

Vulnerability Bounds and Leakage Resilience of Blinded Cryptography under Timing Attacks

Boris Köpf
MPI-SWS
bkoepf@mpi-sws.org

Geoffrey Smith
Florida International University
smithg@cis.fiu.edu

Abstract—We establish formal bounds for the number of min-entropy bits that can be extracted in a timing attack against a cryptosystem that is protected by blinding, the state-of-the-art countermeasure against timing attacks. Compared with existing bounds, our bounds are both tighter and of greater operational significance, in that they directly address the key’s one-guess vulnerability. Moreover, we show that any semantically secure public-key cryptosystem remains semantically secure in the presence of timing attacks, if the implementation is protected by blinding and bucketing. This result shows that, by considering (and justifying) more optimistic models of leakage than recent proposals for leakage-resilient cryptosystems, one can achieve provable resistance against side-channel attacks for standard cryptographic primitives.

I. INTRODUCTION

Side-channel attacks against cryptographic algorithms recover keys from information that is revealed by the algorithm’s physical execution; they pose a serious threat to the security of today’s cryptosystems. Timing attacks [15] are especially significant, because they can be carried out remotely (e.g. over networks [6]), which opens the door to a potentially large number of attackers.

Previous countermeasures against timing attacks were either not fully practical or not backed up by formal security guarantees. In [17], a countermeasure—input blinding and bucketing—has been proposed that at the same time is practical and offers formal security guarantees. However, the security guarantees in [17] are given in terms of a bound on the number of Shannon-bits of the secret key that are leaked. They are not fully satisfactory for two reasons.

First, bounds based on Shannon-entropy yield operational guarantees for the expected effort for recovering secrets by brute-force search. However, examples show that even if the *expected* effort is large, the probability of guessing the secret with a small number of guesses can be unacceptably high [25].

Second, the bounds from [17] only capture the information that is contained in the side-channel observa-

tions. They do not take into consideration adversaries that can leverage additional information, e.g. public keys and chosen plaintext-ciphertext pairs. It is unclear whether the combination of a small amount of secret key leakage together with such additional information could be used to break the system, as was shown for the symmetric-key case in [14]. Public-key cryptosystems that can deal with partial leakage of the secret key are emerging [2], [21], but they rely on sophisticated constructions and are of unclear practicality.

In this paper, we address both of these open problems: First, we show that the countermeasure from [17] delivers bounds on the *min-entropy* leakage of the secret key, which yields strong operational guarantees about the *vulnerability* of the key to being found in a small number of guesses. Moreover, our new bounds also improve on the bounds from [17] in terms of tightness.

Second, we show that any cryptosystem that is secure against adaptive chosen-ciphertext attacks (i.e. that satisfies IND-CCA2) is also secure against combined timing- and adaptive chosen-ciphertext attacks, given that the implementation is protected using the countermeasure from [17]. This implies that side-channel resistance can be achieved using a *standard cryptosystem* implemented with the countermeasure, rather than requiring special leakage-resilient cryptographic primitives. We achieve this result by using a more optimistic leakage model than the ones considered in recent work on leakage-resilient cryptography [2], [10], [20], [21], and showing that this model is justified by the countermeasure.

The technical development of the information-theoretic part of our work is based on modeling an n -observation blinded timing attack using a *channel matrix* C_n^{ta} , whose $[s, (o_1, \dots, o_n)]$ entry gives the conditional probability of observing a sequence of execution times (o_1, \dots, o_n) , given that the secret key is s . The number of rows and the number of columns of C_n^{ta} are both very large. But we show that C_n^{ta} can be *factored* into the product of two channel matrices $T_n U_n$ so that

the *inner dimension* (the number of columns of T_n and the number of rows of U_n) is small. We then show that the number of leaked min-entropy bits is bounded from above by the logarithm of this inner dimension, which leads to our bounds on the vulnerability of the secret key.

For the cryptographic part of our work, we first extend IND-CCA2 to a notion of security that assumes adversaries that have access to a channel C that captures individual execution times of an implementation with blinding and bucketing applied. We show that a polynomial number of outputs of this channel can be equivalently replaced by a single output of a channel T with polynomially bounded range. We further show that any efficient adversary A that breaks the security of the cryptosystem using T can be turned into an efficient adversary A' that violates vanilla IND-CCA2: A' first samples A 's advantage for distinguishing ciphertexts, for all possible outputs of T , and then runs A on the output with the largest estimated advantage, thereby violating IND-CCA2. Finally, we outline some connections between the information-theoretic and the cryptographic aspects of our work.

In summary, our contribution is twofold. First, we provide bounds for the information that is leaked by the timing behavior of a blinded implementation. These bounds improve over existing bounds by providing stronger operational security guarantees. Second, we show that standard cryptographic security guarantees remain valid in the presence of timing leaks that are mitigated by blinding and bucketing. This result, by shifting the burden of proof from the cryptographic primitive to its implementation, opens a new perspective for obtaining practical and leakage-resilient cryptosystems.

The rest of this paper is structured as follows. In Section II we recall the basics of min-entropy leakage, timing attacks, blinded cryptography, and bucketing. In Section III, we present new results on min-entropy channel capacity and use them to bound the min-entropy leakage of the n -observation timing attack on blinded cryptography. In Section IV, we present our cryptographic reduction, which shows that a semantically-secure cryptosystem remains semantically secure when adversaries are given access to timing information, provided that it is protected by blinding and bucketing. In Section V, we discuss our results and suggest some future directions. Finally, Sections VI and VII discuss related work and conclude.

II. PRELIMINARIES

A. Channels and min-entropy

In this section, we briefly recall the motivation and definitions of the *min-entropy* measure of information leakage proposed in [25].

A *channel* is a triple $(\mathcal{S}, \mathcal{O}, C)$, where \mathcal{S} is a finite set of secret input values, \mathcal{O} is a finite set of observable output values, and C is an $|\mathcal{S}| \times |\mathcal{O}|$ matrix, called the *channel matrix*, such that $C[s, o] = P[o|s]$, the conditional probability of obtaining output o given that the input is s . Note that each row of C sums to 1. We will informally refer to channel $(\mathcal{S}, \mathcal{O}, C)$ simply as C .

Any *a priori* distribution P_S on \mathcal{S} determines a random variable S . Moreover, P_S and C determine a joint probability P_\wedge on $\mathcal{S} \times \mathcal{O}$, where $P_\wedge[s, o] = P_S[s]P[o|s]$, and a marginal distribution P_O on \mathcal{O} , where $P_O[o] = \sum_{s \in \mathcal{S}} P_\wedge[s, o]$, giving a random variable O .

We *quantify* the amount of information that flows from S to O by considering an adversary \mathcal{A} who wishes to guess the value of S . It is natural to measure information leakage by comparing \mathcal{A} 's “uncertainty” about S before and after seeing the value of O , using the equation

$$\text{leakage} = \text{initial uncertainty} - \text{remaining uncertainty}.$$

We define “uncertainty” by considering the *vulnerability* of S to being guessed correctly *in one try* by \mathcal{A} ; if we make the worst-case assumption that \mathcal{A} knows P_S and C , then the *a priori* vulnerability is

$$V(S) = \max_{s \in \mathcal{S}} P_S[s]$$

and the *a posteriori* vulnerability is

$$\begin{aligned} V(S|O) &= \sum_{o \in \mathcal{O}} P_O[o] \max_{s \in \mathcal{S}} P[s|o] \\ &= \sum_{o \in \mathcal{O}} \max_{s \in \mathcal{S}} P_\wedge[s, o] \\ &= \sum_{o \in \mathcal{O}} \max_{s \in \mathcal{S}} (P_S[s]P[o|s]). \end{aligned}$$

We convert from vulnerability to uncertainty by taking the negative logarithm, giving Rényi's *min-entropy* [24]. Our definitions, then, are

- initial uncertainty: $H_\infty(S) = -\log V(S)$
- remaining uncertainty: $H_\infty(S|O) = -\log V(S|O)$

(Note however that there is no universally agreed-upon definition of conditional min-entropy $H_\infty(S|O)$ in the literature [7].) Finally, we define the *min-entropy*

leakage from S to O , denoted \mathcal{L}_{SO} , to be

$$\begin{aligned}\mathcal{L}_{SO} &= H_\infty(S) - H_\infty(S|O) \\ &= -\log V(S) - (-\log V(S|O)) \\ &= \log \frac{V(S|O)}{V(S)}.\end{aligned}$$

Thus min-entropy leakage is the logarithm of the factor by which knowledge of O increases the one-guess vulnerability of S .

We remark that if \mathcal{A} is allowed to make *multiple* guesses, then focusing on one-guess vulnerability might seem inadequate. But, in many cases, useful bounds can be obtained simply by observing that allowing i guesses at most increases the vulnerability by a factor of i . Thus if we write V_i for i -guess vulnerability, we have $V_i(S) \leq iV(S)$ and $V_i(S|O) \leq iV(S|O)$.

Most of the literature on quantitative information flow (for example, [8] and [18]) has defined “uncertainty” using *Shannon entropy* (rather than min-entropy), so that leakage from S to O is defined as *mutual information*: $\mathcal{L}_{SO} = H(S) - H(S|O) = I(S; O)$.

Measuring leakage using Shannon entropy gives different operational guarantees than those given by min-entropy. Massey’s *guessing entropy* bound [19] shows that conditional Shannon entropy $H(S|O)$ gives strong bounds on the *expected number of guesses* required to find S by brute-force search, given O . However, examples in [25] show that this expected number of guesses can be arbitrarily high even when S is highly vulnerable to being guessed in one try. In contrast, conditional min-entropy $H_\infty(S|O)$ satisfies

$$V(S|O) = 2^{-H_\infty(S|O)},$$

so it provides strong operational guarantees about the vulnerability of S , given O .

B. Timing attacks and input blinding

In this section, we briefly review the model of timing attacks against blinded cryptography from [17] and show how it can be cast as a channel according to our definition.

We consider implementations of cryptographic algorithms for which the execution time can be captured by a function $f: \mathcal{S} \times \mathcal{M} \rightarrow \mathcal{O}$ such that $f(s, m)$ gives the time required to decrypt message m using secret key s . Here \mathcal{S} is the finite set of secret keys, \mathcal{M} is the finite set of messages, and \mathcal{O} is the set of possible timings (e.g. the range of clock ticks between the worst-case and the best-case execution times). We assume that the adversary has full knowledge of the implementation; in particular we assume that f is known to the adversary.

This deterministic timing model captures special-purpose implementations in synchronous hardware, as well as software implementations on simple microprocessors without state, e.g. without caches and pipelines. For implementations on microprocessors with caches and pipelines (but without interferences such as interrupts and preemptions), the execution time is not necessarily a function of the implementation’s inputs. However, if one forces the processor into a fixed initial state before each execution, the execution time becomes deterministic and modeling the implementation as a function is valid.

In a timing attack against f , the adversary gathers n side-channel observations $o_1 = f(s, m_1), \dots, o_n = f(s, m_n)$ for deducing s or narrowing down its possible values. All known timing attacks against RSA decryption require that the adversary be able to obtain a large number of pairs (m_i, o_i) of ciphertexts and corresponding execution times. *Input blinding* tries to defend against timing attacks by *randomizing* each ciphertext m_i before decryption; it decorrelates the ciphertexts from the execution times and thus renders all known timing attacks ineffective.

Example 1. Consider an RSA decryption $x = m^s \bmod N$, where m is a ciphertext chosen by the adversary, x is the plaintext, $N = p \cdot q$ is the modulus, and s with $s \cdot e = 1 \bmod \varphi(N)$ is the secret key. In the *blinding* phase, one picks a random r that is relatively prime to N and computes $m \cdot r^e \bmod N$. The result of the decryption is $(m \cdot r^e)^s = x \cdot r \bmod N$, which yields x after *unblinding*, i.e., after multiplication with $r^{-1} \bmod N$.

Input blinding techniques are available for many common cryptographic algorithms, including ElGamal and Diffie-Hellman; the mathematical details of these techniques depend on the algebraic properties of the individual cryptosystems.

Looking at Example 1, observe that if the values of the masks r are independent and uniformly distributed, then the blinded inputs to the RSA decryption are also independent and uniformly distributed. We therefore model blinding in general by assuming that when the adversary observes n blinded timings, $f(s, m_1), \dots, f(s, m_n)$, then the messages m_i being decrypted are chosen randomly and independently, according to some probability distribution P_M . It follows that the information leaked in an n -observation blinded timing attack can be captured by a channel $(\mathcal{S}, \mathcal{O}^n, C_n^{ta})$, where $C_n^{ta}[s, (o_1, \dots, o_n)]$ gives the conditional probability of observing the timings (o_1, \dots, o_n) given that the secret key is s .

C. Bucketing

Finally, we describe *bucketing*. Intuitively, a larger set \mathcal{O} of possible timing observations will result in a larger amount of information leakage in a timing attack. This intuition has been formalized in [17], and has led to the proposal of combining blinding with another countermeasure, *bucketing*. Bucketing is the discretization of a system's execution times such that the results of the computation are returned at only a small number of fixed points in time. Bucketing reduces the size of \mathcal{O} , but at some cost to the system's performance, since results must be delayed to the next bucket time.

However, experiments with a 1024-bit RSA implementation [17] show that the performance overhead introduced by bucketing is very small compared to that introduced by a constant-time implementation: for a bucketing of size 5, the performance overhead with respect to an implementation without bucketing is less than 0.7%; for a bucketing of only 2 buckets, this overhead is still less than 3%. In contrast, a bucketing of size 1 (i.e. a constant-time implementation) leads to a performance overhead of more than 36%. (Note that this comparison ignores the overhead introduced by the blinding and unblinding steps. If the involved factors can be precomputed, however, this overhead will be small.) In the following we will typically assume that \mathcal{O} is a set of small size.

III. INFORMATION-THEORETIC BOUNDS

In this section, we will develop techniques that allow us to derive bounds on the maximum min-entropy leakage of the n -observation timing channel C_n^{ta} .

A. Min-entropy channel capacity

We begin by developing the theory of min-entropy channel capacity for an arbitrary channel $(\mathcal{S}, \mathcal{O}, C)$. The channel capacity is the maximum leakage of the channel over all possible *a priori* distributions on \mathcal{S} .

Definition 1. *The min-entropy channel capacity of $(\mathcal{S}, \mathcal{O}, C)$, denoted $\mathcal{ML}(C)$, is the supremum, over all *a priori* distributions on \mathcal{S} , of the min-entropy leakage from \mathcal{S} to \mathcal{O} :*

$$\mathcal{ML}(C) = \sup_{P_S \in \mathcal{D}(\mathcal{S})} \mathcal{L}_{SO}.$$

(Recall from Section II that the random variable S is implicitly associated with the *a priori* distribution P_S .)

It turns out to be quite simple to calculate the min-entropy channel capacity, as it is just the logarithm of the sum of the column maximums of C :

Theorem 1. *The min-entropy channel capacity of C is*

$$\mathcal{ML}(C) = \log \sum_{o \in \mathcal{O}} \max_{s \in \mathcal{S}} P[o|s],$$

and it is realized by a uniform distribution on \mathcal{S} .

Proof: Using the formulas in Section II we have

$$\begin{aligned} \mathcal{L}_{SO} &= \log \frac{V(S|O)}{V(S)} \\ &= \log \frac{\sum_{o \in \mathcal{O}} \max_{s \in \mathcal{S}} (P_S[s] P[o|s])}{\max_{s \in \mathcal{S}} P_S[s]} \\ &\leq \log \frac{\sum_{o \in \mathcal{O}} \max_{s \in \mathcal{S}} P[o|s] (\max_{s \in \mathcal{S}} P_S[s])}{\max_{s \in \mathcal{S}} P_S[s]} \\ &= \log \sum_{o \in \mathcal{O}} \max_{s \in \mathcal{S}} P[o|s] \end{aligned}$$

The upper bound is realized iff the following condition holds: for all $o \in \mathcal{O}$, there exists $s^* \in \mathcal{S}$ such that $P[o|s^*] = \max_{s \in \mathcal{S}} P[o|s]$ and $P_S[s^*] = \max_{s \in \mathcal{S}} P_S[s]$. This condition holds when S is uniformly distributed. It can hold for nonuniform distributions as well, provided that some proper subset of the rows of C includes at least one maximum from each column. ■

Note that Theorem 1 appears already in [5] (though without the observation about nonuniform distributions).

Theorem 1 shows that min-entropy leakage is maximized by a uniform *a priori* distribution P_S . We remark that min-entropy leakage is strongly dependent on the *a priori* distribution. As a trivial example, if the *a priori* vulnerability of S is 1, then there is nothing to leak, so the leakage is 0 regardless of C . More interestingly, consider the following $(n+1) \times n$ channel matrix C_0 :

$$\begin{pmatrix} \frac{1}{n} & \frac{1}{n} & \frac{1}{n} & \cdots & \frac{1}{n} \\ 1 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & & \vdots \\ 0 & 0 & 0 & \cdots & 1 \end{pmatrix}$$

Assume an *a priori* distribution in which s_0 has probability $\frac{1}{2}$ and s_1 through s_n each have probability $\frac{1}{2n}$. Then $V(S) = \frac{1}{2}$. But $V(S|O) = \frac{1}{2}$ as well, since on output o_i we know that the input was either s_0 or s_i , each with probability $\frac{1}{2}$. It follows that the min-entropy leakage under this *a priori* distribution is 0, since C_0 does not then improve the adversary's one-guess probability of guessing S . (Interestingly, if we consider *two-guess* vulnerability V_2 , then we see that $V_2(S) = \frac{n+1}{2n}$ but $V_2(S|O) = 1$.) If, instead, we consider min-entropy channel capacity, then we see that the n column maximums are all 1, giving a capacity of $\log n$, showing that C_0 can indeed leak a great deal.

If we focus on min-entropy channel capacity, then we can easily see that a capacity of 0 implies that there is no leakage at all, because the channel satisfies *probabilistic noninterference*, which means that the output is completely independent of the input:

Corollary 1. *The min-entropy channel capacity of C is 0 iff the rows of C are identical.*

Proof: By Theorem 1, the channel capacity of C is 0 iff the sum of its column maximums is 1. But each row of C sums to 1, so after considering the first row of C the sum of the column maximums is already 1. If any subsequent row differs at all from the earlier rows, it increases the sum of the column maximums above 1, which means that the capacity is positive. Conversely, if the rows of C are identical, then the sum of the column maximums is 1, which means that the capacity is 0. ■

Now we turn to the main result of this section, which concerns the *factorization* of channel matrices. Recall that matrices C_1 and C_2 can be multiplied provided that the number of columns of C_1 is the same as the number of rows of C_2 —we refer to this common number as the *inner dimension*. We have the following theorem, which shows that a factorization of a channel matrix implies a bound on its min-entropy channel capacity:

Theorem 2. *If channel matrix $C = C_1C_2$, where C_1 and C_2 are channel matrices with inner dimension p , then the min-entropy channel capacity of C is at most $\log p$.*

Proof: The key observation is that the $[s, o]$ entry of C is a *convex combination* of the elements of column o of C_2 ; the coefficients are given by row s of C_1 . Letting set \mathcal{T} index the inner dimension, we have

$$C[s, o] = \sum_{t \in \mathcal{T}} C_1[s, t]C_2[t, o].$$

Hence

$$C[s, o] \leq \max_{t \in \mathcal{T}} C_2[t, o],$$

which implies that the sum of the column maximums of C is at most the sum of the column maximums of C_2 . But the sum of the column maximums of C_2 is at most p , since C_2 has p rows, each of which sums to 1. Hence, by Theorem 1, the min-entropy channel capacity of C is at most $\log p$. ■

As a corollary, we can get easy upper bounds based on the number of rows or columns of C :

Corollary 2. *The min-entropy channel capacity of C is at most the logarithm of the number of rows of C , and at most the logarithm of the number of columns of C .*

Proof: If C has q rows and r columns, then we can trivially use both q and r as the inner dimension in a factorization of C using an identity matrix of appropriate dimension: $C = I_{q \times q}C$ and $C = CI_{r \times r}$. ■

Theorem 2 has an intuitive interpretation. We can view a channel matrix C as a pipe, where the number of rows of C is the size of the input end of the pipe, and the number of columns of C is the size of the output end of the pipe. In the case where $C = C_1C_2$, where the inner dimension p is small, we can view the pipe as being narrow in the middle, which prevents it from leaking very much.

B. Min-entropy channel capacity of the n -observation timing attack on blinded cryptography

We now return our attention to the n -observation timing attack on blinded cryptography considered in [17]. Recall from Section II that this attack can be seen as a channel from \mathcal{S} to \mathcal{O}^n that is described by a channel matrix C_n^{ta} , whose $[s, (o_1, \dots, o_n)]$ entry gives the conditional probability of observing the timings (o_1, \dots, o_n) , given that the secret key is s . What is the value of this entry? Under our assumptions, the conditional probability of observing the timing o_i at position i can be calculated by collecting all messages m such that $f(s, m) = o_i$, and summing their probabilities, which are given by distribution P_M . The product of those probabilities over all i gives the probability of observing the timings (o_1, \dots, o_n) . Thus we have

$$C_n^{ta}[s, (o_1, \dots, o_n)] = \prod_{i=1}^n \sum_m \{P_M[m] \mid f(s, m) = o_i\}. \quad (1)$$

Having determined the channel matrix C_n^{ta} , we now wish to determine upper bounds on its min-entropy channel capacity. Our bounds are based on a simple observation about the structure of C_n^{ta} . Looking at Equation 1, we see immediately that if two sequences of timings (o_1, \dots, o_n) and (o'_1, \dots, o'_n) are *permutations* of each other, then for any secret key s ,

$$C_n^{ta}[s, (o_1, \dots, o_n)] = C_n^{ta}[s, (o'_1, \dots, o'_n)]$$

since permuting a sequence of timings just reorders the product on the righthand side of Equation 1. This implies that columns (o_1, \dots, o_n) and (o'_1, \dots, o'_n) of C_n^{ta} are identical. We now argue that this structural property of C_n^{ta} allows us to factor it with a small inner dimension.

We first recall the information-theoretic method of *types*; see for example [9]. The *type* $t_{\bar{o}}$ of a sequence $\bar{o} = (o_1, \dots, o_n)$ is the sequence of *counts* of the

number of occurrences of each element of \mathcal{O} within \bar{o} . We let $\mathcal{T}_n = \{t_{\bar{o}} \mid \bar{o} \in \mathcal{O}^n\}$ be the set of types of sequences in \mathcal{O}^n . Also, we let $|t_{\bar{o}}|$ denote the number of sequences $\bar{o}' \in \mathcal{O}^n$ such that $t_{\bar{o}'} = t_{\bar{o}}$. Using types, we can rephrase our structural property of C_n^{ta} :

Lemma 1. *If \bar{o} and \bar{o}' have the same type, then columns \bar{o} and \bar{o}' of C_n^{ta} are identical.*

Proof: Two sequences in \mathcal{O}^n have the same type iff they are permutations of each other. ■

(We remark that Lemma 1 remains true even if we generalize from a deterministic timing function to a probabilistic one, namely $f : \mathcal{S} \times \mathcal{M} \rightarrow \mathcal{D}(\mathcal{O})$, where $f(s, m)$ gives a *distribution* on \mathcal{O} . This implies that our min-entropy leakage bounds below still hold in this more general timing model.)

Now we can present our factorization:

Lemma 2. *Channel matrix C_n^{ta} can be factored into the product of two channel matrices with inner dimension $|\mathcal{T}_n|$.*

Proof: Let matrix T_n , indexed by \mathcal{S} and \mathcal{T}_n , be defined by

$$T_n[s, t_{\bar{o}}] = |t_{\bar{o}}| C_n^{ta}[s, \bar{o}].$$

(Notice that $T_n[s, t_{\bar{o}}]$ is well defined, since if $t_{\bar{o}} = t_{\bar{o}'}$, then $C_n^{ta}[s, \bar{o}] = C_n^{ta}[s, \bar{o}']$.) Let matrix U_n , indexed by \mathcal{T}_n and \mathcal{O}^n , be defined by

$$U_n[t_{\bar{o}}, \bar{o}'] = \begin{cases} \frac{1}{|t_{\bar{o}}|}, & \text{if } t_{\bar{o}'} = t_{\bar{o}} \\ 0, & \text{otherwise.} \end{cases}$$

(Note that both T_n and U_n are channel matrices, since each of their rows sums to 1.) It is easy to see that $C_n^{ta} = T_n U_n$. We have

$$\begin{aligned} (T_n U_n)[s, \bar{o}'] &= \sum_{t_{\bar{o}} \in \mathcal{T}_n} T_n[s, t_{\bar{o}}] U_n[t_{\bar{o}}, \bar{o}'] \\ &= T_n[s, t_{\bar{o}'}] U_n[t_{\bar{o}'}, \bar{o}'] \\ &= |t_{\bar{o}'}| C_n^{ta}[s, \bar{o}'] \frac{1}{|t_{\bar{o}'}|} \\ &= C_n^{ta}[s, \bar{o}'] \end{aligned}$$

(The second equality above follows from the fact that $U_n[t_{\bar{o}}, \bar{o}'] = 0$ whenever $t_{\bar{o}} \neq t_{\bar{o}'}$.) Thus we have shown that channel $(\mathcal{S}, \mathcal{O}^n, C_n^{ta})$ is the *composition* of channels $(\mathcal{S}, \mathcal{T}_n, T_n)$ and $(\mathcal{T}_n, \mathcal{O}^n, U_n)$. ■

Together, these results now allow us to bound the amount of min-entropy leakage of the n -observation timing attack on blinded cryptography:

Theorem 3. *The min-entropy channel capacity of the channel matrix C_n^{ta} is at most $\log |\mathcal{T}_n|$.*

Proof: Follows immediately from Theorem 2 and Lemma 2. ■

We emphasize that Theorem 1 of [17] establishes the *same* upper bound, $\log |\mathcal{T}_n|$, as we do. But there the bound is on the *Shannon* channel capacity, rather than the *min-entropy* channel capacity, of C_n^{ta} . This makes a big difference in the operational security guarantees that are provided. Specifically, Massey’s lower bound on guessing entropy [19] shows that the upper bound on Shannon channel capacity in [17] implies a strong lower bound on the *expected* number of guesses required to guess the secret key S after the C_n^{ta} attack. But such a bound does not rule out the possibility that the adversary could, nevertheless, have a large probability of guessing S in a small number of tries. In contrast, our bounds on min-entropy channel capacity *do* allow us to rule out such a threat. We will illustrate this concretely, once we have analyzed $|\mathcal{T}_n|$.

As shown in Lemma 1 of [17], it is easy to see that

$$|\mathcal{T}_n| \leq (n+1)^{|\mathcal{O}|}, \quad (2)$$

since each of the $|\mathcal{O}|$ possible timing values must occur between 0 and n times in any sequence (o_1, \dots, o_n) . This gives an immediate corollary:

Corollary 3. *The min-entropy channel capacity of C_n^{ta} is at most $|\mathcal{O}| \log(n+1)$.*

We can tighten this upper bound somewhat by calculating $|\mathcal{T}_n|$ more carefully:

Lemma 3. $|\mathcal{T}_n| = \binom{n+|\mathcal{O}|-1}{n}$.

Proof: Counting the size of \mathcal{T}_n can be viewed as an “Occupancy Problem” as discussed in Section II.5 of Feller [11]. We want to know in how many ways we can place n indistinguishable “balls” (the timing observations) into $|\mathcal{O}|$ “bins” (the possibilities for each observation). In general, the number of ways of putting n indistinguishable balls into b bins turns out to be the binomial coefficient

$$\binom{n+b-1}{n}.$$

To see this, note that each such placement can be represented as a string of n stars (representing the balls) with $b-1$ bars inserted (representing the boundaries between the bins). For example, with $n=5$ and $b=4$, the string

$$**|*||**$$

represents the case when we put 2 balls in the first bin, 1 ball in the second bin, 0 balls in the third bin, and 2 balls

in the fourth bin. If the symbols were all distinguishable, then the number of such strings would be $(n + b - 1)!$. But since the n stars and $b - 1$ bars are indistinguishable, then the total number of strings is

$$\frac{(n + b - 1)!}{n!(b - 1)!},$$

which is equal to the above binomial coefficient. ■

Using Lemma 3, we get the following tighter bound:

Corollary 4. *The min-entropy channel capacity of C_n^{ta} is at most $\log \binom{n + |\mathcal{O}| - 1}{n}$.*

We now are ready to discuss the operational security guarantees provided by Theorem 3 and Lemma 3. The crucial fact is that min-entropy channel capacity is an upper bound on min-entropy leakage, which in turn is the logarithm of the factor by which *one-guess vulnerability* is increased by the attack. Let us consider a concrete scenario. Suppose that the secret key S has *a priori* vulnerability of 2^{-500} , and that an adversary does an n -observation blinded timing attack with $n = 2^{40}$ and $|\mathcal{O}| = 5$ (thanks to bucketing).

Using the bound in Corollary 4, we see that the min-entropy leakage is at most

$$\log \left(\frac{2^{40} + 4}{2^{40}} \right) \approx 155.4$$

This means that the expected *a posteriori* vulnerability of S is at most $2^{-500} \cdot 2^{155.4} = 2^{-344.6}$. Hence the adversary's expected probability of guessing the key in one try is at most $2^{-344.6}$, and of guessing the key in 2^{20} tries is at most $2^{-344.6} \cdot 2^{20} = 2^{-324.6}$.

In contrast, the cruder bound in Corollary 3 (used by [17]) here gives a leakage of $5 \log(2^{40} + 1) \approx 200$. So Lemma 3 considerably improves the leakage bound, from 200 to 155.4 bits. (Of course the more important difference from [17] is that our leakage bound is for min-entropy leakage, rather than Shannon entropy leakage.)

In conclusion, blinding and bucketing give strong guarantees about the vulnerability of the secret key S to timing attacks. In the next section, we broaden our perspective by considering an adversary that *combines* a timing attack with an adaptive chosen-ciphertext attack.

IV. CRYPTOGRAPHIC SECURITY IN THE PRESENCE OF SIDE-CHANNELS

In this section, we show that any public-key cryptosystem that is secure against adaptive chosen-ciphertext attacks (e.g. RSA-OAEP [12]) is also secure against adversaries that can additionally measure the

execution time of the decryption algorithm, provided that the cryptosystem's implementation is protected by blinding and bucketing.

To prove this result, we first extend the notion of security against adaptive chosen-ciphertext adversaries (IND-CCA2) [23] to a notion of security that incorporates adversaries that additionally have access to a channel that leaks information about the secret key. We assume that the range of the channel is bounded and that the outputs of the channel are independent. As discussed in Section II, such a channel captures the timing behavior of an implementation with blinding and bucketing applied. We then show that a polynomial number of independent outputs of this channel can be equivalently replaced by a single output of a channel with polynomially bounded range. Finally, we prove that this polynomially bounded channel can be entirely eliminated, showing that the side-channel information does not constitute a significant advantage for the adversary.

A. A notion of security for blinded cryptography

We begin by recalling and extending basic definitions about public-key cryptography.

Definition 2. *A public key encryption scheme Π is a triple $\Pi = (\mathcal{G}, \mathcal{E}, \mathcal{D})$, where \mathcal{G} is the key generation algorithm, \mathcal{E} the encryption algorithm and \mathcal{D} the decryption algorithm.*

Indistinguishability under adaptive chosen-ciphertext attacks (IND-CCA2) [23] is a strong notion of semantic security for public-key cryptosystems. It captures that no realistic adversary can distinguish between the encryptions of two self-chosen messages with non-negligible probability, even if it has access to a decryption oracle. We model such an adversary as two probabilistic polynomial time algorithms A_1 and A_2 , where A_1 outputs the challenge messages msg_0 and msg_1 , together with a string *state*, which is used for communication between A_1 and A_2 . A_2 is given a challenge ciphertext c , which is the encryption of either msg_0 or msg_1 , and it outputs a guess of which it was. We define two decryption oracles. The first oracle \mathcal{D}_{sk} decrypts every ciphertext it is queried on. The second oracle \mathcal{D}'_{sk} behaves as \mathcal{D}_{sk} , except that it returns a default value for every attempt to decrypt the challenge ciphertext c .

Definition 3 (IND-CCA2). *A public-key encryption scheme $\Pi = (\mathcal{G}, \mathcal{E}, \mathcal{D})$ is secure against adaptive chosen-ciphertext attacks if, for any probabilistic polynomial-time adversary $A = (A_1, A_2)$, it holds that*

$$\text{Adv}_A^{\text{CCA2}}(k) = |P[\text{Ex}_A^{\text{CCA2}}(k)] - \frac{1}{2}|$$

is negligible in k , where $\text{Ex}_A^{\text{CCA2}}(k)$ is the event that the following game outputs true.

- 1) $(sk, pk) \leftarrow \mathcal{G}(1^k)$
- 2) $(msg_0, msg_1, state) \leftarrow A_1^{\mathcal{D}_{sk}}(pk)$ such that $|msg_0| = |msg_1|$
- 3) $b \leftarrow \{0, 1\}$
- 4) $c \leftarrow \mathcal{E}_{pk}(msg_b)$
- 5) $b' \leftarrow A_2^{\mathcal{D}'_{sk}}(c, state)$
- 6) Output $b' = b$

We extend IND-CCA2 by additionally giving the adversary access to the output of a channel that reveals information about the secret key sk . In Section II, a channel $(\mathcal{S}, \mathcal{O}, C)$ represents a conditional probability distribution of observations for given secret keys, which is defined in terms of the channel matrix C . In this section, we generalize channels to accommodate for the derivation of asymptotic security guarantees. More precisely, we consider families of channels $(\mathcal{S}(k), \mathcal{O}(k), C(k))$, where k is the security parameter. Whenever convenient (e.g. to emphasize that a channel serves as an oracle of arity 0 for the adversary) we denote a channel as a family of random variables $C(k) = \{O_{sk}(k) \mid sk \in \mathcal{S}(k)\}$, where $O_{sk}(k)$ is distributed according to $P[O(k) \mid S(k) = sk]$. We assume that, for a given secret key sk , the adversary can access independent repetitions of $O_{sk}(k)$. For a public-key cryptosystem, the distribution of secret keys is given by the key generation algorithm \mathcal{G} . As is standard, we often leave the security parameter k implicit. The following definition formally captures our novel notion of security.

Definition 4 (IND-CCA2-C). *A public-key encryption scheme $\Pi = (\mathcal{G}, \mathcal{E}, \mathcal{D})$ is secure against adaptive chosen-ciphertext attacks in the presence of the channel $C = \{O_{sk}\}$ if, for any probabilistic polynomial-time adversary $A = (A_1, A_2)$, it holds that*

$$\text{Adv}_A^{\text{CCA2C}}(k) = |P[\text{Ex}_A^{\text{CCA2C}}(k)] - \frac{1}{2}|$$

is negligible in k , where $\text{Ex}_A^{\text{CCA2C}}(k)$ is the event that the following game outputs true.

- 1) $(sk, pk) \leftarrow \mathcal{G}(1^k)$.
- 2) $(msg_0, msg_1, state) \leftarrow A_1^{\mathcal{D}_{sk, O_{sk}}}(pk)$ such that $|msg_0| = |msg_1|$
- 3) $b \leftarrow \{0, 1\}$
- 4) $c \leftarrow \mathcal{E}_{pk}(msg_b)$.
- 5) $b' \leftarrow A_2^{\mathcal{D}'_{sk, O_{sk}}}(c, state)$
- 6) Output $b' = b$

Finally, we say that a public key cryptosystem satisfies IND-CCA2-C1 if it satisfies IND-CCA2-C when A is given only a single access to O_{sk} .

Note that IND-CCA2-C and IND-CCA2-C1 both specialize to IND-CCA2 if the output of the channel is independent of the input, i.e. if the channel satisfies probabilistic noninterference.

B. Channels of bounded range

In this section we introduce channels with bounded range. Channels with bounded range model the timing behavior of blinded implementations with bucketing applied. The following definition captures that the maximal number of possible observations b is fixed for all security parameters k , while the set of possible observations may vary for each k .

Definition 5. *A channel $(\mathcal{S}(k), \mathcal{O}(k), C(k))$ has bounded range b if, for every k , $|\mathcal{O}(k)| \leq b$.*

Example 2. The channel $C_1^{ta}(k)$ defined in Equation 1 in Section III captures one timing measurement of a blinded implementation. $C_1^{ta}(k)$ is a channel of bounded range b , if additionally a bucketing of size b is applied.

Note that, from a theoretical point of view, it would be desirable to parameterize b by k to allow the number of buckets to grow with the security parameter. However, as discussed in Section II-C, bucketings of very small size compared to the key length already lead to highly efficient implementations. Thus, modeling the timing behavior of blinded implementations with bucketing as channels with bounded range is reasonable, even for our asymptotic considerations.

For a channel $C = \{O_{sk}\}$ with bounded range b and a polynomial p , we define the channel $T_p^C = \{T_{sk}\}$ as follows. For a fixed k , the output of $T_{sk}(k)$ is obtained by taking the outcomes of $p(k)$ independent repetitions $\bar{o} = o_1, \dots, o_{p(k)}$ of $O_{sk}(k)$, and sorting them with respect to a fixed order on $\mathcal{O}(k)$. This ordered sequence is a canonical representative of the type of the sequence \bar{o} . The range of the variables $T_{sk}(k)$ is independent of sk , and we denote it by $\mathcal{T}_{p(k)}(k)$.

Observe that the channel $T_{p(k)}^C(k)$ corresponds to the factor $T_{p(k)}(k)$ of the factorization of the channel $C_{p(k)}^{ta}(k)$ defined in Lemma 2. According to Equation 2 of Section III, $|\mathcal{T}_{p(k)}(k)|$ is upper-bounded by the polynomial expression $(p(k) + 1)^b$. Moreover, note that the set $\mathcal{T}_{p(k)}(k)$ can be efficiently enumerated.

Lemma 4. *Let Π be a public-key cryptosystem and C a channel with bounded range. If, for all polynomials p , Π satisfies IND-CCA2-C1 in the presence of T_p^C , then Π also satisfies IND-CCA2-C in the presence of C .*

Proof: Let A be an adversary that violates IND-CCA2-C given $C = \{O_{sk}\}$, and whose running time

is bounded by the polynomial p . Then there is also an adversary A' that violates IND-CCA2-C1 given $T_p^C = \{T_{sk}\}$. The adversary A' performs a single call to T_{sk} , saves the result to a dedicated tape, and applies a random permutation to the tape content. A' then runs A , where each call of A to O_{sk} is replaced by reading a new symbol from the dedicated tape. Since the running time (and hence the number of calls to O_{sk}) of A is bounded by p , the tape contains sufficient observations. As all sequences of observations of the same type are equally probable, the content of the additional tape is indistinguishable from independent calls to O_{sk} . A' hence produces output that is indistinguishable from that of A and violates IND-CCA2-C. ■

Observe that the application of a random permutation to the output of T_p^C corresponds to the application of the factor U_p of the channel C_p^{ta} defined in Lemma 2 of Section III. We will shed further light on this connection in Section V.

C. Replacing channels by sampling

We now show that a single access to the channel T_p^C does not constitute a substantial advantage for the adversary.

Lemma 5. *Let Π be a public-key cryptosystem and C a channel with bounded range. If Π satisfies IND-CCA2 then, for all polynomials p , Π also satisfies IND-CCA2-C1 in the presence of T_p^C .*

For the proof of Lemma 5, we assume an adversary A that violates IND-CCA2-C1 in the presence of $T_p^C = \{T_{sk}\}$ for some polynomial p , and we show that this gives an adversary A' that violates IND-CCA2. How can A' do without the oracle T_{sk} ? The idea is that A' can enumerate *all* of the (polynomially-many) possible oracle values $t \in \mathcal{T}_p$ and can estimate A 's advantage given t by repeatedly running A for each such t . More precisely, for each t , A' runs steps 2) through 6) of the Ex_A^{CCA2C} game m times, answering A 's decryption queries using \mathcal{D}_{sk} and answering A 's oracle query using t . Note that A' itself randomly selects the challenge bit and performs the encryption of the corresponding message in steps 3) and 4) of the game. For each t , A' computes the relative frequency with which A wins, and chooses t^* to be the value that gives the largest estimated advantage. A' then completes its own Ex_A^{CCA2} game by simulating A with oracle result t^* .

Notice that A 's advantage with t^* might be substantially smaller than with a correct value from T_{sk} . But, by the Chernoff inequality, the probability that A' picks a “bad” t^* decreases exponentially quickly in m , the number of samples made. By choosing m appropriately,

then, we can show that A' runs in polynomial time and achieves non-negligible advantage, violating IND-CCA2.

Before we present the formal proof of Lemma 5, we give some auxiliary definitions and lemmas. To this end, assume a fixed keypair (sk, pk) and an IND-CCA2-C1 adversary $A = (A_1, A_2)$. For $t \in \mathcal{T}_p$, we define $G(t)$ as the event that the following game outputs true.

- 1) $(msg_0, msg_1, state) \leftarrow A_1^{(\mathcal{D}_{sk,t})}(pk)$ such that $|msg_0| = |msg_1|$
- 2) $b \leftarrow \{0, 1\}$
- 3) $c \leftarrow \mathcal{E}_{pk}(msg_b)$
- 4) $b' \leftarrow A_2^{(\mathcal{D}_{sk,t})}(c, state)$
- 5) Output $b' = b$

We write $G(T)$ in cases where G obtains its argument from a random variable T . Note that $G(T_{sk})$ corresponds to the event Ex_A^{CCA2C} with the fixed keypair (sk, pk) and the restriction that the adversary may access the channel T_{sk} only once.

We further define the random variable \hat{T}_m that captures the sampling process as previously described by

$$\hat{T}_m = \text{argmax}_{t \in \mathcal{T}_p} F_m[G(t)] ,$$

where $F_m[G(t)]$ denotes the relative frequency with which the event $G(t)$ occurs in m independent repetitions. For a fixed $\epsilon > 0$, we define the event S by

$$S \equiv \forall t \in \mathcal{T}_p : |F_m[G(t)] - P[G(t)]| \leq \epsilon .$$

The event S captures that, for all $t \in \mathcal{T}_p$, the sampling process was sufficiently precise. The following lemma states that the probability of imprecise sampling (i.e., that of $\neg S$) decreases exponentially with the number m of samples made.

Lemma 6. $P[\neg S] \leq e^{-2\epsilon^2 m} |\mathcal{T}_p|$

Proof: For i.i.d. random variables X_i ($1 \leq i \leq m$) with $X_i \in \{0, 1\}$ and $P[X_i = 1] = p$, the Chernoff inequality (see e.g. [27]) states that

$$P\left[\left|\frac{1}{m} \sum_{i=1}^m X_i - p\right| > \epsilon\right] \leq e^{-2\epsilon^2 m} . \quad (3)$$

Instantiating X_i by $G(t)$ in (3), we obtain

$$P[|F_m[G(t)] - P[G(t)]| > \epsilon] \leq e^{-2\epsilon^2 m}$$

for each $t \in \mathcal{T}_p$. Thus the probability of $\neg S$ is bounded from above by $e^{-2\epsilon^2 m} |\mathcal{T}_p|$. ■

We further show that if S and $\hat{T}_m = t^*$ hold, the adversary's advantage when given t^* will not be significantly lower than the adversary's advantage when given the output of the oracle T_{sk} .

Lemma 7. For all $t^* \in \mathcal{T}_p$ we have

$$S \wedge \hat{T}_m = t^* \Rightarrow P[G(t^*)] \geq P[G\langle T_{sk} \rangle] - 2\epsilon .$$

Proof: By definition, the event S entails the following events

$$P[G(t^*)] \geq F_m[G(t^*)] - \epsilon \quad (4)$$

$$F_m[G(t)] \geq P[G(t)] - \epsilon \quad (5)$$

for all $t^*, t \in \mathcal{T}_p$. Intuitively, (4) says that the actual probability of winning with t^* is not much worse than the observed relative frequency; (5) says that the observed relative frequency of winning with t is not much worse than the actual probability. We further have that $\hat{T}_m = t^*$ implies

$$F_m[G(t^*)] \geq F_m[G(t)] . \quad (6)$$

Combining (4), (5), and (6) we conclude that $S \wedge \hat{T}_m = t^*$ entails

$$\forall t \in \mathcal{T}_p : P[G(t^*)] \geq P[G(t)] - 2\epsilon . \quad (7)$$

The lemma then follows by taking the convex combination over $P[T_{sk} = t]$, for all $t \in \mathcal{T}_p$, on the right hand side of (7). ■

We are now ready to give the formal justification of Lemma 5.

Proof of Lemma 5: Recall that, for a fixed keypair (sk, pk) , the event $G\langle T_{sk} \rangle$ corresponds to Ex_A^{CCA2C} where A can make a single access to T_{sk} , and observe that $G\langle \hat{T}_m \rangle$ corresponds to $\text{Ex}_{A'}^{CCA2}$. We obtain the following connection between the probabilities of both events.

$$\begin{aligned} P[G\langle \hat{T}_m \rangle] &= \sum_{t^* \in \mathcal{T}} P[G(t^*)]P[\hat{T}_m = t^*] \\ &\geq \sum_{t^* \in \mathcal{T}} P[G(t^*)]P[\hat{T}_m = t^* \wedge S] \\ &\stackrel{(*)}{\geq} \sum_{t^* \in \mathcal{T}} (P[G\langle T_{sk} \rangle] - 2\epsilon)P[\hat{T}_m = t^* \wedge S] \\ &= (P[G\langle T_{sk} \rangle] - 2\epsilon)P[S] \\ &= (P[G\langle T_{sk} \rangle] - 2\epsilon)(1 - P[\neg S]) \\ &\geq P[G\langle T_{sk} \rangle] - P[\neg S] - 2\epsilon . \end{aligned}$$

For (*), observe that Lemma 7 implies

$$P[\hat{T}_m = t^* \wedge S] > 0 \Rightarrow P[G(t^*)] \geq P[G\langle T_{sk} \rangle] - 2\epsilon .$$

To see this, note that, for a fixed t^* , the statement on the right hand side of the implication of Lemma 7 is true whenever there is an elementary event for which

the left hand side is true. $P[\hat{T}_m = t^* \wedge S] > 0$ ensures that such an event exists.

Using Lemma 6 to bound $P[\neg S]$, we further obtain

$$P[G\langle \hat{T}_m \rangle] \geq P[G\langle T_{sk} \rangle] - e^{-2\epsilon^2 m} |\mathcal{T}_p| - 2\epsilon .$$

Up to this point we have assumed a fixed keypair (sk, pk) . We now take the expected value over all keypairs generated by $\mathcal{G}(1^k)$ and make the security parameter k explicit in our notation. We obtain

$$\begin{aligned} P[\text{Ex}_{A'}^{CCA2}(k)] \\ \geq P[\text{Ex}_A^{CCA2C}(k) - e^{-2\epsilon^2 m} |\mathcal{T}_{p(k)}(k)| - 2\epsilon] , \end{aligned}$$

where $|\mathcal{T}_{p(k)}(k)|$ is polynomially bounded. Now $|P[\text{Ex}_A^{CCA2C}(k)] - \frac{1}{2}|$ is non-negligible by assumption, i.e. there is a polynomial q such that

$$|P[\text{Ex}_A^{CCA2C}(k)] - \frac{1}{2}| \geq \frac{1}{q(k)}$$

for infinitely many k . Instantiating ϵ with $\frac{1}{4q(k)}$ and m with $16q(k)^2 k$, we obtain

$$\text{Adv}_{A'}^{CCA2}(k) \geq \frac{1}{2q(k)} - e^{-k} |\mathcal{T}_{p(k)}(k)| ,$$

which is non-negligible and concludes this proof. ■

D. Cryptographic security in the presence of channels

We are now ready to give the main theorem of this section.

Theorem 4. Any public-key cryptosystem that is IND-CCA2 secure is also IND-CCA2-C secure in the presence of a channel with bounded range.

The proof of Theorem 4 follows directly from Lemmas 4 and 5. Furthermore, observe that our proofs are valid without modification for weaker notions of security, such as IND-CCA and IND-CPA. We omit a formalization of the corresponding definitions and statements for better readability.


We conclude this section with the following corollary of Theorem 4 applied to the channel C_1^{ta} (see Example 2 in this section), which states that the security of a cryptosystem is preserved under timing attacks, provided that the cryptosystem's implementation is protected by blinding and bucketing.

Corollary 5. Any public-key cryptosystem that is secure against adaptive chosen-ciphertext attacks is also secure against combined timing- and adaptive chosen-ciphertext attacks, given that the implementation is protected by input blinding and bucketing.

Note that the statement of Corollary 5 relies on assumptions about the implementation of the cryptosystem, as discussed in Section II-B.

V. DISCUSSION AND FUTURE WORK

Our presentations of the information-theoretic bounds in Section III and the cryptographic guarantees in Section IV are largely independent. However, there are close connections between the two parts. As pointed out in Section IV, the factorization of the channel matrix $C^{ta} = TU$ presented in Lemma 2 corresponds to the two core parts of the reduction proof presented in Section IV. The outputs of the matrix T correspond to the outputs of the channel T^C from Section IV, which are representatives of the types of sequences of observations. Furthermore, the channel U implicitly appears as part of the adversary A' defined in the proof of Lemma 4: A' receives the output of T^C as an input and applies a random permutation, thereby obtaining the output of U . This illustrates that it could be possible to generalize the results in Section IV to arbitrary factorizations $C = C_1C_2$ of channel matrices. However, such a generalization would require us to introduce additional assumptions about the efficient enumerability of the inner dimension of the factorization and the efficient computability of the distribution associated with C_2 , together with a parameterization in the security parameter. We have avoided these technicalities by choosing a more direct way of presenting our results. Ignoring these additional assumptions for a moment, the connections between Sections III and IV can be oversimplified as follows:



As future work, we plan to investigate whether the arrow on the left-hand side can be inverted, i.e. whether a given bound on the amount of leaked information implies factorizations of the channel matrix with bounded inner dimension. Such a result could enable us to more directly derive cryptographic security guarantees from the outcomes of a quantitative information-flow analysis.

From a practical perspective, it would be interesting to derive concrete security guarantees for implementations that are protected by blinding and bucketing. Such guarantees would allow one to reason about the required keylength for achieving security in practice, which is beyond the scope of the asymptotic guarantees derived in this paper.

VI. RELATED WORK

Our results are based on the model of side-channels from [16] and the model of unknown-message attacks from [3]. The proposal for combining bucketing and

blinding, together with bounds on the number of Shannon bits that are leaked is due to [17]. We extend these bounds by giving corresponding bounds for the min-entropy/vulnerability leakage, which have been shown to deliver stronger operational security guarantees [25].

With respect to the theory of min-entropy channel capacity, Braun et al. [5] consider “multiplicative leakage” $V(S|O)/V(S)$ (which is essentially the same as our min-entropy leakage) and “additive leakage” $V(S|O) - V(S)$. They show that channel capacity under multiplicative leakage is realized by a uniform *a priori* distribution, but under additive leakage it is realized by some “corner point” distribution that gives probability $\frac{1}{k}$ to k of the elements of \mathcal{S} and probability 0 to the rest.

Several approaches in language-based security use type systems to detect [13], eliminate [1], [4], or mitigate [22] timing leaks. Applying these language-based approaches to cryptographic algorithms requires restrictive programming and precise knowledge about the time consumption of the individual instructions on the underlying machine. The combination of bucketing and blinding weakens these requirements at the price of partial information leakage. In this paper, we show that this leakage is so small that it does not compromise the security of standard cryptographic primitives.

Standaert et al. propose a framework for the evaluation of side-channel attacks [26], where they use two different metrics for the evaluation of systems: an information-theoretic metric and a security metric. The information-theoretic metric captures the amount of information (in Shannon entropy) that a non-adaptive chosen-message adversary can obtain from a side-channel and is not given a direct interpretation in terms of security. The security metric characterizes the security of a system in terms of the success rate of applying a given key recovery strategy to the measurement data. In this way, an analysis with the model of [26] yields assertions about the effectiveness of a particular kind of attack, but not necessarily universal bounds.

Naor and Segev [21] present a generic construction for leakage-resilient public-key cryptosystems from universal hash functions. They show that their construction leads to a variant of the Cramer-Shoup cryptosystem that is CCA2-secure in the presence of leakage of up to $1/6$ (adversarially chosen) bits of the secret key. In our model, only the number of side-channel observations can be chosen by the adversary, and the amount of leaked information corresponds to a number of bits that is logarithmic in the number of measurements. With this more optimistic model of leakage, which is justified by

the use of blinding as a countermeasure, we show that any CCA2-secure cryptosystem is leakage-resilient.

Alwen et al. [2] propose leakage-resilient public-key cryptosystems for a model of leakage that allows the adversary to learn a large (potentially gigabytes) but bounded amount of information about the secret key. Clearly, this model of leakage requires the use of secret keys that are even larger than the amount of leaked information. The challenge with this model of leakage is to devise cryptosystems that operate efficiently on such large keys.

VII. CONCLUSIONS

We established formal bounds for the number of min-entropy bits that can be extracted in a timing attack against a cryptosystem that is protected by blinding and bucketing. Compared with the bounds derived in [17], our bounds are both tighter and of greater operational significance, in that they directly address the probability of guessing the key in one attempt. Moreover, we have shown that any CCA2-secure public-key cryptosystem remains CCA2-secure in the presence of timing attacks, if the implementation is protected by blinding and bucketing. This suggests that, by combining security guarantees for cryptographic primitives with guarantees for their implementation, one can achieve cryptosystems that are leakage-resilient and practical at the same time.

Acknowledgments: We thank Dominique Unruh for many helpful discussions. Geoffrey Smith was partially supported by the National Science Foundation under grant CNS-0831114.

REFERENCES

- [1] Johan Agat. Transforming out Timing Leaks. In *Proc. 27th ACM Symposium on Principles of Programming Languages (POPL 2000)*, pages 40–53. ACM, 2000.
- [2] Joël Alwen, Yevgeniy Dodis, and Daniel Wichs. Leakage-Resilient Public-Key Cryptography in the Bounded-Retrieval Model. In *Proc. Advances in Cryptology (CRYPTO 2009)*, volume 5677, pages 36–54. Springer, Lecture Notes in Computer Science.
- [3] Michael Backes and Boris Köpf. Formally Bounding the Side-Channel Leakage in Unknown-Message Attacks. In *Proc. 13th European Symposium on Research in Computer Security (ESORICS 2008)*, volume 5283 of *Lecture Notes in Computer Science*, pages 517–532. Springer, 2008.
- [4] Gilles Barthe, Tamara Rezk, and Martijn Warnier. Preventing Timing Leaks Through Transactional Branching Instructions. In *Proc. 3rd Workshop on Quantitative Aspects of Programming Languages (QAPL 2006)*, Electronic Notes in Theoretical Computer Science (ENTCS), pages 33–55. Elsevier, 2005.
- [5] Christelle Braun, Konstantinos Chatzikokolakis, and Catuscia Palamidessi. Quantitative notions of leakage for one-try attacks. In *Proc. of the 25th Conference on Mathematical Foundations of Programming Semantics (MFPS 2009)*, volume 249 of *ENTCS*, pages 75–91. Elsevier, 2009.
- [6] David Brumley and Dan Boneh. Remote timing attacks are practical. *Computer Networks*, 48(5):701–716, 2005.
- [7] Christian Cachin. Entropy Measures and Unconditional Security in Cryptography. Dissertation No. 12187. ETH Zürich, 1997.
- [8] Konstantinos Chatzikokolakis, Catuscia Palamidessi, and Prakash Panangaden. Anonymity protocols as noisy channels. *Information and Computation*, 206:378–401, 2008.
- [9] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory, Second Edition*. John Wiley & Sons, Inc., 2006.
- [10] Stefan Dziembowski and Krzysztof Pietrzak. Leakage-Resilient Cryptography. In *Proc. 49th IEEE Symposium Symposium on Foundations of Computer Science (FOCS 2008)*, pages 293–302. IEEE Computer Society, 2008.
- [11] William Feller. *An Introduction to Probability Theory and Its Applications*, volume I. John Wiley & Sons, Inc., Third edition, 1968.
- [12] Eiichiro Fujisaki, Tatsuaki Okamoto, David Pointcheval, and Jacques Stern. RSA-OAEP Is Secure under the RSA Assumption. *Journal of Cryptology*, 17(2):81–104, 2004.
- [13] Daniel Hedin and David Sands. Timing Aware Information Flow Security for a JavaCard-like Bytecode. *Electronic Notes in Theoretical Computer Science (ENTCS)*, 141(1):163–182, 2005.
- [14] John Kelsey, Bruce Schneier, David Wagner, and Chris Hall. Side Channel Cryptanalysis of Product Ciphers. *Journal of Computer Security (JCS)*, 8(2–3):141–158, 2000.
- [15] Paul Kocher. Timing Attacks on Implementations of Diffie-Hellman, RSA, DSS, and Other Systems. In *Proc. Advances in Cryptology (CRYPTO 1996)*, volume 1109 of *Lecture Notes in Computer Science*, pages 104–113. Springer, 1996.
- [16] Boris Köpf and David Basin. An Information-Theoretic Model for Adaptive Side-Channel Attacks. In *Proc. 14th ACM Conference on Computer and Communication Security (CCS 2007)*, pages 286–296. ACM, 2007.
- [17] Boris Köpf and Markus Dürmuth. A Provably Secure And Efficient Countermeasure Against Timing Attacks. In *Proc. 22nd IEEE Computer Security Foundations Symposium (CSF 2009)*, pages 324–335. IEEE Computer Society, 2009.

- [18] Pasquale Malacaria. Assessing security threats of looping constructs. In *Proc. 34th ACM Symposium on Principles of Programming Languages (POPL 2007)*, pages 225–235, Nice, France, January 2007.
- [19] James L. Massey. Guessing and Entropy. In *Proc. 1994 IEEE Symposium on Information Theory (ISIT 1994)*, page 204. IEEE Computer Society, 1994.
- [20] Silvio Micali and Leonid Reyzin. Physically Observable Cryptography (Extended Abstract). In *Proc. First Theory of Cryptography Conference (TCC 2004)*, volume 2951 of *Lecture Notes in Computer Science*, pages 278–296. Springer, 2004.
- [21] Moni Naor and Gil Segev. Public-Key Cryptosystems Resilient to Key Leakage. In *Proc. Advances in Cryptology (CRYPTO 2009)*, volume 5677 of *Lecture Notes in Computer Science*, pages 18–35. Springer, 2009.
- [22] Alessandra Di Piero, Chris Hankin, and Herbert Wiklicky. Quantifying timing leaks and cost optimisation. In *Proc. 10th International Conference on Information and Communications Security (ICICS 2008)*, volume 5308 of *Lecture Notes in Computer Science*, pages 81–96. Springer, 2008.
- [23] Charles Rackoff and Daniel R. Simon. Non-Interactive Zero-Knowledge Proof of Knowledge and Chosen Ciphertext Attack. In *Proc. Advances in Cryptology (CRYPTO 1991)*, volume 576 of *Lecture Notes in Computer Science*, pages 433–444. Springer, 1992.
- [24] Alfréd Rényi. On measures of entropy and information. In *Proceedings of the 4th Berkeley Symposium on Mathematics, Statistics and Probability 1960*, pages 547–561, 1961.
- [25] Geoffrey Smith. On the Foundations of Quantitative Information Flow. In *Proc. 13th International Conference on Foundations of Software Science and Computation Structures (FoSSaCS 2009)*, volume 5504 of *Lecture Notes in Computer Science*, pages 288–302. Springer, 2009.
- [26] Francois-Xavier Standaert, Tal Malkin, and Moti Yung. A Unified Framework for the Analysis of Side-Channel Key Recovery Attacks. In *Proc. 28th International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT 2009)*, volume 5479 of *Lecture Notes in Computer Science*, pages 443–461. Springer, 2009.
- [27] Mathukumalli Vidjasagar. *Learning and Generalization*. Springer, second edition, 2002.