrity architecture. In *DATE*, pages 1172–1177, 2016.
- [9] Catherine Easdon. Undocumented cpu behavior: analyzing undocumented opcodes on intel x86-64, 2020.
- [10] Davide Ferraretto and Graziano Pravadelli. Efficient fault injection in QEMU. In *16th Latin-American Test Symposium, LATS 2015, Puerto Vallarta, Mexico, March 25-27, 2015*, pages 1–6, 2015.
- [11] Olga Goloubeva, Maurizio Rebaudengo, Matteo Sonza Reorda, and Massimo Violante. Soft-Error Detection Using Control Flow Assertions. In *18th IEEE International Symposium on Defect and Fault-Tolerance in VLSI Systems (DFT 2003), 3-5 November 2003, Boston, MA, USA, Proceedings*, pages 581–588, 2003.
- [12] Google Project Zero. Exploiting the DRAM Rowhammer bug to gain Kernel privileges. <https://googleprojectzero.blogspot.com/2015/03/exploiting-dram-rowhammer-bug-to-gain.html>, 2015. Accessed: 2022-11-07.
- [13] James Gratchoff, Niek Timmers, Albert Spruyt, and Lukasz Chmielewski. Proving the wild jungle jump. Technical report, Technical report, University of Amsterdam, 2015.
- [14] J. Alex Halderman, Seth D. Schoen, Nadia Heninger, William Clarkson, William Paul, Joseph A. Calandrino, Ariel J. Feldman, Jacob Appelbaum, and Edward W. Felten. Lest we remember: cold-boot attacks on encryption keys. *Commun. ACM*, 52:91–98, 2009.
- [15] Karine Heydemann, Jean-François Lalande, and Pascal Berthomé. Formally verified software countermeasures for control-flow integrity of smart card C code. *Comput. Secur.*, 85:202–224, 2019.
- [16] Andrea Höller, Armin Krieg, Tobias Rauter, Johannes Iber, and Christian Kreiner. QEMU-Based Fault Injection for a System-Level Analysis of Software Countermeasures Against Fault Attacks. In *DSD*, pages 530–533, 2015.
- [17] Intel. *Intel® Hardware Shield – Intel® Total Memory Encryption*, 2021.
- [18] Intel. *Intel® 64 and IA-32 Architectures Software Developer’s Manual*, 04 2022. Volume 3.
- [19] Intel. *Architecture Specification: Intel® Trust Domain Extensions (Intel® TDX) Module*, 06 2022.
- [20] Intel. *Intel® Architecture Memory Encryption Technologies*, 08 2022. Revision 1.4.
- [21] Volodymyr Kuznetsov, Laszlo Szekeres, Mathias Payer, George Candea, R. Sekar, and Dawn Song. Code-Pointer Integrity. In *OSDI*, pages 147–163, 2014.
- [22] Jean-François Lalande, Karine Heydemann, and Pascal Berthomé. Software Countermeasures for Control Flow Integrity of Smart Card C Codes. In *ESORICS*, volume 8713 of *LNCS*, pages 200–218, 2014.
- [23] Chris Lattner and Vikram S. Adve. LLVM: A Compilation Framework for Lifelong Program Analysis & Transformation. In *CGO*, pages 75–88, 2004.
- [24] Hans Liljestrand, Thomas Nyman, Kui Wang, Carlos China Perez, Jan-Erik Ekberg, and N. Asokan. PAC it up: Towards Pointer Integrity using ARM Pointer Authentication. In *USENIX Security Symposium*, pages 177–194, 2019.
- [25] Jason Lowe-Power, Abdul Mutaal Ahmad, Ayaz Akram, Mohammad Alian, Rico Amslinger, Matteo Andreozzi, Adrià Armejach, Nils Asmussen, Srikant Bharadwaj, Gabe Black, Gedare Bloom, Bobby R. Bruce, Daniel Rodrigues Carvalho, Jerónimo Castrillón, Lihong Chen, Nicolas Derumigny, Stephan Diestelhorst, Wendy Elsasser, Marjan Fari-borz, Amin Farmahini Farahani, Pouya Fotouhi, Ryan Gambord, Jayneel Gandhi, Dibakar Gope, Thomas Grass, Bagus Hanindhito, Andreas Hansson, Swapnil Haria, Austin Harris, Timothy Hayes, Adrian Herrera, Matthew Horsnell, Syed Ali Raza Jafri, Radhika Jagtap, Hanhwi Jang, Reiley Jeyapaul, Timothy M. Jones, Matthias Jung, Subash Kannoth, Hamidreza Khaleghzadeh, Yuetsu Kodama, Tushar Krishna, Tommaso Marinelli, Christian Menard, Andrea Mondelli, Tiago Mück, Omar Naji, Krishnendra Nathella, Hoa Nguyen, Nikos Nikolieris, Lena E. Olson, Marc S. Orr, Binh Pham, Pablo Prieto, Trivikram Reddy, Alec Roelke, Mahyar Samani, Andreas Sandberg, Javier Setoain, Boris Shingarov, Matthew D. Sinclair, Tuan Ta, Rahul Thakur, Giacomo Travaglini, Michael Upton, Nilay Vaish, Ilias Vougioukas, Zhengrong Wang, Norbert Wehn, Christian Weis, David A. Wood, Hongil Yoon, and Eder F. Zulian. The gem5 Simulator: Version 20.0+. *CoRR*, abs/2007.03152, 2020.
- [26] Ali José Mashtizadeh, Andrea Bittau, Dan Boneh, and David Mazières. CCFI: Cryptographically Enforced Control Flow Integrity. In *CCS*, pages 941–951, 2015.
- [27] Kit Murdock, David F. Oswald, Flavio D. Garcia, Jo Van Bulck, Daniel Gruss, and Frank Piessens. Plundervolt: Software-based Fault Injection Attacks against Intel SGX. In *S&P*, pages 1466–1482, 2020.
- [28] Pascal Nasahl, Miguel Osorio, Pirmin Vogel, Michael Schaffner, Timothy Trippel, Dominic Rizzo, and Stefan Mangard. SYNFI: Pre-Silicon Fault Analysis of an Open-Source Secure Element. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2022:56–87, 2022.
- [29] Pascal Nasahl, Robert Schilling, and Stefan Mangard. Protecting Indirect Branches Against Fault Attacks Using ARM Pointer Authentication. In *HOST [1]*, pages 68–79.
- [30] Pascal Nasahl, Robert Schilling, Mario Werner, Jan Hoogerbrugge, Marcel Medwed, and Stefan Mangard. CrypTag: Thwarting Physical and Logical Memory Vulnerabilities using Cryptographically Colored Memory. In *ASIA CCS ’21: ACM Asia Conference on Computer and Communications Security, Virtual Event, Hong Kong, June 7-11, 2021*, pages 200–212, 2021.
- [31] Pascal Nasahl and Niek Timmers. Attacking AUTOSAR using software and hardware attacks. *escar USA*, 2019.
- [32] Shoei Nashimoto, Daisuke Suzuki, Rei Ueno, and Naofumi Homma. Bypassing Isolated Execution on RISC-V with Fault Injection. *IACR Cryptol. ePrint Arch.*, page 1193, 2020.
- [33] Nahmsuk Oh, Philip P. Shirvani, and Edward J. McCluskey. Control-flow checking by software signatures. *IEEE Trans. Reliab.*, 51:111–122, 2002.
- [34] Gabriele Paoloni. How to benchmark code execution times on intel ia-32 and ia-64 instruction set architectures. *Intel Corporation*, 123:170, 2010.
- [35] David Patterson, Jeremy Bennett, Cesare Garlati Palmer Dabbelt, G. S. Madhusudan, and Trevor Mudge. Embench: Open benchmarks for embedded platforms. <https://www.embench.org/>.
- [36] Pengfei Qiu, Dongsheng Wang, Yongqiang Lyu, and Gang Qu. VoltJockey: Breaching TrustZone by Software-Controlled Voltage Manipulation over Multi-core Frequencies. In *CCS*, pages 195–209, 2019.
- [37] Pengfei Qiu, Dongsheng Wang, Yongqiang Lyu, and Gang Qu. VoltJockey: Breaking SGX by Software-Controlled Voltage-Induced Hardware Faults. In *Asian Hardware Oriented Security and Trust Symposium, AsianHOST 2019, Xi’an, China, December 16-17, 2019*, pages 1–6, 2019.
- [38] George A. Reis, Jonathan Chang, Neil Vachharajani, Ram Rangan, and David I. August. SWIFT: Software Implemented Fault Tolerance. In *CGO*, pages 243–254, 2005.
- [39] Jan Richter-Brockmann, Aein Rezaei Shahmirzadi, Pascal Sasdrich, Amir Moradi, and Tim Güneysu. FIVER - Robust Verification of Countermeasures against Fault Injections. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2021:447–473, 2021.
- [40] Robert Schilling, Pascal Nasahl, and Stefan Mangard. FIPAC: Thwarting Fault- and Software-Induced Control-Flow Attacks with ARM Pointer Authentication. In *COSADE*, volume 13211 of *LNCS*, pages 100–124, 2022.
- [41] Robert Schilling, Pascal Nasahl, Stefan Weiglhofer, and Stefan Mangard.

- SecWalk: Protecting Page Table Walks Against Fault Attacks. In *HOST* [1], pages 56–67.
- [42] Robert Schilling, Mario Werner, and Stefan Mangard. Securing conditional branches in the presence of fault attacks. In *DATE*, pages 1586–1591, 2018.
- [43] Mohammad Shokrollah-Shirazi and Seyed Ghassem Miremadi. FPGA-Based Fault Injection into Synthesizable Verilog HDL Models. In *Second International Conference on Secure System Integration and Reliability Improvement, SSIRI 2008, July 14-17, 2008, Yokohama, Japan*, pages 143–149, 2008.
- [44] Standard Performance Evaluation Corporation. SPEC CPU@ 2017. <https://www.spec.org/cpu2017/>, 2022. Accessed: 2022-11-07.
- [45] Stefan Steinegger, David Schrammel, Samuel Weiser, Pascal Nasahl, and Stefan Mangard. SERVAS! Secure Enclaves via RISC-V Authentication Shield. In *ESORICS*, volume 12973 of *LNCS*, pages 370–391, 2021.
- [46] Laszlo Szekeres, Mathias Payer, Tao Wei, and Dawn Song. SoK: Eternal War in Memory. In *S&P*, pages 48–62, 2013.
- [47] Adrian Tang, Simha Sethumadhavan, and Salvatore J. Stolfo. CLKSCREW: Exposing the Perils of Security-Oblivious Energy Management. In *USENIX Security Symposium*, pages 1057–1074, 2017.
- [48] Niek Timmers and Cristofaro Mune. Escalating Privileges in Linux Using Voltage Fault Injection. In *FDTC*, pages 1–8, 2017.
- [49] Niek Timmers, Albert Spruyt, and Marc Witteman. Controlling PC on ARM Using Fault Injection. In *FDTC*, pages 25–35, 2016.
- [50] Aurélien Vasselle, Hugues Thiebauld, Quentin Maouhoub, Adèle Morisset, and Sébastien Ermeneux. Laser-Induced Fault Injection on Smartphone Bypassing the Secure Boot-Extended Version. *IEEE Trans. Computers*, 69:1449–1459, 2020.
- [51] Rajesh Venkatasubramanian, John P. Hayes, and Brian T. Murray. Low-Cost On-Line Fault Detection Using Control Flow Assertions. In *9th IEEE International On-Line Testing Symposium (IOLTS 2003), 7-9 July 2003, Kos Island, Greece*, pages 137–143, 2003.
- [52] Mario Werner, Thomas Unterluggauer, David Schaffenrath, and Stefan Mangard. Sponge-Based Control-Flow Protection for IoT Devices. In *EURO S&P*, pages 214–226, 2018.
- [53] Mario Werner, Erich Wenger, and Stefan Mangard. Protecting the Control Flow of Embedded Processors against Fault Attacks. In *CARDIS*, volume 9514 of *LNCS*, pages 161–176, 2015.
- [54] Kent D. Wilken and John Paul Shen. Continuous Signature Monitoring: Efficient Concurrent-Detection of Processor Control Errors. In *Proceedings International Test Conference 1988, Washington, D.C., USA, September 1988*, pages 914–925, 1988.
- [55] Kent D. Wilken and John Paul Shen. Continuous signature monitoring: low-cost concurrent detection of processor control errors. *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.*, 9:629–641, 1990.