



Quantitative Mitigation of Timing Side Channels

Saeid Tizpaz-Niari^(✉), Pavol Černý,
and Ashutosh Trivedi

University of Colorado Boulder, Boulder, USA
Saeid.TizpazNiari@colorado.edu



Abstract. Timing side channels pose a significant threat to the security and privacy of software applications. We propose an approach for *mitigating* this problem by decreasing the strength of the side channels as measured by entropy-based objectives, such as min-guess entropy. Our goal is to minimize the information leaks while guaranteeing a user-specified maximal acceptable performance overhead. We dub the decision version of this problem *Shannon mitigation*, and consider two variants, *deterministic* and *stochastic*. First, we show that the deterministic variant is NP-hard. However, we give a polynomial algorithm that finds an optimal solution from a restricted set. Second, for the stochastic variant, we develop an approach that uses optimization techniques specific to the entropy-based objective used. For instance, for min-guess entropy, we used mixed integer-linear programming. We apply the algorithm to a threat model where the attacker gets to make *functional observations*, that is, where she observes the running time of the program for the same secret value combined with different public input values. Existing mitigation approaches do not give confidentiality or performance guarantees for this threat model. We evaluate our tool SCHMIT on a number of micro-benchmarks and real-world applications with different entropy-based objectives. In contrast to the existing mitigation approaches, we show that in the functional-observation threat model, SCHMIT is scalable and able to maximize confidentiality under the performance overhead bound.

1 Introduction

Information leaks through timing side channels remain a challenging problem [13, 16, 24, 29, 35, 37, 47]. A program leaks secret information through timing side channels if an attacker can deduce secret values (or their properties) by observing response times. We consider the problem of mitigating timing side channels. Unlike elimination techniques [7, 31, 46] that aim to completely remove timing leaks without considering the performance penalty, the goal of mitigation techniques [10, 26, 48] is to weaken the leaks, while keeping the penalty low.

We define the *Shannon mitigation* problem that decides whether there is a mitigation policy to achieve a lower bound on a given security entropy-based

measure while respecting an upper bound on the performance overhead. Consider an example where the program-under-analysis has a secret variable with seven possible values, and has three different timing behaviors, each forming a cluster of secret values. It takes 1 second if the secret value is 1, it takes 5 seconds if the secret is between 2 and 5, and it takes 10 seconds if the secret value is 6 or 7. The *entropy-based measure* quantifies the remaining uncertainty about the secret after timing observations. Min-guess entropy [11, 25, 41] for this program is 1, because if the observed execution time is 1, the attacker guesses the secret in one try. A *mitigation policy* involves merging some timing clusters by introducing delays. A good solution might be to introduce a 9 second delay if the secret is 1, which merges two timing clusters. But, this might be disallowed by the budget on the performance overhead. Therefore, another solution must be found, such as introducing a 4 seconds delay when the secret is one.

We develop two variants of the Shannon mitigation problem: *deterministic* and *stochastic*. The mitigation policy of the deterministic variant requires us to move all secret values associated to an observation to another observation, while the policy of the stochastic variant allows us to move only a portion of secret values in an observation to another one. We show that the deterministic variant of the Shannon mitigation problem is intractable and propose a dynamic programming algorithm to approximate the optimal solution for the problem by searching through a restricted set of solutions. We develop an algorithm that reduces the problem in the stochastic variant to a well-known optimization problem that depends on the entropy-based measure. For instance, with min-guess entropy, the optimization problem is mixed integer-linear programming.

We consider a threat model where an attacker knows the public inputs (known-message attacks [26]), and furthermore, where the public input changes much more often than the secret inputs (for instance, secrets such as bank account numbers do not change often). As a result, for each secret, the attacker observes a timing function of the public inputs. We call this model *functional observations* of timing side channels.

We develop our tool SCHMIT that has three components: side channel discovery [45], search for the mitigation policy, and the policy enforcement. The side channel discovery builds the functional observations [45] and measures the entropy of secret set after the observations. The mitigation policy component includes the implementation of the dynamic programming and optimization algorithms. The enforcement component is a monitoring system that uses the program internals and functional observations to enforce the policy at runtime. To summarize, we make the following contributions:

- We formalize the *Shannon mitigation* problem with two variants and show that the complexity of finding deterministic mitigation policy is NP-hard.
- We describe two algorithms for synthesizing the mitigation policy: one is based on dynamic programming for the deterministic variant, that is in polynomial time and results in an approximate solution, and the other one solves the stochastic variant of the problem with optimization techniques.

- We consider a threat model that results in functional observations. On a set of micro-benchmarks, we show that existing mitigation techniques are not secure and efficient for this threat model.
- We evaluate our approach on five real-world Java applications. We show that SCHMIT is scalable in synthesizing mitigation policy within a few seconds and significantly improves the security (entropy) of the applications.

```

Example(int high, int low) {
int t_high = high, t_low = low;
while (t_high > 0) {
  if (t_high % 2 == 1) {
    while (t_low > 0) {
      if (t_low % 2 == 1) {
        res += compute(t_low, t_high);}
        t_low = t_low >> 1;}}
    t_high = t_high >> 1;}
return res;}

```




Fig. 1. (a) The example used in Sect. 2. (b) The timing functions for each secret value of the program.

2 Overview

First, we describe the threat model considered in this paper. Second, we describe our approach on a running example. Third, we compare the results of SCHMIT with the existing mitigation techniques [10, 26, 48] and show that SCHMIT achieves the highest entropy (i.e., best mitigation) for all three entropy objectives.

Threat Model. We assume that the attacker has access to the source code and the mitigation model, and she can sample the run-time of the application arbitrarily many times on her own machine. During an attack, she intends to guess a fixed secret of the target machine by observing the mitigated running time. Since we consider the attack models where the attacker knows the public inputs and the secret inputs are less volatile than public inputs, her observations are functional observations, where for each secret value, she learns a function from the public inputs to the running time.

Example 2.1. Consider the program shown in Fig. 1(a). It takes secret and public values as inputs. The running time depends on the number of set bits in both secret and public inputs. We assume that secret and public inputs can be between 1 and 1023. Figure 1(b) shows the running time of different secret values as timing functions, i.e., functions from the public inputs to the running time.

Side channel discovery. One can use existing tools to find the initial functional observations [44,45]. In Example 2.1, functional observations are $\mathcal{F} = \langle y, 2y, \dots, 10y \rangle$ where y is a variable whose value is the number of set bits in the public input. The corresponding secret classes after this observation is $\mathcal{S}_{\mathcal{F}} = \langle 1_1, 1_2, 1_3, \dots, 1_{10} \rangle$ where 1_n shows a set of secret values that have n set bits. The sizes of classes are $B = \{10, 45, 120, 210, 252, 210, 120, 45, 10, 1\}$. We use L_1 -norm as metric to calculate the distance between the functional observations \mathcal{F} . This distance (penalty) matrix specifies extra performance overhead to move from one functional observation to another. With the assumption of uniform distributions over the secret input, Shannon entropy, guessing entropy, and the min-guessing entropy are 7.3, 90.1, and 1.0, respectively. These entropies are defined in Sect. 3 and measure the remaining entropy of the secret set after the observations. We aim to maximize the entropy measures, while keeping the performance overhead below a threshold, say 60% for this example.

Mitigation with Schmit. We use our tool SCHMIT to mitigate timing leaks of Example 2.1. The mitigation policy for the Shannon entropy objective is shown in Fig. 2(a). The policy results in two classes of observations. The policy requires to move functional observations $\langle y, 2y, \dots, 5y \rangle$ to $\langle 6y \rangle$ and all other observations $\langle 7y, 8y, 9y \rangle$ to $\langle 10y \rangle$. To enforce this policy, we use a monitoring system at runtime. The monitoring system uses a decision tree model of the initial functional observations. The decision tree model characterizes each functional observation with associated program internals such as method calls or basic block invocations [43,44]. The decision tree model for the Example 2.1 is shown in Fig. 2(b). The monitoring system records program internals and matches it with the decision tree model to detect the current functional observation. Then, it adds delays, if necessary, to the execution time in order to enforce the mitigation policy. With this method, the mitigated functional observation is $\mathcal{G} = \langle 6y, 10y \rangle$ and the secret




Fig. 2. (a) Mitigation policy calculation with deterministic algorithm (left). The observations $x1$ and $x2$ stands for all observations from $C_2 - C_5$ and from $C_8 - C_9$, resp.; (b) Leaned discriminant decision tree (center): it characterizes the functional clusters of Fig. 1(b) with internals of the program in Fig. 1(a); and (c) observations (right) after the mitigation by SCHMIT results in two classes of observations.

class is $\mathcal{S}_G = \langle \{1_1, 1_2, 1_3, 1_4, 1_5, 1_6\}, \{1_7, 1_8, 1_9, 1_{10}\} \rangle$ as shown in Fig. 2 (c). The performance overhead of this mitigation is 43.1%. The Shannon, guessing, and min-guess entropies have improved to 9.7, 459.6, and 193.5, respectively.

Comparison with state of the art. We compare our mitigation results to black-box mitigation scheme [10] and bucketing [26]. *Black-box double scheme technique.* We use the double scheme technique [10] to mitigate the leaks of Example 2.1. This mitigation uses a prediction model to release events at scheduled times. Let us consider the prediction for releasing the event i at N -th epoch with $S(N, i) = \max(inp_i, S(N, i-1)) + p(N)$, where inp_i is the time arrival of the i -th request, $S(N, i-1)$ is the prediction for the request $i-1$, and $p(N) = 2^{N-1}$ models the basis for the prediction scheme at N -th epoch. We assume that the request are the same type and the sequence of public input requests for each secret are received in the beginning of epoch $N = 1$. Figure 3(a) shows the functional observations after applying the predictive mitigation. With this mitigation, the classes of observations are $\mathcal{S}_G = \langle 1_1, \{1_2, 1_3\}, \{1_4, 1_5, 1_6, 1_7\}, \{1_8, 1_9, 1_{10}\} \rangle$. The number of classes of observations is reduced from 10 to 4. The performance overhead is 39.9%. The Shannon, guessing, and min-guess entropies have increased to 9.00, 321.5, and 5.5, respectively. *Bucketing.* We consider the mitigation approach with buckets [26]. For Example 2.1, if the attacker does not know the public input (unknown-message attacks [26]), the observations are $\{1.1, 2.1, 3.3, \dots, 9.9, 10.9, \dots, 109.5\}$ as shown in Fig. 3(b). We apply the bucketing algorithm in [26] for this observations, and it finds two buckets $\{37.5, 109.5\}$ shown with the red lines in Fig. 3(b). The bucketing mitigation requires to move the observations to the closet bucket. Without functional observations, there are 2 classes of observations. However, with functional observations, there are more than 2 observations. Figure 3(c) shows how the pattern of observations are leaking through functional side channels. There are 7 classes of observations: $\mathcal{S}_G = \langle \{1_1, 1_2, 1_3\}, \{1_4\}, \{1_5\}, \{1_6\}, \{1_7\}, \{1_8\}, \{1_9\}, \{1_{10}\} \rangle$. The Shannon, guessing, and min-guess entropies are 7.63, 102.3, and 1.0, respectively.




Fig. 3. (a) The execution time after mitigation using the double scheme technique [10]. There are four classes of functional observations after the mitigation. (b) Mitigation with bucketing [26]. All observations require to move to the closet red line. (c) Functional observations distinguish 7 classes of observations after mitigating with bucketing.

Overall, SCHMIT achieves the higher entropy measures for all three objectives under the performance overhead of 60%.

3 Preliminaries

For a finite set Q , we use $|Q|$ for its cardinality. A *discrete probability distribution*, or just *distribution*, over a set Q is a function $d : Q \rightarrow [0, 1]$ such that $\sum_{q \in Q} d(q) = 1$. Let $\mathcal{D}(Q)$ denote the set of all discrete distributions over Q . We say a distribution $d \in \mathcal{D}(Q)$ is a *point distribution* if $d(q)=1$ for a $q \in Q$. Similarly, a distribution $d \in \mathcal{D}(Q)$ is *uniform* if $d(q)=1/|Q|$ for all $q \in Q$.

Definition 1 (Timing Model). *The timing model of a program \mathcal{P} is a tuple $[\mathcal{P}] = (X, Y, \mathcal{S}, \delta)$ where $X = \{x_1, \dots, x_n\}$ is the set of secret-input variables, $Y = \{y_1, \dots, y_m\}$ is the set of public-input variables, $\mathcal{S} \subseteq \mathbb{R}^n$ is a finite set of secret-inputs, and $\delta : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}_{\geq 0}$ is the execution-time function of the program over the secret and public inputs.*

We assume that the adversary knows the program and wishes to learn the value of the secret input. To do so, for some fixed secret value $s \in \mathcal{S}$, the adversary can invoke the program to estimate (to an arbitrary precision) the execution time of the program. If the set of public inputs is empty, i.e. $m = 0$, the adversary can only make *scalar observations* of the execution time corresponding to a secret value. In the more general setting, however, the adversary can arrange his observations in a functional form by estimating an approximation of the *timing function* $\delta(s) : \mathbb{R}^m \rightarrow \mathbb{R}_{\geq 0}$ of the program.

A *functional observation* of the program \mathcal{P} for a secret input $s \in \mathcal{S}$ is the function $\delta(s) : \mathbb{R}^m \rightarrow \mathbb{R}_{\geq 0}$ defined as $\mathbf{y} \in \mathbb{R}^m \mapsto \delta(s, \mathbf{y})$. Let $\mathcal{F} \subseteq [\mathbb{R}^m \rightarrow \mathbb{R}_{\geq 0}]$ be the finite set of all functional observations of the program \mathcal{P} . We define an order \prec over the functional observations \mathcal{F} : for $f, g \in \mathcal{F}$ we say that $f \prec g$ if $f(y) \leq g(y)$ for all $y \in \mathbb{R}^m$.

The set \mathcal{F} characterizes an equivalence relation $\equiv_{\mathcal{F}}$, namely secrets with equivalent functional observations, over the set \mathcal{S} , defined as following: $s \equiv_{\mathcal{F}} s'$ if there is an $f \in \mathcal{F}$ such that $\delta(s) = \delta(s') = f$. Let $\mathcal{S}_{\mathcal{F}} = \langle S_1, S_2, \dots, S_k \rangle$ be the quotient space of \mathcal{S} characterized by the observations $\mathcal{F} = \langle f_1, f_2, \dots, f_k \rangle$. We write \mathcal{S}_f for the secret set $S \in \mathcal{S}_{\mathcal{F}}$ corresponding to the observations $f \in \mathcal{F}$. Let $\mathcal{B} = \langle B_1, B_2, \dots, B_k \rangle$ be the size of observational equivalence class in $\mathcal{S}_{\mathcal{F}}$, i.e. $B_i = |\mathcal{S}_{f_i}|$ for $f_i \in \mathcal{F}$ and let $B = |\mathcal{S}| = \sum_{i=1}^k B_i$.

Shannon entropy, guessing entropy, and min-guess entropy are three prevalent information metrics to quantify information leaks in programs. Köpf and Basin [25] characterize expressions for various information-theoretic measures on information leaks when there is a uniform distribution on \mathcal{S} given below.

Proposition 1 (Köpf and Basin [25]). *Let $\mathcal{F} = \langle f_1, \dots, f_k \rangle$ be a set of observations and let \mathcal{S} be the set of secret values. Let $\mathcal{B} = \langle B_1, \dots, B_k \rangle$ be the*

corresponding size of secret set in each class of observation and $B = \sum_{i=1}^k B_i$. Assuming a uniform distribution on \mathcal{S} , entropies can be characterized as:

1. **Shannon Entropy:** $SE(\mathcal{S}|\mathcal{F}) \stackrel{\text{def}}{=} (\frac{1}{B}) \sum_{1 \leq i \leq k} B_i \log_2(B_i)$,
2. **Guessing Entropy:** $GE(\mathcal{S}|\mathcal{F}) \stackrel{\text{def}}{=} (\frac{1}{2B}) \sum_{1 \leq i \leq k} B_i^2 + \frac{1}{2}$, and
3. **Min-Guess Entropy:** $mGE(\mathcal{S}|\mathcal{F}) \stackrel{\text{def}}{=} \min_{1 \leq i \leq k} \{(B_i + 1)/2\}$.

4 Shannon Mitigation Problem

Our goal is to mitigate the information leakage due to the timing side channels by adding synthetic delays to the program. An aggressive, but commonly-used, mitigation strategy aims to eliminate the side channels by adding delays such that every secret value yields a common functional observation. However, this strategy may often be impractical as it may result in unacceptable performance degradations of the response time. Assuming a well-known penalty function associated with the performance degradation, we study the problem of maximizing entropy while respecting a bound on the performance degradation. We dub the decision version of this problem Shannon mitigation.

Adding synthetic delays to execution-time of the program, so as to mask the side-channel, can give rise to new functional observations that correspond to upper-envelopes of various combinations of original observations. Let $\mathcal{F} = \langle f_1, f_2, \dots, f_k \rangle$ be the set of functional observations. For $I \subseteq 1, 2, \dots, k$, let $f_I = \mathbf{y} \in \mathbb{R}^m \mapsto \sup_{i \in I} f_i(\mathbf{y})$ be the functional observation corresponding to upper-envelope of the functional observations in the set I . Let $\mathcal{G}(\mathcal{F}) = \{f_I : I \neq \emptyset \subseteq \{1, 2, \dots, k\}\}$ be the set of all possible functional observations resulting from the upper-envelope calculations. To change the observation of a secret value with functional observation f_i to a new observation f_I (we assume that $i \in I$), we need to add delay function $f_I^i : \mathbf{y} \in \mathbb{R}^m \mapsto f_I(\mathbf{y}) - f_i(\mathbf{y})$.

Mitigation Policies. Let $\mathcal{G} \subseteq \mathcal{G}(\mathcal{F})$ be a set of admissible post-mitigation observations. A *mitigation policy* is a function $\mu : \mathcal{F} \rightarrow \mathcal{D}(\mathcal{G})$ that for each secret $s \in \mathcal{S}_f$ suggests the probability distribution $\mu(f)$ over the functional observations. We say that a mitigation policy is *deterministic* if for all $f \in \mathcal{F}$ we have that $\mu(f)$ is a point distribution. Abusing notations, we represent a deterministic mitigation policy as a function $\mu : \mathcal{F} \rightarrow \mathcal{G}$. The semantics of a mitigation policy recommends to a program analyst a probability $\mu(f)(g)$ to elevate a secret input $s \in \mathcal{S}_f$ from the observational class f to the class $g \in \mathcal{G}$ by adding $\max\{0, g(p) - f(p)\}$ units delay to the corresponding execution-time $\delta(s, p)$ for all $p \in Y$. We assume that the mitigation policies respect the order, i.e. for every mitigation policy μ and for all $f \in \mathcal{F}$ and $g \in \mathcal{G}$, we have that $\mu(f)(g) > 0$ implies that $f \prec g$. Let $M_{(\mathcal{F} \rightarrow \mathcal{G})}$ be the set of mitigation policies from the set of observational clusters \mathcal{F} into the clusters \mathcal{G} .

For the functional observations $\mathcal{F} = \langle f_1, \dots, f_k \rangle$ and a mitigation policy $\mu \in M_{(\mathcal{F} \rightarrow \mathcal{G})}$, the resulting observation set $\mathcal{F}[\mu] \subseteq \mathcal{G}$ is defined as:

$$\mathcal{F}[\mu] = \{g \in \mathcal{G} : \text{there exists } f \in \mathcal{F} \text{ such that } \mu(f)(g) > 0\}.$$

Since the mitigation policy is stochastic, we use average sizes of resulting observations to represent fitness of a mitigation policy. For $\mathcal{F}[\mu] = \langle g_1, g_2, \dots, g_\ell \rangle$, we define their expected class sizes $\mathcal{B}_\mu = \langle C_1, C_2, \dots, C_\ell \rangle$ as $C_i = \sum_{j=1}^i \mu(f_j)(f_i) \cdot B_j$ (observe that $\sum_{i=1}^\ell C_i = B$). Assuming a uniform distribution on \mathcal{S} , various entropies for the expected class size after applying a policy $\mu \in M_{(\mathcal{F} \rightarrow \mathcal{G})}$ can be characterized by the following expressions:

1. **Shannon Entropy:** $\text{SE}(\mathcal{S}|\mathcal{F}, \mu) \stackrel{\text{def}}{=} (\frac{1}{B}) \sum_{1 \leq i \leq \ell} C_i \log_2(C_i)$,
2. **Guessing Entropy:** $\text{GE}(\mathcal{S}|\mathcal{F}, \mu) \stackrel{\text{def}}{=} (\frac{1}{2B}) \sum_{1 \leq i \leq \ell} C_i^2 + \frac{1}{2}$, and
3. **Min-Guess Entropy:** $\text{mGE}(\mathcal{S}|\mathcal{F}, \mu) \stackrel{\text{def}}{=} \min_{1 \leq i \leq \ell} \{(C_i + 1)/2\}$.

We note that the above definitions do not represent the expected entropies, but rather entropies corresponding to the expected cluster sizes. However, the three quantities provide bounds on the expected entropies after applying μ . Since Shannon and Min-Guess entropies are concave functions, from Jensen's inequality, we get that $\text{SE}(\mathcal{S}|\mathcal{F}, \mu)$ and $\text{mGE}(\mathcal{S}|\mathcal{F}, \mu)$ are upper bounds on expected Shannon and Min-Guess entropies. Similarly, $\text{GE}(\mathcal{S}|\mathcal{F}, \mu)$, being a convex function, give a lower bound on expected guessing entropy.

We are interested in maximizing the entropy while respecting constraints on the overall performance of the system. We formalize the notion of performance by introducing performance penalties: there is a function $\pi : \mathcal{F} \times \mathcal{G} \rightarrow \mathbb{R}_{\geq 0}$ such that elevating from the observation $f \in \mathcal{F}$ to the functional observation $g \in \mathcal{G}$ adds an extra $\pi(f, g)$ performance overheads to the program. The expected performance penalty associated with a policy μ , $\pi(\mu)$, is defined as the probabilistically weighted sum of the penalties, i.e. $\sum_{f \in \mathcal{F}, g \in \mathcal{G}: f \prec g} |\mathcal{S}_f| \cdot \mu(f)(g) \cdot \pi(f, g)$. Now, we introduce our key decision problem.

Definition 2 (Shannon Mitigation). *Given a set of functional observations $\mathcal{F} = \langle f_1, \dots, f_k \rangle$, a set of admissible post-mitigation observations $\mathcal{G} \subseteq \mathcal{G}(\mathcal{F})$, set of secrets \mathcal{S} , a penalty function $\pi : \mathcal{F} \times \mathcal{G} \rightarrow \mathbb{R}_{\geq 0}$, a performance penalty upper bound $\Delta \in \mathbb{R}_{\geq 0}$, and an entropy lower-bound $E \in \mathbb{R}_{\geq 0}$, the Shannon mitigation problem $\text{SHAN}_{\mathcal{E}}(\mathcal{F}, \mathcal{G}, \mathcal{S}, \pi, E, \Delta)$, for a given entropy measure $\mathcal{E} \in \{\text{SE}, \text{GE}, \text{mGE}\}$, is to decide whether there exists a mitigation policy $\mu \in M_{(\mathcal{F} \rightarrow \mathcal{G})}$ such that $\mathcal{E}(\mathcal{S}|\mathcal{F}, \mu) \geq E$ and $\pi(\mu) \leq \Delta$. We define the deterministic Shannon mitigation variant where the goal is to find a deterministic such policy.*

5 Algorithms for Shannon Mitigation Problem

5.1 Deterministic Shannon Mitigation

We first establish the intractability of the deterministic variant.

Theorem 1. *Deterministic Shannon mitigation problem is NP-complete.*

Proof. It is easy to see that the deterministic Shannon mitigation problem is in NP: one can guess a certificate as a deterministic mitigation policy $\mu \in M_{(\mathcal{F} \rightarrow \mathcal{G})}$

and can verify in polynomial time that it satisfies the entropy and overhead constraints. Next, we sketch the hardness proof for the min-guess entropy measure by providing a reduction from the *two-way partitioning* problem [28]. For the Shannon entropy and guess entropy measures, a reduction can be established from the Shannon capacity problem [18] and the Euclidean sum-of-squares clustering problem [8], respectively.

Given a set $A = \{a_1, a_2, \dots, a_k\}$ of integer values, the two-way partitioning problem is to decide whether there is a partition $A_1 \uplus A_2 = A$ into two sets A_1 and A_2 with equal sums, i.e. $\sum_{a \in A_1} a = \sum_{a \in A_2} a$. W.l.o.g assume that $a_i \leq a_j$ for $i \leq j$. We reduce this problem to a deterministic Shannon mitigation problem $\text{SHAN}_{\text{mGE}}(\mathcal{F}_A, \mathcal{G}_A, \mathcal{S}_A, \pi_A, E_A, \Delta_A)$ with k clusters $\mathcal{F}_A = \mathcal{G}_A = \langle f_1, f_2, \dots, f_k \rangle$ with the secret set $\mathcal{S}_A = \langle S_1, S_2, \dots, S_k \rangle$ such that $|S_i| = a_i$. If $\sum_{1 \leq i \leq k} a_i$ is odd then the solution to the two-way partitioning instance is trivially **no**. Otherwise, let $E_A = (1/2) \sum_{1 \leq i \leq k} a_i$. Notice that any deterministic mitigation strategy that achieves min-guess entropy larger than or equal to E_A must have at most two clusters. On the other hand, the best min-guess entropy value can be achieved by having just a single cluster. To avoid this and force getting two clusters corresponding to the two partitions of a solution to the two-way partitions problem instance A , we introduce performance penalties such that merging more than $k - 2$ clusters is disallowed by keeping performance penalty $\pi_A(f, g) = 1$ and performance overhead $\Delta_A = k - 2$. It is straightforward to verify that an instance of the resulting min-guess entropy problem has a **yes** answer if and only if the two-way partitioning instance does. \square

Since the deterministic Shannon mitigation problem is intractable, we design an approximate solution for the problem. Note that the problem is hard even if we only use existing functional observations for mitigation, i.e., $\mathcal{G} = \mathcal{F}$. Therefore, we consider this case for the approximate solution. Furthermore, we assume the following *sequential dominance* restriction on a deterministic policy μ : for $f, g \in \mathcal{F}$ if $f \prec g$ then either $\mu(f) \prec g$ or $\mu(f) = \mu(g)$. In other words, for any given $f \prec g$, f can not be moved to a higher cluster than g without having g be moved to that cluster. For example, Fig. 4(a) shows Shannon mitigation problem with four functional observations and all possible mitigation policies (we represent $\mu(f_i)(f_j)$ with $\mu(i, j)$). Figure 4(b) satisfies the sequential dominance restriction, while Fig. 4(c) does not.

The search for the deterministic policies satisfying the sequential dominance restriction can be performed efficiently using dynamic programming by effective use of intermediate results' memorizations.

Algorithm (1) provides a pseudocode for the dynamic programming solution to find a deterministic mitigation policy satisfying the sequential dominance. The key idea is to start with considering policies that produce a single cluster for subclasses P_i of the problem with the observation from $\langle f_1, \dots, f_i \rangle$, and then compute policies producing one additional cluster in each step by utilizing the previously computed sub-problems and keeping track of the performance penalties. The algorithm terminates as soon as the solution of the current step respects the performance bound. The complexity of the algorithm is $O(k^3)$.




Fig. 4. (a). Example of Shannon mitigation problem with all possible mitigation policies for 4 classes of observations. (b,c) Two examples of the mitigation policies that results in 2 and 3 classes of observations.

5.2 Stochastic Shannon Mitigation Algorithm

Next, we solve the (stochastic) Shannon mitigation problem by posing it as an optimization problem. Consider the stochastic Shannon mitigation problem $\text{SHAN}_{\mathcal{E}}(\mathcal{F}, \mathcal{G} = \mathcal{F}, \mathcal{S}_{\mathcal{F}}, \pi, E, \Delta)$ with a stochastic policy $\mu : \mathcal{F} \rightarrow \mathcal{D}(\mathcal{G})$ and

Algorithm 1. APPROXIMATE DETERMINISTIC SHANNON MITIGATION

Input: The Shannon entropy problem $\text{SHAN}_{MGE}(\mathcal{F}, \mathcal{G} = \mathcal{F}, \mathcal{S}_{\mathcal{F}}, \pi, E, \Delta)$

Output: The entropy table (T).

```

1 for  $i = 1$  to  $k$  do
2    $T(i, 1) = \mathcal{E}(\bigcup_{j=1}^i S_j)$ 
3   if  $\sum_{1 \leq j \leq i} \pi(j, i)(B_j/B) \leq \Delta$  then  $\Pi(i, 1) = \sum_{1 \leq j \leq i} \pi(j, i)(B_j/B)$ 
4   else  $\Pi(i, 1) = \infty$ 
5 if  $\Pi(k, 1) < \infty$  then return  $T$ ;
6 for  $r = 2$  to  $k$  do
7   for  $i = 1$  to  $k$  do
8      $\Omega(i, r) = \{j : 1 \leq j < i \text{ and } \Pi(j, r-1) + \sum_{j < q \leq i} \pi(q, i)(B_q/B) \leq \Delta\}$ 
9     if  $\Omega \neq \emptyset$  then  $T(i, r) = \max_{j \in \Omega(i, r)} \left( \min(T(j, r-1), \mathcal{E}(\bigcup_{q=j+1}^i S_q)) \right)$ 
10    else  $T(i, r) = -\infty$ 
11    Let  $j$  be the index that maximizes  $T(i, r)$ 
12    if  $\Omega \neq \emptyset$  then  $\Pi(i, r) = (\Pi(j, r-1) + \sum_{j < q \leq i} \pi(q, i)(B_q/B))$ 
13    else  $\Pi(i, r) = \infty$ 
14  if  $\Pi(k, r) < \infty$  then return  $T$ ;
15 return  $T$ ;
    
```

$\mathcal{S}_{\mathcal{F}} = \langle S_1, S_2, \dots, S_k \rangle$. The following program characterizes the optimization problem that solves the Shannon mitigation problem with stochastic policy.

Maximize \mathcal{E} , subject to:

1. $0 \leq \mu(f_i)(f_j) \leq 1$ for $1 \leq i \leq j \leq k$
2. $\sum_{i \leq j \leq k} \mu(f_i)(f_j) = 1$ for all $1 \leq i \leq k$.
3. $\sum_{i=1}^k \sum_{j=i}^k |S_i| \cdot \mu(f_i)(f_j) \cdot \pi(f_i, f_j) \leq \Delta$.
4. $C_j = \sum_{i=1}^j |S_i| \cdot \mu(f_i)(f_j)$ for $1 \leq j \leq k$.

Here, the objective function \mathcal{E} is one of the following functions:

1. **Guessing Entropy** $\mathcal{E}_{GE} = \sum_{j=1}^k C_j^2$
2. **Min-Guess Entropy** $\mathcal{E}_{MGE} = \min_{1 \leq j \leq k} \{C_j \mid C_j > 0\}$
3. **Shannon Entropy** $\mathcal{E}_{SE} = \sum_{j=1}^k C_j \cdot \log_2(C_j)$

The linear constraints for the problem are defined as the following. The condition (1) and (2) express that μ provides a probability distributions, condition (3) provides restrictions regarding the performance constraint, and the condition (4) is the entropy specific constraint. The objective function of the optimization problem is defined based on the entropy criteria from \mathcal{E} . For the simplicity, we omit the constant terms from the objective function definitions. For the guessing entropy, the problem is an instance of linearly constrained quadratic optimization problem [33]. The problem with Shannon entropy is a non-linear optimization problem [12]. Finally, the optimization problem with min-guess entropy is an instance of mixed integer programming [32]. We evaluate the scalability of these solvers empirically in Sect. 6 and leave the exact complexity as an open problem. We show that the min-guess entropy objective function can be efficiently solved with the branch and bound algorithms [36]. Figure 4(b,c) show two instantiations of the mitigation policies that are possible for the stochastic mitigation.

6 Implementation Details

A. Environmental Setups. All timing measurements are conducted on an Intel NUC5i5RYH. We switch off JIT Compilation and run each experiment multiple times and use the mean running time. This helps to reduce the effects of environmental factors such as the Garbage Collections. All other analyses are conducted on an Intel i5-2.7 GHz machine.

B. Implementation of Side Channel Discovery. We use the technique presented in [45] for the side channel discovery. The technique applies the functional

data analysis [38] to create B-spline basis and fit functions to the vector of timing observations for each secret value. Then, the technique applies the functional data clustering [21] to obtain K classes of observations. We use the number of secret values in a cluster as the class size metric and the L_1 distance norm between the clusters as the penalty function.

C. Implementation of Mitigation Policy Algorithms. For the stochastic optimization, we encode the Shannon entropy and guessing entropy with linear constraints in Scipy [22]. Since the objective functions are non-linear (for the Shannon entropy) and quadratic (for the guessing entropy), Scipy uses sequential least square programming (SLSQP) [34] to maximize the objectives. For the stochastic optimization with the min-guess entropy, we encode the problem in Gurobi [19] as a mixed-integer programming (MIP) problem [32]. Gurobi solves the problem efficiently with branch-and-bound algorithms [1]. We use Java to implement the dynamic programming.

D. Implementation of Enforcement. The enforcement of mitigation policy is implemented in two steps. *First*, we use the initial timing functions and characterize them with program internal properties such as basic block calls. To do so, we use the decision tree learning approach presented in [45]. The decision tree model characterizes each functional observations with properties of program internals. *Second*, given the policy of mitigation, we enforce the mitigation policy with a monitoring system implemented on top of the Javassist [15] library. The monitoring system uses the decision tree model and matches the properties enabled during an execution with the tree model (detection of the current cluster). Then, it adds extra delays, based on the mitigation policy, to the current execution-time and enforces the mitigation policy. Note that the dynamic monitoring can result in a few micro-second delays. For the programs with timing differences in the order of micro-seconds, we transform source code using the decision tree model. The transformation requires manual efforts to modify and compile the new program. But, it adds negligible delays.

E. Micro-benchmark Results. Our goal is to compare different mitigation methods in terms of their security and performance. We examine the computation time of our tool SCHMIT in calculating the mitigation policies. See appendix for the relationships between performance bounds and entropy measures.

Applications: Mod_Exp applications [30] are instances of square-and-multiply modular exponentiation ($R = y^k \text{ mod } n$) used for secret key operations in RSA [39]. Branch_and_Loop series consist of 6 applications where each application has conditions over secret values and runs a linear loop over the public values. The running time of the applications depend on the slope of the linear loops determined by the secret input.

Computation time comparisons: Fig. 5 shows the computation time for Branch_and_Loop applications (the applications are ordered in x-axis based on the discovered number of observational classes). For the min-guess entropy, we observe that both stochastic and dynamic programming approaches are efficient and fast as shown in Fig. 5(a). For the Shannon and guessing entropies,

Table 1. Micro-benchmark results. M.E and B.L stand for Mod_Exp and Branch_and_Loop applications. Legend: #S: no. of secret values, #P: no. of public values, Δ : Upper bound over performance penalty, ϵ : clustering parameter, #K: classes of observations before mitigation, #K_X: classes of observations after mitigation with X technique, mGE: Min-guess entropy before mitigation, mGE_X: Min-guess entropy after mitigation with X, O_X: Performance overhead added after mitigation with X.

App(s)	Initial Characteristics				Double Scheme				Bucketing			SCHMIT (Determin.)			SCHMIT (Stoch.)			
	#S	#P	Δ	ϵ	#K	mGE	#K _{DS}	mGE _{DS}	O _{DS} (%)	#K _B	mGE _B	O _B (%)	K _D	#mGE _D	O _D (%)	#K _S	mGE _S	O _S (%)
M.E.1	32	32	0.5	1.0	1	16.5	1	16.5	0.0	1	16.5	0.0	1	16.5	0.0	1	16.5	0.0
M.E.2	64	64	0.5	1.0	2	16.5	1	32.5	5,221	1	32.5	27.6	1	32.5	21.4	1	32.5	21.4
M.E.3	128	128	0.5	2.0	2	32.5	1	64.5	5,407	1	64.5	33.9	1	64.5	22.7	1	64.5	22.7
M.E.4	256	256	0.5	2.0	4	10.5	1	128.5	6,679	1	128.5	30.7	1	128.5	28.3	1	128.5	28.3
M.E.5	512	512	0.5	5.0	23	1.0	1	256.5	7,294	2	128.5	50.0	1	256.5	31.0	1	253.0	30.3
M.E.6	1,024	1,024	0.5	8.0	40	1.0	1	512.5	7,822	20	1.0	34.5	2	27.5	46.7	5	85.5	50.0
B.L.1	25	50	0.5	10.0	4	3.0	3	3.0	73.0	3	3.0	17.5	2	5.5	26.1	2	6.5	34.9
B.L.2	50	50	0.5	10.0	8	3.0	4	3.0	61.3	5	3.0	21.9	2	10.5	45.3	2	13.0	45.3
B.L.3	100	50	0.5	20.0	16	3.0	4	8.0	42.4	8	3.0	33.4	2	20.5	48.3	2	21.5	50
B.L.4	200	50	0.5	20.0	32	3.0	6	3.0	36.9	16	3.0	28.7	2	48.0	48.7	2	50.5	49.7
B.L.5	400	50	0.5	20.0	64	3.0	8	3.0	35.4	32	3.0	27.2	3	65.5	32.0	2	100.5	50.0
B.L.6	800	50	0.5	20.0	125	3.0	12	8.0	37.8	29	3.0	52.5	3	133.0	34.6	2	200.5	49.6

the dynamic programming is scalable, while the stochastic mitigation is computationally expensive beyond 60 classes of observations as shown in Fig. 5(b,c).

Mitigation Algorithm Comparisons: Table 1 shows micro-benchmark results that compare the four mitigation algorithms with the two program series. Double scheme mitigation technique [10] does not provide guarantees on the performance overhead, and we can see that it is increased by more than 75 times for `mod_exp_6`. Double scheme method reduces the number of classes of observations. However, we observe that this mitigation has difficulty improving the min-guess entropy. Second, Bucketing algorithm [26] can guarantee the performance overhead, but it is not an effective method to improve the security of functional observations, see the examples `mod_exp_6` and `Branch_and_Loop_6`. Third, in the algorithms, SCHMIT guarantees the performance to be below a certain bound, while it results in the highest entropy values. In most cases, the stochastic optimization technique achieves the highest min-entropy value. Here, we show the results with min-guess entropy measure. Also, we have strong evidences to show that SCHMIT achieves higher Shannon and guessing entropies. For example, in `B.L_5`, the initial Shannon entropy has improved from 2.72 to 6.62, 4.1, 7.56, and 7.28 for the double scheme, the bucketing, the stochastic, and the deterministic algorithms, respectively.




Fig. 5. Computation time for synthesizing mitigation policy over `Branch_and_Loop` applications. Computation time for min-guess entropy (a) takes only few seconds. Computation time for the Shannon entropy (b) and guessing entropy (c) are expensive using Stochastic optimization. We set time-out to be 10 hours.

7 Case Study

Research Question. Does SCHMIT scale well and improve the security of applications (entropy measures) within the given performance bounds?

Methodology. We use the deterministic and stochastic algorithms for mitigating the leaks. We show our results for the min-guess entropy, but other entropy measures can be applied as well. Since the task is to mitigate existing leakages, we assume that the secret and public inputs are given.

Objects of Study. We consider four real-world applications:

In the inset table, we show the basic characteristics of these benchmarks.

Application	Num methods	Num secret	Num public	ϵ	Initial clusters	Initial. Min-guess
GabFeed	573	1,105	65	6.50	34	1.0
Jetty	63	800	635	0.1	20	4.5
Java Verbal Expressions	61	2,000	10	0.02	9	50.5
Password Checker	6	20	2,620	0.05	6	1.0

GabFeed is a chat server with 573 methods [4]. There is a side channel in the authentication part of the application where the application takes users' public keys and its own private key, and generating a common key [14]. The vulnerability leaks the number of set bits in the secret key. Initial functional observations are shown in Fig. 6a. There are 34 clusters and min-guess entropy is 1. We aim to maximize the min-guess entropy under the performance overhead of 50%.

Jetty. We mitigate the side channels in `util.security` package of Eclipse Jetty web server. The package has `Credential` class which had a timing side channel. This vulnerability was analyzed in [14] and fixed initially in [6]. Then, the developers noticed that the implementation in [6] can still leak information and fixed this issue with a new implementation in [5]. However, this new implementation is still leaking information [45]. We apply SCHMIT to mitigate this timing side channels. Initial functional observations is shown in Fig. 6d. There are 20 classes of observations and the initial min-guess entropy is 4.5. We aim to maximize the min-guess entropy under the performance overhead of 50%.

Java Verbal Expressions is a library with 61 methods that construct regular expressions [2]. There is a timing side channel in the library similar to password comparison vulnerability [3] if the library has secret inputs. In this case, starting from the initial character of a candidate expression, if the character matches with the regular expression, it slightly takes more time to respond the request than otherwise. This vulnerability can leak all the regular expressions. We consider regular expressions to have a maximum size of 9. There are 9 classes of observations and the initial min-guess entropy is 50.5. We aim to maximize the min-guess entropy under the performance overhead of 50%.

Password Checker. We consider the password matching example from loginBad program [9]. The password stored in the server is secret, and the user's guess is a public input. We consider 20 secret (lengths at most 6) and 2,620 public inputs. There are 6 different clusters, and the initial min-guess entropy is 1.

Findings for GabFeed. With the stochastic algorithm, SCHMIT calculates the mitigation policy that results in 4 clusters. This policy improves the min-guess entropy from 1 to 138.5 and adds an overhead of 42.8%. With deterministic algorithm, SCHMIT returns 3 clusters. The performance overhead is 49.7% and the min-guess entropy improves from 1 to 106. The user chooses the deterministic policy and enforces the mitigation. We apply CART decision tree learning and characterizes the classes of observations with GabFeed method calls as shown in

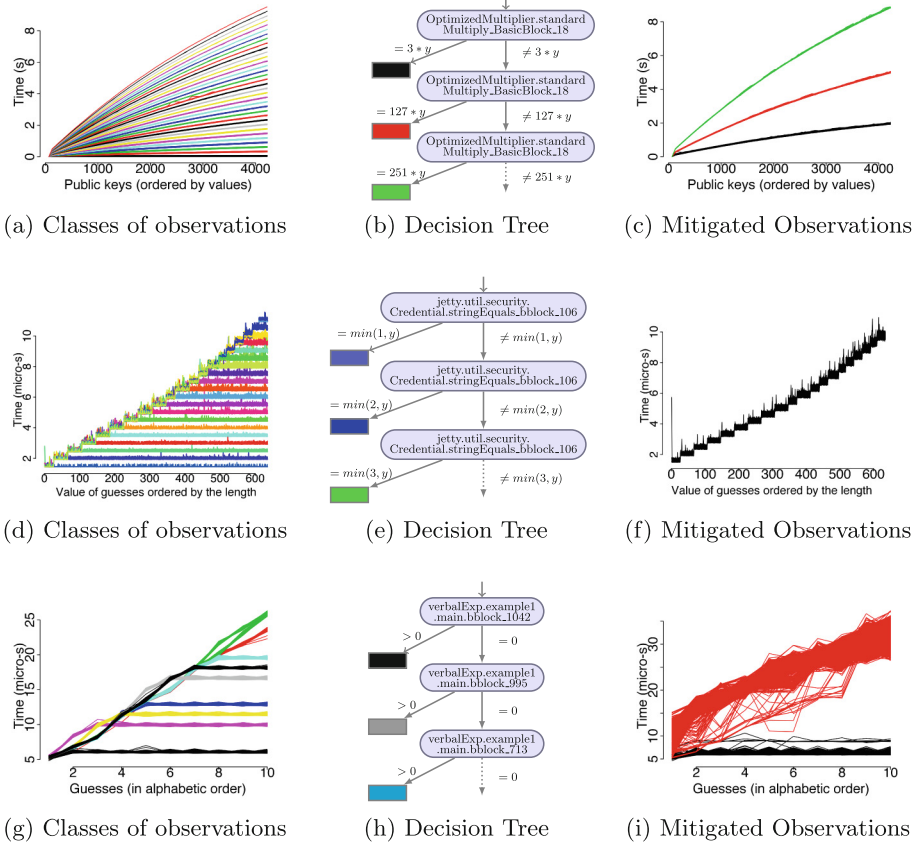


Fig. 6. Initial functional observations, decision tree, and the mitigated observations from left to right for Gabfeed, Jetty, and Verbal Expressions from top to bottom.

Fig. 6b. The monitoring system uses the decision tree model and automatically detects the current class of observation. Then, it adds extra delays based on the mitigation policy to enforce it. The results of the mitigation is shown in Fig. 6c. Answer for our research question. *Scalability:* It takes about 1 second to calculate the stochastic and the deterministic policies. *Security:* Stochastic and deterministic variants improve the min-guess entropy more than 100 times under the given performance overhead of 50%, respectively.

Findings for Jetty. The stochastic algorithm and the deterministic algorithm find the same policy that results in 1 cluster with 39.6% performance overhead. The min-guess entropy improves from 4.5 to 400.5. For the enforcement, SCHMIT first uses the initial clusterings and specifies their characteristics with program internals that result in the decision tree model shown in Fig. 6e. Since the response time is in the order of micro-seconds, we transform the source code using the decision tree model by adding extra counter variables. The results of

the mitigation is shown in Fig. 6f. *Scalability*: It takes less than 1 second to calculate the policies for both algorithms. *Security*: Stochastic and deterministic variants improve the min-guess entropy 89 times under the given performance overhead.

Findings for Java Verbal Expressions. For the stochastic algorithm, the policy results in 2 clusters, and the min-guess entropy has improved to 500.5. The performance overhead is 36%. For the dynamic programming, the policy results in 2 clusters. This adds 28% of performance overhead, while it improves the min-guess entropy from 50.5 to 450.5. The user chooses to use the deterministic policy for the mitigation. For the mitigation, we transform the source code using the decision tree model and add the extra delays based on the mitigation policy.

Findings for Password Matching. Both the deterministic and the stochastic algorithms result in finding a policy with 2 clusters where the min-guess entropy has improved from 1 to 5.5 with the performance overhead of 19.6%. For the mitigation, we transform the source code using the decision tree model and add extra delays based on the mitigation policy if necessary.

8 Related Work

Quantitative theory of information have been widely used to measure how much information is being leaked with side-channel observations [11, 20, 25, 41]. Mitigation techniques increase the remaining entropy of secret sets leaked through the side channels, while considering the performance [10, 23, 26, 40, 48, 49].

Köpf and Dürmuth [26] use a bucketing algorithm to partition programs' observations into intervals. With the unknown-message threat model, Köpf and Dürmuth [26] propose a dynamic programming algorithm to find the optimal number of possible observations under a performance penalty. The works [10, 48] introduce different black-box schemes to mitigate leaks. In particular, Askarov et al. [10] show the quantizing time techniques, which permit events to release at scheduled constant slots, have the worst case leakage if the slot is not filled with events. Instead, they introduce the double scheme method that has a schedule of predictions like the quantizing approach, but if the event source fails to deliver events at the predicted time, the failure results in generating a new schedule in which the interval between predictions is doubled. We compare our mitigation technique with both algorithms throughout this paper.

Elimination of timing side channels is a common technique to guarantee the confidentiality of software [7, 17, 27, 30, 31, 46]. The work [46] aims to eliminate side channels using static analysis enhanced with various techniques to keep the performance overheads low without guaranteeing the amounts of overhead. In contrast, we use dynamic analysis and allow a small amount of information to leak, but we guarantee an upper-bound on the performance overhead.

Machine learning techniques have been used for explaining timing differences between traces [42–44]. Tizpaz-Niari et al. [44] consider performance issues in softwares. They also cluster execution times of programs and then explain what

program properties distinguish the different functional clusters. We adopt their techniques for our security problem.

Acknowledgements. The authors would like to thank Mayur Naik for shepherding our paper and providing useful suggestions. This research was supported by DARPA under agreement FA8750-15-2-0096.

References

1. Branch and bound algorithm for mip problems. <http://www.gurobi.com/resources/getting-started/mip-basics>
2. Verbal expressions library. <https://github.com/VerbalExpressions/JavaVerbalExpressions>
3. Timing attack in google keyczar library (2009). <https://rdist.root.org/2009/05/28/timing-attack-in-google-keyczar-library/>
4. Gabfeed application (2016). https://github.com/Apogee-Research/STAC/tree/master/Engagement_Challenges/Engagement_2/gabfeed_1
5. Timing side-channel on the length of password in eclipse jetty May 2017. <https://github.com/eclipse/jetty.project/commit/2baa1abe4b1c380a30deacca1ed-367466a1a62ea>
6. Timing side-channel on the password in eclipse jetty May 2017. <https://github.com/eclipse/jetty.project/commit/f3751d70787fd8ab93932a51c60514c2eb37cb58>
7. Agat, J.: Transforming out timing leaks. In: Proceedings of the 27th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, pp. 40–53. ACM (2000)
8. Aloise, D., Deshpande, A., Hansen, P., Popat, P.: Np-hardness of euclidean sum-of-squares clustering. *Mach. Learn.* **75**(2), 245–248 (2009)
9. Antonopoulos, T., Gazzillo, P., Hicks, M., Koskinen, E., Terauchi, T., Wei, S.: Decomposition instead of self-composition for proving the absence of timing channels. In: PLDI, pp. 362–375. ACM (2017)
10. Askarov, A., Zhang, D., Myers, A.C.: Predictive black-box mitigation of timing channels. In: Proceedings of the 17th ACM Conference on Computer and Communications Security, pp. 297–307. ACM (2010)
11. Backes, M., Köpf, B., Rybalchenko, A.: Automatic discovery and quantification of information leaks. In: 2009 30th IEEE Symposium on Security and Privacy, pp. 141–153. IEEE (2009)
12. Bertsekas, D.P.: Nonlinear programming. Athena Scientific, 2016. Tech. rep., ISBN 978-1-886529-05-2
13. Brumley, D., Boneh, D.: Remote timing attacks are practical. *Comput. Netw.* **48**(5), 701–716 (2005)
14. Chen, J., Feng, Y., Dillig, I.: Precise detection of side-channel vulnerabilities using quantitative cartesian hoare logic. In: CCS, pp. 875–890 (2017)
15. Chiba, S.: Javassist - a reflection-based programming wizard for java. In: Proceedings of OOPSLA 1998 Workshop on Reflective Programming in C++ and Java, vol. 174 (1998)
16. Dhem, J.-F., Koeune, F., Leroux, P.-A., Mestré, P., Quisquater, J.-J., Willems, J.-L.: A practical implementation of the timing attack. In: Quisquater, J.-J., Schneier, B. (eds.) CARDIS 1998. LNCS, vol. 1820, pp. 167–182. Springer, Heidelberg (2000). <https://doi.org/10.1007/10721064.15>

17. Eldib, H., Wang, C.: Synthesis of masking countermeasures against side channel attacks. In: Biere, A., Bloem, R. (eds.) CAV 2014. LNCS, vol. 8559, pp. 114–130. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-08867-9_8
18. Fallgren, M.: On the complexity of maximizing the minimum shannon capacity in wireless networks by joint channel assignment and power allocation. In: 2010 IEEE 18th International Workshop on Quality of Service (IWQoS), pp. 1–7 (2010)
19. Gurobi, L.: Optimization: Gurobi optimizer reference manual (2018). <http://www.gurobi.com>
20. Heusser, J., Malacaria, P.: Quantifying information leaks in software. In: Proceedings of the 26th Annual Computer Security Applications Conference, pp. 261–269. ACM (2010)
21. Jacques, J., Preda, C.: Functional data clustering: a survey. *Adv. Data Anal. Classif.* **8**(3), 231–255 (2014)
22. Jones, E., Oliphant, T., Peterson, P., et al.: SciPy: open source scientific tools for Python (2001). <http://www.scipy.org/>
23. Kadloor, S., Kiyavash, N., Venkitasubramaniam, P.: Mitigating timing based information leakage in shared schedulers. In: 2012 Proceedings IEEE Infocom, pp. 1044–1052. IEEE (2012)
24. Kocher, P.C.: Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In: Kobitz, N. (ed.) CRYPTO 1996. LNCS, vol. 1109, pp. 104–113. Springer, Heidelberg (1996). https://doi.org/10.1007/3-540-68697-5_9
25. Köpf, B., Basin, D.: An information-theoretic model for adaptive side-channel attacks. In: Proceedings of the 14th ACM Conference on Computer and Communications Security, pp. 286–296. CCS 2007, ACM, New York (2007)
26. Köpf, B., Dürmuth, M.: A provably secure and efficient countermeasure against timing attacks. In: 22nd IEEE Computer Security Foundations Symposium, 2009, CSF 2009, pp. 324–335. IEEE (2009)
27. Köpf, B., Mantel, H.: Transformational typing and unification for automatically correcting insecure programs. *Int. J. Inf. Secur.* **6**(2–3), 107–131 (2007)
28. Korf, R.E.: A complete anytime algorithm for number partitioning. *AI* **106**, 181–203 (1998)
29. Lamson, B.W.: A note on the confinement problem. *Commun. ACM* **16**(10), 613–615 (1973)
30. Mantel, H., Starostin, A.: Transforming out timing leaks, more or less. In: Pernul, G., Ryan, P.Y.A., Weippl, E. (eds.) ESORICS 2015. LNCS, vol. 9326, pp. 447–467. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24174-6_23
31. Molnar, D., Piotrowski, M., Schultz, D., Wagner, D.: The program counter security model: automatic detection and removal of control-flow side channel attacks. In: Won, D.H., Kim, S. (eds.) ICISC 2005. LNCS, vol. 3935, pp. 156–168. Springer, Heidelberg (2006). https://doi.org/10.1007/11734727_14
32. Nemhauser, G.L., Wolsey, L.A.: Integer programming and combinatorial optimization. In: Nemhauser, G.L., Savelsbergh, M.W.P., Sigismondi, G.S. (1992). *Constraint Classification for Mixed Integer Programming Formulations*. COAL Bulletin, vol. 20, pp. 8–12. Wiley, Chichester (1988)
33. Nocedal, J., Wright, S.J.: *Numerical Optimization 2nd* (2006)
34. Nocedal, J., Wright, S.J.: *Sequential Quadratic Programming*. Springer, New York (2006)
35. Padlipsky, M., Snow, D., Karger, P.: Limitations of End-to-End Encryption in Secure Computer Networks. Tech. rep, MITRE CORP BEDFORD MA (1978)
36. Papadimitriou, C.H., Steiglitz, K.: *Combinatorial Optimization: Algorithms and Complexity*. Courier Corporation, North Chelmsford (1998)

37. Phan, Q.S., Bang, L., Pasareanu, C.S., Malacaria, P., Bultan, T.: Synthesis of adaptive side-channel attacks. In: 2017 IEEE 30th Computer Security Foundations Symposium (CSF), pp. 328–342. IEEE (2017)
38. Ramsay, J., Hooker, G., Graves, S.: *Functional Data Analysis with R and MATLAB*. Springer Science & Business Media, Berlin (2009)
39. Rivest, R.L., Shamir, A., Adleman, L.: A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM* **21**(2), 120–126 (1978)
40. Schinzel, S.: An efficient mitigation method for timing side channels on the web. In: 2nd International Workshop on Constructive Side-Channel Analysis and Secure Design (COSADE) (2011)
41. Smith, G.: On the foundations of quantitative information flow. In: de Alfaro, L. (ed.) *FoSSaCS 2009*. LNCS, vol. 5504, pp. 288–302. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-00596-1_21
42. Song, L., Lu, S.: Statistical debugging for real-world performance problems. In: *Proceedings of the 2014 ACM International Conference on Object Oriented Programming Systems Languages & Applications*, pp. 561–578. OOPSLA 2014 (2014). <https://doi.org/10.1145/2660193.2660234>
43. Tizpaz-Niari, S., Černý, P., Chang, B.-Y.E., Sankaranarayanan, S., Trivedi, A.: Discriminating traces with time. In: Legay, A., Margaria, T. (eds.) *TACAS 2017*. LNCS, vol. 10206, pp. 21–37. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-662-54580-5_2
44. Tizpaz-Niari, S., Černý, P., Chang, B.E., Trivedi, A.: Differential performance debugging with discriminant regression trees. In: 32nd AAAI Conference on Artificial Intelligence (AAAI), pp. 2468–2475 (2018)
45. Tizpaz-Niari, S., Černý, P., Trivedi, A.: Data-driven debugging for functional side channels. arXiv preprint. [arXiv:1808.10502](https://arxiv.org/abs/1808.10502) (2018)
46. Wu, M., Guo, S., Schaumont, P., Wang, C.: Eliminating timing side-channel leaks using program repair. In: *Proceedings of the 27th ACM SIGSOFT International Symposium on Software Testing and Analysis*, pp. 15–26. ACM (2018)
47. Yarom, Y., Genkin, D., Heninger, N.: Cachebleed: a timing attack on openssl constant-time rsa. *J. Cryptographic Eng.* **7**(2), 99–112 (2017)
48. Zhang, D., Askarov, A., Myers, A.C.: Predictive mitigation of timing channels in interactive systems. In: *Proceedings of the 18th ACM Conference on Computer and Communications Security*, pp. 563–574. ACM (2011)
49. Zhang, D., Askarov, A., Myers, A.C.: Language-based control and mitigation of timing channels. *PLDI* **47**(6), 99–110 (2012)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

