

CacheAudit: A Tool for the Static Analysis of Cache Side Channels

Goran Doychev¹, Dominik Feld², Boris Köpf¹, Laurent Mauborgne¹, and Jan Reineke²

¹IMDEA Software Institute

²Saarland University

Abstract

We present CacheAudit, a versatile framework for the automatic, static analysis of cache side channels. CacheAudit takes as input a program binary and a cache configuration, and it derives formal, quantitative security guarantees for a comprehensive set of side-channel adversaries, namely those based on observing cache states, traces of hits and misses, and execution times.

Our technical contributions include novel abstractions to efficiently compute precise overapproximations of the possible side-channel observations for each of these adversaries. These approximations then yield upper bounds on the information that is revealed. In case studies we apply CacheAudit to binary executables of algorithms for symmetric encryption and sorting, obtaining the first formal proofs of security for implementations with countermeasures such as preloading and data-independent memory access patterns.

1 Introduction

Side-channel attacks recover secret inputs to programs from non-functional characteristics of computations, such as time [30], power [31], or memory consumption [26]. Typical goals of side-channel attacks are the recovery of cryptographic keys and private information about users.

Processor caches are a particularly rich source of side-channels because their behavior can be monitored in various ways, which is demonstrated by three documented classes of side-channel attacks: (1) In *time-based attacks* [30, 12] the adversary monitors the overall execution time of a victim, which is correlated with the number of cache hits and misses during execution. Time-based attacks are especially daunting because they can be carried out remotely over the network [7]. (2) In *access-based attacks* [39, 38, 23] the adversary probes the cache state by timing its own accesses to memory. Access-based attacks require that attacker and victim share the

same hardware platform, which is common in the cloud and has already been exploited [40, 48]. (3) In *trace-based attacks* [6] the adversary monitors the sequence of cache hits and misses. This can be achieved, e.g., by monitoring the CPU’s power consumption and is particularly relevant for embedded systems.

A number of proposals have been made for countering cache-based side-channel attacks. Some proposals focus entirely on modifications of the hardware platform; they either solve the problem for specific algorithms such as AES [2] or require modifications to the platform [45] that are so significant that their rapid adoption seems unlikely. The bulk of proposals relies on controlling the interactions between the software and the hardware layers, either through the operating system [23, 47], the client application [12, 38, 16], or both [28]. Reasoning about these interactions can be tricky and error-prone because it relies on the specifics of the binary code and the microarchitecture.

In this paper we present CacheAudit, a tool for the static exploration of the interactions of a program with the cache. CacheAudit takes as input a program binary and a cache configuration and delivers formal security guarantees that cover all possible executions of the corresponding system. The security guarantees are quantitative upper bounds on the amount of information that is contained in the side-channel observations of timing-, access-, and trace-based adversaries, respectively. CacheAudit can be used to formally analyze the effect on the leakage of software countermeasures and cache configurations, such as preloading of tables or increasing the cache’s line size. The design of CacheAudit is modular and facilitates the extension with any cache model for which efficient abstractions are in place. The current implementation of CacheAudit supports caches with LRU, FIFO, and PLRU replacement strategies.

We demonstrate the scope of CacheAudit in case studies where we analyze the side-channel leakage of representative algorithms for symmetric encryption and sort-

ing. We highlight the following two results: (1) For the reference implementation of the Salsa20 [11] stream cipher (which was designed to be resilient to cache side-channel attacks) CacheAudit can formally prove non-leakage on the basis of the binary executable, for all adversary models and replacement strategies. (2) For a library implementation of AES 128 [3], CacheAudit confirms that the preloading of tables significantly improves the security of the executable: for most adversary models and replacement strategies, we can in fact prove non-leakage of the executable, whenever the tables fit entirely into the cache. However, for access-based adversaries and LRU caches, CacheAudit reports small, non-zero bounds. And indeed, with LRU (as opposed to, e.g., FIFO), the *ordering* of blocks within a cache set reveals information about the victim’s final memory accesses.

On a technical level, we build on the fact that the amount of leaked information corresponds to the number of possible side-channel observations, which can be over-approximated by abstract interpretation and model counting [33]. To realize CacheAudit based on this insight, we propose three novel abstract domains that keep track of the observations of access-based, time-based, and trace-based adversaries, respectively. In particular:

1. We propose an abstract domain that tracks relational information about the memory blocks that may be cached. Opposed to existing abstract domains used in worst-case execution time analysis [21], our novel domain can retain analysis precision in the presence of array accesses to unknown positions.

2. We propose an abstract domain that tracks the traces of cache hits and misses that may occur during execution. We use a technique based on prefix trees and hash consing to compactly represent such sets of traces, and to count their number.

3. We propose an abstract domain that tracks the possible execution times of a program. This domain captures timing variations due to control flow and caches by associating hits and misses with their respective latencies and adding the execution time of the respective commands. We formalize the connection of these domains in an abstract interpretation [17] framework that captures the relationship between microarchitectural state and program code. We use this framework to formally prove the correctness of the derived upper bounds on the leakage to the corresponding side-channel adversaries.

In summary, our main contributions are both theoretical and practical: On a theoretical level, we define novel abstract domains that are suitable for the analysis of cache side channels, for a comprehensive set of adversaries. On a practical level, we build CacheAudit, the first tool for the automatic, quantitative information-flow analysis of cache side-channels, and we show how it can be used to derive formal security guarantees from

binary executables of sorting algorithms and state-of-the-art cryptosystems.

Outline The remainder of the paper is structured as follows. In Section 2, we illustrate the power of CacheAudit on a simple example program. In Section 3 we define the semantics and side channels of programs. We describe the analysis framework, the design of CacheAudit, and the novel abstract domains in Sections 4, 5 and 6, respectively. We present experimental results in Section 7, before we discuss prior work and conclude in Sections 8 and 9.

2 Illustrative Example

In this section, we illustrate on a simple example program the kind of guarantees CacheAudit can derive. Namely, we consider an implementation of BubbleSort that receives its input in an array *a* of length *n*. We assume that the contents of *a* are secret and we aim to deduce how much information a cache side-channel adversary can learn about the relative ordering of the elements of *a*.

```

1 void BubbleSort(int a[], int n)
2 {
3     int i, j, temp;
4     for (i = 0; i < (n - 1); ++i)
5         for (j = 0; j < n - 1 - i; ++j)
6             if (a[j] > a[j+1])
7                 {
8                     temp = a[j+1];
9                     a[j+1] = a[j];
10                    a[j] = temp;
11                }
12 }
```

To begin with, observe that the conditional swap in lines 6–11 is executed exactly $\frac{n(n-1)}{2}$ times. A *trace-based* adversary that can observe, for each instruction, whether it corresponds to a cache hit or a miss is likely to be able to distinguish between the two alternative paths in the conditional swap, hence we expect this adversary to be able to distinguish between $2^{\frac{n(n-1)}{2}}$ execution traces. A *timing-based* adversary who can observe the overall execution time is likely to be able to distinguish between $\frac{n(n-1)}{2} + 1$ possible execution times, corresponding to the number of times the swap has been carried out. For an *access-based* adversary who can probe the final cache state upon termination, the situation is more subtle: evaluating the guard in line 6 requires accessing both *a*[*j*] and *a*[*j*+1], which implies that both will be present in the cache when the swap in lines 8–10 is carried out. Assuming we begin with an empty cache, we expect that there is only one possible final cache state.

CacheAudit enables us to perform such analyses (for a particular n) formally and automatically, based on actual x86 binary executables and different cache types. CacheAudit achieves this by tracking compact representations of supersets of possible cache states and traces of hits and misses, and by counting the corresponding number of elements. For the above example, CacheAudit was able to precisely confirm the intuitive bounds, for a random selection of n in $\{2, \dots, 64\}$.

In terms of security, the number of possible observations corresponds to the factor by which the cache observation increases the probability of correctly guessing the secret ordering of inputs. Hence, for $n = 32$ and a uniform distribution on this order (i.e. an initial probability of $\frac{1}{32!} = 3.8 \cdot 10^{-36}$), the bounds derived by CacheAudit imply that the probability of determining the correct input order from the side-channel observation is 1 for a trace-based adversary, $3.7 \cdot 10^{-33}$ for a time-based adversary, and remains unchanged for an access-based adversary.

3 Caches, Programs, and Side Channels

3.1 A Primer on Caches

Caches are fast but small memories that store a subset of the main memory’s contents to bridge the latency gap between the CPU and main memory. To profit from spatial locality and to reduce management overhead, main memory is logically partitioned into a set of *memory blocks* \mathcal{B} . Blocks are cached as a whole in cache lines of the same size.

When accessing a memory block, the cache logic has to determine whether the block is stored in the cache (“cache hit”) or not (“cache miss”). To enable an efficient look-up, each block can only be stored in a small number of cache lines. For this purpose, caches are partitioned into equally-sized *cache sets*. The size of a cache set is called the *associativity* k of the cache. There is a function *set* that determines the cache set a memory block maps to.

Since the cache is much smaller than main memory, a *replacement policy* must decide which memory block to replace upon a cache miss. Usually, replacement policies treat sets independently, so that accesses to one set do not influence replacement decisions in other sets. Well-known replacement policies in this class are least-recently used (LRU), used in various Freescale processors such as the MPC603E and the TriCore17xx; pseudo-LRU (PLRU), a cost-efficient variant of LRU, used in the Freescale MPC750 family and multiple Intel microarchitectures; first-in first-out (FIFO), also known as ROUND ROBIN, used in several ARM and Freescale processors such as the ARM922 and the Freescale MPC7450 family.

A more comprehensive overview can be found in [22].

3.2 Programs and Computations

A program $P = (\Sigma, I, F, \mathcal{E}, \mathcal{T})$ consists of the following components

- Σ - a set of *states*
- $I \subseteq \Sigma$ - a set of *initial* states
- $F \subseteq \Sigma$ - a set of *final* states
- \mathcal{E} - a set of *events*
- $\mathcal{T} \subseteq \Sigma \times \mathcal{E} \times \Sigma$ - a *transition relation*

A *computation* of P is an alternating sequence of states and events $\sigma_0 e_0 \sigma_1 e_1 \dots \sigma_n$ such that $\sigma_0 \in I$ and that for all $i \in \{0, \dots, n-1\}$, $(\sigma_i, e_i, \sigma_{i+1}) \in \mathcal{T}$. The set of all computations is the *trace collecting semantics* $Col(P) \subseteq Traces$ of a program, where *Traces* denotes the set of all alternating sequences of states and events. When considering terminating programs, the trace collecting semantics can be formally defined as the least fixpoint of the *next* operator containing I :

$$Col(P) = \text{lfp}_T^{\subseteq} \lambda S. S \cup \text{next}(S),$$

where *next* describes the effect of one computation step:

$$\text{next}(S) = \{t. \sigma_n e_n \sigma_{n+1} \mid t. \sigma_n \in S \wedge (\sigma_n, e_n, \sigma_{n+1}) \in \mathcal{T}\}$$

3.3 Cache Updates and Cache Effects

For reasoning about cache side channels, we consider a semantics in which the cache is part of the program state. Namely, the state will consist of logical memories (representing values of memory locations and registers) in \mathcal{M} and a cache state in \mathcal{C} , i.e., $\Sigma = \mathcal{M} \times \mathcal{C}$.

The *memory update* $upd_{\mathcal{M}}$ is a function $upd_{\mathcal{M}}: \mathcal{M} \rightarrow \mathcal{M}$ that is determined solely by the instruction set semantics. The memory update has effects on the cache that are described by a function $eff_{\mathcal{M}}: \mathcal{M} \rightarrow \mathcal{E}_{\mathcal{M}}$. The memory effect is an argument to the *cache update* function $upd_{\mathcal{C}}: \mathcal{C} \times \mathcal{E}_{\mathcal{M}} \rightarrow \mathcal{C}$.

In the setting of this paper, $eff_{\mathcal{M}}$ determines which block of main memory is accessed, which is required to compute the cache update $upd_{\mathcal{C}}$, i.e., $\mathcal{E}_{\mathcal{M}} = \mathcal{B} \cup \{\perp\}$, where \perp denotes that no memory block is accessed.

We formally describe $upd_{\mathcal{C}}$ only for the LRU strategy. For formalizations of other strategies, see [22]. Upon a cache miss, LRU replaces the least-recently-used memory block. To this end, it tracks the ages of memory blocks within each cache set, where the youngest block has age 0 and the oldest cached block has age $k-1$. Thus, the state of the cache can be modeled as a function that assigns an age to each memory block, where

non-cached blocks are assigned age k :

$$\mathcal{C} := \{c \in \mathcal{B} \rightarrow A \mid \forall a, b \in \mathcal{B} : a \neq b \Rightarrow ((\text{set}(a) = \text{set}(b)) \Rightarrow (c(a) \neq c(b) \vee c(a) = c(b) = k))\},$$

where $A := \{0, \dots, k-1, k\}$ is the set of ages. The constraint encodes that no two blocks in the same cache set can have the same age. For readability we omit the additional constraint that blocks of non-zero age are preceded by other blocks, i.e. that cache sets do not contain “holes”.

The cache update for LRU is then given by

$$\text{upd}_{\mathcal{C}}(c, b) := \lambda b' \in \mathcal{B}. \begin{cases} 0 & : b' = b \\ c(b') & : \text{set}(b') \neq \text{set}(b) \\ c(b') + 1 & : \text{set}(b') = \text{set}(b) \wedge c(b') < c(b) \\ c(b') & : \text{set}(b') = \text{set}(b) \wedge c(b') \geq c(b) \end{cases}$$

In the setting of this paper, the events \mathcal{E} consist of cache hits and misses, which are described by the cache effect $\text{eff}_{\mathcal{C}} : \mathcal{C} \times \mathcal{B} \rightarrow \mathcal{E}$:

$$\text{eff}_{\mathcal{C}}(c, m) := \begin{cases} \text{hit} & : c(m) < k \\ \text{miss} & : \text{else} \end{cases}$$

Both $\text{upd}_{\mathcal{C}}$ and $\text{eff}_{\mathcal{C}}$ are naturally extended to the case where no memory access occurs. Then, the cache state remains unchanged and the cache effect is \perp . So $\mathcal{E} = \{\text{hit}, \text{miss}, \perp\}$.

With this, we can now connect the components and obtain the global transition relation $\mathcal{T} \subseteq \Sigma \times \mathcal{E} \times \Sigma$ by

$$\mathcal{T} = \{((m_1, c_1), e, (m_2, c_2)) \mid m_2 = \text{upd}_{\mathcal{M}}(m_1) \wedge c_2 = \text{upd}_{\mathcal{C}}(c_1, \text{eff}_{\mathcal{M}}(m_1)) \wedge e = \text{eff}_{\mathcal{C}}(c_1, \text{eff}_{\mathcal{M}}(m_1))\},$$

which formally captures the asymmetric relationship between caches, logical memories, and events.

3.4 Side Channels

For a deterministic, terminating program P , the transition relation is a function, and the program can be modeled as a mapping $P : I \rightarrow \text{Traces}$.

We model an adversary’s view on the computations of P as a function $\text{view} : \text{Traces} \rightarrow \mathcal{O}$ mapping traces to a finite set of observations \mathcal{O} . The composition $C = \text{view} \circ P : I \rightarrow \mathcal{O}$ defines a mapping from initial states to observations, which we call a *channel* of P . Whenever view is determined by the cache and event components of traces, we call C a *side channel* of P .

We next define views corresponding to the observations of access-based, trace-based, and timing-based side-channel adversaries.

The view of an *access-based* adversary that shares the memory space with the victim is defined by

$$\text{view}^{\text{acc}} : (m_0, c_0)e_0 \dots e_{n-1}(m_n, c_n) \mapsto c_n$$

and captures that the adversary can determine (by probing) which memory blocks are contained in the cache upon termination of the victim. An adversary that does *not* share the memory space with the victim can only observe how many blocks the victim has loaded (by probing how many of its own blocks have been evicted), but not which. We denote this view by $\text{view}^{\text{accd}}$. The view of a *trace-based* adversary is defined by

$$\text{view}^{\text{tr}} : \sigma_0 e_0 \dots e_{n-1} \sigma_n \mapsto e_0 \dots e_{n-1}$$

and captures that the adversary can determine for each instruction whether it results in a hit, miss, or does not access memory. The view of a *time-based* adversary is defined by

$$\text{view}^{\text{time}} : \sigma_0 e_0 \dots e_{n-1} \sigma_n \mapsto t_{\text{hit}} \cdot |\{i \mid e_i = \text{hit}\}| + t_{\text{miss}} \cdot |\{i \mid e_i = \text{miss}\}| + t_{\perp} \cdot |\{i \mid e_i = \perp\}|,$$

and captures that the adversary can determine the overall execution time of the program. Here, t_{hit} , t_{miss} , and t_{\perp} are the execution times (e.g. in clock cycles) of instructions that imply cache hits, cache misses, or no memory accesses at all. While the view of the time-based adversary as defined above is rather simplistic, e.g. disregarding effects of pipelining and out-of-order execution, notice that our semantics and our tool can be extended to cater for a more fine-grained, instruction- and context-dependent modeling of execution times. We denote the side channels corresponding to the four views by C^{acc} , C^{accd} , C^{tr} , and C^{time} , respectively.

3.5 Quantification of Side Channels

We characterize the security of a channel $C : I \rightarrow \mathcal{O}$ as the difficulty of guessing the secret input from the channel output.

Formally, we model the choice of a secret input by a random variable X with $\text{ran}(X) \subseteq I$ and the corresponding observation by a random variable $C(X)$ with $\text{ran}(C(X)) \subseteq \mathcal{O}$. We model the attacker as another random variable \hat{X} . The goal of the attacker is to estimate the value of X , i.e. it is successful if $\hat{X} = X$. We make the assumption that the attacker does not have information about the value of X beyond what is contained in $C(X)$, which we formalize as the requirement that $X \rightarrow C(X) \rightarrow \hat{X}$ form a Markov chain. The following theorem expresses a security guarantee in terms of an upper bound on the attacker’s success probability.

Theorem 1. *Let $X \rightarrow C(X) \rightarrow \hat{X}$ be a Markov chain. Then*

$$P(X = \hat{X}) \leq \max_{\sigma \in I} P(X = \sigma) \cdot |\text{ran}(C)|$$

For the interpretation of the statement observe that if the adversary has no information about the value of X (i.e., if \hat{X} and X are statistically independent), its success probability is bounded by the probability of the most likely value of X , i.e. $P(X = \hat{X}) \leq \max_{\sigma \in I} P(X = \sigma)$, where equality can be achieved. Theorem 1 hence states that the size of the active range of C is an upper bound on the factor by which this probability is increased when the attacker sees $C(X)$ and is, in that sense, an upper bound for the amount of information leaked by C . We will often give bounds on $|\text{ran}(C)|$ on a log-scale, in which case they represent upper bounds on the number of leaked bits.

For a formal connection to traditional (entropy-based) presentations of quantitative information-flow analysis [42] and a proof of Theorem 1, see Appendix B.

3.6 Adversarially-chosen Cache States

We sometimes assume that initial states are pairs consisting of *high* and *low* components, i.e. $I = I_{hi} \times I_{lo}$, where only the high component is meant to be kept secret and the low component may be provided by the adversary, a common setting in information-flow analysis [41]. In this case, a program and a view define a *family* of channels $C_{\sigma_{lo}} : I_{hi} \rightarrow O$, one for each low component $\sigma_{lo} \in I_{lo}$.

A particularly interesting instance is the decomposition into secret memory $I_{hi} = \mathcal{M}$ and adversarially chosen cache $I_{lo} = \mathcal{C}$. While bounds for the corresponding channel can be derived by considering all possible initial cache states, corresponding analyses suffer from poor precision. The following lemma enables us to derive bounds for the general case, based on the empty cache state. Notice that it only holds for access-driven adversaries.

Lemma 1. *For any $adv \in \{acc, accd\}$ and initial cache state $c \in \mathcal{C}$, in any of those two cases:*

- *The strategy is LRU*
- *or the strategy is PLRU or FIFO and no block in c is accessed during program execution,*

we have:

$$\left| \text{ran}(C_{\emptyset}^{adv}) \right| = \left| \text{ran}(C_c^{adv}) \right|,$$

where \emptyset is a shorthand for the empty cache state.

This lemma was proved in [32] for the LRU case with initial cache state not containing any block of the victim, and since the proof is based on the fact that memory blocks in the cache do not affect the position of memory

blocks that are accessed during computation when the two sets of memory blocks are disjoint, the proof can be straightforwardly extended to the FIFO and PLRU cases. That means that results on the empty initial cache state hold in all cases when the attacker cannot put memory blocks from the victim process in the cache. For LRU, we can remove the restriction on the initial cache state because the position in the cache of any block accessed during a given computation does not depend on its initial position, and any other memory block can be considered as in a disjoint memory space.

4 Automatic Quantification of Cache Side Channels

Theorem 1 enables the quantification of cache side channels by determining their range. However, computing this range exactly is practically infeasible in most cases. Abstract interpretation [17] overcomes this fundamental problem by resorting to an approximation of the state space and the transition relation. In addition, abstract interpretation allows us to design abstractions in a modular way, so that the soundness proofs can easily be split into independent lemmas and the tool we built from the framework described in this section is easily extensible and parametric.

4.1 Sound Abstraction of Leakage

As we have a fixpoint characterization of the collecting semantics, it is easy to frame a sound static analysis using abstract interpretation: we design an abstract domain Traces^\sharp where the meaning of abstract elements is given by a function $\gamma : \text{Traces}^\sharp \rightarrow \mathcal{P}(\text{Traces})$. If we have an abstract transfer function $\text{next}^\sharp : \text{Traces}^\sharp \rightarrow \text{Traces}^\sharp$ such that the following local soundness condition holds:

$$\forall a \in \text{Traces}^\sharp : \gamma(\text{next}^\sharp(a)) \supseteq \text{next}(\gamma(a)), \quad (1)$$

then any post-fixpoint of next^\sharp that is greater than an abstraction of the inputs I is a sound over-approximation of the collecting semantics. We use $\text{Col}(P)^\sharp$ to denote any such post-fixpoint.

Theorem 2 (Local soundness implies global soundness, from [17]).

$$(1) \Rightarrow \text{Col}(P) \subseteq \gamma(\text{Col}(P)^\sharp).$$

An immediate consequence of Theorem 2 is the following statement that shows how a sound abstract analysis can be used to derive upper bounds for the leaked information (see Theorem 1).

Theorem 3 (Upper bounds on leakage). *For $adv \in \{tr, time, acc\}$, we have*

$$\left| view^{adv} \left(\gamma \left(Col(P)^\sharp \right) \right) \right| \geq \left| ran(C^{adv}) \right|.$$

4.2 Abstraction Using a Control Flow Graph

In order to simplify the problem and to come up with tractable and modular abstractions, we design independent abstractions for cache states, memory, and sequences of events.

- \mathcal{M}^\sharp abstracts memory and $\gamma_{\mathcal{M}} : \mathcal{M}^\sharp \rightarrow \mathcal{P}(\mathcal{M})$ formalizes its meaning.
- \mathcal{C}^\sharp abstracts cache configurations and $\gamma_{\mathcal{C}} : \mathcal{C}^\sharp \rightarrow \mathcal{P}(\mathcal{C})$ formalizes its meaning.
- \mathcal{E}^\sharp abstracts sequences of events and $\gamma_{\mathcal{E}} : \mathcal{E}^\sharp \rightarrow \mathcal{P}(\mathcal{E}^*)$ formalizes its meaning.

But, since cache updates and events depend on memory state, independent analyses would be too imprecise. Consequently, in order to maintain some of the relations, we link the three abstract domains for memory state, caches, and events through a finite set of labels L so that our abstract domain is

$$Traces^\sharp = L \rightarrow \mathcal{M}^\sharp \times \mathcal{C}^\sharp \times \mathcal{E}^\sharp.$$

We will write $a^{\mathcal{M}}(l)$, $a^{\mathcal{C}}(l)$ and $a^{\mathcal{E}}(l)$ for the first, second and third components of an abstract element $a(l)$.

Labels roughly correspond to nodes in a control flow graph in classical data flow analyses. One could simply use program locations as labels. But with our setting, we can use more general labels, allowing a more fine-grained analysis where we can distinguish values of flags or results of previous tests, in the spirit of [35]. To capture that, we associate a meaning with these labels via a function $\gamma_L : L \rightarrow \mathcal{P}(Traces)$. If the labels are program locations, then $\gamma_L(l)$ is the set of traces ending in a state in location l . The analogy with control flow graphs can be extended to edges of that graph: using the *next* function, we define the successors and predecessors of a location l as: $succ(l) = \{k \mid next(\gamma_L(l)) \cap \gamma_L(k) \neq \emptyset\}$, and $pred(l) = \{k \mid next(\gamma_L(k)) \cap \gamma_L(l) \neq \emptyset\}$.

Then we can describe our trace abstract domain as a reduced cardinal power [18] of L and the reduced cardinal product of \mathcal{M}^\sharp , \mathcal{C}^\sharp and \mathcal{E}^\sharp , with the meaning function:

$$\begin{aligned} \gamma(a) &= \{\sigma_0 e_0 \sigma_1 \dots \sigma_n \in Traces \mid \forall i \leq n, \forall l \in L : \\ &\quad \sigma_0 e_0 \sigma_1 \dots \sigma_i \in \gamma_L(l) \Rightarrow \\ &\quad \sigma_i^{\mathcal{M}} \in \gamma_{\mathcal{M}}(a^{\mathcal{M}}(l)) \wedge \sigma_i^{\mathcal{C}} \in \gamma_{\mathcal{C}}(a^{\mathcal{C}}(l)) \\ &\quad \wedge e_0 \dots e_{i-1} \in \gamma_{\mathcal{E}}(a^{\mathcal{E}}(l))\} \end{aligned} \quad (2)$$

and as usual for these reduced products, the abstract transfer function $next^\sharp$ will be decomposed into:

$$next^\sharp(a) = \lambda l. (next_{\mathcal{M}^\sharp}(a, l), next_{\mathcal{C}^\sharp}(a, l), next_{\mathcal{E}^\sharp}(a, l)). \quad (3)$$

Each next function on abstract memory, cache and events implements partial reductions [19] using effects abstracting the concrete effects defined in the previous section.

4.3 Local Soundness

Since we have a reduced composition of abstract domains, we can use the soundness theorems of [18]: it is sufficient to show soundness of each individual abstract domain to show the soundness of the entire analysis. The abstract *next*^{sharp} operation is implemented using local update functions describing the change from one label to the next when applying the *next* function. For each label $k \in L$, and for each $l \in succ(k)$:

- Abstract update function $upd_{\mathcal{M}^\sharp, (k, l)} : \mathcal{M}^\sharp \rightarrow \mathcal{M}^\sharp$, and
- abstract memory effect $eff_{\mathcal{M}^\sharp, (k, l)} : \mathcal{M}^\sharp \rightarrow \mathcal{P}(\mathcal{E}_{\mathcal{M}})$.

For the cache domain we do not need separate functions for each pair (k, l) , as the cache update only depends on the accessed block which is delivered by the abstract memory effect. Similarly, the update of the event domain only depends on the effect computed by the abstract cache effect. Thus, we have:

- Abstract cache update $upd_{\mathcal{C}^\sharp} : \mathcal{C}^\sharp \times \mathcal{P}(\mathcal{E}_{\mathcal{M}}) \rightarrow \mathcal{C}^\sharp$,
- abstract cache effect $eff_{\mathcal{C}^\sharp} : \mathcal{C}^\sharp \times \mathcal{P}(\mathcal{E}_{\mathcal{M}}) \rightarrow \mathcal{P}(\mathcal{E}_{\mathcal{C}})$,
- and, for the events, an abstract update function $upd_{\mathcal{E}^\sharp} : \mathcal{E}^\sharp \times \mathcal{P}(\mathcal{E}_{\mathcal{C}}) \rightarrow \mathcal{E}^\sharp$.

Once we have these functions, we can approximate the effect of the *next* function on each label l , using the abstract values associated with the labels that can lead to l , $pred(l)$. This yields the following abstract transfer functions, where $\sqcup^{\mathcal{M}^\sharp}, \sqcup^{\mathcal{C}^\sharp}, \sqcup^{\mathcal{E}^\sharp}$ refer to the join functions of the respective domains:

- For the memory domain:

$$next_{\mathcal{M}^\sharp}(a, l) = \bigsqcup_{k \in pred(l)}^{ \mathcal{M}^\sharp } upd_{\mathcal{M}^\sharp, (k, l)}(a^{\mathcal{M}}(k))$$

- For the cache domain:

$$next_{\mathcal{C}^\sharp}(a, l) = \bigsqcup_{k \in pred(l)}^{ \mathcal{C}^\sharp } upd_{\mathcal{C}^\sharp} \left(a^{\mathcal{C}}(k), eff_{\mathcal{M}^\sharp, (k, l)}(a^{\mathcal{M}}(k)) \right)$$

- For the events:

$$\begin{aligned} next_{\mathcal{E}^\sharp}(a, l) &= \\ &\bigsqcup_{k \in pred(l)}^{ \mathcal{E}^\sharp } upd_{\mathcal{E}^\sharp} \left(a(k)^{\mathcal{E}}, eff_{\mathcal{C}^\sharp} \left(a^{\mathcal{C}}(k), eff_{\mathcal{M}^\sharp, (k, l)}(a^{\mathcal{M}}(k)) \right) \right) \end{aligned}$$

Now from Equations 1, 2, and 3, we can derive local conditions in each domain that are sufficient to guarantee local soundness for the whole analysis:

Definition 1 (Local soundness of abstract domains). *The abstract domains are locally sound if the abstract joins are over-approximations of unions, and if for any function $f^\sharp \in \{\text{upd}_{\mathcal{M}^\sharp, (k,l)}, \text{eff}_{\mathcal{M}^\sharp, (k,l)}, \text{upd}_{\mathcal{C}^\sharp}, \text{eff}_{\mathcal{C}^\sharp}, \text{upd}_{\mathcal{E}^\sharp}\}$ approximating concrete function $f \in \{\text{upd}_{\mathcal{M}}, \text{eff}_{\mathcal{M}}, \text{upd}_{\mathcal{C}}, \text{eff}_{\mathcal{C}}, \text{next}\}$ and corresponding meaning function γ_f , we have for any abstract value x :*

$$\gamma_f(f^\sharp(x)) \supseteq f(\gamma_f(x)).$$

For example, for the cache abstract domain, we have the following local soundness conditions:

$$\begin{aligned} \forall c^\sharp \in \mathcal{C}^\sharp, M \in \mathcal{P}(\mathcal{E}_{\mathcal{M}}): \\ \gamma_C(\text{upd}_{\mathcal{C}^\sharp}(c^\sharp, M)) \supseteq \text{upd}_{\mathcal{C}}(\gamma_C(c^\sharp), M), \\ \text{eff}_{\mathcal{C}^\sharp}(c^\sharp, M) \supseteq \text{eff}_{\mathcal{C}}(\gamma_C(c^\sharp), M), \end{aligned}$$

$$\forall \mathcal{G}^\sharp \subseteq \mathcal{C}^\sharp: \gamma_C\left(\bigsqcup_{G^\sharp \in \mathcal{G}^\sharp} G^\sharp\right) \supseteq \bigcup_{G^\sharp \in \mathcal{G}^\sharp} \gamma_C(G^\sharp).$$

Lemma 2 (Local Soundness Conditions). *If local soundness on the abstract memory, cache, and events domains are satisfied, then the corresponding next^\sharp function satisfies local soundness.*

Due to the above lemma, abstract domains for the three aspects memory, cache, and events can be separately developed and proved correct. We exploit this fact in this paper, and we plan to develop further abstractions in the future, targeting different classes of adversaries or improving precision.

4.4 Soundness of Delivered Bounds

We implemented the framework described above in a tool named CacheAudit. Thanks to the previous results, CacheAudit provides the following guarantees.

Theorem 4. *The bounds derived by CacheAudit soundly over-approximate $|\text{ran}(C^{\text{adv}})|$, for $\text{adv} \in \{\text{tr}, \text{acc}, \text{time}\}$, and hence correspond to upper bounds on the maximal amount of leaked information.*

The statement is an immediate consequence of combining Lemma 2 with Theorems 2 and 3, under the assumption that all involved abstract domains satisfy local soundness conditions, and that the corresponding counting procedures are correct. We formally prove the validity of these assumptions only for our novel relational and trace domains, see Section 6. For the other domains, corresponding proofs are either standard (e.g. the value domain) or out of scope of this submission.

5 Tool Design and Implementation

In this section we describe the architecture and implementation of CacheAudit.

We take advantage of the compositionality of the framework described in Section 4 and use a generic iterator module to compute fixpoints, where we rely on independent modules for the abstract domains that correspond to the components of the next^\sharp operation. Figure 1 depicts the overall architecture of CacheAudit, with the individual modules described below. CacheAudit is implemented in currently about 7.5 KLOC of OCaml [5], which we plan to make publicly available.

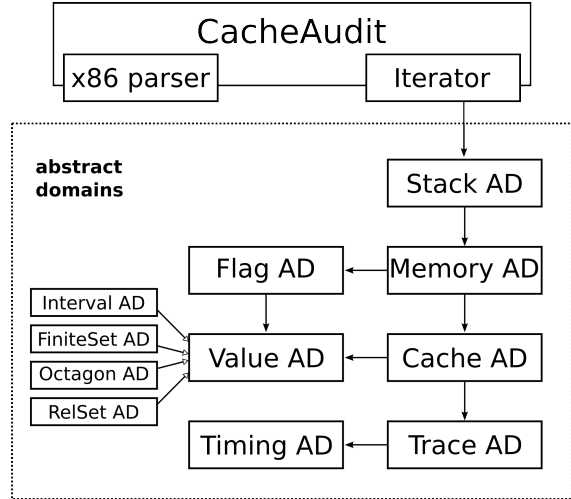


Figure 1: The architecture of CacheAudit. The solid boxes represent modules. Black-headed arrows mean that the module at the head is an argument of the module at the tail. White-headed arrows represent is-a relationships.

5.1 Control Flow Reconstruction

The first stage of the analysis is similar to a compiler front end. The main challenge is that we directly analyze x86 executables with no explicit control flow graph, which we need for guiding the fixpoint computation.

For the parsing phase, we rely on Chlipala’s parser for x86 executables [14], which we extended to a set of instructions that is sufficient for our case studies (but not yet complete). For the control-flow reconstruction, we consider only programs without dynamically computed jump and call targets, which is why it suffices to identify the basic blocks and link them according to the corresponding branching conditions and (static) branch targets. We plan to integrate more sophisticated techniques for control-flow reconstruction [29] in the future.

5.2 Iterator

The iterator module is responsible for the computation of the $next^\sharp$ operator and of the approximation of its fix-point using adequate iteration strategies [18]. Our analysis uses an *iterative* strategy, i.e., it stabilizes components of the abstract control flow graph according to a weak topological ordering, which we compute using Bourdoncle’s algorithm [13].

The iterator also implements parts of the reduced cardinal power, based on the labels computed according to the control-flow graph: Each label is associated with an initial abstract state. The analysis computes the effect of the commands executed from that label to its successors on the initial abstract state, and propagates the resulting final states using the abstract domains described below. In order to increase precision, we expand locations using loop unfolding, so that we have a number of different initial and final abstract states for each label inside loops, depending on a parameter describing the number of loop unfoldings we want to perform. Most of our examples (such as the cryptographic algorithms) require only a small, constant number of loop iterations, so that we can choose unfolding parameters that avoid joining states stemming from different iterations.

5.3 Abstract Domains

As described in Section 4, we decompose the abstract domain used by the iterator into mostly independent abstract domains describing different aspects of the concrete semantics.

Value Abstract Domains A value abstract domain represents sets of mapping from variables to (integer) values. Value abstract domains are used by the cache abstract domain to represent ages of blocks in the cache, and by the flag abstract domain to represent values stored at the addresses used in the program. We implemented different value abstract domains, such as the interval domain, an exact finite sets domain (where the sets become intervals when they are growing too large) and a relational set domain (as described in Section 6.1).

Flag Abstract Domain In x86 binaries, there are no high level guards: instead, most operations modify flags which are then queried in conditional branches. In order to deal precisely with such branches, we need to record relational informations between the values of variables and the values of these flags. To that purpose, for each operation that modifies the flags, we compute an over-approximation of the values of the arguments that may lead to a particular flag combination. The flag abstract domain represents an abstract state as a mapping from

values of flags to an element of the value abstract domain. When the analysis reaches a conditional branch, it can identify which combination of flag values corresponds to the branch and select the appropriate abstract values.

Memory Abstract Domain The memory abstract domain associates memory addresses and registers with variables and translates machine instructions into the corresponding operations on those variables, which are represented using flag abstract domains as described above. One important aspect for efficiency is that variables corresponding to addresses are created dynamically during the analysis whenever they are needed. The memory abstract domain further records all accesses to main memory using a cache abstract domain, as described below.

Stack Abstract Domain Operations on the stack are handled by a dedicated stack abstract domain. In this way the memory abstract domain does not have to deal with stack operations such as procedure calls, for which special techniques can be implemented to achieve precise interprocedural analysis.

Cache Abstract Domain Cache abstract domains only keep information about the cache, using a representation of sets of mappings from blocks to ages in the cache. This is implemented using an instance of value abstract domains. Effects from the memory domain are passed to the cache abstract domain at the memory abstract domain level, so that the cache domain knows which addresses are touched during computation. The cache abstract domain passes information about the presence or absence of cache hits and misses to the trace abstract domain, which we present in Section 6.2. The timings are then obtained as an abstraction from the traces.

6 Abstract Domains for Cache Adversaries

6.1 Cache State Domains

Abstractions of cache states are at the heart of analyses for all three cache adversaries considered in this paper. Thus, precise abstraction of cache states is crucial to determine tight leakage bounds.

Intuition behind Relational Sets The current state-of-the-art abstraction for LRU replacement by Ferdinand et al. [21] maintains an upper and a lower bound on the age of every memory block. This abstraction was developed with the sole goal of classifying memory accesses as cache hits or cache misses. In contrast, our goal is to develop abstractions whose concretizations are

small, as these will yield better bounds on the maximal leakage of a channel. To this end, we propose a new domain called *relational sets* that improves previous work along two dimensions:

1. Instead of *intervals* of ages of memory blocks, we maintain *sets* of ages or memory blocks.
2. Instead of maintaining *independent* information about the age of each memory blocks, we record the *relation* between ages of different memory blocks.

In addition to increasing precision, moving from intervals to sets allows us to analyze caches with FIFO and PLRU replacement. Interval-based analysis of FIFO and PLRU has been shown to be rather imprecise in the context of worst-case execution time analysis.

Motivating Example Consider the following method, which performs a table lookup based on a secret input, as it may occur in e.g. an AES implementation:

```

unsigned int A[size];

int getElement(int secret) {
    if (secret < size)
        return A[secret];
}

```

Assume we want to determine the possible cache states after one invocation of `getElement`. As the value of `secret` is unknown to the analysis, every memory location of the array might be accessed.

Assuming the array was not cached before the invocation of `getElement`, the interval-based domain by Ferdinand et al. [21] determines a lower bound of 0 and an upper bound of k on the age of each array element.

By tracking sets instead of intervals of ages for each memory block, we would get 0 and k as possible ages of each array element.

Both non-relational domains, however, are not powerful enough to infer or even express the fact, that *only one* of the array’s memory blocks has been accessed, and can thus be cached. Therefore, the number of possible cache states represented by non-relational abstractions grows exponentially in the size of the array, while the actual number of possible cache states only grows linearly.

A relational domain, tracking the possible ages of, e.g., pairs of memory blocks, would indeed yield a linear growth in the number of possible cache states. For each pair of array elements, it would be able to infer that only one of the two blocks may be cached. From this, it follows, that only one of all of the array elements may be cached.

Figure 2 shows experimental results for the example program with three domains: the interval domain (*IV*),

size	8	16	32	64	128	256
LRU/ <i>IV</i>	1	2	4	8	16	32
LRU/ <i>Set</i>	1	2	4	8	16	32
LRU/ <i>Rel</i>	1	1.58	2.32	3.17	4.01	5.04

Figure 2: Bounds on the number of leaked bits about the parameter *secret* for varying array sizes. The cache parameters are fixed, with a block size of 32 bytes, associativity 4 and cache size 4 KB.

and two instances of the *relational sets* domain, tracking sets of ages of individual blocks (*Set*) and sets of ages of pairs of blocks (*Rel*), respectively.

We do not see an improvement of *sets* over *intervals* in this particular example, as the information that a block has either age 0 or age k can be inferred from the intervals in the counting procedure. This is because the considered arrays are small and thus no two array elements map to the same cache set. We will see in the case studies, however, that *sets* alone often improve over *intervals*.

Formalization of Relational Sets Precision and analysis cost of the *relational sets* domain can be controlled by a parameter \mathcal{N} . The parameter \mathcal{N} is a set of sets of memory block: $\mathcal{N} \subseteq \mathcal{P}(\mathcal{B})$.

The idea is to precisely track the possible combinations of ages of memory blocks for each set of blocks contained in \mathcal{N} . For example, if we want to track information for individual blocks and for pairs of blocks, we would choose \mathcal{N} to be

$$\mathcal{N}_{rel} = \{N \subseteq \mathcal{B} \mid 1 \leq |N| \leq 2\}.$$

On the other hand, for a more efficient, yet still more precise analysis than intervals, we may opt for a non-relational analysis, tracking the sets of ages of individual blocks:

$$\mathcal{N}_{sets} = \{N \subseteq \mathcal{B} \mid |N| = 1\}.$$

Each choice of \mathcal{N} induces a different domain $\mathcal{C}_{\mathcal{N}}^{\sharp}$ of abstract cache states. In fact, the two domains *Set* and *Rel* from Figure 2 in the motivating example are instances of the relational sets domain with $\mathcal{N} = \mathcal{N}_{sets}$ and $\mathcal{N} = \mathcal{N}_{rel}$, respectively.

For each set of memory blocks N in \mathcal{N} , an abstract cache state collects the possible “partial” cache states when limiting attention to the ages of the blocks in N . Such a partial cache state can be represented by a function $f : N \rightarrow A$, similarly to concrete cache states, as defined in Section 3.1. Thus, in an abstract cache state \mathcal{C}^{\sharp} , for each $N \in \mathcal{N}$, $\mathcal{C}^{\sharp}(N)$ contains a set of such functions. Formally, the domain of abstract cache states $\mathcal{C}_{\mathcal{N}}^{\sharp}$ is de-

defined as follows:

$$\mathcal{C}_{\mathcal{N}}^{\sharp} := \left\{ \mathcal{C}^{\sharp} : \mathcal{N} \rightarrow \mathcal{P}(\mathcal{B} \rightarrow A) \mid \forall N \in \mathcal{N} : \mathcal{C}^{\sharp}(N) \in \mathcal{P}(N \rightarrow A) \right\}.$$

By $\mathcal{B} \rightarrow A$ we denote the set of partial functions from \mathcal{B} to A . In the following, we assume a fixed set \mathcal{N} , and thus omit the subscript \mathcal{N} , unless we need to explicitly refer to a different set.

The meaning of abstract cache states is formalized by a concretization function $\gamma_{\mathcal{C}} : \mathcal{C}^{\sharp} \rightarrow \mathcal{P}(\mathcal{C})$. To define $\gamma_{\mathcal{C}}$, we first need helper functions $\gamma_{\mathcal{C},N} : \mathcal{P}(N \rightarrow A) \rightarrow \mathcal{P}(\mathcal{C})$ that formalize the meaning of the set of partial cache states \mathcal{C}^{\sharp} stored for a particular set of blocks N :

$$\gamma_{\mathcal{C},N}(c^{\sharp}) := \left\{ c \in \mathcal{C} \mid c|_N \in c^{\sharp} \right\},$$

where $c|_N$ denotes the restriction of the function c to the domain N , which is a subset of the original domain \mathcal{B} of c . Then, the concretization function $\gamma_{\mathcal{C}}$ is

$$\gamma_{\mathcal{C}}(\mathcal{C}^{\sharp}) := \bigcap_{N \in \mathcal{N}} \gamma_{\mathcal{C},N}(\mathcal{C}^{\sharp}(N)).$$

Thus, a concrete state is represented by an abstract state if its restriction to each of the sets in \mathcal{N} is contained in the respective sets of the abstract state.

We can also define an abstraction function $\alpha_{\mathcal{C}} : \mathcal{P}(\mathcal{C}) \rightarrow \mathcal{C}^{\sharp}$, assigning to each set of concrete states its unique best abstract representation:

$$\alpha_{\mathcal{C}}(\mathcal{C}) := \lambda N \in \mathcal{N}. \{c|_N \mid c \in \mathcal{C}\}.$$

We need a join function $\sqcup^{\mathcal{C}^{\sharp}} : \mathcal{P}(\mathcal{C}^{\sharp}) \rightarrow \mathcal{C}^{\sharp}$ on abstract states to combine analysis information from e.g. different program paths. This simply amounts to a pointwise union:

$$\sqcup^{\mathcal{C}^{\sharp}} D := \lambda N \in \mathcal{N}. \bigcup_{\mathcal{C}^{\sharp} \in D} \mathcal{C}^{\sharp}(N).$$

Note that the join function induces a partial order $\sqsubseteq^{\mathcal{C}^{\sharp}}$ on abstract states, as $a \sqsubseteq^{\mathcal{C}^{\sharp}} b \Leftrightarrow b = a \sqcup^{\mathcal{C}^{\sharp}} b$. It is fairly easy to see that both the abstraction and the concretization function are monotone functions. In fact, $(\mathcal{P}(\mathcal{C}), \subseteq) \xleftrightarrow[\alpha_{\mathcal{C}}]{\gamma_{\mathcal{C}}} (\mathcal{C}^{\sharp}, \sqsubseteq^{\mathcal{C}^{\sharp}})$ forms a *Galois connection*. Thus, we can compose abstraction, concretization, and cache update function to obtain the best abstract transformer $upd_{\mathcal{C}^{\sharp}} : \mathcal{C}^{\sharp} \times \mathcal{P}(\mathcal{E}_{\mathcal{M}}) \rightarrow \mathcal{C}^{\sharp}$:

$$upd_{\mathcal{C}^{\sharp}}(\mathcal{C}^{\sharp}, M) := \alpha_{\mathcal{C}}(upd_{\mathcal{C}}(\gamma_{\mathcal{C}}(\mathcal{C}^{\sharp}), M)),$$

where $upd_{\mathcal{C}}$ is lifted to sets of concrete states and sets of memory blocks in the expected way. Due to the

monotony of $\alpha_{\mathcal{C}}$ and the lifted version of $upd_{\mathcal{C}}$, we also have that

$$upd_{\mathcal{C}^{\sharp}}(\mathcal{C}^{\sharp}, M) = \bigsqcup_{b \in M}^{C^{\sharp}} \alpha_{\mathcal{C}}(upd_{\mathcal{C}}(\gamma_{\mathcal{C}}(\mathcal{C}^{\sharp}), b)).$$

Similarly, we can define the best abstract effects:

$$\begin{aligned} eff_{\mathcal{C}^{\sharp}} & : \mathcal{C}^{\sharp} \times \mathcal{P}(\mathcal{E}_{\mathcal{M}}) \rightarrow \mathcal{P}(\mathcal{E}_{\mathcal{C}}), \\ eff_{\mathcal{C}^{\sharp}}(\mathcal{C}^{\sharp}, M) & := eff_{\mathcal{C}}(\gamma_{\mathcal{C}}(\mathcal{C}^{\sharp}), M), \end{aligned}$$

where $eff_{\mathcal{C}}$ is lifted to sets of concrete states and sets of memory blocks.

Lemma 3 (Local soundness of best abstract transformer). *By construction [17], the best abstract transformer and the best abstract effects satisfy the local soundness condition stated in Definition 1.*

Implementation of the Best Abstract Transformer for Relational Sets

The definition of the best abstract transformer as described above suggests an implementation that explicitly computes the concretization, applies the concrete cache update function, and abstracts the resulting set of concrete cache states. While the concretization of any abstract cache state is finite in our case, it may be extremely large. Thus, such a naive approach will not work in practice.

In an efficient implementation of the abstract transformer, we would like to separately update the information $\mathcal{C}^{\sharp}(N)$ associated with each set N in \mathcal{N} without constructing complete concrete cache states. Examining the concrete cache update function $upd_{\mathcal{C}}$, defined in Section 3.1, notice that it is possible to do this, if the accessed block b belongs to N . In that case, even the third and fourth condition can be precisely evaluated using the information available “locally” in $\mathcal{C}^{\sharp}(N)$. If the accessed block b maps to the same cache set as other blocks contained in N , but b itself is not part of N , then the third and fourth condition in the concrete cache update function cannot be evaluated. The idea is thus to *partially concretize* abstract states, by adding the accessed block b to each set N in \mathcal{N} . To this end, we define the following set: $\mathcal{N}_b = \{N \cup \{b\} \mid N \in \mathcal{N}\}$.

The partial concretization function $\gamma_{\mathcal{N}_b}$ translates states from $\mathcal{C}_{\mathcal{N}}^{\sharp}$ into the slightly more concrete domain $\mathcal{C}_{\mathcal{N}_b}^{\sharp}$:

$$\gamma_{\mathcal{N}_b}(\mathcal{C}^{\sharp}) := \lambda N \in \mathcal{N}_b. \left\{ c|_N \mid c \in \gamma_{\mathcal{C}}(\mathcal{C}^{\sharp}) \right\}.$$

This can be implemented rather efficiently, by enumerating all extensions of the functions found in $\mathcal{C}^{\sharp}(N)$ and checking whether or not they agree with the constraints for all other N' in \mathcal{N} . There is a corresponding partial

abstraction function $\alpha_{\mathcal{N}_b}$ that simply drops information about block b if it is not contained in N .

With partial concretization and partial abstraction in place, we can realize the best abstract transformer as follows:

$$\text{upd}_{C^\sharp}(C^\sharp, M) = \bigsqcup_{b \in M} \alpha_{\mathcal{N}_b}(\text{upd}_{C^\sharp, \mathcal{N}_b}(\gamma_{\mathcal{N}_b}(C^\sharp), b)),$$

where $\text{upd}_{C^\sharp, \mathcal{N}_b} : C_{\mathcal{N}_b}^\sharp \rightarrow C_{\mathcal{N}_b}^\sharp$ applies upd_C^N , defined below, to each $c^\sharp \in C^\sharp(N)$:

$$\text{upd}_{C^\sharp, \mathcal{N}_b}(C^\sharp, b) := \lambda N \in \mathcal{N}_b. \left\{ \text{upd}_C^N(c^\sharp, b) \mid c^\sharp \in C^\sharp(N) \right\}.$$

upd_C^N closely resembles its concrete counterpart, as all of the conditions required to update the age of a block can be precisely evaluated in c^\sharp :

$$\text{upd}_C^N(c^\sharp, b) := \lambda b' \in N. \begin{cases} 0 & : b' = b \\ c^\sharp(b') & : \text{set}(b') \neq \text{set}(b) \\ c^\sharp(b') + 1 & : \text{set}(b') = \text{set}(b) \wedge c^\sharp(b') < c^\sharp(b) \\ c^\sharp(b') & : \text{set}(b') = \text{set}(b) \wedge c^\sharp(b') \geq c^\sharp(b) \end{cases}$$

In fact, the concrete update function upd_C is a special case of upd_C^N for $N = \mathcal{B}$. For FIFO and PLRU, we can construct update functions on partially concretized states in a similar fashion.

Theorem 5 (Local soundness using partial concretization). *The abstract transformer based on partial concretization is equal to the best abstract transformer, i.e.*

$$\forall C^\sharp \in C^\sharp, M \in \mathcal{P}(\mathcal{E}_M) : \bigsqcup_{b \in M} \alpha_{\mathcal{N}_b}(\text{upd}_{C^\sharp, \mathcal{N}_b}(\gamma_{\mathcal{N}_b}(C^\sharp), b)) = \bigsqcup_{b \in M} \alpha_C(\text{upd}_C(\gamma_C(C^\sharp), b)).$$

By Lemma 3 this implies the local soundness of the abstract transformer based on partial concretization.

Lazy Representation of Relational Sets In an abstract state C^\sharp , there are usually some sets $N \in \mathcal{N}$ for which $C^\sharp(N)$ is redundant w.r.t. the subsets of N . In such cases, the concretization of $C^\sharp(N)$ is equal to the intersection of the concretizations of its subsets:

$$\gamma_{C, N}(C^\sharp(N)) = \bigcap_{\substack{N' \subseteq N \\ N' \in \mathcal{N}}} \gamma_{C, N'}(C^\sharp(N')).$$

Our implementation detects such cases and omits the explicit representation of analysis information for redundant sets. As a consequence, the space and time complexity of the implementation is roughly related to the amount of non-redundant relational information.

Counting for Relational Sets In order to bound the maximal leakage of the cache side channel, we need to determine the number of concrete cache states represented by an abstract cache state. A naive approach is to explicitly compute an abstract cache state’s concretization. This is feasible only if the leakage is very small.

Our approach is to divide the counting problem into several independent counting problems for distinct subsets of memory blocks. One could separately count the number of states of each cache set, and multiply the results. However, this may lead to imprecise results if states of different cache sets depend on each other. We call two memory blocks a and b *dependent*, if they map to the same cache set, or if there is a non-redundant set $N \in \mathcal{N}$ containing a and b . We compute the finest partition of the set of memory blocks such that dependent memory blocks belong to the same part of the partition. As memory blocks in different partitions are “independent”, we can separately determine the number of partial cache states for each part without losing any precision when we multiply the results.

6.2 A Trace Domain

We devise an abstract domain for keeping track of the sets of event traces that may occur during the execution of a program. Following the way events are computed in the concrete, namely as a function from cache states and memory effects (see Section 3.3), the abstract cache domain provides abstract cache effects.

In our current implementation of CacheAudit, we use an exact representation for sets of event traces: we can represent any finite set of event traces, and assuming an incoming set of traces \mathcal{S} and a set of cache effects E , we compute the resulting event set precisely as follows:

$$\text{upd}_{\mathcal{E}^\sharp}(\mathcal{S}, E) = \{\sigma.e \mid \sigma \in \mathcal{S} \wedge e \in E\}$$

Then soundness is obvious, since the abstract operation is the same as its concrete counterpart. Due to loop unfolding, we do not require widenings, even though the domain contains infinite ascending chains (see Section 5.2).

Lemma 4. *The trace domain is locally sound.*

Representation for Sets of Event Traces We represent sets of finite event traces corresponding to a particular program location by an acyclic graph with vertices V , a dedicated root $r \in V$, and a node labeling $\ell : V \rightarrow \mathcal{P}(\mathcal{E}) \cup \{\perp\}$. In this graph, every node $v \in V$ represents a set of traces $\gamma(v) \in \mathcal{P}(\mathcal{E}^*)$ in the following way:

1. For the root r , $\gamma(r) = \{\varepsilon\}$

2. For v with $L(v) = \sqcup$ and predecessors u_1, \dots, u_n , $\gamma(v) = \bigcup_{i=1}^n \gamma(u_i)$.
3. For v with $L(v) \neq \sqcup$ and predecessors u_1, \dots, u_n , $\gamma(v) = \{t.u \mid u \in L(v) \wedge t \in \bigcup_{i=1}^n \gamma(u_i)\}$

Intuitively, every $v \in V$ represents a set of event traces, namely the sequences of labels of paths from r to v .

In the context of *CacheAudit*, we need to implement two operations on this data structure, namely (1) the join $\sqcup^{\mathcal{E}^\#}$ of two sets of traces and the (2) addition $upd_{\mathcal{E}^\#}(\mathcal{S}, E)$ of a cache event to a particular set of traces.

For the join of two sets of traces represented by v and w , we add a new vertex u with label \sqcup and add edges from v and w to u .

For the extension of a set of traces represented by a vertex v by a set of cache events E , we first check whether v already has a child w labeled with E . If so, we use w as a representation of the extended set of traces. If not, we add a new vertex u with label E and add an edge (u, v) . In this way we make maximal use of sharing and obtain a prefix DAG.

The correctness of the representation follows by construction. In our *CacheAudit*, we use hash consing for efficiently building the prefix dag.

Counting Sets of Traces The following algorithm $count_{tr}$ overapproximates the number of traces that are represented by a given graph.

1. For the root r , $count_{tr}(r) = 1$
2. For v with $L(v) = \sqcup$ and predecessors u_1, \dots, u_n , $count_{tr}(v) = \sum_{i=1}^n count_{tr}(u_i)$
3. For v with $L(v) \neq \sqcup$ and predecessors u_1, \dots, u_n , $count_{tr}(v) = |L(v)| \cdot \sum_{i=1}^n count_{tr}(u_i)$

The soundness of this counting, i.e. the fact that $|\gamma(v)| \leq count_{tr}(v)$, follows by construction. Notice that this counting procedure is precise if the labels represent singleton events, because then every trace is uniquely represented in the graph. However, the precision dramatically decreases with larger sets of labels. In our case, labels contain at most two events and the counting is sufficiently precise.

Counting Timing Variations We currently model execution time as a simple abstraction of traces, see Section 3. In particular, timing is computed from a traces over $\mathcal{E} = \{hit, miss, \perp\}$ by multiplying the number of occurrences of each event by the time they consume: t_{hit} , t_{miss} , and t_{\perp} , respectively. The following algorithm $count_{time}$ overapproximates the set of timing behaviors that are represented by a given graph.

1. For the root r , $count_{time}(r) = \{0\}$
2. For v with $L(v) = \sqcup$ and predecessors u_1, \dots, u_n , $count_{time}(v) = \bigcup_{i=1}^n count_{time}(u_i)$
3. For v with $L(v) \neq \sqcup$ and predecessors u_1, \dots, u_n ,

$$count_{time}(v) = \left\{ t_x + t \mid x \in L(v) \wedge t \in \bigcup_{i=1}^n count_{time}(u_i) \right\}$$

The soundness of $count_{time}$, i.e. the fact that it delivers a superset of the number of possible timing behaviors follows by construction.

7 Case Studies

In this section we demonstrate the capabilities of *CacheAudit* in case studies where we use it to analyze the cache side channels of algorithms for sorting and symmetric encryption. All results are based on the automatic analysis of corresponding 32-bit x86 Linux executables that we compiled using *gcc* with disabled stack canaries and without any compiler optimizations.

7.1 AES 128

We analyze the AES implementation from the PolarSSL library [3] with keys of 128 bits, where we consider the implementation with and without preloading of tables, for all attacker models, different replacement strategies, associativities, and line sizes. All results are presented as upper bounds of the leakage in *bits*; for their interpretation see Theorem 1. In some cases, *CacheAudit* reports upper bounds that exceed the key size (128 bits), which corresponds to an imprecision of the static analysis. We opted against truncating to 128 bits to illustrate the degree of imprecision. The full data of our analysis are given in Appendix C. Here, we highlight some of our findings.

- *Preloading* almost consistently leads to better security guarantees in all scenarios, see e.g. Figure 3. However, the effect becomes clearly more apparent for cache sizes beyond 8KB, which is explained by the PolarSSL AES tables exceeding the size of the 4KB cache by 256B. For cache sizes that are larger than the preloaded tables, we can prove noninterference for C^{acc} and FIFO, C^{accd} and LRU, and for C^{tr} and C^{time} on LRU, FIFO, and PLRU. For C^{acc} with shared memory spaces and LRU, this result does *not* hold because the adversary can obtain information about the order of memory blocks in the cache.




Figure 3: Effect of attacker model and preloading (PL) on the security guarantee, for the LRU replacement strategy. The horizontal axis gives the cache size, and the vertical axis gives the leakage in bits.




Figure 4: Effect of cache line size on the security guarantee, for C^{acc} and C^{accd} , and LRU replacement strategy without preloading. Different curves correspond to different cache line sizes. The horizontal axis gives the cache size, and the vertical axis gives the leakage bits.

- A larger *line size* consistently leads to better security guarantees for access-based adversaries, see e.g. Figure 4. This follows because more array indices map to a line which decreases the resolution of the attacker’s observations.
- In terms of *replacement strategies*, we consistently derive the lowest bounds for LRU, followed by PLRU and FIFO (see, e.g. Figures 9 and 10), where the only exception is the case of C^{acc} and preloading (see Figure 5). In this case FIFO is more secure because with LRU the adversary can obtain information about the ordering of memory blocks in the cache.
- In terms of *cache size*, we consistently derive better bounds for larger caches, with the exception of C^{accd} . For this adversary model the bounds increase because larger caches correspond to distributing the table to more sets, which increases its possibilities to observe variations. The guarantees we obtain for C^{accd} and C^{acc} converge for caches of 4 ways and sizes beyond 16KB, see e.g. Figure 4. This is due to the fact that each cache set can contain at most one unique block of the 4KB table. In that way, the ability of C^{acc} to observe ordering of blocks within a set does not give it power.
- In terms of *precision*, set-based analyses consistently match or improve over the bounds delivered by




Figure 5: Effect of replacement strategy on the security guarantee, for C^{acc} and C^{accd} , with and without preloading (PL). The horizontal axis gives the cache size, and the vertical axis gives the leakage in bits.

interval-based analyses, see Figure 7 in the appendix. Notice that the improvement is given in bits, i.e. on a logarithmic scale. The analysis time for the examples was typically below 60s and peaked at 365s for AES without preloading and 4KB cache.

Comparison to [32]: The PolarSSL AES implementation has already been analyzed in [32] with respect to access-based adversaries and LRU replacement. The results obtained by CacheAudit go beyond that analysis in that we derive bounds w.r.t. access-based, trace-based, and time-based adversaries, for LRU, FIFO, and PLRU strategies. For access-based adversaries and LRU, the bounds we derive are lower than those in [32]; in particular, for C^{accd} we derive bounds of zero for implementations with preloading for *all* caches sizes that are larger than the AES tables—which is obtained in [32] only for caches of 128KB. While these results are obtained for different platforms (x86 vs ARM) and are hence not directly comparable, they do suggest a significant increase in precision. In contrast to [32], this is achieved without any code instrumentation.

7.2 Salsa20

Salsa20 is a stream cipher by Bernstein [11]. Internally, the cipher uses XOR, 32-bit addition mod 2^{32} , and constant-distance rotation operation on an internal state of 16 32-bit words. The lack of key-dependent memory lookups intends to avoid cache side channels in software implementations of the cipher. With CacheAudit we could formally confirm this intuition by automated analysis of the reference implementation of the Salsa20 encryption, which includes a function call to a hash function. Specifically, we analyze the leakage of the encryption operation on an arbitrary 512-byte message for C^{acc} , C^{tr} , and C^{time} , FIFO and LRU strategies, on 4KB caches with line size of 32B, where we consistently obtain up-

per bounds of 0 for the leakage. The time required for analyzing each of the cases was below 11s.

7.3 Sorting Algorithms

In this section we use CacheAudit to establish bounds on the cache side channels of different sorting algorithms. This case study is inspired by an early investigation of secure sorting algorithms [9]. While the authors of [9] consider only time-based adversaries and noninterference as a security property, CacheAudit allows us to give quantitative answers for a comprehensive set of side-channel adversaries, based on the binary executables and concrete cache models.

As examples, we use the implementations of BubbleSort, InsertionSort, and SelectionSort from [4], which are given in Section 2 and Appendix A, respectively, where we use integer arrays of lengths from 8 to 64.

The results of our analysis are summarized in Figure 6. In the following we highlight some of our findings.

- We obtain the same bounds for BubbleSort and SelectionSort, which is explained by the similar structure of their control flow. A detailed explanation of those bounds is given in Section 2. InsertionSort has a different control flow structure, which is reflected by our data. In particular InsertionSort has only $n!$ possible execution traces due to premature abortion of the inner loop, which leads to better bounds w.r.t. trace-based adversaries. However, InsertionSort leaks more information to timing-based adversaries, because the number of iterations in the inner loop varies and thus fewer executions have the same timing.
- For access-based adversaries we obtain zero bounds for all algorithms. For trace-based adversaries, the derived bounds do not imply meaningful security guarantees: the bounds reported for InsertionSort are in the order of $\log_2(n!)$, which corresponds to the maximum information contained in the ordering of the elements; the bounds reported for the other sorting algorithm exceed this maximum, which is caused by the imprecision of the static analysis.
- We performed an analysis of the sorting algorithms for smaller (256B) and larger (64KB) cache sizes and obtained the exact same bounds as in Figure 6, with the exception of the case of arrays of 64 entries and 256B caches: there the leakage increases because the arrays do not fit entirely into the cache due to their misalignment with the memory blocks.

7.4 Discussion and Outlook

A number of comments are in order when interpreting the bounds delivered by CacheAudit. First, we obtained all of the bounds for an *empty* initial cache. As described

in Section 3.6, for access-based adversaries they immediately extend to bounds for arbitrary initial cache states, as long as the victim does not access any block that is contained in it. This is relevant, e.g. for an adversary who can fill the initial cache state only with lines from its own disjoint memory space. For LRU, our bounds extend to arbitrary initial cache states without further restriction.

Second, while CacheAudit relies on more accurate models of cache and timing than any information-flow analysis we are aware of, there are several timing-relevant features of hardware it does not capture (and make assertions about) at this point yet, including pipelines, TLBs, and multiple levels of caches.

Third, for the case of AES and Salsa20, the derived bounds hold for the leakage about the key in *one* execution, with respect to any payload. For the case of zero leakage (i.e., noninterference), the bounds trivially extend to bounds for multiple executions and imply strong security guarantees. For the case of non-zero leakage, the bounds can add up when repeatedly running the victim process with a fixed key and varying payload, leading to a decrease in security guarantees. One of our prime targets for future work is to derive security guarantees that hold for multiple executions of the victim process. One possibility is to employ leakage-resilient cryptosystems [20, 46], where our work can be used to bound the range of the leakage functions.

8 Related Work

The work most closely related to ours is [32]. There, the authors quantify cache side channels by connecting a commercial, closed-source tool for the static analysis of worst-case execution times [1] to an algorithm for counting concretizations of abstract cache states. The application of the tool to side-channel analysis is limited to access-based adversaries and requires heavy code instrumentation. In contrast, CacheAudit provides tailored abstract domains for all kinds of cache side-channel adversaries, different replacement strategies, and is modular and open for further extensions. Furthermore, the bounds delivered by CacheAudit are significantly tighter than those reported in [32], see Section 7.

Zhang et al. [47] propose an approach for mitigating timing side channels that is based on contracts between software and hardware. The contract is enforced on the software side using a type system, and on the hardware side, e.g., by using dedicated hardware such as partitioned caches. The analysis ensures that an adversary cannot obtain any information by observing public parts of the memory; any confidential information the adversary obtains must be via timing, which is controlled using dedicated mitigate commands. Tiwari et al. [44] sketch a novel microarchitecture that facilitates information-flow

array length	8			16			32			64		
	C^{tr}	C^{time}	C^{acc}	C^{tr}	C^{time}	C^{acc}	C^{tr}	C^{time}	C^{acc}	C^{tr}	C^{time}	C^{acc}
BubbleSort	28	4.86	0	120	6.92	0	496	8.96	0	2016	11	0
InsertionSort	15.23	6.91	0	44.3	10.15	0	117.7	13.3	0	296	15.8	0
SelectionSort	28	4.86	0	120	6.92	0	496	8.96	0	2016	11	0

Figure 6: The table illustrates the security guarantees derived by *CacheAudit* for the implementations of BubbleSort, SelectionSort, and InsertionSort, for trace-based, timing-based, and access-based adversaries, for LRU caches of 4KB and line sizes of 32B.

tracking by design, where they use noninterference as a baseline confidentiality property. Other mitigation techniques include coding guidelines [16] for thwarting cache attacks on x86 CPUs, or novel cache architectures that are resistant to cache side-channel attacks [45]. The goal of our approach is orthogonal to those approaches in that we focus on the *analysis* of microarchitectural side channels rather than on their mitigation. Our approach does not rely on a specific platform; rather it can be applied to any language and hardware architecture, for which abstractions are in place.

Kim et al. put forward StealthMem [28], a system-level defense against cache-timing attacks in virtualized environments. The core of StealthMem is a software-based mechanism that locks pages of a virtual machine into the cache and avoids that they are evicted by other VMs. StealthMem can be seen as a lightweight variant of flushing/preloading countermeasures. As future work, we plan to use our tool to derive formal, quantitative guarantees for programs using StealthMem.

For the case of AES, there are efficient software implementations that avoid the use of data caches by bit-slicing [27]. Furthermore, a model for statistical estimation of the effectiveness of AES cache attacks based on sizes of cache lines and lookup tables has been presented in [43]. In contrast, our analysis technique applies to arbitrary programs.

Technically, our work builds on methods from quantitative information-flow analysis (QIF) [15], where the automation by reduction to counting problems appears in [10, 37, 25, 36], and the connection to abstract interpretation in [33]. Finally, our work goes beyond language-based approaches that consider caching [8, 24] in that we rely on more realistic models of caches and aim for more permissive, quantitative guarantees.

9 Conclusions

We presented CacheAudit, the first automatic tool for the static derivation of formal, quantitative security guarantees against cache side-channel attacks. We demonstrate the usefulness of CacheAudit by establishing the first formal proofs of security of software-based countermeasures for a comprehensive set of adversaries and based on executable code.

The open architecture of CacheAudit makes it an ideal platform for future research on microarchitectural side channels. In particular, we are currently investigating the derivation of security guarantees for concurrent adversaries. Progress along those lines will provide a handle for extending our security guarantees to the operating system level. We will further investigate abstractions for hardware features such as pipelines, out-of-order execution, and leakage-resilient cache designs, with the goal of providing broad tool support for reasoning about side-channels arising at the hardware/software interface.

Acknowledgments We thank Ignacio Echeverría for helping with the implementation.

References

- [1] AbsInt aiT Worst-Case Execution Time Analyzers. <http://www.absint.com/a3/>.
- [2] Intel Advanced Encryption Standard (AES) Instructions Set. <http://software.intel.com/file/24917>.
- [3] PolarSSL. <http://polarssl.org/>.
- [4] Sorting algorithms. <http://www.codebeach.com/2008/09/sorting-algorithms-in-c.html>.
- [5] The Caml Language. <http://caml.inria.fr/>.
- [6] O. Aciıçmez and Ç. K. Koç. Trace-Driven Cache Attacks on AES. In *ICICS*, volume 4307 of *LNCS*, pages 112–121. Springer, 2006.
- [7] O. Aciıçmez, W. Schindler, and Ç. K. Koç. Cache based remote timing attack on the AES. In *CT-RSA*, volume 4377 of *LNCS*, pages 271–286. Springer, 2007.
- [8] J. Agat. Transforming out Timing Leaks. In *POPL 2000*, pages 40–53. ACM, 2000.
- [9] J. Agat and D. Sands. On confidentiality and algorithms. In *IEEE Symposium on Security and Privacy (SSP)*, pages 64–77. IEEE Computer Society, 2001.

- [10] M. Backes, B. Köpf, and A. Rybalchenko. Automatic Discovery and Quantification of Information Leaks. In *SSP*, pages 141–153. IEEE, 2009.
- [11] D. Bernstein. Salsa20. <http://cr.yp.to/snuffle.html>.
- [12] D. J. Bernstein. Cache-timing attacks on AES. Technical report, 2005.
- [13] F. Bourdoncle. Efficient chaotic iteration strategies with widenings. In *Proc. International Conference on Formal Methods in Programming and their Applications (FMPA '93)*, pages 128–141. Springer, 1993.
- [14] A. Chlipala. Modular development of certified program verifiers with a proof assistant. In *ICFP*, pages 160–171. ACM, 2006.
- [15] D. Clark, S. Hunt, and P. Malacaria. A static analysis for quantifying information flow in a simple imperative language. *JCS*, 15(3):321–371, 2007.
- [16] B. Coppens, I. Verbauwhede, K. D. Bosschere, and B. D. Sutter. Practical mitigations for timing-based side-channel attacks on modern x86 processors. In *SSP*, pages 45–60. IEEE, 2009.
- [17] P. Cousot and R. Cousot. Abstract interpretation: a unified lattice model for static analysis of programs by construction of approximation of fixpoints. In *Proc. of the 4th ACM Symposium of Principles of Programming Languages (POPL 1977)*, pages 238–252, 1977.
- [18] P. Cousot and R. Cousot. Systematic design of program analysis frameworks. In *Proc. of the 6th ACM Symposium of Principles of Programming Languages (POPL 1979)*, pages 269–282, 1979.
- [19] P. Cousot, R. Cousot, and L. Mauborgne. Theories, solvers and static analysis by abstract interpretation. *Journal of the ACM*, 59(6):31, 2012.
- [20] S. Dziembowski and K. Pietrzak. Leakage-Resilient Cryptography. In *FOCS*, pages 293–302. IEEE, 2008.
- [21] C. Ferdinand, F. Martin, R. Wilhelm, and M. Alt. Cache behavior prediction by abstract interpretation. *Science of Computer Programming*, 35(2):163–189, 1999.
- [22] D. Grund. *Static Cache Analysis for Real-Time Systems – LRU, FIFO, PLRU*. PhD thesis, Saarland University, 2012.
- [23] D. Gullasch, E. Bangerter, and S. Krenn. Cache Games - Bringing Access-Based Cache Attacks on AES to Practice. In *SSP*, pages 490–505. IEEE, 2011.
- [24] D. Hedin and D. Sands. Timing Aware Information Flow Security for a JavaCard-like Bytecode. *ENTCS*, 141(1):163–182, 2005.
- [25] J. Heusser and P. Malacaria. Quantifying information leaks in software. In *ACSAC*, pages 261–269. ACM, 2010.
- [26] S. Jana and V. Shmatikov. Memento: Learning secrets from process footprints. In *IEEE Symposium on Security and Privacy (SSP)*, pages 143–157. IEEE Computer Society, 2012.
- [27] E. Käsper and P. Schwabe. Faster and timing-attack resistant AES-GCM. In *CHES*, volume 5747 of *LNCS*, pages 1–17, 2009.
- [28] T. Kim, M. Peinado, and G. Mainar-Ruiz. Stealthmem: System-level protection against cache-based side channel attacks in the cloud. In *19th USENIX Security Symposium*. USENIX Association, 2012.
- [29] J. Kinder, F. Zuleger, and H. Veith. An abstract interpretation-based framework for control flow reconstruction from binaries. In *VMCAI*, pages 214–228. Springer, 2009.
- [30] P. Kocher. Timing Attacks on Implementations of Diffie-Hellman, RSA, DSS, and Other Systems. In *CRYPTO*, volume 1109 of *LNCS*, pages 104–113. Springer, 1996.
- [31] P. Kocher, J. Jaffe, and B. Jun. Differential Power Analysis. In *Proc. Advances in Cryptology (CRYPTO 1999)*, volume 1666 of *LNCS*, pages 388–397. Springer, 1999.
- [32] B. Köpf, L. Mauborgne, and M. Ochoa. Automatic Quantification of Cache Side-Channels. In *Proc. 24th International Conference on Computer Aided Verification (CAV)*, pages 564–580. Springer, 2012.
- [33] B. Köpf and A. Rybalchenko. Approximation and Randomization for Quantitative Information-Flow Analysis. In *CSF*, pages 3–14. IEEE, 2010.
- [34] B. Köpf and G. Smith. Vulnerability Bounds and Leakage Resilience of Blinded Cryptography under Timing Attacks. In *CSF*, pages 44–56. IEEE, 2010.
- [35] L. Mauborgne and X. Rival. Trace partitioning in abstract interpretation based static analyzers. In *ESOP*, volume 3444 of *LNCS*, pages 5–20. Springer, 2005.

- [36] Z. Meng and G. Smith. Calculating bounds on information leakage using two-bit patterns. In *PLAS*. ACM, 2011.
- [37] J. Newsome, S. McCamant, and D. Song. Measuring channel capacity to distinguish undue influence. In *PLAS*, pages 73–85. ACM, 2009.
- [38] D. A. Osvik, A. Shamir, and E. Tromer. Cache Attacks and Countermeasures: the Case of AES. In *CT-RSA*, volume 3860 of *LNCS*, pages 1–20. Springer, 2006.
- [39] C. Percival. Cache missing for fun and profit. In *BSDCan*, 2005.
- [40] T. Ristenpart, E. Tromer, H. Shacham, and S. Savage. Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds. In *ACM Conference on Computer and Communications Security*, pages 199–212. ACM, 2009.
- [41] A. Sabelfeld and A. C. Myers. Language-based Information-Flow Security. *IEEE Journal on Selected Areas in Communications*, 21(1):5–19, 2003.
- [42] G. Smith. On the Foundations of Quantitative Information Flow. In *FoSSaCS*, volume 5504 of *LNCS*, pages 288–302. Springer, 2009.
- [43] K. Tiri, O. Aciçmez, M. Neve, and F. Andersen. An Analytical Model for Time-Driven Cache Attacks. In *FSE*, volume 4593 of *LNCS*, pages 399–413. Springer, 2007.
- [44] M. Tiwari, J. Oberg, X. Li, J. Valamehr, T. E. Levin, B. Hardekopf, R. Kastner, F. T. Chong, and T. Sherwood. Crafting a usable microkernel, processor, and I/O system with strict and provable information flow security. In *ISCA*, pages 189–200. ACM, 2011.
- [45] Z. Wang and R. B. Lee. New cache designs for thwarting software cache-based side channel attacks. In *ISCA*, pages 494–505. ACM, 2007.
- [46] Y. Yu, F.-X. Standaert, O. Pereira, and M. Yung. Practical leakage-resilient pseudorandom generators. In *ACM Conference on Computer and Communications Security (CCS)*, pages 141–151. ACM, 2010.
- [47] D. Zhang, A. Askarov, and A. C. Myers. Language-based control and mitigation of timing channels. In *PLDI*, pages 99–110. ACM, 2012.
- [48] Y. Zhang, A. Juels, M. K. Reiter, and T. Ristenpart. Cross-VM Side Channels and Their Use to Extract Private Keys. In *CCS, to appear*. ACM, 2012.

A Example Code

Selection Sort

```
void SelectionSort(int a[], int array_size)
{
    int i;
    for (i = 0; i < array_size - 1; ++i)
    {
        int j, min, temp;
        min = i;
        for (j = i+1; j < array_size; ++j)
        {
            if (a[j] < a[min])
                min = j;
        }
        temp = a[i];
        a[i] = a[min];
        a[min] = temp;
    }
}
```

Insertion Sort

```
void InsertionSort(int a[], int array_size)
{
    int i, j, index;
    for (i = 1; i < array_size; ++i)
    {
        index = a[i];
        for (j = i; j > 0 && a[j-1] > index; j--)
            a[j] = a[j-1];

        a[j] = index;
    }
}
```

B Proofs

Proof of Theorem 1. Let $X \rightarrow Y \rightarrow \hat{X}$ be a Markov chain. We show that

$$P(X = \hat{X}) \leq \frac{1}{2^{H(X|Y)}},$$

where

$$H(X|Y) = -\log_2 \sum_y P(Y = y) \max_x P(X = x|Y = y)$$

is the standard definition of conditional min-entropy [42]:

$$\begin{aligned}
P(X = \hat{X}) &= \sum_{y,x} P(X = \hat{X} = x \wedge Y = y) \\
&= \sum_y P(Y = y) \sum_x P(X = \hat{X} = x | Y = y) \\
&\stackrel{(*)}{=} \sum_y P(Y = y) \sum_x P(X = x | Y = y) P(\hat{X} = x | Y = y) \\
&\leq \sum_y P(Y = y) \max_x P(X = x | Y = y) \\
&= 2^{-H(X|Y)},
\end{aligned}$$

where (*) follows from the Markov property.

From [34] it is known that the reduction in min-entropy (defined by $H(X) = -\log_2 \max_x P(X = x)$) of X when Y is known is bounded by the size of the range of Y , i.e. $H(X) - H(X|Y) \leq \log_2 |\text{ran}(Y)|$, from which we conclude

$$P(X = \hat{X}) \leq \max_x P(X = x) |\text{ran}(Y)| .$$

□

C Comprehensive results of case studies

C^{acc}	4KB	8KB	16KB	32KB	$\geq 64KB$	4KB	8KB	16KB	32KB	$\geq 64KB$
LRU/Set	100.8	81.7	70.7	70.7	68.0	96.4	42.4	8.0	3.0	0.0
LRU/IV	102.9	84.4	72.0	70.7	68.0	99.3	44.7	8.0	3.0	0.0
FIFO/Set	108.0	86.3	73.0	73.0	68.0	66.3	0.0	0.0	0.0	0.0
FIFO/IV	109.9	89.0	74.3	73.0	68.0	68.2	0.0	0.0	0.0	0.0

(a) AES without preloading (b) AES with preloading

Figure 7: The table illustrates the effect of the replacement strategy on the security guarantees we derive for access-based adversaries (C^{acc}), for different cache sizes. The rows correspond to replacement strategies/abstract domains: *Set* corresponds to our novel set-based domain, and *IV* corresponds to the interval-based domain of [21]. The entries in the table denote the security guarantee in bits, with a 64-byte cache line size.

C^{accd}	4KB	8KB	16KB	32KB	$\geq 64KB$	4KB	8KB	16KB	32KB	$\geq 64KB$
LRU	35.9	53.3	67.6	69.2	68.0	24.3	0.0	0.0	0.0	0.0
FIFO	35.9	53.3	67.6	69.2	68.0	22.3	0.0	0.0	0.0	0.0

(a) AES without preloading (b) AES with preloading

Figure 8: Effect of the replacement strategy the security guarantee for the C^{accd} -adversary, for different cache sizes.

C^{tr}	4KB	8KB	16KB	32KB	$\geq 64KB$	4KB	8KB	16KB	32KB	$\geq 64KB$
LRU	193.0	160.0	160.0	160.0	155.0	202.0	0.0	0.0	0.0	0.0
FIFO	354.0	273.0	226.0	226.0	155.0	362.0	0.0	0.0	0.0	0.0
PLRU	212.0	160.0	160.0	160.0	155.0	219.0	0.0	0.0	0.0	0.0

(a) AES without preloading (b) AES with preloading

Figure 9: Effect of the replacement strategy the security guarantee for the C^{tr} -adversary, for different cache sizes.

C^{time}	4KB	8KB	16KB	32KB	$\geq 64KB$	4KB	8KB	16KB	32KB	$\geq 64KB$
LRU	7.6	7.4	7.4	7.4	7.3	7.7	0.0	0.0	0.0	0.0
FIFO	8.5	8.1	7.9	7.9	7.3	8.6	0.0	0.0	0.0	0.0
PLRU	7.8	7.4	7.4	7.4	7.3	7.8	0.0	0.0	0.0	0.0

(a) AES without preloading

C^{time}	4KB	8KB	16KB	32KB	$\geq 64KB$	4KB	8KB	16KB	32KB	$\geq 64KB$
LRU	7.7	0.0	0.0	0.0	0.0	7.7	0.0	0.0	0.0	0.0
FIFO	8.6	0.0	0.0	0.0	0.0	8.6	0.0	0.0	0.0	0.0
PLRU	7.8	0.0	0.0	0.0	0.0	7.8	0.0	0.0	0.0	0.0

(b) AES with preloading

Figure 10: Effect of the replacement strategy the security guarantee for the C^{time} -adversary, for different cache sizes.

C^{acc}	4KB	8KB	16KB	32KB	$\geq 64KB$	4KB	8KB	16KB	32KB	$\geq 64KB$
32B	204.6	165.0	142.3	142.3	136.0	185.3	81.6	14.0	6.0	0.0
64B	100.8	81.7	70.7	70.7	68.0	96.4	42.4	8.0	3.0	0.0
128B	52.5	42.6	36.7	36.4	35.0	48.5	21.2	4.0	1.0	0.0

(a) AES without preloading

C^{acc}	4KB	8KB	16KB	32KB	$\geq 64KB$	4KB	8KB	16KB	32KB	$\geq 64KB$
32B	40.2	0.0	0.0	0.0	0.0	40.2	0.0	0.0	0.0	0.0
64B	24.3	0.0	0.0	0.0	0.0	24.3	0.0	0.0	0.0	0.0
128B	13.0	0.0	0.0	0.0	0.0	13.0	0.0	0.0	0.0	0.0

(b) AES with preloading

Figure 11: Effect of the cache line size the security guarantee for the C^{acc} -adversary, for different cache sizes. The results were obtained for the LRU replacement strategy.

C^{accd}	4KB	8KB	16KB	32KB	$\geq 64KB$	4KB	8KB	16KB	32KB	$\geq 64KB$
32B	72.7	106.9	135.7	139.0	136.0	40.2	0.0	0.0	0.0	0.0
64B	35.9	53.3	67.6	69.2	68.0	24.3	0.0	0.0	0.0	0.0
128B	18.0	27.1	34.4	35.6	35.0	13.0	0.0	0.0	0.0	0.0

(a) AES without preloading

C^{accd}	4KB	8KB	16KB	32KB	$\geq 64KB$	4KB	8KB	16KB	32KB	$\geq 64KB$
32B	40.2	0.0	0.0	0.0	0.0	40.2	0.0	0.0	0.0	0.0
64B	24.3	0.0	0.0	0.0	0.0	24.3	0.0	0.0	0.0	0.0
128B	13.0	0.0	0.0	0.0	0.0	13.0	0.0	0.0	0.0	0.0

(b) AES with preloading

Figure 12: Effect of the cache line size the security guarantee for the C^{accd} -adversary, for different cache sizes. The results were obtained for the LRU replacement strategy.

C^{tr}	4KB	8KB	16KB	32KB	$\geq 64KB$	4KB	8KB	16KB	32KB	$\geq 64KB$
32B	216.0	160.0	160.0	160.0	155.0	224.0	0.0	0.0	0.0	0.0
64B	193.0	160.0	160.0	160.0	155.0	202.0	0.0	0.0	0.0	0.0
128B	198.0	164.0	163.0	163.0	159.0	203.0	0.0	0.0	0.0	0.0

(a) AES without preloading

C^{tr}	4KB	8KB	16KB	32KB	$\geq 64KB$	4KB	8KB	16KB	32KB	$\geq 64KB$
32B	224.0	0.0	0.0	0.0	0.0	224.0	0.0	0.0	0.0	0.0
64B	202.0	0.0	0.0	0.0	0.0	202.0	0.0	0.0	0.0	0.0
128B	203.0	0.0	0.0	0.0	0.0	203.0	0.0	0.0	0.0	0.0

(b) AES with preloading

Figure 13: Effect of the cache line size the security guarantee for the C^{tr} -adversary, for different cache sizes. The results were obtained for the LRU replacement strategy.

C^{time}	4KB	8KB	16KB	32KB	$\geq 64KB$	4KB	8KB	16KB	32KB	$\geq 64KB$
32B	7.8	7.4	7.4	7.4	7.3	7.9	0.0	0.0	0.0	0.0
64B	7.6	7.4	7.4	7.4	7.3	7.7	0.0	0.0	0.0	0.0
128B	7.7	7.4	7.4	7.4	7.4	7.7	0.0	0.0	0.0	0.0

(a) AES without preloading

C^{time}	4KB	8KB	16KB	32KB	$\geq 64KB$	4KB	8KB	16KB	32KB	$\geq 64KB$
32B	7.9	0.0	0.0	0.0	0.0	7.9	0.0	0.0	0.0	0.0
64B	7.7	0.0	0.0	0.0	0.0	7.7	0.0	0.0	0.0	0.0
128B	7.7	0.0	0.0	0.0	0.0	7.7	0.0	0.0	0.0	0.0

(b) AES with preloading

Figure 14: Effect of the cache line size the security guarantee for the C^{time} -adversary, for different cache sizes. The results were obtained for the LRU replacement strategy.

C^{acc}	4KB	8KB	16KB	32KB	64KB	$\geq 128KB$	4KB	8KB	16KB	32KB	$\geq 64KB$
1-way	72.0	70.7	68.0	68.0	68.0	68.0	11.0	0.0	0.0	0.0	0.0
2-way	82.8	70.7	70.7	68.0	68.0	68.0	59.7	8.0	3.0	0.0	0.0
4-way	100.8	81.7	70.7	70.7	68.0	68.0	96.4	42.4	8.0	3.0	0.0
8-way	79.83	100.4	81.7	70.0	70.0	68.0	83.0	n/a	42.4	8.0	3.0

(a) AES without preloading

C^{acc}	4KB	8KB	16KB	32KB	64KB	$\geq 128KB$	4KB	8KB	16KB	32KB	$\geq 64KB$
1-way	11.0	0.0	0.0	0.0	0.0	0.0	11.0	0.0	0.0	0.0	0.0
2-way	59.7	8.0	3.0	0.0	0.0	0.0	59.7	8.0	3.0	0.0	0.0
4-way	96.4	42.4	8.0	3.0	0.0	0.0	96.4	42.4	8.0	3.0	0.0
8-way	83.0	n/a	42.4	8.0	3.0	0.0	83.0	n/a	42.4	8.0	3.0

(b) AES with preloading

Figure 15: Effect of the cache associativity on information leakage (in bits) for the C^{acc} -adversary, for different cache sizes. The results were obtained for the LRU replacement strategy.

C^{accd}	4KB	8KB	16KB	32KB	64KB	$\geq 128KB$	4KB	8KB	16KB	32KB	$\geq 64KB$
1-way	67.6	69.2	68.0	68.0	68.0	68.0	16.7	8.8	0.0	0.0	0.0
2-way	53.3	67.6	69.2	68.0	68.0	68.0	22.4	0.0	0.0	0.0	0.0
4-way	35.9	53.3	67.6	69.2	68.0	68.0	24.3	0.0	0.0	0.0	0.0
8-way	17.0	35.6	53.3	67.0	68.6	68.0	17.3	n/a	0.0	0.0	0.0

(a) AES without preloading

(b) AES with preloading

Figure 16: Effect of the cache associativity on information leakage (in bits) for the C^{accd} -adversary, for different cache sizes. The results were obtained for the LRU replacement strategy.

C^{tr}	4KB	8KB	16KB	32KB	64KB	$\geq 128KB$	4KB	8KB	16KB	32KB	$\geq 64KB$
1-way	241.0	241.0	155.0	155.0	155.0	155.0	182.0	130.0	0.0	0.0	0.0
2-way	201.0	160.0	160.0	155.0	155.0	155.0	210.0	0.0	0.0	0.0	0.0
4-way	193.0	160.0	160.0	160.0	155.0	155.0	202.0	0.0	0.0	0.0	0.0
8-way	160.0	160.0	160.0	156.0	156.0	155.0	164.0	0.0	0.0	0.0	0.0

(a) AES without preloading

(b) AES with preloading

Figure 17: Effect of the cache associativity on information leakage (in bits) for the C^{tr} -adversary, for different cache sizes. The results were obtained for the LRU replacement strategy, with a 64-byte cache line size.

C^{time}	4KB	8KB	16KB	32KB	$\geq 64KB$	4KB	8KB	16KB	32KB	$\geq 64KB$
1-way	8.0	8.0	7.3	7.3	7.3	7.6	7.1	0.0	0.0	0.0
2-way	7.7	7.4	7.4	7.3	7.3	7.8	0.0	0.0	0.0	0.0
4-way	7.6	7.4	7.4	7.4	7.3	7.7	0.0	0.0	0.0	0.0
8-way	7.4	7.4	7.4	7.3	7.3	7.4	0.0	0.0	0.0	0.0

(a) AES without preloading

(b) AES with preloading

Figure 18: Effect of the cache associativity on information leakage (in bits) for the C^{time} -adversary, for different cache sizes. The results were obtained for the LRU replacement strategy, with a 64-byte cache line size.