Symbolic Verification of Cache Side-channel Freedom

Sudipta Chattopadhyay
Singapore University of Technology and Design

Abhik Roychoudhury
National University of Singapore

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 is to allow symbolic reasoning over counterexample traces and to combine 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 seconds. Besides, in all except one case, CACHEFIX synthesizes *all* patches within 20 minutes to ensure cache side-channel freedom of the respective routines during execution.

1 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. Firstly, 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.* non-spurious) counterexample is produced. Secondly, 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 behaviour 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 v on input variables and a sequence of actions that needs to be applied when the program is processed with inputs satisfying v. 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 [8, 14]. 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 [15, 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.

We organize the remainder of the paper as follows. After providing an overview of CACHEFIX (Section 2), 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 behaviour 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 4.3). In Section 4.4, we show the generalization of our CACHEFIX approach over timing-based attacks [14] and trace-based attacks [8].
(3) We discuss a systematic exploration of counterexamples to synthesize patches and to shield program executions against cache timing attacks (Section 5). We provide theoretical guarantees that such patch synthesis converges towards 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 non-spurious counterexamples within 75 seconds on average. Besides, in most cases, CACHEFIX generated all patches within

20 minutes to ensure cache side-channel freedom during execution. Our implementation and all experimental data are publicly available.

2 OVERVIEW

In this section, we demonstrate the general insight behind our approach through examples. We consider the simple code fragments in Figure 1(a)-(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 Figure 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 Figure 1(a)-(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 behaviour can be formulated via a well-defined set of predicates. To this end, let us assume $set(r_i)$ captures the cache set accessed by instruction r_i and $tag(r_i)$ captures the accessed cache tag by the same instruction. Consider the instruction r_3 in Figure 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:

$$\begin{aligned}\Gamma(r_3) \equiv \neg\ & 0 \leq key \leq 127 \wedge \rho_{13}^{set} \wedge \neg\rho_{13}^{tag} \\ \wedge \neg\ & key \geq 128 \wedge \rho_{23}^{set} \wedge \neg\rho_{23}^{tag}\end{aligned} \tag{1}$$

where $\rho_{ji}^{tag} \equiv tag(r_j) \neq tag(r_i)$ and $\rho_{ji}^{set} \equiv set(r_j) = set(r_i)$. 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 $miss_3 = 1$ (i.e. r_3 suffers a cache miss) and $miss_3 = 0$ (i.e. r_3 is a cache hit), otherwise. Formally, we set the cache behaviour of r_3 as follows:

$$\Gamma(r_3) \Leftrightarrow (miss_3 = 1)\,; \quad \neg\Gamma(r_3) \Leftrightarrow (miss_3 = 0) \tag{2}$$

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

In general, we note that the cache behaviour of the program in Figure 1(a), i.e., the cache behaviours of $r_1, \ldots, r_4$; can be formulated accurately via the following set of predicates related to cache semantics:

$$Pred_{cache} = \{\rho_{ji}^{set} \cup \rho_{ji}^{tag} \mid 1 \leq j < i \leq 4\} \tag{3}$$

The size of $Pred_{cache}$ depends on the number of memory-related instructions. However, $|Pred_{cache}|$ 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 Figure 1(a) when all feasible traces exhibit the same number of cache misses. Hence, such a property φ can be formulated as the non-existence of two traces tr_1 and tr_2 as follows:

$$\varphi \equiv \nexists tr_1, \nexists tr_2\ s.t. \left(\sum_{i=1}^{4} miss_i^{(tr_1)} \neq \sum_{i=1}^{4} miss_i^{(tr_2)}\right) \tag{4}$$

where $miss_i^{(tr)}$ 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 $Pred_{cache}$ (cf. Equation 3). In other words, even if some predicates in $Pred_{cache}$ 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 $Pred_{init} = \{\rho_{13}^{set}, \rho_{13}^{tag}, \rho_{23}^{set}, \rho_{23}^{tag}\}$. The rest of the predicates in $Pred_{cache} \setminus Pred_{init}$ are set to *unknown* value. With this configuration at hand, CacheFix returns counterexample traces tr_1 and tr_2 (cf. Equation 4) to reflect that φ does not hold for the program in Figure 1(a). In particular, the following traces are returned:

$$tr_1 \equiv \langle miss_1 = miss_3 = 1, miss_2 = miss_4 = 0\rangle$$
$$tr_2 \equiv \langle miss_1 = 0, miss_2 = miss_3 = miss_4 = 1\rangle$$

Given tr_1 and tr_2, we check whether any of them are spurious. To this end, we reconstruct the logical condition (cf. Equation 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 $Pred_{cache}$.

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:

$$\mathcal{U} \equiv \neg\rho_{34}^{set} \vee \rho_{34}^{tag} \tag{5}$$

Intuitively, with the initial abstraction $Pred_{init}$, 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:

$$Pred_{cur} = \{\rho_{13}^{set}, \rho_{13}^{tag}, \rho_{23}^{set}, \rho_{23}^{tag}, \rho_{34}^{set}, \rho_{34}^{tag}\}$$

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

$$Pred_{unknown} = \{\rho_{12}^{set}, \rho_{12}^{tag}, \rho_{14}^{set}, \rho_{14}^{tag}, \rho_{24}^{set}, \rho_{24}^{tag}\}$$

Therefore, it was possible to verify φ by tracking only half of the predicates in $Pred_{cache}$. 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 $Pred_{unknown}$ were not required for the verification process, as neither r_1 nor r_2 influences the cache behaviour 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 [9]. 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

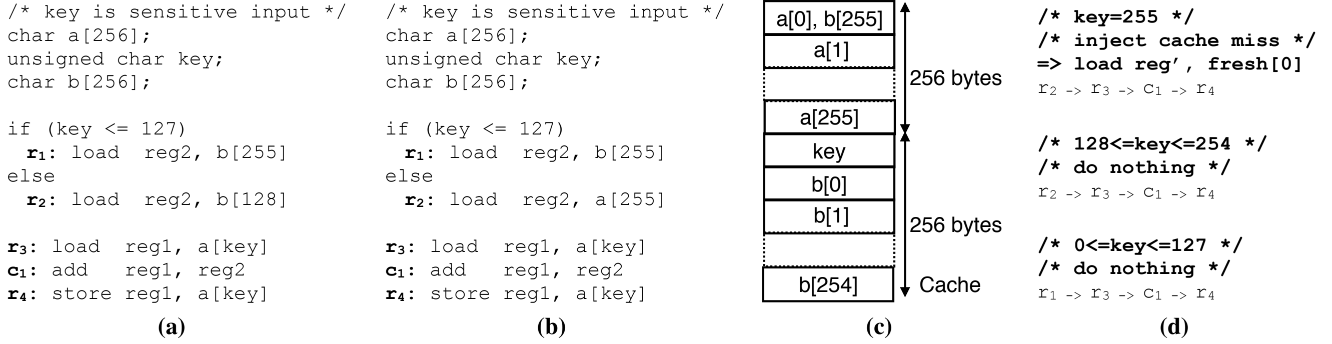

Figure 1: A code fragment (a) satisfying cache side-channel freedom, (b) violating cache side-channel freedom. (c) Mapping of variables into the cache. (d) Runtime actions and the execution order for the program in Figure 1(b) to ensure cache side-channel freedom.

behaviour. Besides, the example in Figure 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 behaviour. We accomplish this by leveraging our verification results.

We consider the example in Figure 1(b) and let us assume that we start with the initial abstraction $Pred_{init} = \{\rho_{13}^{set}, \rho_{13}^{tag}, \rho_{23}^{set}, \rho_{23}^{tag}\}$. CacheFix returns the following counterexample while verifying φ (cf. Equation 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 $miss_1, \ldots, miss_4$ 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 Figure 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 Figure 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 Figure 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**. Figure 2 outlines the overall workflow of CacheFix. The abstraction refinement process is guaranteed to converge towards the most precise abstraction of cache semantics to (dis)prove the cache side-channel freedom. Moreover, as observed in Figure 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.

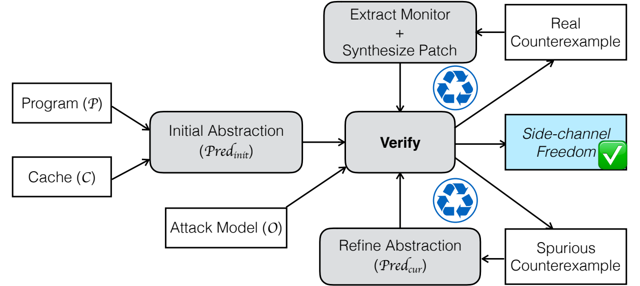

Figure 2: Workflow of our symbolic verification and patching

3 THREAT AND SYSTEM MODEL

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:

- $\mathcal{O}_{time} : \{h, m\}^* \rightarrow \mathbb{N}$. $\mathcal{O}_{time}$ maps each execution trace to the number of cache misses suffered by the same. This attack model imitates cache timing attacks [14].
- $\mathcal{O}_{trace} : \{h, m\}^* \rightarrow \{0, 1\}^*$. $\mathcal{O}_{trace}$ 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 [8].

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 [8, 14]. We currently do not handle more advanced attacks on shared caches [32]. First-level caches can either be partitioned (instruction vs. 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.

4 ABSTRACTION REFINEMENT

Notations. We represent cache via a triple $\langle 2^{\mathcal{S}}, 2^{\mathcal{B}}, \mathcal{A} \rangle$ where $2^{\mathcal{S}}$, $2^{\mathcal{B}}$ and $\mathcal{A}$ capture the number of cache sets, cache line size and cache associativity, respectively. We use $set(r_i)$ and $tag(r_i)$ 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 ($miss_i = 1$) or a hit ($miss_i = 0$). For instructions r_i and r_j, we have $j < i$ if and only if r_j was (symbolically) executed before r_i.

4.1 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 behaviour might be affected by all instructions executing prior to r_i. Concretely, the cache behaviour of r_i can be accurately predicted based on the set of logical predicates $Pred_{set}$ and $Pred_{tag}$ as follows:

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

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

$$Pred_{set} = \bigcup_{i=1} Pred^i_{set}; \quad Pred_{tag} = \bigcup_{i=1} Pred^i_{tag} \quad (7)$$

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

$$\begin{aligned} Pred_{init} = \bigcup_{i=1} & \; p \mid p \in Pred_{tag} \cup Pred_{set} \; \wedge |\sigma(r_i)| = 1 \; \wedge \\ & \forall k \in [1, i). \; |\sigma(r_k)| = 1 \wedge guard_k \Rightarrow true\} \end{aligned} \quad (8)$$

$\sigma(r_i)$ 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 $Pred_{init} = \phi$. 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: firstly, the set $Pred_{init}$ can be computed efficiently during symbolic execution. Secondly, as the predicates in $Pred_{init}$ 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 C, attack model O
Output: Successful verification or a concrete counterexample

1: /* Ψ is a formula representation of $\mathcal{P}$ */
2: /* $Pred$ is cache-semantics-related predicates */
3: /* Γ determines cache behaviour of all instructions */
4: $(\Psi, Pred, \Gamma)$:= ExecuteSymbolic($\mathcal{P}$, C)
5: /* Formulate initial abstraction (cf. Equation 8) */
6: $Pred_{cur}$:=$Pred_{init}$:= GetInitialAbstraction($Pred$)
7: /* Rewrite Ψ with initial abstraction */
8: Rewrite(Ψ, $Pred_{init}$)
9: /* Formulate cache side-channel freedom property */
10: φ := GetProperty(O)
11: /* Invoke symbolic verification to check $\Psi \wedge \neg\varphi$ */
12: (res, tr_1, tr_2) := Verify(Ψ, φ)
13: **while** (res=$false$) $\wedge$ (tr_1 or tr_2 is spurious) **do**
14: /* Extract unsatisfiable core from tr_1 and/or tr_2 */
15: $\mathcal{U}$:= UnsatCore(tr_1, tr_2, Γ)
16: /* Refine abstractions and repeat verification */
17: $Pred_{cur}$:= Refine($Pred_{init}$, $\mathcal{U}$, $Pred$)
18: Rewrite(Ψ, $Pred_{cur}$)
19: (res, tr_1, tr_2) := Verify(Ψ, φ)
20: $Pred_{init}$:= $Pred_{cur}$
21: **end while**
22: return res

4.2 Abstract domain refinement

We use the mapping $\Gamma : \{r_1, r_2, \ldots, r_N\} \rightarrow \{true, false\}$ 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:

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

$\Gamma(r_i)$ depends on predicates in $Pred^i_{set} \cup Pred^i_{tag}$ and the cache configuration. We show the formulation of $\Gamma(r_i)$ in Section 4.3.

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 $\mathcal{P}$ (*cf.* procedure ExecuteSymbolic) and systematically translating the cache and program semantics of each instruction into a set of constraints (*cf.* 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 $\Psi(r_i) \equiv \psi_1 \vee \psi_2 \vee \ldots \vee \psi_{j-1} \vee \psi_j$, 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 (*cf.* lines 13-14) and program semantics (*cf.* lines 18) of

Procedure 2 Symbolically Tracking Program and Cache States

```
1:  /* symbolically execute 𝒫 with cache configuration C*/
2:  procedure EXECUTESYMBOLIC(𝒫, C)
3:      i := 1; Ψ := true; Pred_set := Pred_tag := ϕ
4:      r_i := GETNEXTINSTRUCTION(𝒫)
5:      while r_i ≠ exit do
6:          if r_i is memory-related instruction then
7:              /* Collect predicates for cache semantics */
8:              Pred_set ∪ = Pred^i_set; Pred_tag ∪ = Pred^i_tag
9:              /* Γ(r_i) determines cache behaviour of r_i */
10:             Formulate Γ(r_i) /* see Section 4.3 */
11:             Γ ∪ = {Γ(r_i)}
12:             /* Integrate cache semantics within Ψ */
13:             Ψ := CONVERT(Ψ, Γ(r_i) ⇔ (miss_i = 1))
14:             Ψ := CONVERT(Ψ, ¬Γ(r_i) ⇔ (miss_i = 0))
15:         end if
16:         /* Integrate program semantics of r_i within Ψ */
17:         /* φ(r_i) is a predicate capturing r_i semantics */
18:         Ψ := CONVERT(Ψ, φ(r_i))
19:         i := i + 1
20:         r_i := GETNEXTINSTRUCTION(𝒫)
21:     end while
22:     return (Ψ, Pred_set ∪ Pred_tag, Γ)
23: end procedure
```

r_i. For instance, to integrate cache semantics of a memory-related instruction r_i, $\Psi(r_i)$ is converted to $\Psi(r_i) \wedge (\Gamma(r_i) \Leftrightarrow (miss_i = 1)) \wedge (\neg\Gamma(r_i) \Leftrightarrow (miss_i = 0))$. 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 (SSA) 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 (cf. Equation 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 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 non-existence of any two traces tr_1 and tr_2 that have different number of cache misses (cf. Equation 4). In other words, if the following formula is satisfied with more than one valuations for $\sum_{i=1}^{N} miss_i$, then side-channel freedom is violated:

$$\Psi \wedge \left(\left(\sum_{i=1}^{N} miss_i \right) \geq 0 \right) \tag{10}$$

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 $trace \in \{tr_1, tr_2\}$. Recall from Equation 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:

$$\Gamma_{trace} = \bigwedge_{i=1}^{N} \begin{cases} \Gamma(r_i), & \text{if } miss_i^{(trace)} = 1; \\ \neg\Gamma(r_i), & \text{if } miss_i^{(trace)} = 0; \end{cases} \tag{11}$$

In Equation 11, $miss_i^{(trace)}$ 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 $Pred_{init}$ 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 $Pred_{tag} \cup Pred_{set}$. Finally, we refine the abstraction (*cf.* procedure REFINE in Algorithm 1) to $Pred_{cur}$ by including all predicates in the unsatisfiable core as follows:

$$Pred_{cur} := Pred_{init} \cup \{p^+ \mid p^+ \in \mathsf{UnsatCore}(\Gamma_{trace}) \setminus Pred_{init}\} \tag{12}$$

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

4.3 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 appendix). To simplify the formulation, we will use the following abbreviations for the rest of the section:

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

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 $r \in \{r_1, r_2, \ldots, r_{i-1}\}$ access different memory blocks than the memory block accessed by r_i. This is captured as follows:

$$\Theta_i^{cold} \equiv \bigwedge_{j \in [1,i)} \neg\rho_{ji}^{set} \vee \rho_{ji}^{tag} \vee \neg guard_j \tag{14}$$

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:

$\phi_{\mathbf{ji}}^{\mathbf{cnf,dir}}$: 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:

$$\phi_{ji}^{cnf,dir} \equiv \rho_{ji}^{tag} \wedge \rho_{ji}^{set} \tag{15}$$

$\phi_{\mathbf{ji}}^{\mathbf{rel,dir}}$: 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 $r_{1\ldots3}$ 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. $\phi_{ji}^{rel,dir}$ is formally captured as follows:

$$\phi_{ji}^{rel,dir} \equiv \bigwedge_{j<k<i} \rho_{ki}^{tag} \vee \neg\rho_{ki}^{set} \vee \neg guard_k \tag{16}$$

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

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

$$\Theta_i^{cnf,dir} \equiv \bigvee_{j \in [1,i)} \phi_{ji}^{cnf,dir} \wedge \phi_{ji}^{rel,dir} \wedge guard_j \tag{17}$$

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:

$$\boxed{\Gamma(r_i) \equiv guard_i \wedge \left(\Theta_i^{cold} \vee \Theta_i^{cnf,dir}\right)} \tag{18}$$

4.4 Property for cache side-channel freedom

In this paper, we instantiate our checker for timing-based attacks [14] and trace-based attacks [8] 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:

$$\boxed{\Psi \wedge \left(\left(\sum_{i=1}^{N} miss_i\right) \geq 0\right)_{sol\left(\sum_{i=1}^{N} miss_i\right)} \leq 1} \tag{19}$$

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

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

$$\xi(r_i) = \begin{cases} 1, & \text{if } guard_i \wedge (miss_i = 1) \text{ holds;} \\ 0, & \text{if } guard_i \wedge (miss_i = 0) \text{ holds;} \end{cases} \tag{20}$$

The following verification goal ensures side-channel freedom:

$$\boxed{\Psi \wedge \left(|dom(\xi)| \geq 0 \;\wedge\; \left(\|_{r_i \in dom(\xi)}\, \xi(r_i)\right) \geq 0\right)_{sol(X)} \leq 1} \tag{21}$$

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

5 RUNTIME MONITORING

CacheFix produces the first real counterexample when it discovers two traces with different observations (w.r.t. attack model 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:

- We analyze a counterexample trace tr and extract the symbolic condition for which the same trace would be generated,
- We systematically explore unique counterexamples with the objective to reduce the uncertainty to guess sensitive inputs,
- 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.

Analyzing a 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.

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

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

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

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 (*cf.* Equation 20 for ξ):

$$\Psi \wedge \neg \left(\left(\sum_{i=1}^{N} miss_i \right) \neq o \right);$$
$$\Psi \wedge \neg \left(\begin{array}{c} |dom(\xi)| \neq \left[|dom(\xi)| \right]_o \\ \vee \;\; \|_{r_i \in dom(\xi)}\, \xi(r_i) \neq \left[\|_{r_i \in dom(\xi)}\, \xi(r_i) \right]_o \end{array} \right) \quad (23)$$

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 Equation 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 Equation 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 Equation 23 becomes unsatisfiable.

To explore a different equivalence class of observation than that of observation o, CacheFix negates the verification goal. For O_{time}, as an example, the verification goal is changed as follows:

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

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

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.

$O_{\mathbf{time}}$. Assume $\Omega = \{\langle v_1, o_1 \rangle, \langle v_2, o_2 \rangle, \ldots, \langle v_k, o_k \rangle\}$ 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:

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

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

$O_{\mathbf{trace}}$. 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:

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

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 Figure 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):

$$tr_1'' \equiv \langle \mathbf{miss}, miss, hit, hit \rangle; \;\; tr_2'' \equiv \langle miss, miss, hit, \mathbf{hit} \rangle$$

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

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

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

Practical consideration. In practice, the injection of a cache miss can be performed via accessing a fresh memory block (*cf.* Figure 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 re-accessing 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.

***Properties guaranteed by* CacheFix**. CacheFix satisfies the following crucial properties (proofs are included in the appendix) 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 O_{time} or 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 (cf. Equation 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 $\mathcal{G}_{cap}^{init}$, $\mathcal{G}_{shn}^{init}$ and $\mathcal{G}_{min}^{init}$, 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 $\mathcal{G}_{cap}^{init}$ (respectively, $\mathcal{G}_{shn}^{init}$ and $\mathcal{G}_{min}^{init}$) even in the presence of attacks captured via O_{time} and O_{trace}.*

6 IMPLEMENTATION AND EVALUATION

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 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. Secondly, 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.

Subject programs and cache. We have chosen security-critical subjects from `OpenSSL` [5], `GDK` [3], arithmetic routines from `libfixedtimefixedpoint` [4] and elliptic curve routines from `FourQlib` [2] to evaluate CACHEFIX (cf. Table 1). 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.

Efficiency of checking. Table 1 captures a summary of our evaluation for CACHEFIX. The outcome of this evaluation is either a successful verification (✓) or a non-spurious counterexample (✗). CACHEFIX accomplished the verification tasks for all subjects only within a few minutes. The maximum time taken by our checker was 390 seconds 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 behaviour 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 1 clearly demonstrates the effectiveness of our abstraction refinement process. Specifically, for `AES`, `DES` and `fix_pow`, the checker does not terminate in *ten minutes* 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 1.

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` (*cf.* `validate` in Table 1).

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

Overhead from monitors. We evaluated the time taken by CACHEFIX for counterexample exploration and patch generation (*cf.* Section 5). For each generated patch, we have also evaluated the overhead induced by the same at runtime (cf. Table 2).

Table 2 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 O_{time} and 371 actions for 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 (cf. "% actions" in Table 2) with respect to the total

Table 1: Summary of CACHEFIX Evaluation. Timeout is set to ten minutes. $Time$ 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)	$Pred_{set} \cup Pred_{tag}$	O_{time}				O_{trace}			
				$\lvert Pred_{cur} \rvert$	Result	$Time$ (secs)	$Time_{all}$ (secs)	$\lvert Pred_{cur} \rvert$	Result	$Time$ (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

Table 2: 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	O_{time}						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%
eccpt_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%

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 behaviour 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 (cf. Table 2). Each explored equivalence class was primarily attributed to a unique program path. Nevertheless, due to more involved computations (*cf.* Section 5), the overhead of CACHEFIX in attack model O_{trace} is higher than the overhead in attack model O_{time}. As observed from Table 2, CACHEFIX discovers significantly more equivalence classes w.r.t. attack model O_{trace} as compared to the number of equivalence classes w.r.t. attack model O_{time}. This implies AES is more vulnerable to O_{trace} as compared to O_{time}. Excluding DES subject to O_{trace}, our exploration terminates in all scenarios within 20 mins.

> *To explore counterexample and generate patches,* CacheFix *takes* $2.27x$ *more time with* O_{trace}, *as compared to* 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.*

Sensitivity w.r.t. 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 appendix). 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. $Pred_{set} \cup Pred_{tag}$) 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.

7 REVIEW OF PRIOR WORKS

Earlier works on cache analysis are based on abstract interpretation [35] and its combination with model checking [18], 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 [8, 14, 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 [15, 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 [9, 12]. Existing works based on symbolic execution [11, 34], taint analysis [20, 33] and verifying timing-channel freedom [10] 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 [13, 16, 17] are complementary to CacheFix. These works have the flavour of testing and and they do not provide capabilities to ensure cache side-channel freedom.

8 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. Firstly, 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. Secondly, 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 seconds. Moreover, in most scenarios, CacheFix generated patches within 20 minutes 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: **(blinded)**

REFERENCES

[1] [n. d.]. CBMC: Bounded Model Checking for Software. ([n. d.]). http://www.cprover.org/cbmc/ (Date last accessed 23-October-2017).

[2] [n. d.]. FourQLib Library. ([n. d.]). https://github.com/Microsoft/FourQlib/ (Date last accessed 20-October-2017).

[3] [n. d.]. GDK Library. ([n. d.]). https://developer.gnome.org/gdk3/3.22/ (Date last accessed 20-October-2017).

[4] [n. d.]. A library for doing constant-time fixed-point numeric operations. ([n. d.]). https://github.com/kmowery/libfixedtimefixedpoint/ (Date last accessed 20-October-2017).

[5] [n. d.]. OpenSSL Library. ([n. d.]). https://github.com/openssl/openssl/ (Date last accessed 20-October-2017).

[6] [n. d.]. SSW Library. ([n. d.]). https://github.com/mengyao/Complete-Striped-Smith-Waterman-Library (Date last accessed 23-October-2017).

[7] [n. d.]. Symbolic Verification of Cache Side-channel Freedom. ([n. d.]). https://github.com/esweek2018/emsoft2018/blob/master/cachefix_supplement.pdf.

[8] Onur Acıiçmez and Çetin Kaya Koç. 2006. Trace-driven cache attacks on AES. In *Information and Communications Security*. Springer.

[9] José Bacelar Almeida, Manuel Barbosa, Gilles Barthe, François Dupressoir, and Michael Emmi. 2016. Verifying Constant-Time Implementations. In *USENIX*. 53–70.

[10] Timos Antonopoulos, Paul Gazzillo, Michael Hicks, Eric Koskinen, Tachio Terauchi, and Shiyi Wei. 2017. Decomposition instead of self-composition for proving the absence of timing channels. In *PLDI*. 362–375.

[11] Michael Backes, Boris Köpf, and Andrey Rybalchenko. 2009. Automatic Discovery and Quantification of Information Leaks. In *IEEE S&P*. 141–153.

[12] Gilles Barthe, Gustavo Betarte, Juan Diego Campo, Carlos Daniel Luna, and David Pichardie. 2014. System-level Non-interference for Constant-time Cryptography. In *CCS*. 1267–1279.

[13] Tiyash Basu and Sudipta Chattopadhyay. 2017. Testing Cache Side-Channel Leakage. In *ICST Workshops*. 51–60.

[14] Daniel J Bernstein. 2005. Cache-timing attacks on AES. (2005).

[15] Pablo Cañones, Boris Köpf, and Jan Reineke. 2017. Security Analysis of Cache Replacement Policies. In *POST*. 189–209.

[16] Sudipta Chattopadhyay. 2017. Directed Automated Memory Performance Testing. In *TACAS*. 38–55.

[17] Sudipta Chattopadhyay, Moritz Beck, Ahmed Rezine, and Andreas Zeller. 2017. Quantifying the information leak in cache attacks via symbolic execution. In *MEMOCODE*. 25–35.

[18] Sudipta Chattopadhyay and Abhik Roychoudhury. 2013. Scalable and precise refinement of cache timing analysis via path-sensitive verification. *Real-Time Systems* 49, 4 (2013), 517–562.

[19] Edmund M. Clarke, Armin Biere, Richard Raimi, and Yunshan Zhu. 2001. Bounded Model Checking Using Satisfiability Solving. *Formal Methods in System Design* 19, 1 (2001), 7–34.

[20] James Clause, Wanchun Li, and Alessandro Orso. 2007. Dytan: a generic dynamic taint analysis framework. In *ISSTA*. 196–206.

[21] Stephen Crane, Andrei Homescu, Stefan Brunthaler, Per Larsen, and Michael Franz. 2015. Thwarting Cache Side-Channel Attacks Through Dynamic Software Diversity.. In *NDSS*.

[22] Goran Doychev, Boris Köpf, Laurent Mauborgne, and Jan Reineke. 2015. CacheAudit: a tool for the static analysis of cache side channels. *TISSEC* 18, 1 (2015), 4.

[23] Moritz Lipp et al. 2018. Meltdown. *ArXiv e-prints* (2018). arXiv:1801.01207

[24] Paul Kocher et al. 2018. Spectre Attacks: Exploiting Speculative Execution. *ArXiv e-prints* (2018). arXiv:1801.01203

[25] Qian Ge, Yuval Yarom, David Cock, and Gernot Heiser. 2016. A Survey of Microarchitectural Timing Attacks and Countermeasures on Contemporary Hardware. In *Cryptology ePrint Archive*. https://eprint.iacr.org/2016/613.pdf/.

[26] Daniel Gruss, Raphael Spreitzer, and Stefan Mangard. 2015. Cache Template Attacks: Automating Attacks on Inclusive Last-Level Caches. In *USENIX Security*.

[27] Roberto Guanciale, Hamed Nemati, Christoph Baumann, and Mads Dam. 2016. Cache Storage Channels: Alias-Driven Attacks and Verified Countermeasures. In *IEEE Symposium on Security and Privacy*. 38–55.

[28] David Gullasch, Endre Bangerter, and Stephan Krenn. 2011. Cache games–bringing access-based cache attacks on AES to practice. In *IEEE Symposium on Security and Privacy*. IEEE.

[29] Boris Köpf and David A. Basin. 2007. An information-theoretic model for adaptive side-channel attacks. In *CCS*. 286–296.

[30] Boris Köpf, Laurent Mauborgne, and Martín Ochoa. 2012. Automatic quantification of cache side-channels. In *CAV*. Springer.

[31] Moritz Lipp, Daniel Gruss, Raphael Spreitzer, Clémentine Maurice, and Stefan Mangard. 2016. ARMageddon: Cache Attacks on Mobile Devices. In *USENIX Security Symposium*. 549–564.

[32] Fangfei Liu, Yuval Yarom, Qian Ge, Gernot Heiser, and Ruby B. Lee. 2015. Last-Level Cache Side-Channel Attacks are Practical. In *IEEE Symposium on Security and Privacy*. 605–622.

[33] James Newsome, Stephen McCamant, and Dawn Song. 2009. Measuring channel capacity to distinguish undue influence. In *PLAS*. 73–85.

[34] Corina S. Pasareanu, Quoc-Sang Phan, and Pasquale Malacaria. 2016. Multi-run side-channel analysis using Symbolic Execution and Max-SMT. In *CSF*.

[35] Henrik Theiling, Christian Ferdinand, and Reinhard Wilhelm. 2000. Fast and precise WCET prediction by separated cache and path analyses. *Real-Time Systems* 18, 2-3 (2000).

[36] Zhenghong Wang and Ruby B. Lee. 2007. New Cache Designs for Thwarting Software Cache-based Side Channel Attacks. In *ISCA*. 494–505.

[37] Yuval Yarom and Katrina Falkner. 2014. FLUSH+RELOAD: A High Resolution, Low Noise, L3 Cache Side-Channel Attack. In *USENIX Security Symposium*. 719–732.

[38] Yuval Yarom, Daniel Genkin, and Nadia Heninger. 2017. CacheBleed: a timing attack on OpenSSL constant-time RSA. *J. Cryptographic Engineering* 7, 2 (2017), 99–112.

APPENDIX

The appendix includes additional cache models (e.g. LRU and FIFO) incorporated within CacheFix, the theoretical guarantees and additional experimental results.

Theoretical Guarantees

In this section, we include the detailed proof of the properties satisfied by CacheFix.

Property 3. ***(Monotonicity)*** *Consider a victim program $\mathcal{P}$ with sensitive input $\mathcal{K}$. Given attack models O_{time} or 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 (cf. Equation 25-26) employed at runtime.*

Proof. Consider the generic attack model $O : \{h, m\}^* \rightarrow \mathbb{X}$ that maps each trace to an element in the countable set $\mathbb{X}$. For a victim program $\mathcal{P}$, assume $TR \subseteq \{h, m\}^*$ is the set of all execution traces. After one round of patch synthesis, assume $TR' \subseteq \{h, m\}^*$ is the set of all execution traces in $\mathcal{P}$ when the synthesized patches are applied at runtime. By construction, each round of patch synthesis merges two equivalence classes of observations (*cf.* Algorithm 3), hence, making them indistinguishable by the attacker O. As a result, the following relationship holds:

$$|O(TR)| = |O(TR')| + 1 \tag{27}$$

Channel capacity $\mathcal{G}_{cap}^{\mathcal{P}}$ equals to $\log |O(TR)|$ for the original program $\mathcal{P}$, but it reduces to $\log |O(TR')|$ when the synthesized patches are applied. We conclude the proof as the same argument holds for any round of patch synthesis. □

Property 4. ***(Convergence)*** *Let us assume 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 $\mathcal{G}_{cap}^{init}$, $\mathcal{G}_{shn}^{init}$ and $\mathcal{G}_{min}^{init}$, via channel capacity, Shannon entropy and Min entropy, respectively. If our checker terminates and all synthesized patches are applied at runtime, then our framework guarantees that the channel capacity (respectively, Shannon entropy and Min entropy) will remain $\mathcal{G}_{cap}^{init}$ (respectively, $\mathcal{G}_{shn}^{init}$ and $\mathcal{G}_{min}^{init}$) even in the presence of attacks captured via O_{time} and O_{trace}.*

Proof. Consider the generic attack model $O : \{h, m\}^* \rightarrow \mathbb{X}$, mapping each execution trace to an element in the countable set $\mathbb{X}$. We assume a victim program $\mathcal{P}$ that exhibits a set of execution traces $TR \subseteq \{h, m\}^*$. From Equation 27, we know that $|O(TR)|$ decreases with each round of patch synthesis. Given that our checker terminates, we obtain the program $\mathcal{P}$, together with a set of synthesized patches that are applied when $\mathcal{P}$ executes. Assume $TR^f \in \{h, m\}^*$ is the set of execution traces obtained from $\mathcal{P}$ when all patches are systematically applied. Clearly, $\left|O\left(TR^f\right)\right| = 1$.

The channel capacity of $\mathcal{P}$, upon the termination of our checker, is $\log \left|O\left(TR^f\right)\right| = \log 1 = 0$. This concludes that the channel capacity does not change even in the presence of attacks O_{time} and O_{trace}.

For a given distribution λ of sensitive input $\mathcal{K}$, Shannon entropy $\mathcal{G}_{shn}^{init}$ is computed as follows:

$$\mathcal{G}_{shn}^{init}(\lambda) = -\sum_{\mathcal{K} \in \mathbb{K}} \lambda(\mathcal{K}) \log_2 \lambda(\mathcal{K}) \tag{28}$$

where $\mathbb{K}$ captures the domain of sensitive input $\mathcal{K}$. For a given equivalence class of observation $o \in O\left(TR^f\right)$, the remaining uncertainty is computed as follows:

$$\mathcal{G}_{shn}^{final}(\lambda_o) = -\sum_{\mathcal{K} \in \mathbb{K}} \lambda_o(\mathcal{K}) \log_2 \lambda_o(\mathcal{K}) \tag{29}$$

$\lambda_o(\mathcal{K})$ captures the probability that the sensitive input is $\mathcal{K}$, given the observation o is made by the attacker. Finally, to evaluate the remaining uncertainty of the patched program version, $\mathcal{G}_{shn}^{final}(\lambda_o)$ is averaged over all equivalence class of observations as follows:

$$\mathcal{G}_{shn}^{final}\left(\lambda_{O\left(TR^f\right)}\right) = \sum_{o \in O\left(TR^f\right)} pr(o)\, \mathcal{G}_{shn}^{final}(\lambda_o) \tag{30}$$

where $pr(o)$ captures the probability of the observation $o \in O\left(TR^f\right)$. However, we have $\left|O\left(TR^f\right)\right| = 1$. Hence, for any $o \in O\left(TR^f\right)$, we get $pr(o) = 1$ and $\lambda_o(\mathcal{K}) = \lambda(\mathcal{K})$. Plugging these observations into Equation 30 and Equation 29, we get the following:

$$\boxed{\begin{aligned} \mathcal{G}_{shn}^{final}\left(\lambda_{O\left(TR^f\right)}\right) &= \sum_{o \in O\left(TR^f\right)} \mathcal{G}_{shn}^{final}(\lambda_o) \\ &= -\sum_{\mathcal{K} \in \mathbb{K}} \lambda(\mathcal{K}) \log_2 \lambda(\mathcal{K}) \\ &= \mathcal{G}_{shn}^{init}(\lambda) \end{aligned}} \tag{31}$$

Finally, for a given distribution λ of sensitive input $\mathcal{K}$, the min entropy $\mathcal{G}_{min}^{init}$ is computed as follows:

$$\mathcal{G}_{min}^{init}(\lambda) = -\log_2 \max_{\mathcal{K} \in \mathbb{K}} \lambda(\mathcal{K}) \tag{32}$$

Therefore, min entropy captures the best strategy of an attacker, that is, to choose the most probable secret.

Similar to Shannon entropy, for a given equivalence class of observation $o \in O(TR^f)$, the remaining uncertainty is computed as follows:

$$\mathcal{G}^{init}_{min}(\lambda_o) = -\log_2 \max_{\mathcal{K}\in\mathbb{K}} \lambda_o(\mathcal{K}) \tag{33}$$

$\lambda_o(\mathcal{K})$ captures the probability that the sensitive input is $\mathcal{K}$, given the observation o is made by the attacker.

Finally, we obtain the min entropy of the patched program version via the following relation:

$$\mathcal{G}^{final}_{min}\left(\lambda_{O(TR^f)}\right) = -\log_2 \left(\sum_{o\in O(TR^f)} pr(o) \max_{\mathcal{K}\in\mathbb{K}} \lambda_o(\mathcal{K}) \right) \tag{34}$$

Since $pr(o) = 1$ and $\lambda_o(\mathcal{K}) = \lambda(\mathcal{K})$ for any $o \in O(TR^f)$, we get the following from Equation 34 and Equation 33:

$$\boxed{\begin{aligned} \mathcal{G}^{final}_{min}\left(\lambda_{O(TR^f)}\right) &= -\log_2 \left(\sum_{o\in O(TR^f)} \max_{\mathcal{K}\in\mathbb{K}} \lambda_o(\mathcal{K}) \right) \\ &= -\log_2 \max_{\mathcal{K}\in\mathbb{K}} \lambda(\mathcal{K}) \\ &= \mathcal{G}^{init}_{min}(\lambda) \end{aligned}} \tag{35}$$

Equation 31 and Equation 35 conclude this proof.

□

Modeling LRU and FIFO cache semantics

To formulate the conditions for conflict misses in set-associative caches, it is necessary to understand the notion of cache conflict. We use the following definition of cache conflict to formulate $\Gamma(r_i)$:

DEFINITION 1. ***(Cache conflict)*** r_j generates a cache conflict to r_i if and only if $1 \le j < i$, $\sigma(r_j) \ne \sigma(r_i)$ and the execution of r_j can change the relative position of $\sigma(r_i)$ within the $set(r_i)$-state immediately before instruction r_i.

Recall that $\sigma(r_i)$ captures the memory block accessed at r_i and $set(r_i)$ captures the cache set accessed by r_i. The state of a cache set is an ordered $\mathcal{A}$-tuple – capturing the relative positions of all memory blocks within the respective cache set. For instance, $\langle m_1, m_2\rangle$ captures the state of a two-associative cache set. The rightmost memory block (*i.e.* m_2) captures the first memory block to be evicted from the cache set if a block $m \notin \{m_1, m_2\}$ is accessed and mapped to the same cache set.

Challenges with LRU policy. To illustrate the unique challenges related to set-associative caches, let us consider the following sequence of memory accesses in a two-way associative cache and with LRU replacement policy: $(r_1 : m_1) \rightarrow (r_2 : m_2) \rightarrow (r_3 : m_2) \rightarrow (r_4 : m_1)$. We assume both m_1 and m_2 are mapped to the same cache set. If the cache is empty before r_1, r_4 will still incur a *cache hit*. This is because, r_4 suffers cache conflict only once, from the memory block m_2. To incorporate the aforementioned phenomenon into our cache semantics, we only count cache conflicts from the closest access to a given memory block. Therefore, in our example, we count cache conflicts to r_4 from r_3 and discard the cache conflict from r_2. Formally, we introduce the following additional condition for instruction r_j to inflict a cache conflict to instruction r_i.

$\boldsymbol{\phi^{eqv,lru}_{ji}}$: No instruction between r_j and r_i accesses the same memory block as r_j. This is to ensure that r_j is the closest to r_i in terms of accessing the memory block $\sigma(r_j)$. We capture $\phi^{eqv,lru}_{ji}$ formally as follows:

$$\phi^{eqv,lru}_{ji} \equiv \bigwedge_{j<k<i} \rho^{tag}_{jk} \vee \neg\rho^{set}_{jk} \vee \neg guard_k \tag{36}$$

Hence, r_j inflicts a unique cache conflict to r_i only if $\phi^{eqv,lru}_{ji}$, $\phi^{cnf,lru}_{ji} \equiv \phi^{cnf,dir}_{ji}$, $\phi^{rel,lru}_{ji} \equiv \phi^{rel,dir}_{ji}$ are all satisfiable.

Challenges with FIFO policy. Unlike LRU replacement policy, the cache state does not change for a cache hit in FIFO replacement policy. For example, consider the following sequence of memory accesses in a two-way associative FIFO cache: $(r_1 : m_1) \rightarrow (r_2 : m_2) \rightarrow (r_3 : m_1) \rightarrow (r_4 : m_1)$. Let us assume m_1, m_2 map to the same cache set and the cache is empty before r_1. In this example, r_2 generates a cache conflict to r_4 even though m_1 is accessed between r_2 and r_4. This is because r_3 is a cache hit and it does not change cache states.

In general, to formulate $\Gamma(r_i)$, we need to know whether any instruction r_j, prior to r_i, was a cache miss. This, in turn, is captured via $\Gamma(r_j)$. Concretely, r_j generates a unique cache conflict to r_i if all the following conditions are satisfied.

$\boldsymbol{\phi^{cnf,fifo}_{ji}}$: If r_j accesses the same cache set as r_i, but accesses a different cache-tag as compared to r_i and r_j suffers a cache miss. This is formalized as follows:

$$\phi^{cnf,fifo}_{ji} \equiv \rho^{tag}_{ji} \wedge \rho^{set}_{ji} \wedge \Gamma(r_j) \tag{37}$$

$\boldsymbol{\phi^{rel,fifo}_{ji}}$: No cache miss between r_j and r_i access the same memory block as r_i. $\phi^{rel,fifo}_{ji}$ ensures that the relative position of the memory block $\sigma(r_i)$ within $set(r_i)$ was not reset between r_j and r_i. $\phi^{rel,fifo}_{ji}$ is formalized as follows:

$$\phi^{rel,fifo}_{ji} \equiv \bigwedge_{j<k<i} \rho^{tag}_{ki} \vee \neg\rho^{set}_{ki} \vee \neg guard_k \vee \neg\Gamma(r_k) \tag{38}$$

$\boldsymbol{\phi^{eqv,fifo}_{ji}}$: No cache miss between r_j and r_i access the same memory block as r_j. We note that $\phi^{eqv,fifo}_{ji}$ ensures that r_j is the closest cache miss to r_i accessing the memory block $\sigma(r_j)$. This, in turn, ensures that we count the cache conflict from memory block $\sigma(r_j)$ to instruction r_i *only once*. We formulate $\phi^{eqv,fifo}_{ji}$ as follows:

$$\phi^{eqv,fifo}_{ji} \equiv \bigwedge_{j<k<i} \rho^{tag}_{jk} \vee \neg\rho^{set}_{jk} \vee \neg guard_k \vee \neg\Gamma(r_k) \tag{39}$$

Formulating cache conflict in set-associative caches. With the intuition mentioned in the preceding paragraphs, we formalize the unique cache conflict from r_j to r_i via the following logical conditions:

$$\Theta^{+,x}_{j,i} \equiv \left(\phi^{cnf,x}_{ji} \wedge \phi^{rel,x}_{ji} \wedge \phi^{eqv,x}_{ji} \wedge guard_j\right) \Rightarrow \eta_{ji} = 1 \tag{40}$$

$$\Theta_{j,i}^{-,x} \equiv \neg\phi_{ji}^{cnf,x} \vee \neg\phi_{ji}^{rel,x} \vee \neg\phi_{ji}^{eqv,x} \vee \neg guard_j \Rightarrow \eta_{ji} = 0 \tag{41}$$

where $x = \{lru, fifo\}$. Concretely, η_{ji} is set to 1 if r_j creates a unique cache conflict to r_i and η_{ji} is set to 0 otherwise.

Computing $\Gamma(r_i)$ for Set-associative Caches. To formulate $\Gamma(r_i)$ for set-associative caches, we need to check whether the number of unique cache conflicts to r_i exceeds the associativity ($\mathcal{A}$) of the cache. Based on this intuition, we formalize $\Gamma(r_i)$ for set-associative caches as follows:

$$\Gamma(r_i) \equiv guard_i \wedge \left(\Theta_i^{cold} \vee \left(\sum_{j\in[1,i)} \eta_{ji} \geq \mathcal{A}\right)\right) \tag{42}$$

We note that $\sum_{j\in[1,i)} \eta_{ji}$ accurately counts the number of unique cache conflicts to the instruction r_i (*cf.* Equation 40-Equation 41). Hence, the condition $\left(\sum_{j\in[1,i)} \eta_{ji} \geq \mathcal{A}\right)$ precisely captures whether $\sigma(r_i)$ is replaced from the cache before r_i is executed. If r_i does not suffer a cold miss and $\left(\sum_{j\in[1,i)} \eta_{ji} < \mathcal{A}\right)$, then r_i will be a cache hit when executed, as captured by the condition $\neg\Gamma(r_i)$.

Detailed Runtime Monitoring

Algorithm 3 outlines the overall process. The procedure MONITORING takes the following inputs:

- Ψ: A symbolic representation of program and cache semantics with the current level of abstraction,
- $\mathcal{O}$ and φ: The model of the attacker ($\mathcal{O}$) and a property φ initially capturing cache side-channel freedom w.r.t. $\mathcal{O}$,
- *Pred* and $Pred_{cur}$: Cache semantics related predicates (*Pred*) and the current level of abstraction ($Pred_{cur}$), and
- Γ: Symbolic conditions to determine the cache behaviour.

Additional Experimental Results

Sensitivity w.r.t. cache. Figure 4 outlines the evaluation for routines that violate side-channel freedom and for attack model $\mathcal{O}_{time}$. Nevertheless, the conclusion holds for all routines and attack models. Figure 4 captures the number of equivalence classes explored (hence, the number of patches generated *cf.* Algorithm 3) with respect to time. We make the following crucial observations from Figure 4. Firstly, the scalability of our checker is stable across a variety of cache configurations. This is because we encode cache semantics within a program via symbolic constraints on cache conflict. The size of these constraints depends on the number of memory-related instructions, but its size is not heavily influenced by the size of the cache. Secondly, the number of equivalence classes of observations does not vary significantly across cache configurations. Indeed, the number of equivalence classes may even increase (hence, increased channel capacity) with a bigger cache size (*e.g.* in DES and AES). However, for all cache configurations, CACHEFIX generated all the patches that need to be applied for making the respective programs cache side-channel free.

The scalability of CacheFix *is stable across a variety of cache configurations. Moreover, in all cache configurations,* CacheFix *generated all required patches to ensure the cache side-channel freedom w.r.t. $\mathcal{O}_{time}$.*

Algorithm 3 Monitor extraction and instrumentation

1: **procedure** MONITORING(Ψ, $\mathcal{O}$, φ, *Pred*, $Pred_{cur}$, Γ)
2: /* If φ captures side-channel freedom, then $trace$ is
3: any of the two traces constituting the counterexample */
4: $(res, trace)$:= VERIFY(Ψ, φ)
5: **while** $(res{=}false) \wedge (trace \neq$ spurious$)$ **do**
6: /* Extract observation from $trace$ */
7: o := GETOBSERVATION($trace$)
8: /* Extract monitor from $trace$ */
9: $v_o := v$:= EXTRACTMONITOR($trace$)
10: /* Refine Ψ to find unique counterexamples */
11: $\Psi := \Psi \wedge \neg v$
12: /* Refine φ to find all traces exhibiting o */
13: φ := REFINEOBJECTIVE(φ, o)
14: /* Check the unsatisfiability of $\Psi \wedge \neg\varphi$ */
15: $(res', trace')$:= VERIFY(Ψ, φ)
16: **while** $(res'{=}false)$ **do**
17: **if** $(trace' \neq$ spurious$)$ **then**
18: v := EXTRACTMONITOR($trace'$)
19: /* Combine monitors with observation o */
20: $v_o := v_o \vee v$
21: /* Refine Ψ for unique counterexamples */
22: $\Psi := \Psi \wedge \neg v$
23: /* Check the unsatisfiability of $\Psi \wedge \neg\varphi$ */
24: $(res', trace')$:= VERIFY(Ψ, φ)
25: **else**
26: /* Refine abstraction to repeat verification */
27: ABSREFINE(Ψ, *Pred*, $Pred_{cur}$, $trace'$, Γ)
28: $(res', trace')$:= VERIFY(Ψ, φ)
29: **end if**
30: **end while**
31: Let Ω holds the set of monitor, observation pairs
32: $\Omega \cup:= \{\langle v_o, o\rangle\}$
33: /* Instrument patches for monitor v_o */
34: INSTRUMENTPATCH(Ω, $\mathcal{O}$)
35: /* Refine objective to find new observations */
36: $\varphi := \neg\varphi$
37: /* Check the unsatisfiability of $\Psi \wedge \neg\varphi$ */
38: $(res, trace)$:= VERIFY(Ψ, φ)
39: **end while**
40: /* Program is still not side-channel free */
41: /* Refine abstraction to repeat verification loop */
42: **if** $(res{=}false)$ **then**
43: ABSREFINE(Ψ, *Pred*, $Pred_{cur}$, $trace$, Γ)
44: MONITORING(Ψ, $\mathcal{O}$, φ, *Pred*, $Pred_{cur}$, Γ)
45: **end if**
46: **end procedure**
47: **procedure** ABSREFINE(Ψ, *Pred*, $Pred_{cur}$, $trace$, Γ)
48: /* Extract unsatisfiable core */
49: $\mathcal{U}$:= UNSATCORE($trace$, Γ)
50: /* Refine abstractions (see Section 4) */
51: $Pred_{cur}$:= REFINE($Pred_{cur}$, $\mathcal{U}$, *Pred*)
52: /* Rewrite Ψ with the refined abstraction */
53: REWRITE(Ψ, $Pred_{cur}$)
54: **end procedure**

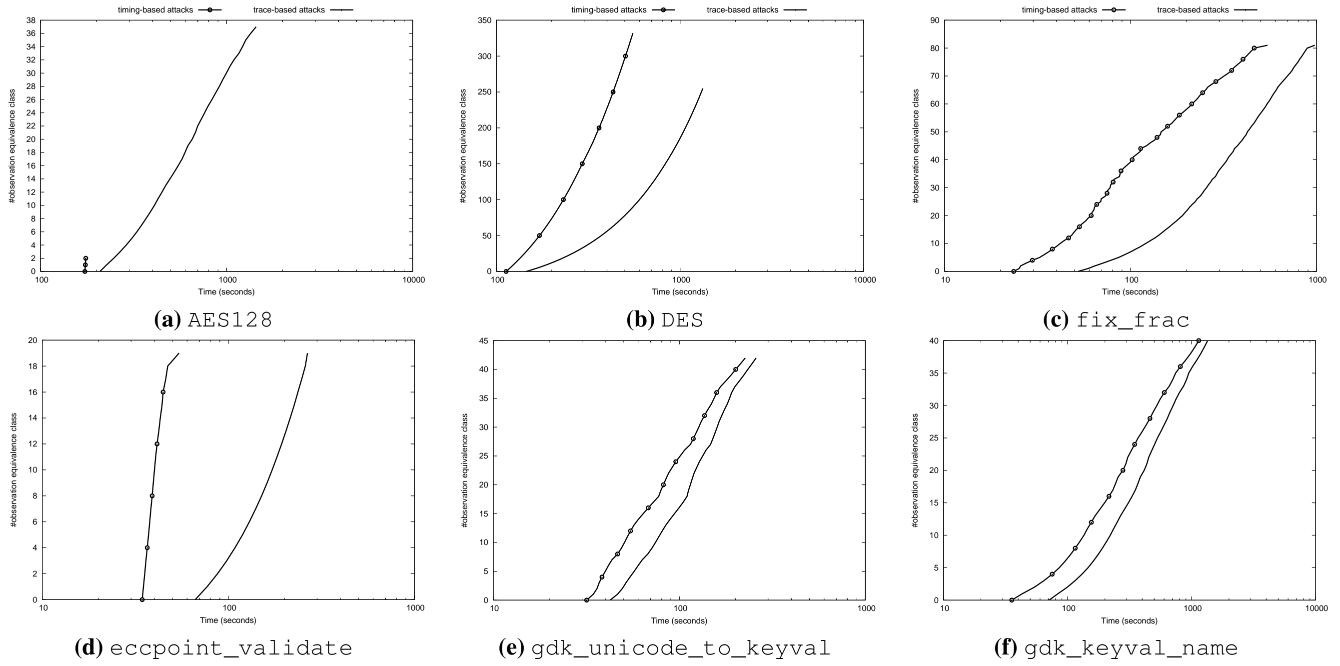

Figure 3: Overhead of counterexample exploration and patch synthesis

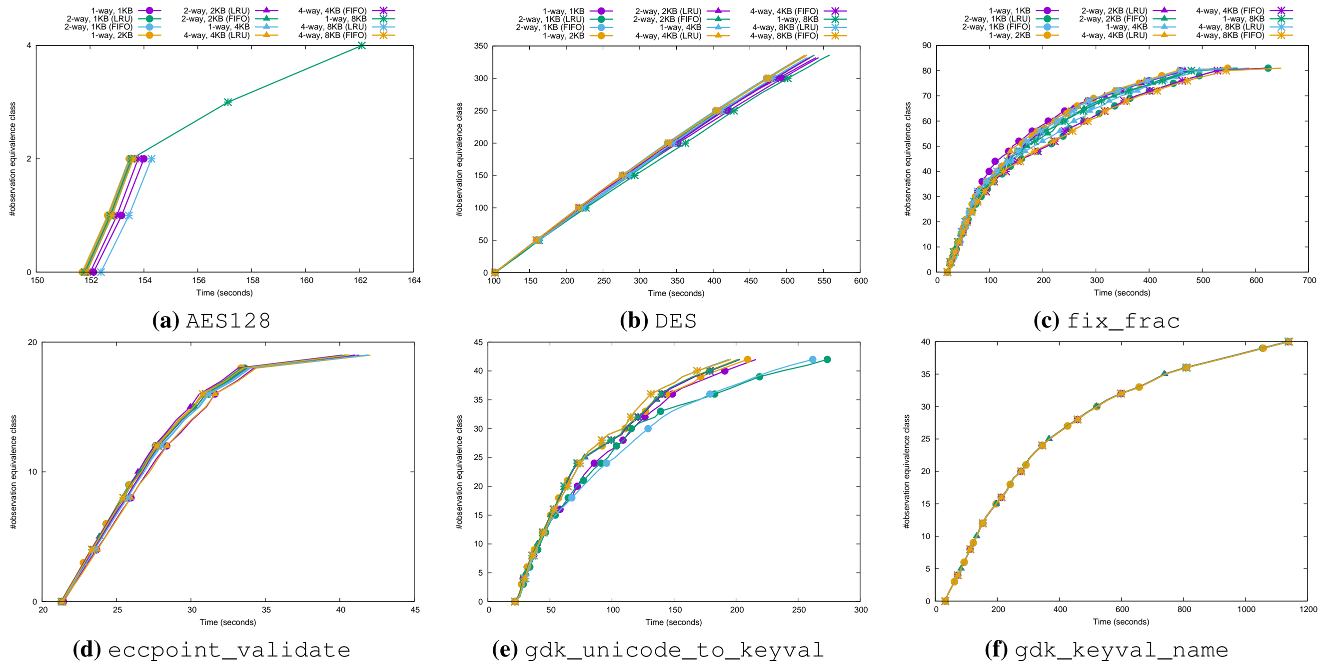

Figure 4: CACHEFIX sensitivity w.r.t. cache