

Symbolic Verification of Cache Side-Channel Freedom

Sudipta Chattopadhyay^{ID} and Abhik Roychoudhury

Abstract—Cache timing attacks allow third-party observers to retrieve sensitive information from program executions. But, is it possible to automatically check the vulnerability of a program against cache timing attacks and then, automatically shield program executions against these attacks? For a given program, a cache configuration and an attack model, our CACHEFIX framework either verifies the cache side-channel freedom of the program or synthesizes a series of patches to ensure cache side-channel freedom during program execution. At the core of our framework is a novel symbolic verification technique based on automated abstraction refinement of cache semantics. The power of such a framework allows symbolic reasoning over counterexample traces and combines it with runtime monitoring for eliminating cache side channels during program execution. Our evaluation with routines from OpenSSL, libfixedtimefixedpoint, GDK, and FourQlib libraries reveals that our CACHEFIX approach (dis)proves cache side-channel freedom within an average of 75 s. In nearly all test cases, CACHEFIX synthesizes all patches within 20 min to ensure cache side-channel freedom of the respective routines during execution.

Index Terms—Computational and artificial intelligence, computer security, formal verification, software engineering.

I. INTRODUCTION

CACHE timing attacks [23], [24] are among the most critical *side-channel attacks* [25] that retrieve sensitive information from program executions. Recent cache attacks [31] further show that cache side-channel attacks are practical even in commodity embedded processors, such as in ARM-based embedded platforms [31]. The basic idea of a cache timing attack is to observe the timing of cache hits and misses for a program execution. Subsequently, the attacker use such timing to guess the sensitive input via which the respective program was activated.

Given the practical relevance, it is crucial to verify whether a given program (e.g., an encryption routine) satisfies *cache*

side-channel freedom, meaning the program is *not vulnerable* to cache timing attacks. However, verification of such a property is challenging for several reasons. First, the verification of cache side-channel freedom requires a systematic integration of cache semantics within the program semantics. This, in turn, is based on the derivation of a suitable abstraction of cache semantics. Our proposed CACHEFIX approach automatically builds such an abstraction and systematically refines it until a proof of cache side-channel freedom is obtained or a real (i.e., nonspurious) counterexample is produced. Second, proving cache side-channel freedom of a program requires reasoning over multiple execution traces. To this end, we propose a symbolic verification technique within our CACHEFIX framework. Concretely, we capture the cache behavior of a program via symbolic constraints over program inputs. Then, we leverage recent advances on satisfiability modulo theory (SMT) and constraint solving to (dis)prove the cache side-channel freedom of a program.

An appealing feature of our CACHEFIX approach is to employ symbolic reasoning over the real counterexample traces. To this end, we systematically explore real counterexample traces and apply such symbolic reasoning to synthesize patches. Each synthesized patch captures a symbolic condition ν on input variables and a sequence of actions that needs to be applied when the program is processed with inputs satisfying ν . The application of a patch is guaranteed to reduce the channel capacity of the program under inspection. Moreover, if our checker terminates, then our CACHEFIX approach guarantees to synthesize all patches that *completely shields the program* against cache timing attacks [7], [13]. Intuitively, our CACHEFIX approach can start with a program \mathcal{P} vulnerable to cache timing attack. Then, it leverages a systematic combination of symbolic verification and runtime monitoring to execute \mathcal{P} with cache side-channel freedom.

It is the precision and the novel mechanism implemented within CACHEFIX that set us apart from the state of the art. Existing works on analyzing cache side channels [14], [22], [30] are incapable to automatically build and refine abstractions for cache semantics. Besides, these works are not directly applicable when the underlying program *does not* satisfy cache side-channel freedom. Given an arbitrary program, our CACHEFIX approach generates proofs of its cache side-channel freedom or generates input(s) that manifest the violation of cache side-channel freedom. Moreover, our symbolic reasoning framework provides capabilities to systematically synthesize patches and completely eliminate cache side channels during program execution.

Manuscript received April 3, 2018; revised June 8, 2018; accepted July 2, 2018. Date of current version October 18, 2018. This work was supported in part by SUTD under Grant SRIS17123, and in part by the National Research Foundation, Prime Minister's Office, Singapore, under Its National Cybersecurity Research and Development Program, administered by the National Cybersecurity Research and Development Directorate under Award NRF2014NCR-NCR001-21. (Corresponding author: Sudipta Chattopadhyay.)

S. Chattopadhyay is with the Information Systems Technology and Design, Singapore University of Technology and Design, Singapore (e-mail: sudipta_chattopadhyay@sutd.edu.sg).

A. Roychoudhury is with the School of Computing, National University of Singapore, Singapore (e-mail: abhik@comp.nus.edu.sg).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TCAD.2018.2858402

We organize the remainder of this paper as follows. After providing an overview of CACHEFIX (Section II), we make the following contributions.

- 1) We present CACHEFIX, a novel symbolic verification framework to check the cache side-channel freedom of an arbitrary program. To the best of our knowledge, this is the first application of automated abstraction refinement and symbolic verification to check the cache behavior of a program.
- 2) We instantiate our CACHEFIX approach with direct-mapped caches, as well as with set-associative caches with least recently used (LRU) and first-in-first-out (FIFO) policy (Section IV-C). In Section IV-D, we show the generalization of our CACHEFIX approach over timing-based attacks [13] and trace-based attacks [7].
- 3) We discuss a systematic exploration of counterexamples to synthesize patches and to shield program executions against cache timing attacks (Section V). We provide theoretical guarantees that such patch synthesis converges toward completely eliminating cache side channels during execution.
- 4) We provide an implementation of CACHEFIX and evaluate it with 25 routines from OpenSSL, GDK, FourQlib, and libfixedtimefixedpoint libraries. Our evaluation reveals that CACHEFIX can establish proof or generate nonspurious counterexamples within 75 s on average. Besides, in most cases, CACHEFIX generated all patches within 20 min to ensure cache side-channel freedom during execution. Our implementation and all experimental data are publicly available.

II. OVERVIEW

In this section, we demonstrate the general insight behind our approach through examples. We consider the simple code fragments in Fig. 1(a) and (b), where `key` is a sensitive input. In this example, we will assume a direct-mapped cache having a size of 512 bytes. For the sake of brevity, we also assume that `key` is stored in a register and accessing `key` does not involve the cache. The mapping of different program variables into the cache appears in Fig. 1(c). Finally, we assume the presence of an attacker who observes the number of cache misses in the victim program. For such an attacker, examples in Fig. 1(a) and (b) satisfy cache side-channel freedom if and only if the number of cache misses suffered is independent of `key`.

Why Symbolic Verification?: Cache side-channel freedom of a program critically depends on how it interacts with the cache. We make an observation that the program cache behavior can be formulated via a well-defined set of predicates. To this end, let us assume $\text{set}(r_i)$ captures the cache set accessed by instruction r_i and $\text{tag}(r_i)$ captures the accessed cache tag by the same instruction. Consider the instruction r_3 in Fig. 1(a). We introduce a symbolic variable miss_3 , which we intend to set to one if r_3 suffers a cache miss and to set to zero otherwise. We observe that miss_3 depends on the following

logical condition:

$$\Gamma(r_3) \equiv \neg \left(0 \leq \text{key} \leq 127 \wedge \rho_{13}^{\text{set}} \wedge \neg \rho_{13}^{\text{tag}} \right) \wedge \neg \left(\text{key} \geq 128 \wedge \rho_{23}^{\text{set}} \wedge \neg \rho_{23}^{\text{tag}} \right) \quad (1)$$

where $\rho_{ji}^{\text{tag}} \equiv (\text{tag}(r_j) \neq \text{tag}(r_i))$ and $\rho_{ji}^{\text{set}} \equiv (\text{set}(r_j) = \text{set}(r_i))$. Intuitively, $\Gamma(r_3)$ checks whether both r_1 and r_2 , if executed, load different memory blocks than the one accessed by r_3 . Therefore, if $\Gamma(r_3)$ is evaluated to true, then $\text{miss}_3 = 1$ (i.e., r_3 suffers a cache miss) and $\text{miss}_3 = 0$ (i.e., r_3 is a cache hit), otherwise. Formally, we set the cache behavior of r_3 as follows:

$$\Gamma(r_3) \Leftrightarrow (\text{miss}_3 = 1); \quad \neg \Gamma(r_3) \Leftrightarrow (\text{miss}_3 = 0). \quad (2)$$

The style of encoding, as shown in (2), facilitates the usage of state-of-the-art solvers for verifying cache side-channel freedom.

In general, we note that the cache behavior of the program in Fig. 1(a), i.e., the cache behaviors of r_1, \dots, r_4 ; can be formulated accurately via the following set of predicates related to cache semantics:

$$\text{Pred}_{\text{cache}} = \left\{ \rho_{ji}^{\text{set}} \cup \rho_{ji}^{\text{tag}} \mid 1 \leq j < i \leq 4 \right\}. \quad (3)$$

The size of $\text{Pred}_{\text{cache}}$ depends on the number of memory-related instructions. However, $|\text{Pred}_{\text{cache}}|$ does not vary with the cache size.

Key Insight in Abstraction Refinement: If the attacker observes the number of cache misses, then the cache side-channel freedom holds for the program in Fig. 1(a) when all feasible traces exhibit the same number of cache misses. Hence, such a property φ can be formulated as the nonexistence of two traces tr_1 and tr_2 as follows:

$$\varphi \equiv \exists \text{tr}_1, \exists \text{tr}_2 \text{ s.t. } \left(\sum_{i=1}^4 \text{miss}_i^{(\text{tr}_1)} \neq \sum_{i=1}^4 \text{miss}_i^{(\text{tr}_2)} \right) \quad (4)$$

where $\text{miss}_i^{(\text{tr})}$ captures the valuation of miss_i in trace tr .

Our key insight is that to establish a proof of φ (or its lack thereof), it is not necessary to accurately track the values of all predicates in $\text{Pred}_{\text{cache}}$ [see (3)]. In other words, even if some predicates in $\text{Pred}_{\text{cache}}$ have unknown values, it might be possible to (dis)prove φ . This phenomenon occurs due to the inherent design principle of caches and we exploit this in our abstraction refinement process.

To realize our hypothesis, we first start with an initial set of predicates (possibly empty) whose values are accurately tracked during verification. In this example, let us assume that we start with an initial set of predicates $\text{Pred}_{\text{init}} = \{\rho_{13}^{\text{set}}, \rho_{13}^{\text{tag}}, \rho_{23}^{\text{set}}, \rho_{23}^{\text{tag}}\}$. The rest of the predicates in $\text{Pred}_{\text{cache}} \setminus \text{Pred}_{\text{init}}$ are set to unknown value. With this configuration at hand, CACHEFIX returns counterexample traces tr_1 and tr_2 [see (4)] to reflect that φ does not hold for the program in Fig. 1(a). In particular, the following traces are returned:

$$\begin{aligned} \text{tr}_1 &\equiv \langle \text{miss}_1 = \text{miss}_3 = 1, \text{miss}_2 = \text{miss}_4 = 0 \rangle \\ \text{tr}_2 &\equiv \langle \text{miss}_1 = 0, \text{miss}_2 = \text{miss}_3 = \text{miss}_4 = 1 \rangle. \end{aligned}$$

Given tr_1 and tr_2 , we check whether any of them are spurious. To this end, we reconstruct the logical condition

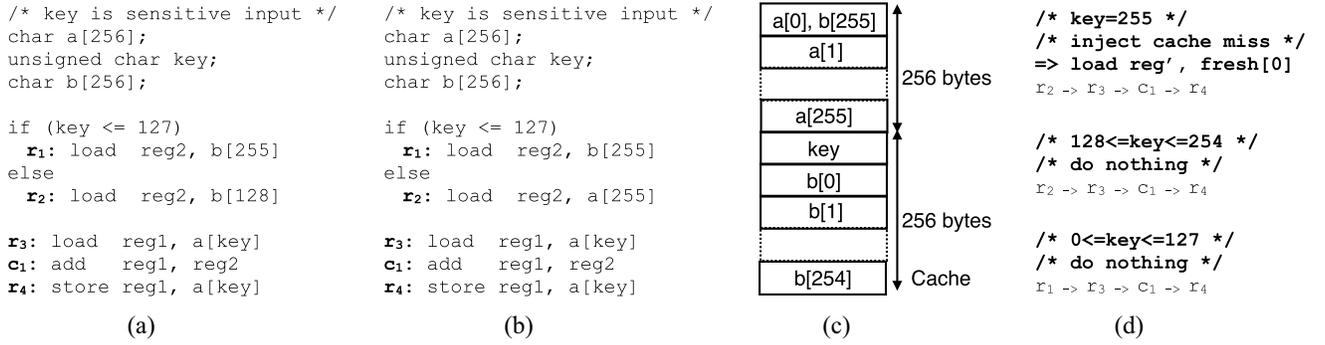


Fig. 1. Code fragment (a) satisfying cache side-channel freedom and (b) violating cache side-channel freedom. (c) Mapping of variables into the cache. (d) Runtime actions and the execution order for the program in Fig. 1(b) to ensure cache side-channel freedom.

[see (2)] that led to the specific valuations of $miss_i$ variables in a trace. For instance in trace tr_2 , such a logical condition is captured via $\neg\Gamma(r_1) \wedge \bigwedge_{i \in [2,4]} \Gamma(r_i)$. It turns out that $\neg\Gamma(r_1) \wedge \bigwedge_{i \in [2,4]} \Gamma(r_i)$ is *unsatisfiable*, making tr_2 spurious. This happened due to the incompleteness in tracking the predicates $\text{Pred}_{\text{cache}}$.

To systematically augment the set of predicates and rerun our verification process, we extract the unsatisfiable core from $\neg\Gamma(r_1) \wedge \bigwedge_{i \in [2,4]} \Gamma(r_i)$. Specifically, we get the following unsatisfiable core:

$$\mathcal{U} \equiv \neg\rho_{34}^{\text{set}} \vee \rho_{34}^{\text{tag}}. \quad (5)$$

Intuitively, with the initial abstraction $\text{Pred}_{\text{init}}$, our checker CACHEFIX failed to observe that r_3 and r_4 access the same memory block, hence, \mathcal{U} is unsatisfiable. We then augment our initial set of predicates with the predicates in \mathcal{U} and therefore, refining the abstraction as follows:

$$\text{Pred}_{\text{cur}} = \left\{ \rho_{13}^{\text{set}}, \rho_{13}^{\text{tag}}, \rho_{23}^{\text{set}}, \rho_{23}^{\text{tag}}, \rho_{34}^{\text{set}}, \rho_{34}^{\text{tag}} \right\}.$$

CACHEFIX successfully verifies the cache side-channel freedom of the program in Fig. 1(a) with the set of predicates Pred_{cur} . We note that the predicates in $\text{Pred}_{\text{cache}} \setminus \text{Pred}_{\text{cur}} \neq \phi$. In particular, we still have unknown values assigned to the following set of predicates:

$$\text{Pred}_{\text{unknown}} = \left\{ \rho_{12}^{\text{set}}, \rho_{12}^{\text{tag}}, \rho_{14}^{\text{set}}, \rho_{14}^{\text{tag}}, \rho_{24}^{\text{set}}, \rho_{24}^{\text{tag}} \right\}.$$

Therefore, it was possible to verify φ by tracking only half of the predicates in $\text{Pred}_{\text{cache}}$. Intuitively, ρ_{12}^{set} and ρ_{12}^{tag} were not needed to be tracked as r_1 and r_2 cannot appear in a single trace, as captured via the program semantics. In contrast, the rest of the predicates in $\text{Pred}_{\text{unknown}}$ were not required for the verification process, as neither r_1 nor r_2 influences the cache behavior of r_4 —it is influenced completely by r_3 .

Key Insight in Fixing: In general, the state-of-the-art in fixing cache side-channel is to revert to constant-time programming style [8]. Constant-time programming style imposes heavy burden on a programmer to follow certain programming patterns, such as to ensure the absence of input-dependent branches and input-dependent memory accesses. Yet, most programs do not exhibit constant-time behavior. Besides, the example in Fig. 1(a) shows that an application can still have

constant cache-timing, despite not following the constant-time programming style. Using our CACHEFIX approach, we observe that it is not necessary to always write constant-time programs. Instead, the executions of such programs can be manipulated to exhibit constant time behavior. We accomplish this by leveraging our verification results.

We consider the example in Fig. 1(b) and let us assume that we start with the initial abstraction $\text{Pred}_{\text{init}} = \{\rho_{13}^{\text{set}}, \rho_{13}^{\text{tag}}, \rho_{23}^{\text{set}}, \rho_{23}^{\text{tag}}\}$. CACHEFIX returns the following counterexample while verifying φ [see (4)]:

$$tr_1 \equiv \langle miss_1 = miss_3 = 1, miss_2 = miss_4 = 0 \rangle$$

$$tr_2 \equiv \langle miss_2 = 1, miss_1 = miss_3 = miss_4 = 0 \rangle.$$

If we reconstruct the logical condition that led to the specific valuations of $miss_1, \dots, miss_4$ in tr_1 and tr_2 , then we get the symbolic formulas $\Gamma(r_1) \wedge \neg\Gamma(r_2) \wedge \Gamma(r_3) \wedge \neg\Gamma(r_4)$ and $\neg\Gamma(r_1) \wedge \Gamma(r_2) \wedge \neg\Gamma(r_3) \wedge \neg\Gamma(r_4)$, respectively. Both the formulas are satisfiable for the example in Fig. 1(b). Intuitively, this happens due to r_2 , which loads the same memory block as accessed by r_3 only if $key = 255$.

We observe that tr_2 will be equivalent to tr_1 if a cache miss is inserted in the beginning of tr_2 . To this end, we need to know all inputs that lead to tr_2 . Thanks to the symbolic nature of our analysis, we obtain the exact symbolic condition, i.e., $\neg\Gamma(r_1) \wedge \Gamma(r_2) \wedge \neg\Gamma(r_3) \wedge \neg\Gamma(r_4)$ that manifests the trace tr_2 . Therefore, if the program in Fig. 1(b) is executed with any input satisfying $\neg\Gamma(r_1) \wedge \Gamma(r_2) \wedge \neg\Gamma(r_3) \wedge \neg\Gamma(r_4)$, then a cache miss is injected as shown in Fig. 1(d). This ensures the cache side-channel freedom during program execution, as all traces exhibit the same number of cache misses.

Our proposed fixing mechanism is novel that it does not rely on any specific programming style. Moreover, as we generate the fixes by directly leveraging the verification results, we can provide strong security guarantees during program execution.

Overall Workflow of CACHEFIX: Fig. 2 outlines the overall workflow of CACHEFIX. The abstraction refinement process is guaranteed to converge toward the most precise abstraction of cache semantics to (dis)prove the cache side-channel freedom. Moreover, as observed in Fig. 2, our cache side channel fixing is guided by program verification output, enabling us to give cache side channel freedom guarantees about the fixed program.

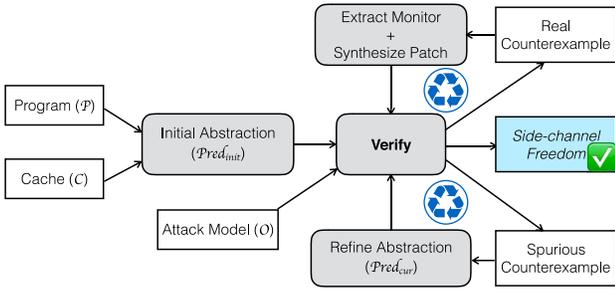


Fig. 2. Workflow of our symbolic verification and patching.

III. THREAT AND SYSTEM MODEL

In the following, we discuss the threat models and the type of processors targeted by CACHEFIX.

A. Threat Model

We assume that an attacker makes observations on the execution traces of victim program \mathcal{P} and the implementation of \mathcal{P} is known to the attacker. Besides, there does not exist any error in the observations made by the attacker. We also assume that an attacker can execute arbitrary user-level code on the processor that runs the victim program. This, in turn, allows the attacker to flush the cache (e.g., via accessing a large array) before the victim routine starts execution. We, however, do not assume that the attacker can access the address space of the victim program \mathcal{P} . We believe the aforementioned assumptions on the attacker are justified, as we aim to verify the cache side-channel freedom of programs against *strong attacker models*.

We capture an execution trace via a sequence of hits (h) and misses (m). Hence, formally we model an attacker as the mapping $\mathcal{O} : \{h, m\}^* \rightarrow \mathbb{X}$, where \mathbb{X} is a countable set. For $tr_1, tr_2 \in \{h, m\}^*$, an attacker can distinguish tr_1 from tr_2 if and only if $\mathcal{O}(tr_1) \neq \mathcal{O}(tr_2)$. In this paper, we instantiate our checker for the following realistic attack models.

- 1) $\mathcal{O}_{\text{time}} : \{h, m\}^* \rightarrow \mathbb{N}$. $\mathcal{O}_{\text{time}}$ maps each execution trace to the number of cache misses suffered by the same. This attack model imitates cache timing attacks [13].
- 2) $\mathcal{O}_{\text{trace}} : \{h, m\}^* \rightarrow \{0, 1\}^*$. $\mathcal{O}_{\text{trace}}$ maps each execution trace to a bitvector (h is mapped to 0 and m is mapped to 1). This attack model imitates trace-based attacks [7].

B. Processor Model

We assume an ARM-style processor with one or more cache levels. However, we consider timing attacks only due to first-level instruction or data caches [7], [13]. We currently do not handle more advanced attacks on shared caches [32]. First-level caches can either be partitioned (instruction versus data) or unified. We assume that set-associative caches have either LRU or FIFO replacement policy. Other deterministic replacement policies can easily be integrated within CACHEFIX via additional symbolic constraints. Finally, our timing model only takes into account the effect of caches. Timing effects due to other micro-architectural features (e.g., pipeline and branch prediction) are currently not handled. For the sake

of brevity, we discuss the timing effects due to memory-related instructions. It is straightforward to integrate the timing effects of computation instructions (e.g., add) into CACHEFIX.

IV. ABSTRACTION REFINEMENT

Notations: We represent cache via a triple $(2^S, 2^B, \mathcal{A})$ where 2^S , 2^B , and \mathcal{A} capture the number of cache sets, cache line size, and cache associativity, respectively. We use $\text{set}(r_i)$ and $\text{tag}(r_i)$ to capture the cache set and cache tag, respectively, accessed by instruction r_i . Additionally, we introduce a symbolic variable miss_i to capture whether r_i was a miss ($\text{miss}_i = 1$) or a hit ($\text{miss}_i = 0$). For instructions r_i and r_j , we have $j < i$ if and only if r_j was (symbolically) executed before r_i .

A. Initial Abstract Domain

We assume that a routine may start execution with any initial cache state, but it does not access memory blocks within the initial state during execution [22]. Hence, for a given instruction r_i , its cache behavior might be affected by all instructions executing prior to r_i . Concretely, the cache behavior of r_i can be accurately predicted based on the set of logical predicates Pred_{set} and Pred_{tag} as follows:

$$\begin{aligned} \text{Pred}_{\text{set}}^i &= \{\text{set}(r_j) = \text{set}(r_i) \mid 1 \leq j < i\} \\ \text{Pred}_{\text{tag}}^i &= \{\text{tag}(r_j) \neq \text{tag}(r_i) \mid 1 \leq j < i\}. \end{aligned} \quad (6)$$

Intuitively, $\text{Pred}_{\text{set}}^i$ captures the set of predicates checking whether any instruction prior to r_i accesses the same cache set as r_i . Similarly, $\text{Pred}_{\text{tag}}^i$ checks whether any instruction prior to r_i has a different cache tag than $\text{tag}(r_i)$. Based on this intuition, the following set of predicates are sufficient to predict the cache behaviors of N memory-related instructions:

$$\text{Pred}_{\text{set}} = \bigcup_{i=1}^N \text{Pred}_{\text{set}}^i; \quad \text{Pred}_{\text{tag}} = \bigcup_{i=1}^N \text{Pred}_{\text{tag}}^i. \quad (7)$$

For the sake of efficiency, however, we launch verification with a smaller set of predicates $\text{Pred}_{\text{init}} \subseteq \text{Pred}_{\text{tag}} \cup \text{Pred}_{\text{set}}$ as follows:

$$\begin{aligned} \text{Pred}_{\text{init}} &= \bigcup_{i=1}^N \{p \mid p \in \text{Pred}_{\text{tag}} \cup \text{Pred}_{\text{set}} \wedge |\sigma(r_i)| = 1 \\ &\quad \wedge \forall k \in [1, i). |\sigma(r_k)| = 1 \wedge \text{guard}_k \Rightarrow \text{true}\}. \end{aligned} \quad (8)$$

$\sigma(r_i)$ captures the set of memory blocks accessed by instruction r_i and guard_k captures the control condition under which r_k is executed. In general, our CACHEFIX approach works even if $\text{Pred}_{\text{init}} = \phi$. However, to accelerate the convergence of CACHEFIX, we start with the predicates whose values can be statically determined (i.e., independent of inputs). Intuitively, we take this approach for two reasons. First, the set $\text{Pred}_{\text{init}}$ can be computed efficiently during symbolic execution. Second, as the predicates in $\text{Pred}_{\text{init}}$ have constant valuation, they reduce the size of the formula to be discharged to the SMT solver. We note that guard_k depends on the program semantics. The abstraction of program semantics is an orthogonal problem and for the sake of brevity, we skip its discussion here.

Algorithm 1 Abstraction Refinement Algorithm

Input: Program \mathcal{P} , cache configuration \mathcal{C} , attack model \mathcal{O}
Output: Successful verification or a concrete counterexample

```

1: /*  $\Psi$  is a formula representation of  $\mathcal{P}$  */
2: /*  $Pred$  is cache-semantic-related predicates */
3: /*  $\Gamma$  determines cache behavior of all instructions */
4:  $(\Psi, Pred, \Gamma) := EXECUTESYMBOLIC(\mathcal{P}, \mathcal{C})$ 
5: /* Formulate initial abstraction (see equation (8)) */
6:  $Pred_{cur} := Pred_{init} := GETINITIALABSTRACTION(Pred)$ 
7: /* Rewrite  $\Psi$  with initial abstraction */
8:  $REWRITE(\Psi, Pred_{init})$ 
9: /* Formulate cache side-channel freedom property */
10:  $\varphi := GETPROPERTY(\mathcal{O})$ 
11: /* Invoke symbolic verification to check  $\Psi \wedge \neg\varphi$  */
12:  $(res, tr_1, tr_2) := VERIFY(\Psi, \varphi)$ 
13: while  $(res=false) \wedge (tr_1 \text{ or } tr_2 \text{ is spurious})$  do
14:   /* Extract unsatisfiable core from  $tr_1$  and/or  $tr_2$  */
15:    $\mathcal{U} := UNSATCORE(tr_1, tr_2, \Gamma)$ 
16:   /* Refine abstractions and repeat verification */
17:    $Pred_{cur} := REFINE(Pred_{init}, \mathcal{U}, Pred)$ 
18:    $REWRITE(\Psi, Pred_{cur})$ 
19:    $(res, tr_1, tr_2) := VERIFY(\Psi, \varphi)$ 
20:    $Pred_{init} := Pred_{cur}$ 
21: end while
22: return  $res$ 

```

B. Abstract Domain Refinement

We use the mapping $\Gamma : \{r_1, r_2, \dots, r_N\} \rightarrow \{\text{true}, \text{false}\}$ to capture the conditions under which r_i was a cache hit (i.e., $miss_i = 0$) or a cache miss (i.e., $miss_i = 1$). In particular, the following holds:

$$\Gamma(r_i) \Leftrightarrow (miss_i = 1); \quad \neg\Gamma(r_i) \Leftrightarrow (miss_i = 0). \quad (9)$$

$\Gamma(r_i)$ depends on predicates in $Pred_{set}^i \cup Pred_{tag}^i$ and the cache configuration. We show the formulation of $\Gamma(r_i)$ in Section IV-C.

EXECUTESYMBOLIC: Algorithm 1 captures the overall verification process based on our systematic abstraction refinement. The symbolic verification engine computes a formula representation Ψ of the program \mathcal{P} . This is accomplished via a symbolic execution on program \mathcal{P} (see procedure **EXECUTESYMBOLIC**) and systematically translating the cache and program semantics of each instruction into a set of constraints (see procedure **CONVERT**).

CONVERT: During the symbolic execution, a set of symbolic states, each capturing a unique execution path reaching an instruction r_i , is maintained. This set of symbolic states can be viewed as a disjunction $\Psi(r_i) \equiv \psi_1 \vee \psi_2 \vee \dots \vee \psi_{j-1} \vee \psi_j$, where $\Psi(r_i) \Rightarrow \Psi$ and each ψ_i symbolically captures a unique execution path leading to instruction r_i . At each instruction r_i , the procedure **CONVERT** translates $\Psi(r_i)$ in such a fashion that $\Psi(r_i)$ integrates both the cache semantics (see lines 13 and 14) and program semantics (see lines 18) of r_i . For instance, to integrate cache semantics of a memory-related instruction r_i , $\Psi(r_i)$ is converted to $\Psi(r_i) \wedge (\Gamma(r_i) \Leftrightarrow$

Procedure 2 Symbolically Tracking Program and Cache States

```

1: /* symbolically execute  $\mathcal{P}$  with cache configuration  $\mathcal{C}$  */
2: procedure EXECUTESYMBOLIC( $\mathcal{P}, \mathcal{C}$ )
3:    $i := 1; \Psi := \text{true}; Pred_{set} := Pred_{tag} := \phi$ 
4:    $r_i := GETNEXTINSTRUCTION(\mathcal{P})$ 
5:   while  $r_i \neq \text{exit}$  do
6:     if  $r_i$  is memory-related instruction then
7:       /* Collect predicates for cache semantics */
8:        $Pred_{set} \cup = Pred_{set}^i; Pred_{tag} \cup = Pred_{tag}^i$ 
9:       /*  $\Gamma(r_i)$  determines cache behavior of  $r_i$  */
10:      Formulate  $\Gamma(r_i)$  /* see Section IV-C */
11:       $\Gamma \cup = \{\Gamma(r_i)\}$ 
12:      /* Integrate cache semantics within  $\Psi$  */
13:       $\Psi := CONVERT(\Psi, \Gamma(r_i) \Leftrightarrow (miss_i = 1))$ 
14:       $\Psi := CONVERT(\Psi, \neg\Gamma(r_i) \Leftrightarrow (miss_i = 0))$ 
15:    end if
16:    /* Integrate program semantics of  $r_i$  within  $\Psi$  */
17:    /*  $\varphi(r_i)$  is a predicate capturing  $r_i$  semantics */
18:     $\Psi := CONVERT(\Psi, \varphi(r_i))$ 
19:     $i := i + 1$ 
20:     $r_i := GETNEXTINSTRUCTION(\mathcal{P})$ 
21:  end while
22:  return  $(\Psi, Pred_{set} \cup Pred_{tag}, \Gamma)$ 
23: end procedure

```

$(miss_i = 1)) \wedge (\neg\Gamma(r_i) \Leftrightarrow (miss_i = 0))$. Similarly, the program semantics of instruction r_i , as captured via $\varphi(r_i)$, is integrated within $\Psi(r_i)$ as $\Psi(r_i) \wedge \varphi(r_i)$. Translating the *program semantics* of each instruction to a set of constraints is a standard technique in any symbolic model checking [19]. Moreover, such a translation is typically carried out on a program in static single assignment form and takes into account both data and control flow. Unlike classic symbolic analysis, however, we consider both the cache semantics and program semantics of an execution path, as explained in the preceding.

GETINITIALABSTRACTION: We start our verification with an initial abstraction of cache semantics [see (8)]. Such an initial abstraction contains a partial set of logical predicates $Pred_{init} \subseteq Pred_{set} \cup Pred_{tag}$. Based on $Pred_{init}$, we rewrite Ψ via the procedure **REWRITE** as follows: we walk through Ψ and look for occurrences of any predicate $p^- \in (Pred_{set} \cup Pred_{tag}) \setminus Pred_{init}$. For any p^- discovered in Ψ , we replace p^- with a fresh symbolic variable V_{p^-} . Intuitively, this means that during the verification process, we assume any truth value for the predicates in $(Pred_{set} \cup Pred_{tag}) \setminus Pred_{init}$. This, in turn, substantially reduces the size of the symbolic formula Ψ and simplifies the verification process.

VERIFY AND GETPROPERTY: The procedure **VERIFY** invokes the solver to check the cache side-channel freedom of \mathcal{P} with respect to attack model \mathcal{O} . The property φ , capturing the cache side-channel freedom, is computed via **GETPROPERTY**. For example, in timing-based attacks, φ is captured via the nonexistence of any two traces tr_1 and tr_2 that have different number of cache misses [see (4)]. In other words, if the following formula is satisfied with more than

one valuations for $\sum_{i=1}^N \text{miss}_i$, then side-channel freedom is violated:

$$\Psi \wedge \left(\left(\sum_{i=1}^N \text{miss}_i \right) \geq 0 \right). \quad (10)$$

Here, N captures the total number of memory-related instructions encountered during the symbolic execution of \mathcal{P} .

REFINE AND REWRITE: When our verification process fails, we check the feasibility of a counterexample trace $\text{trace} \in \{\text{tr}_1, \text{tr}_2\}$. Recall from (9) that r_i is a cache miss if and only if $\Gamma(r_i)$ holds true. We leverage this relation to construct the following formula Γ_{trace} for feasibility checking:

$$\Gamma_{\text{trace}} = \bigwedge_{i=1}^N \begin{cases} \Gamma(r_i), & \text{if } \text{miss}_i^{(\text{trace})} = 1 \\ \neg\Gamma(r_i), & \text{if } \text{miss}_i^{(\text{trace})} = 0. \end{cases} \quad (11)$$

In (11), $\text{miss}_i^{(\text{trace})}$ captures the valuation of symbolic variable miss_i in the counterexample trace. We note that trace is not a spurious counterexample if and only if Γ_{trace} is *satisfiable*, hence, highlighting the violation of cache side-channel freedom.

If Γ_{trace} is *unsatisfiable*, then our initial abstraction $\text{Pred}_{\text{init}}$ was insufficient to (dis)prove the cache side-channel freedom. In order to refine this abstraction, we extract the *unsatisfiable core* from the symbolic formula Γ_{trace} via the procedure UNSATCORE. Such an unsatisfiable core contains a set of CNF clauses $\in \bigcup_{k \in [1, N]} \Gamma(r_k)$. We note each $\Gamma(r_k)$ is a function of the set of predicates $\text{Pred}_{\text{tag}} \cup \text{Pred}_{\text{set}}$. Finally, we refine the abstraction (see procedure REFINE in Algorithm 1) to Pred_{cur} by including all predicates in the unsatisfiable core as follows:

$$\text{Pred}_{\text{cur}} := \text{Pred}_{\text{init}} \cup \{p^+ \mid p^+ \in \text{UnsatCore}(\Gamma_{\text{trace}}) \wedge p^+ \notin \text{Pred}_{\text{init}}\}. \quad (12)$$

With the refined abstraction Pred_{cur} , we rewrite the symbolic formula Ψ (see procedure REWRITE). In particular, we identify the placeholder symbolic variables for predicates in the set $\text{Pred}_{\text{cur}} \setminus \text{Pred}_{\text{init}}$. We rewrite Ψ by replacing these placeholder symbolic variables with the respective predicates in the set $\text{Pred}_{\text{cur}} \setminus \text{Pred}_{\text{init}}$. It is worthwhile to note that the placeholder symbolic variables in $(\text{Pred}_{\text{tag}} \cup \text{Pred}_{\text{set}}) \setminus \text{Pred}_{\text{cur}}$ remain unchanged.

C. Modeling Cache Semantics

For each memory-related instruction r_i , the formulation of $\Gamma(r_i)$ is critical to prove the cache side-channel freedom. The formulation of $\Gamma(r_i)$ depends on the configuration of caches. Due to space constraints, we will only discuss the symbolic model for direct-mapped caches (symbolic models for LRU and FIFO caches are provided in the supplement [18]). To simplify the formulation, we will use the following abbreviations for the rest of the section:

$$\rho_{ij}^{\text{set}} \equiv (\text{set}(r_i) = \text{set}(r_j)); \quad \rho_{ij}^{\text{tag}} \equiv (\text{tag}(r_i) \neq \text{tag}(r_j)). \quad (13)$$

We also distinguish between the following variants of misses.

- 1) *Cold Misses:* Cold misses occur when a memory block is accessed for the first time.

- 2) *Conflict Misses:* All cache misses that are not cold misses are referred to as conflict misses.

Formulating Conditions for Cold Misses: Cold cache misses occur when a memory block is accessed for *the first time during program execution*. In order to check whether r_i suffers a cold miss, we check whether all instructions $r \in \{r_1, r_2, \dots, r_{i-1}\}$ access different memory blocks than the memory block accessed by r_i . This is captured as follows:

$$\Theta_i^{\text{cold}} \equiv \bigwedge_{j \in [1, i)} \left(\neg \rho_{ji}^{\text{set}} \vee \rho_{ji}^{\text{tag}} \vee \neg \text{guard}_j \right). \quad (14)$$

Recall that guard_j captures the control condition under which r_j is executed. Hence, if guard_j is evaluated false for a trace, then r_j does not appear in the respective trace. If Θ_i^{cold} is satisfied, then r_i inevitably suffers a cold cache miss.

Formulating Conditions for Conflict Cache Misses: For direct-mapped caches, an instruction r_i suffers a conflict miss due to an instruction r_j if all of the following conditions are satisfied.

$\phi_{ji}^{\text{cnf,dir}}$: If r_j accesses the same cache set as r_i , however, r_j accesses a different cache tag as compared to r_i . This is formally captured as follows:

$$\phi_{ji}^{\text{cnf,dir}} \equiv \rho_{ji}^{\text{tag}} \wedge \rho_{ji}^{\text{set}}. \quad (15)$$

$\phi_{ji}^{\text{rel,dir}}$: No instruction between r_j and r_i accesses the same memory block as r_i . For instance, consider the memory-block access sequence $(r_1 : m_1) \rightarrow (r_2 : m_2) \rightarrow (r_3 : m_2)$, where both m_1 and m_2 are mapped to the same cache set and $r_{1..3}$ captures the respective memory-related instructions. It is not possible for r_1 to inflict a conflict miss for r_3 , as the memory block m_2 is reloaded by instruction r_2 . $\phi_{ji}^{\text{rel,dir}}$ is formally captured as follows:

$$\phi_{ji}^{\text{rel,dir}} \equiv \bigwedge_{j < k < i} \left(\rho_{ki}^{\text{tag}} \vee \neg \rho_{ki}^{\text{set}} \vee \neg \text{guard}_k \right). \quad (16)$$

Intuitively, $\phi_{ji}^{\text{rel,dir}}$ captures that all instructions between r_j and r_i either access a different memory block than r_i (hence, satisfying $\rho_{ki}^{\text{tag}} \vee \neg \rho_{ki}^{\text{set}}$) or does not appear in the execution trace (hence, satisfying $\neg \text{guard}_k$).

Given the intuition mentioned in the preceding paragraphs, we conclude that r_i suffers a conflict miss if both $\phi_{ji}^{\text{cnf,dir}}$ and $\phi_{ji}^{\text{rel,dir}}$ are satisfied for *any instruction executing prior to r_i* . This is captured in the symbolic condition $\Theta_i^{\text{cnf,dir}}$ as follows:

$$\Theta_i^{\text{cnf,dir}} \equiv \bigvee_{j \in [1, i)} \left(\phi_{ji}^{\text{cnf,dir}} \wedge \phi_{ji}^{\text{rel,dir}} \wedge \text{guard}_j \right). \quad (17)$$

Computing $\Gamma(r_i)$: For direct-mapped caches, r_i can be a cache miss if it is either a cold cache miss or a conflict miss. Hence, $\Gamma(r_i)$ is captured symbolically as follows:

$$\Gamma(r_i) \equiv \text{guard}_i \wedge \left(\Theta_i^{\text{cold}} \vee \Theta_i^{\text{cnf,dir}} \right). \quad (18)$$

D. Property for Cache Side-Channel Freedom

In this paper, we instantiate our checker for timing-based attacks [13] and trace-based attacks [7] as follows.

Timing-Based Attacks: In timing-based attacks, an attacker aims to distinguish traces based on their timing. In our framework, we verify the following property to ensure cache side-channel freedom:

$$\left| \Psi \wedge \left(\left(\sum_{i=1}^N \text{miss}_i \right) \geq 0 \right) \right|_{\text{sol}(\sum_{i=1}^N \text{miss}_i)} \leq 1. \quad (19)$$

N captures the number of symbolically executed, memory-related instructions. $\text{sol}(\sum_{i=1}^N \text{miss}_i)$ captures the number of valuations of $\sum_{i=1}^N \text{miss}_i$. Intuitively, (19) aims to check that the underlying program has exactly one cache behavior, in terms of the total number of cache misses.

Trace-Based Attacks: In trace-based attacks, an attacker monitors the cache behavior of each memory access. We define a partial function $\xi : \{r_1, \dots, r_N\} \rightarrow \{0, 1\}$ as follows:

$$\xi(r_i) = \begin{cases} 1, & \text{if } \text{guard}_i \wedge (\text{miss}_i = 1) \text{ holds} \\ 0, & \text{if } \text{guard}_i \wedge (\text{miss}_i = 0) \text{ holds.} \end{cases} \quad (20)$$

The following verification goal ensures side-channel freedom:

$$\left| \Psi \wedge |\text{dom}(\xi)| \geq 0 \wedge (\|_{r_i \in \text{dom}(\xi)} \xi(r_i)) \geq 0 \right|_{\text{sol}(X)} \leq 1 \quad (21)$$

where $\text{dom}(\xi)$ captures the domain of ξ , $\|$ captures the ordered (with respect to the indexes of r_i) concatenation operation and $X = \langle |\text{dom}(\xi)|, \|_{r_i \in \text{dom}(\xi)} \xi(r_i) \rangle$. Intuitively, we check whether there exists exactly one cache behavior sequence.

V. RUNTIME MONITORING

CACHEFIX produces the first real counterexample when it discovers two traces with different observations (with respect to attack model \mathcal{O}). These traces are then analyzed to compute a set of runtime actions that are guaranteed to reduce the uncertainty to guess sensitive inputs. Overall, our runtime monitoring involves the following crucial steps.

- 1) We analyze a counterexample trace tr and extract the symbolic condition for which the same trace would be generated,
- 2) We systematically explore unique counterexamples with the objective to reduce the uncertainty to guess sensitive inputs,
- 3) We compute a set of runtime actions that need to be applied for improving the cache side-channel freedom.

In the following, we discuss these three steps in more detail.

A. Analyzing Counterexample Trace

Given a real counterexample trace, we extract a symbolic condition that captures all the inputs for which the same counterexample trace can be obtained. Thanks to the symbolic nature of our analysis, CACHEFIX already includes capabilities to extract these monitors as follows:

$$\nu \equiv \bigwedge_{r_i \in \text{trace}} \begin{cases} \Gamma(r_i), & \text{if } \text{miss}_i^{(\text{trace})} = 1 \\ \neg\Gamma(r_i), & \text{if } \text{miss}_i^{(\text{trace})} = 0 \end{cases} \quad (22)$$

where $\text{miss}_i^{(\text{trace})}$ is the valuation of symbolic variable miss_i in trace. We note that $\nu \Rightarrow \neg\Gamma(r_j)$ for any r_j that does not

appear in trace (i.e., $r_j \notin \text{trace}$). Hence, to formulate ν , it was sufficient to consider only the instructions that appear in trace.

Once we extract a monitor ν from counterexample trace, the symbolic system Ψ is refined to $\Psi \wedge \neg\nu$. This is to ensure that we only explore unique counterexample traces.

B. Systematic Exploration of Counterexamples

The order of exploring counterexamples is crucial to satisfy *monotonicity*, i.e., to reduce the channel capacity (a standard metric to quantify the information flow from sensitive input to attacker observation) of \mathcal{P} with each round of patch generation. To this end, CACHEFIX employs a strategy that can be visualized as an exploration of the equivalence classes of observations (e.g., #cache misses), i.e., we explore all counterexamples in the same equivalence class in one shot. In order to find another counterexample exhibiting the same observation as observation o , we modify the verification goal as follows, for timing and trace-based attacks, respectively [see (20) for ξ]:

$$\begin{aligned} & \Psi \wedge \neg \left(\left(\sum_{i=1}^N \text{miss}_i \right) \neq o \right) \\ & \Psi \wedge \neg \left(\bigvee_{\|_{r_i \in \text{dom}(\xi)} \xi(r_i)} \left[\begin{array}{l} |\text{dom}(\xi)| \neq [|\text{dom}(\xi)]_o \\ \|_{r_i \in \text{dom}(\xi)} \xi(r_i) \neq [\|_{r_i \in \text{dom}(\xi)} \xi(r_i)]_o \end{array} \right] \right) \end{aligned} \quad (23)$$

where $[X]_o$ captures the valuation of X with respect to observation o and N is the total number of symbolically executed, memory-related instructions. If (23) is unsatisfiable, then it captures the absence of any more counterexample with observation o . We note that Ψ is automatically refined to avoid discovering duplicate or spurious counterexamples. If (23) is satisfiable, our checker provides another real counterexample with the observation o . We repeat the process until no more real counterexample with the observation o is found, at which point (23) becomes unsatisfiable.

To explore a different equivalence class of observation than that of observation o , CACHEFIX negates the verification goal. For $\mathcal{O}_{\text{time}}$, as an example, the verification goal is changed as follows:

$$\Psi \wedge \neg \left(\left(\sum_{i=1}^N \text{miss}_i \right) = o \right). \quad (24)$$

We note that (24) is satisfiable if and only if there exists an execution trace with observation differing from o .

C. Runtime Actions to Improve Side-Channel Freedom

Our checker maintains the record of all explored observations and the symbolic conditions capturing the equivalence classes of respective observations. At each round of patch (i.e., runtime action) synthesis, we walk through this record and compute the necessary runtime actions for improving cache side-channel freedom.

$\mathcal{O}_{\text{time}}$: Assume $\Omega = \{\langle v_1, o_1 \rangle, \langle v_2, o_2 \rangle, \dots, \langle v_k, o_k \rangle\}$ where each o_i captures a unique number of observed cache misses and v_i symbolically captures all inputs that lead to observation o_i . Our goal is to manipulate executions so that they lead to the

same number of cache misses. To this end, the patch synthesis stage determines the amount of cache misses that needs to be added for each element in Ω . Concretely, the set of runtime actions generated are as follows:

$$\left\langle v_1, \left(\max_{i \in [1, k]} o_i - o_1 \right) \right\rangle, \dots, \left\langle v_k, \left(\max_{i \in [1, k]} o_i - o_k \right) \right\rangle. \quad (25)$$

In practice, when a program is run with input I , we check whether $I \in v_x$ for some $x \in [1, k]$. Subsequently, $(\max_{i \in [1, k]} o_i - o_x)$ cache misses were injected before the program starts executing.

\mathcal{O}_{trace} : During trace-based attacks, the attacker makes an observation on the sequence of cache hits and misses in an execution trace. Therefore, our goal is to manipulate executions in such a fashion that all execution traces lead to the same sequence of cache hits and misses. To accomplish this, each runtime action involves the injection of cache misses or hits before execution, after execution or at an arbitrary point of execution. It also involves invalidating an address in cache. Concretely, this is formalized as follows:

$$\langle v_i, \langle (c_1, a_1), (c_2, a_2), \dots, (c_k, a_k) \rangle \rangle \quad (26)$$

where v_i captures the symbolic input condition where the runtime actions are employed. For any input satisfying v_i , we count the number of instructions executed. If the number of executed instructions reaches c_j , then we perform the action a_j (e.g., injecting hits/misses or invalidating an address in cache), for any $j \in [1, k]$.

As an example, consider a trace-based attack in the example of Fig. 1(b). Our checker will manipulate counterexample traces by injecting cache misses and hits as follows (injected cache hits and misses are highlighted in bold):

$$\text{tr}_1'' \equiv \langle \mathbf{miss}, \text{miss}, \text{hit}, \text{hit} \rangle; \quad \text{tr}_2'' \equiv \langle \text{miss}, \text{miss}, \text{hit}, \mathbf{hit} \rangle$$

Therefore, the following actions are generated to ensure cache side-channel freedom against trace-based attacks:

$$\langle \text{key} = 255, \langle (0, \text{miss}) \rangle \rangle, \langle 0 \leq \text{key} \leq 254, \langle (3, \text{hit}) \rangle \rangle$$

We use string alignment algorithm [6] to make two traces equivalent (via insertion of cache hits/misses or substitution of hits to misses).

D. Practical Consideration

In practice, the injection of a cache miss can be performed via accessing a fresh memory block [see Fig. 1(d)]. However, unless the injection of a cache miss happens to be in the beginning or at the end of execution, the cache needs to be disabled before and enabled after such a cache miss. Consequently, our injection of misses does not affect cache states. In ARM-based processor, this is accomplished via manipulating the C bit of CP15 register. The injection of a cache hit can be performed via tracking the last accessed memory address and reaccessing the same address.

To change a cache hit to a cache miss, the accessed memory address needs to be invalidated in the cache. CACHEFIX symbolically tracks the memory address accessed at each memory-related instruction. When the program is run with

input $I \in v_i$, we concretize all memory addresses with respect to I . Hence, while applying an action that involves cache invalidation, we know the exact memory address that needs to be invalidated. In ARM-based processor, the instruction MCR provides capabilities to invalidate an address in the cache.

We note the preceding manipulations on an execution requires additional registers. We believe this is possible by using some system register or using a locked portion in the cache.

E. Properties Guaranteed by CACHEFIX

CACHEFIX satisfies the following crucial properties (proofs are included in the supplement [18]) on channel capacity, Shannon entropy, and min entropy; which are standard metrics to quantify the information flow from sensitive inputs to the attacker observation.

Property 1 (Monotonicity): Consider a victim program \mathcal{P} with sensitive input \mathcal{K} . Given attack models \mathcal{O}_{time} or \mathcal{O}_{trace} , assume that the channel capacity to quantify the uncertainty of guessing \mathcal{K} is $\mathcal{G}_{cap}^{\mathcal{P}}$. CACHEFIX guarantees that $\mathcal{G}_{cap}^{\mathcal{P}}$ monotonically decreases with each synthesized patch [see (25), (26)] employed at runtime.

Property 2 (Convergence): Consider a victim program \mathcal{P} with sensitive input \mathcal{K} . In the absence of any attacker, assume that the uncertainty to guess \mathcal{K} is $\mathcal{G}_{cap}^{\text{init}}$, $\mathcal{G}_{shn}^{\text{init}}$, and $\mathcal{G}_{min}^{\text{init}}$, via channel capacity, Shannon entropy, and Min entropy, respectively. If CACHEFIX terminates and all synthesized patches are applied at runtime, then the channel capacity (respectively, Shannon entropy and Min entropy) will remain $\mathcal{G}_{cap}^{\text{init}}$ (respectively, $\mathcal{G}_{shn}^{\text{init}}$ and $\mathcal{G}_{min}^{\text{init}}$) even in the presence of attacks captured via \mathcal{O}_{time} and \mathcal{O}_{trace} .

VI. IMPLEMENTATION AND EVALUATION

In this section, we will discuss our implementation setup and key findings from the evaluation.

A. Implementation Setup

The input to CACHEFIX is the target program and a cache configuration. We have implemented CACHEFIX on top of CBMC bounded model checker [1]. It first builds a formula representation of the input program via symbolic execution. Then, it checks the (un)satisfiability of this formula against a specification property. Despite being a bounded model checker, CBMC is used as a classic verification tool in our experiments. In particular for program loops, CBMC first attempts to derive loop bounds automatically. If CBMC fails to derive certain loop bounds, then the respective loop bounds need to be provided manually. Nevertheless, during the verification process, CBMC checks all manually provided loop bounds and the verification fails if any such bound was erroneous. In our experiments, all loop bounds were automatically derived by CBMC. In short, if CACHEFIX successfully verifies a program, then the respective program exhibits cache side-channel freedom for the given cache configuration and targeted attack models.

The implementation of our checker impacts the entire workflow of CBMC. We first modify the symbolic execution engine

TABLE I

SUMMARY OF CACHEFIX EVALUATION. TIMEOUT IS SET TO TEN MINUTES. $Time_{cur}$ CAPTURES THE TIME TAKEN BY CACHEFIX (I.E., # PREDICATES $|Pred_{cur}|$), WHEREAS $Time_{all}$ CAPTURES THE TIME TAKEN WHEN ALL PREDICATES (I.E., # PREDICATES $|Pred_{set} \cup Pred_{tag}|$) ARE CONSIDERED

Library	Routine	Size (LOC)	$ Pred_{set} \cup Pred_{tag} $	\mathcal{O}_{time}				\mathcal{O}_{trace}			
				$ Pred_{cur} $	Result	$Time_{cur}$ (secs)	$Time_{all}$ (secs)	$ Pred_{cur} $	Result	$Time_{cur}$ (secs)	$Time_{all}$ (secs)
OpenSSL [5]	AES128	740	231959	115537	✗	148.73	timeout	115545	✗	175.34	timeout
	DES	2124	205567	95529	✗	105.87	timeout	95639	✗	110.32	timeout
	RC5	1613	50836	25277	✓	26.88	541.35	26277	✓	34.84	585.32
GDK [3]	keyname	712	20827	18178	✗	21.75	120.21	18178	✗	24.32	125.11
	unicode	862	21917	19178	✗	27.49	117.78	19178	✗	32.09	119.56
fixedt [4]	fix_eq	334	71	32	✓	1.70	5.70	32	✓	4.31	6.08
	fix_cmp	400	957	474	✓	2.09	6.12	476	✓	2.08	6.15
	fix_mul	930	33330	16651	✓	22.44	100.32	17016	✓	20.19	121.37
	fix_conv_64	350	211	102	✓	1.69	1.81	102	✓	1.78	1.98
	fix_sqrt	2480	150127	74961	✓	60.54	359.94	87834	✓	67.90	581.45
	fix_exp	1128	101418	50655	✓	53.47	400.11	56194	✓	55.12	416.21
	fix_ln	1140	92113	46025	✓	58.89	114.89	47065	✓	61.11	127.21
	fix_pow	2890	643961	321885	✓	389.97	timeout	323787	✓	377.55	timeout
	fix_ceil	390	266	128	✓	1.70	1.70	128	✓	4.65	4.70
	fix_conv_double	650	1921	953	✓	2.70	2.75	945	✓	4.16	4.20
fix_frac	370	60172	53638	✗	22.14	23.18	51432	✗	22.15	24.58	
FourQ [2]	eccmadd	1550	324661	162058	✓	220.59	437.77	165453	✓	214.16	502.94
	eccnorm	1303	165291	82464	✓	105.97	198.90	83100	✓	114.93	219.09
	pt_setup	1345	3850	1866	✓	10.45	11.66	1901	✓	14.90	14.95
	eccdouble	1364	531085	265219	✓	285.70	558.78	267889	✓	312.31	597.18
	R1_to_R2	1352	85750	42731	✓	73.07	249.12	41014	✓	84.11	398.77
	R1_to_R3	1328	25246	12538	✓	26.95	53.21	12555	✓	24.89	54.45
	R2_to_R4	1322	28415	14105	✓	32.29	55.11	17001	✓	31.12	61.88
	R5_to_R1	1387	12531	5255	✓	22.07	34.05	5278	✓	24.32	52.88
	eccpt_validate	1406	57129	38012	✗	34.48	61.91	37948	✗	34.31	71.54

of CBMC to insert the predicates related to cache semantics. As a result, upon the termination of symbolic execution, the formula representation of the program encodes both the cache semantics and the program semantics. Second, we systematically rewrite this formula based on our abstraction refinement, with the aim of verifying cache side-channel freedom. Finally, we modify the verification engine of CBMC to systematically explore different counterexamples, instrument patches and refining the side-channel freedom properties *on-the-fly*. To manipulate and solve symbolic formulas, we leverage Z3 theorem prover. All reported experiments were performed on an Intel i7 machine, having 8GB RAM and running OSX.

B. Subject Programs and Cache

We have chosen security-critical subjects from OpenSSL library [5], GDK library [3], arithmetic routines from libfixedtimefixedpoint [4], and elliptic curve routines from FourQlib [2] to evaluate CACHEFIX (see Table I). The choice of our subjects is driven by the criticality of the respective routines in developing cryptographic software. To stress test our checker, we include representative routines exhibiting constant cache-timing, as well as routines exhibiting variable cache timing. We set the default cache to be 1KB direct-mapped, with a line size of 32 bytes.

C. Efficiency of Checking

Table I captures a summary of our evaluation for CACHEFIX. The outcome of this evaluation is either a successful verification (✓) or a nonspurious counterexample (✗). CACHEFIX accomplished the verification tasks for all subjects only within a few minutes. The maximum time taken by our checker was 390 s for the routine `fix_pow` – a constant time implementation of powers (x^y). `fix_pow` has complex

memory access patterns, however, its flat structure ensures cache side-channel freedom.

To check the effectiveness of our abstraction refinement process, we compare CACHEFIX with a variant of our checker where all predicates in $Pred_{set} \cup Pred_{tag}$ are considered. Therefore, such a variant does not employ any abstraction refinement, as the set of predicates $Pred_{set} \cup Pred_{tag}$ is sufficient to determine the cache behavior of all instructions in the program. We compare CACHEFIX with this variant in terms of the number of predicates, as well as the verification time. We record the set of predicates $Pred_{cur}$ considered in CACHEFIX when it terminates with a successful verification (✓) or a real counterexample (✗). Table I clearly demonstrates the effectiveness of our abstraction refinement process. Specifically, for AES, DES and `fix_pow`, the checker does not terminate in 10 min when all predicates in $Pred_{set} \cup Pred_{tag}$ are considered during the verification process. In general, the refinement process reduces the number of considered predicates by a factor of 1.81x on average. This leads to a substantial improvement in verification time, as observed from Table I.

The routines chosen from OpenSSL library are single path programs. However, AES and DES exhibit input-dependent memory accesses, hence, violating side-channel freedom. The other routines violate cache side-channel freedom due to input-dependent loop trip counts. For example, routines chosen from the GDK library employ a binary search of the input keystroke over a large table. We note that both libfixedtimefixedpoint and FourQlib libraries include comments involving the security risks in `fix_frac` and `ecc_point_validate` (see validate in Table I).

D. Overhead From Monitors

We evaluated the time taken by CACHEFIX for counterexample exploration and patch generation (see Section V).

TABLE II

OVERHEAD IN GENERATING MONITORS AND DUE TO THE APPLIED RUNTIME ACTIONS. Time CAPTURES THE TIME TO GENERATE ALL PATCHES. THE OVERHEAD (MAXIMUM AND AVERAGE) CAPTURES THE NUMBER OF EXTRA CACHE MISSES, HITS, AND INVALIDATIONS INTRODUCED BY CACHEFIX IN ABSOLUTE TERM (I.E., # ACTIONS) AND WITH RESPECT TO THE TOTAL NUMBER OF INSTRUCTIONS (I.E., % ACTIONS)

Routine	\mathcal{O}_{time}						\mathcal{O}_{trace}					
	#Equivalence class	Time (secs)	Max. overhead		Avg. overhead		#Equivalence class	Time (secs)	Max. overhead		Avg. overhead	
			(# actions)	(% actions)	(# actions)	(% actions)			(# actions)	(% actions)	(# actions)	(% actions)
AES128	3	5	117	0.1%	60	0.05%	38	1271	120	0.1%	90	0.07%
DES	333	444	300	0.3%	150	0.15%	982	12601	371	0.6%	231	0.4%
fix_frac	82	521	80	1.2%	40	0.6%	82	956	119	1.5%	62	1.1%
ecopt_validate	20	22	18	0.09%	9	0.05%	20	234	21	1.1%	11	0.04%
keyname	41	1119	39	2.6%	19	1.2%	41	1324	45	3.2%	28	2.1%
unicode	43	197	41	2.4%	20	1.2%	43	229	49	2.4%	28	1.8%

For each generated patch, we have also evaluated the overhead induced by the same at runtime (see Table II).

CACHEFIX verifies or generates real counterexamples at an average within 70.38 secs w.r.t. attack model \mathcal{O}_{time} and at an average within 74.12 secs w.r.t. \mathcal{O}_{trace} . Moreover, the abstraction refinement process embodied within CACHEFIX substantially reduces the verification time as opposed to when no refinement process was employed.

Table II captures the maximum and average overhead induced by CACHEFIX at runtime. We compute the overhead via the number of additional runtime actions (i.e., cache misses, hits or invalidations) introduced solely via CACHEFIX. In absolute terms, the maximum (average) overhead captures the maximum (average) number of runtime actions induced over all equivalence classes. The maximum overhead was introduced in case of DES – 300 actions for \mathcal{O}_{time} and 371 actions for \mathcal{O}_{trace} . This is primarily due to the difficulty in making a large number of traces equivalent in terms of the number of cache misses and the sequence of hit/miss, respectively. Although the number of actions introduced by CACHEFIX is non-negligible, we note that their effect is minimal on the overall execution. To this end, we execute the program for 100 different inputs in each explored equivalence class and measure the overhead introduced by CACHEFIX (see “% actions” in Table II) with respect to the total number of instructions executed. We observe that the maximum overhead reaches up to 3.2% and the average overhead is up to 2.1%. We believe this overhead is acceptable in the light of cache side-channel freedom guarantees provided by CACHEFIX.

Except AES and DES, the cache behavior of a single program path is independent of program inputs. For the respective subjects, exactly the same number of equivalence classes were explored for both attack models (see Table II). Each explored equivalence class was primarily attributed to a unique program path. Nevertheless, due to more involved computations (see Section V), the overhead of CACHEFIX in attack model \mathcal{O}_{trace} is higher than the overhead in attack model \mathcal{O}_{time} . As observed from Table II, CACHEFIX discovers significantly more equivalence classes with respect to attack model \mathcal{O}_{trace} as compared to the number of equivalence classes with respect to attack model \mathcal{O}_{time} . This implies AES is more vulnerable to \mathcal{O}_{trace} as compared to \mathcal{O}_{time} . Excluding DES subject to \mathcal{O}_{trace} , our exploration terminates in all scenarios within 20 mins.

To explore counterexample and generate patches, CACHEFIX takes 2.27x more time with \mathcal{O}_{trace} , as compared to \mathcal{O}_{time} . Moreover, excluding DES, CACHEFIX explores all equivalence classes of observations within 20 minutes in all scenarios. Finally, the runtime overhead induced by CACHEFIX is only up to 3.2% with respect to the number of executed instructions.

E. Sensitivity With Respect to Cache Configuration

We evaluated CACHEFIX for a variety of cache associativity (1-way, 2-way, and 4-way), cache size (from 1KB to 8KB) and with LRU as well as FIFO replacement policies (detailed experiments are included in the supplement [18]). We observed that the verification time increases marginally (about 7%) when set-associative caches were used instead of direct-mapped caches and does not vary significantly with respect to replacement policy. Finally, we observed changes in the number of equivalence classes of observations for both AES and DES while running these subjects with different replacement policies. However, neither AES nor DES satisfied cache side-channel freedom for any of the cache size and replacement policies tested in our evaluation. The relatively low verification time results from the fact that the total number of predicates (i.e., $\text{Pred}_{set} \cup \text{Pred}_{tag}$) is independent of cache size and replacement policy. Nevertheless, the symbolic encoding for set-associative caches is more involved than direct-mapped caches. This results in an average increase to the number of predicates considered for verification (i.e., Pred_{cur}) by a factor of 1.5x. However, such an increased number of predicates does not translate to significant verification timing for set-associative caches.

VII. REVIEW OF PRIOR WORKS

Earlier works on cache analysis are based on abstract interpretation [35] and its combination with model checking [17], to estimate the worst-case execution time (WCET) of a program. In contrast to these approaches, CACHEFIX automatically builds and refine the abstraction of cache semantics for verifying side-channel freedom. Cache attacks are one of the most critical side-channel attacks [7], [13], [25]–[28], [32], [37], [38]. In contrast to the literature on side-channel attacks, we do not engineer new cache attacks in this paper. Based on a configurable attack model, CACHEFIX verifies and reinstates the cache side-channel freedom of arbitrary programs.

Orthogonal to approaches proposing countermeasures [21], [36], the fixes generated by CACHEFIX is guided by

program verification output. Thus, CACHEFIX can provide cache side-channel freedom guarantees about the fixed program. Moreover, CACHEFIX can be leveraged to formally verify whether existing countermeasures are capable to ensure side-channel freedom.

In contrast to recent approaches on statically analyzing cache side channels [14], [22], [29], [30], our CACHEFIX approach automatically constructs and refines the abstractions for verifying cache side-channel freedom. Moreover, contrary to CACHEFIX, approaches based on static analysis are not directly applicable when the underlying program does not satisfy cache side-channel freedom. CACHEFIX targets verification of arbitrary software programs, over and above constant-time implementations [8], [11]. Existing works based on symbolic execution [10], [34], taint analysis [20], [33], and verifying timing-channel freedom [9] ignore cache attacks. Moreover, these works do not provide capabilities for automatic abstraction refinement and patch synthesis for ensuring side-channel freedom. Finally, in contrast to these works, we show that our CACHEFIX approach scales with routines from real cryptographic libraries.

Finally, recent approaches on testing and quantifying cache side-channel leakage [12], [15], [16] are complementary to CACHEFIX. These works have the flavor of testing and they do not provide capabilities to ensure cache side-channel freedom.

VIII. DISCUSSION

In this paper, we propose CACHEFIX, a novel approach to automatically verify and restore cache side-channel freedom of arbitrary programs. The key novelty in our approach is two fold. First, our CACHEFIX approach automatically builds and refines abstraction of cache semantics. Although targeted to verify cache side-channel freedom, we believe CACHEFIX is applicable to verify other cache timing properties, such as WCET. Second, the core symbolic engine of CACHEFIX systematically combines its reasoning power with runtime monitoring to ensure cache side-channel freedom during program execution. Our evaluation reveals promising results, for 25 routines from several cryptographic libraries, CACHEFIX (dis)proves cache side-channel freedom within an average 75 s. Moreover, in most scenarios, CACHEFIX generated patches within 20 min to ensure cache side-channel freedom during program execution. Despite this result, we believe that CACHEFIX is only an initial step for the automated verification of cache side-channel freedom. In particular, we do not account cache attacks that are more powerful than timing or trace-based attacks. Besides, we do not implement the synthesized patches in a commodity embedded system to check their performance impact. We hope that the community will take this effort forward and push the adoption of formal tools for the evaluation of cache side-channel. For reproducibility and research, our tool and all experimental data are publicly available: <https://bitbucket.org/sudiptac/cache-fix>.

ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewers for their insightful comments.

REFERENCES

- [1] *CBMC: Bounded Model Checking for Software*. Accessed: Oct. 23, 2017. [Online]. Available: <https://www.cprover.org/cbmc/>
- [2] *FourQLib Library*. Accessed: Oct. 20, 2017. [Online]. Available: <https://github.com/Microsoft/FourQLib>
- [3] *GDK Library*. Accessed: Oct. 20, 2017. [Online]. Available: <https://developer.gnome.org/gdk3/>
- [4] *A Library for Doing Constant-Time Fixed-Point Numeric Operations*. Accessed: Oct. 20, 2017. [Online]. Available: <https://github.com/imdea-software/verifying-constant-time>
- [5] *OpenSSL Library*. Accessed: Oct. 20, 2017. [Online]. Available: <https://github.com/openssl/openssl>
- [6] *SSW Library*. Accessed: Oct. 23, 2017. [Online]. Available: <https://github.com/mengyao/Complete-Striped-Smith-Waterman-Library>
- [7] O. Acıncımez and Ç. K. Koç, "Trace-driven cache attacks on AES," in *Information and Communications Security*. Heidelberg, Germany: Springer, 2006.
- [8] J. B. Almeida *et al.*, "Verifying constant-time implementations," in *Proc. USENIX*, 2016, pp. 53–70.
- [9] T. Antonopoulos *et al.*, "Decomposition instead of self-composition for proving the absence of timing channels," in *Proc. PLDI*, 2017, pp. 362–375.
- [10] M. Backes, B. Köpf, and A. Rybalchenko, "Automatic discovery and quantification of information leaks," in *Proc. IEEE S P*, 2009, pp. 141–153.
- [11] G. Barthe, G. Betarte, J. D. Campo, C. D. Luna, and D. Pichardie, "System-level non-interference for constant-time cryptography," in *Proc. CCS*, 2014, pp. 1267–1279.
- [12] T. Basu and S. Chattopadhyay, "Testing cache side-channel leakage," in *Proc. ICST Workshops*, 2017, pp. 51–60.
- [13] D. J. Bernstein, "Cache-timing attacks on AES," Dept. Math. Stat. Comput. Sci., Univ. Illinois at Chicago, Chicago, IL, USA, Rep., 2005. [Online]. Available: <http://palms.ee.princeton.edu/system/files/Cache-timing+attacks+on+AES.pdf>
- [14] P. Cañones, B. Köpf, and J. Reineke, "Security analysis of cache replacement policies," in *Proc. POST*, 2017, pp. 189–209.
- [15] S. Chattopadhyay, "Directed automated memory performance testing," in *Proc. TACAS*, 2017, pp. 38–55.
- [16] S. Chattopadhyay, M. Beck, A. Rezine, and A. Zeller, "Quantifying the information leak in cache attacks via symbolic execution," in *Proc. MEMOCODE*, 2017, pp. 25–35.
- [17] S. Chattopadhyay and A. Roychoudhury, "Scalable and precise refinement of cache timing analysis via path-sensitive verification," *Real Time Syst.*, vol. 49, no. 4, pp. 517–562, 2013.
- [18] S. Chattopadhyay and A. Roychoudhury, "Symbolic verification of cache side-channel freedom," *CoRR*, vol. abs/1807.04701, 2018. [Online]. Available: <http://arxiv.org/abs/1807.04701>
- [19] E. M. Clarke, A. Biere, R. Raimi, and Y. Zhu, "Bounded model checking using satisfiability solving," *Formal Methods Syst. Design*, vol. 19, no. 1, pp. 7–34, 2001.
- [20] J. Clause, W. Li, and A. Orso, "Dytan: A generic dynamic taint analysis framework," in *Proc. ISSTA*, 2007, pp. 196–206.
- [21] S. Crane, A. Homescu, S. Brunthaler, P. Larsen, and M. Franz, "Thwarting cache side-channel attacks through dynamic software diversity," in *Proc. 22nd Annu. Netw. Distrib. Syst. Security Symp. (NDSS)*, San Diego, CA, USA, 2015. [Online]. Available: <https://www.ndss-symposium.org/ndss2015/thwarting-cache-side-channel-attacks-through-dynamic-software-diversity>
- [22] G. Doychev, B. Köpf, L. Mauborgne, and J. Reineke, "CacheAudit: A tool for the static analysis of cache side channels," *TISSEC*, vol. 18, no. 1, 2015, Art. no. 4.
- [23] M. Lipp *et al.*, "Meltdown: Reading kernel memory from user space," in *Proc. 27th USENIX Security Symp. (USENIX Security)*, 2018.
- [24] P. Kocher *et al.*, "Spectre attacks: Exploiting speculative execution," in *Proc. 40th IEEE Symp. Security Privacy (S&P)*, 2019.
- [25] Q. Ge, Y. Yarom, D. Cock, and G. Heiser, "A survey of microarchitectural timing attacks and countermeasures on contemporary hardware," *J. Cryptograph. Eng.*, vol. 8, no. 1, pp. 1–27, 2018.
- [26] D. Gruss, R. Spreitzer, and S. Mangard, "Cache template attacks: Automating attacks on inclusive last-level caches," in *Proc. USENIX Security*, 2015, pp. 897–912.
- [27] R. Guanciale, H. Nemati, C. Baumann, and M. Dam, "Cache storage channels: Alias-driven attacks and verified countermeasures," in *Proc. IEEE Symp. Security Privacy*, 2016, pp. 38–55.
- [28] D. Gullasch, E. Bangerter, and S. Krenn, "Cache games—Bringing access-based cache attacks on AES to practice," in *Proc. IEEE Symp. Security Privacy*, 2011, pp. 490–505.

- [29] B. Köpf and D. A. Basin, "An information-theoretic model for adaptive side-channel attacks," in *Proc. CCS*, 2007, pp. 286–296.
- [30] B. Köpf, L. Mauborgne, and M. Ochoa, "Automatic quantification of cache side-channels," in *Proc. CAV*, 2012, pp. 564–580.
- [31] M. Lipp, D. Gruss, R. Spreitzer, C. Maurice, and S. Mangard, "ARMageddon: Cache attacks on mobile devices," in *Proc. USENIX Security Symp.*, 2016, pp. 549–564.
- [32] F. Liu, Y. Yarom, Q. Ge, G. Heiser, and R. B. Lee, "Last-level cache side-channel attacks are practical," in *Proc. IEEE Symp. Security Privacy*, 2015, pp. 605–622.
- [33] J. Newsome, S. McCamant, and D. Song, "Measuring channel capacity to distinguish undue influence," in *Proc. PLAS*, 2009, pp. 73–85.
- [34] C. S. Pasareanu, Q.-S. Phan, and P. Malacaria, "Multi-run side-channel analysis using symbolic execution and Max-SMT," in *Proc. CSF*, 2016, pp. 387–400.
- [35] H. Theiling, C. Ferdinand, and R. Wilhelm, "Fast and precise WCET prediction by separated cache and path analyses," *Real Time Syst.*, vol. 18, nos. 2–3, pp. 157–179, 2000.
- [36] Z. Wang and R. B. Lee, "New cache designs for thwarting software cache-based side channel attacks," in *Proc. ISCA*, 2007, pp. 494–505.
- [37] Y. Yarom and K. Falkner, "FLUSH+RELOAD: A high resolution, low noise, L3 cache side-channel attack," in *Proc. USENIX Security Symp.*, 2014, pp. 719–732.
- [38] Y. Yarom, D. Genkin, and N. Heninger, "CacheBleed: A timing attack on openssl constant-time RSA," *J. Cryptograph. Eng.*, vol. 7, no. 2, pp. 99–112, 2017.



Sudipta Chattopadhyay received the Ph.D. degree in computer science from the National University of Singapore, Singapore, in 2013.

He is an Assistant Professor with the Information Systems Technology and Design Pillar, Singapore University of Technology and Design, Singapore. In his doctoral dissertation, he researched on Execution-Time Predictability, focussing on Multicore Platforms. He seeks to understand the influence of execution platform on critical software properties, such as performance, energy, robustness, and security. His research interests include program analysis, embedded systems, and compilers.

Mr. Chattopadhyay serves in the review board of the IEEE TRANSACTIONS ON SOFTWARE ENGINEERING.



Abhik Roychoudhury received the Ph.D. degree in computer science from the State University of New York at Stony Brook, Stony Brook, NY, USA, in 2000.

He is a Professor of computer science with the National University of Singapore, Singapore. He is currently leading the TSUNAMI center, a large five-year long targeted research effort funded by the National Research Foundation in the domain of trust-worthy software. His research has been funded by various agencies and companies, including the National Research Foundation, Ministry of Education, A*STAR, Defense Research and Technology Office, DSO National Laboratories, Microsoft, and IBM. He has authored a book entitled *Embedded Systems and Software Validation (Systems-on-Silicon Series)* (Morgan Kaufmann, 2009), which has also been officially translated to Chinese by Tsinghua University Press. His research group has built scalable techniques for testing, debugging, and repair of programs using systematic semantic analysis. The research on automatically repairing programs at a large-scale contributes to the vision of self-healing software. His research interests include software testing and analysis, software security, and trust-worthy software construction.

Dr. Roychoudhury is also the Lead Principal Investigator of the Singapore Cyber-Security Consortium, which is a consortium of over 35 companies in the cyber-security space engaging with academia for research and collaboration. He has served in various capacities in the program committees and organizing committees of various conferences on software engineering, specifically serving as the Program Chair of ACM International Symposium on Software Testing and Analysis in 2016 and the General Chair of ACM SIGSOFT Symposium on Foundations of Software Engineering in 2022. He has served as an Editorial Board Member of the IEEE TRANSACTIONS ON SOFTWARE ENGINEERING from 2014 to 2018. He has been an ACM Distinguished Speaker since 2013.