

A Parallel Repetition Theorem for Leakage Resilience

Zvika Brakerski*

Yael Tauman Kalai†

Abstract

A leakage resilient encryption scheme is one which stays secure even against an attacker that obtains a bounded amount of side information on the secret key (say λ bits of “leakage”). A fundamental question is whether parallel repetition amplifies leakage resilience. Namely, if we secret share our message, and encrypt the shares under two independent keys, will the resulting scheme be resilient to 2λ bits of leakage?

Surprisingly, Lewko and Waters (FOCS 2010) showed that this is false. They gave an example of a public-key encryption scheme that is (CPA) resilient to λ bits of leakage, and yet its 2-repetition is not resilient to even $(1 + \epsilon)\lambda$ bits of leakage. In their counter-example, the repeated schemes share secretly generated public parameters.

In this work, we show that under a reasonable strengthening of the definition of leakage resilience (one that captures known proof techniques for achieving non-trivial leakage resilience), parallel repetition *does* in fact amplify leakage (for CPA security). In particular, if fresh public parameters are used for each copy of the Lewko-Waters scheme, then their negative result does not hold, and leakage is amplified by parallel repetition.

More generally, we show that given t schemes that are resilient to $\lambda_1, \dots, \lambda_t$ bits of leakage, respectfully, their direct product is resilient to $\sum(\lambda_i - 1)$ bits. We present our amplification theorem in a general framework that applies other cryptographic primitives as well.

*Stanford University, zvika@stanford.edu.

†Microsoft Research, yael@microsoft.com.

1 Introduction

In recent years, motivated by a large variety of real-world physical attacks, there has been a major effort by the cryptographic community to construct schemes that are resilient to leakage from the secret keys. This successful line of work gave rise to many constructions of leakage-resilient cryptographic primitives, including stream ciphers [DP08, Pie09], signature schemes [KV09, FKPR10], symmetric and public-key encryption schemes [AGV09, NS09, DKL09, DGK⁺10], as well as more complicated primitives.

A natural question to ask is: Does parallel repetition amplify leakage? More concretely, suppose we are given a public-key encryption scheme \mathcal{E} that remains secure even if λ bits about the secret key are leaked. Is it possible to amplify the leakage-resilience to $t\lambda$ by taking t copies of \mathcal{E} , and encrypting a message m by secret sharing it, and encrypting the i^{th} share using \mathcal{E}_i (we denote the resulting scheme by \mathcal{E}^t)? Using an appropriate definition of parallel repetition, a similar question can be asked for signatures.

Alwen, Dodis, and Wichs [ADW09] and Alwen, Dodis, Naor, Segev, Walfish and Wichs [ADN⁺10] were able to amplify leakage resilience for particular schemes, using the specific properties of these schemes. They raised the fundamental question of whether leakage resilience can *always* be amplified by parallel repetition. They predicted that such a result will be hard or even impossible to prove under the known definitions.

Recently, Lewko and Waters [LW10] gave a striking negative result, giving an example of a public-key encryption scheme that is resilient to λ bits of leakage but whose 2 repetition is not resilient to even $(1 + \epsilon)\lambda$ bits. This was followed by a work of Jain and Pietrzak [JP11] who presented a signature scheme where increasing the number of repetitions does not improve the leakage resilience at all. We elaborate on these negative results (and on how they go hand-in-hand with our positive results) in Section 1.2.

1.1 Our Results

We give positive results, by proving direct product theorems for leakage resilience. In particular, we show that parallel repetition does amplify the leakage resilience (almost) as expected.

The leakage model we consider is based on the “noisy leakage” model of Naor and Segev [NS09].¹ In this model, “legal” leakage functions are poly-size circuits that reduce the min-entropy of the secret key by at most λ . A scheme is said to be λ -leakage resilient if every PPT adversary, that asks for a “legal” leakage function, breaks the scheme with only negligible probability.

In this work, we consider a slightly relaxed leakage model. Instead of requiring the leakage function to *always* reduce the min-entropy of sk by at most λ , we require that it should be hard to break the scheme on those leakage values that do reduce the min-entropy by at most λ . In other words, we consider a *point-wise* definition: We say that a scheme is *point-wise λ -leakage resilient* if for any PPT adversary, that asks for a poly-size leakage function L , the probability that both the leakage value $y \leftarrow L(pk, sk)$ reduces the min-entropy of sk by at most λ , and that $\mathcal{A}(pk, y)$ breaks the scheme, is negligible.

We believe that this leakage model is of independent interest, as it captures our “intent” better: As long as the secret key is left with enough min-entropy, the scheme is secure. Moreover, we

¹While “entropic leakage” may be a more suitable name for this model, we stick with the terminology of [NS09] for historic reasons.

note that all known constructions that are λ -leakage resilient are also point-wise λ -leakage resilient (including [NS09, KV09, DGK⁺10, BG10]). We elaborate on this in Section 4.

At first it may seem that point-wise leakage is equivalent to noisy leakage. However, the difficulty is that it may be hard to determine whether a leakage value $y \leftarrow L(pk, sk)$ indeed reduces the min-entropy of sk by at most λ . If this was efficiently determined, then indeed we would have a reduction between the two models.

For technical reasons (see Section 1.3), we need to further relax our leakage model for our results to go through. We consider two (incomparable) relaxations.

First Relaxation: Almost λ -Leakage. In the first relaxation, instead of requiring that sk has high min-entropy (given pk, y), we require that it is statistically close to a random variable with high min-entropy. A scheme that is secure in this model is said to be *point-wise almost λ -leakage resilient*. We can prove a direct product theorem of any *constant* number of repetitions under this definition.

Theorem 1. *Let $c \in \mathbb{N}$ be a constant, and for every $i \in [c]$, let \mathcal{E}_i be a point-wise almost λ_i -leakage-resilient public-key encryption scheme. Then, $\mathcal{E}_1 \times \dots \times \mathcal{E}_c$ is point-wise almost λ -leakage-resilient, where $\lambda = \sum_{i=1}^c (\lambda_i - 1)$.*

We refer the reader to Section 1.3 and Section 5 for an overview of the proof, and to Appendix C for the formal proof.

Second Relaxation: Leakage with Small Advice. In the second relaxation, we give the adversary an additional logarithmic (in the security parameter) number of bits of (possibly hard to compute) advice (quite surprisingly, we were unable to reduce this model to the point-wise λ -leakage model). A scheme that is secure in this model is said to be *point-wise λ -leakage resilient with logarithmic advice*. We can prove a direct product theorem of any *polynomial* number of repetitions under this definition.

We note that it is not clear what it means to have t different leakage resilient schemes when t is super constant, since there is a different number of schemes for each value of the security parameter. While one can come up with a proper definition (involving a generation algorithm that, for every value of the security parameter, gets i and implements \mathcal{E}_i), for the sake of clarity, we choose to state the theorem below only for parallel repetition of the same scheme.

Theorem 2. *Let $t = t(k)$ be a polynomial in the security parameter. Let \mathcal{E} be a public-key encryption scheme that is point-wise λ -leakage resilient with logarithmic advice. Then \mathcal{E}^t is point-wise $t(\lambda - 1)$ -leakage resilient with logarithmic advice.*

We refer the reader to Section 1.3 for an overview of the proof, and to Section 6 and Appendix D for the formal proof.

The Relation Between our Models. Interestingly, we are not able to show that our relaxations are equivalent to one another, nor to show that they are implied by (plain) point-wise leakage resilience. This is surprising since in the bounded leakage model,² a negligible change in the secret-key distribution, or adding a logarithmic number of hard to compute bits, does not change the

²Where the leakage function's output is required to be bounded by λ bits, as opposed to our requirement that the secret key has high residual entropy.

model. In a nutshell, the reason that this does not carry to our models, is that having high min-entropy is not an efficiently verifiable condition, and that statistical indistinguishability does not preserve min-entropy.

We are able to show, however, that point-wise λ -leakage resilience implies λ -bounded leakage resilience (for the same value of λ), and thus in particular, our relaxed models also imply bounded leakage resilience. We note that proving the above is somewhat nontrivial since we do not want to suffer a degradation in λ . We refer the reader to Section 3.3 for a formal presentation, and to Appendix B for the proof.

Our Models and Current Proof Techniques. We show that for essentially all known schemes that are resilient to non-trivial leakage (i.e. super-logarithmic in the hardness of the underlying problem), amplification of leakage resilience via parallel repetition works. Specifically, this includes the Lewko-Waters counterexample, if the public parameters are chosen independently for each copy of the scheme. In order to do this, we identify a proof template that is used in all leakage resilience proofs, and show that this template is strong enough to prove point-wise leakage resilience, as well as our relaxed notions. See Section 4 for the full details.

The Lewko-Waters counterexample uses its public parameters in a very particular way that makes the argument not go through (see below).

1.2 Prior Work

As we claimed above, all known leakage resilient schemes are proved using the same proof template, and remain secure under our leakage models. This implies that parallel repetition should amplify security for all known schemes, which does not seem to coincide with the negative results of [LW10, JP11]. We explain this alleged discrepancy below.

The Lewko-Waters Counterexample. Lewko and Waters [LW10] construct a public key encryption scheme that is resilient to non-trivial length-bounded leakage, and prove that parallel repetition does not amplify its leakage resilience. However, the copies of their encryption scheme share public parameters: They are all using the same bilinear group. Their scheme, like all other schemes we are aware of, is (computationally indistinguishable from) point-wise leakage resilient and our theorems imply that parallel repetition does amplify its resilience to leakage. This is true so long as the public parameters are generated anew for each copy of the scheme: In our proof, we need to be able to sample key pairs for the scheme in question. Lewko and Waters use the public parameters in an extremely pathological (and clever!) way: The public parameters enable to generate keys for their actual scheme, but not for the computationally indistinguishable scheme where leakage resilience is actually proven. However, if we consider the generation of public parameters as a part of the key generation process, then new key pairs can always be generated, and parallel repetition works.

The Jain-Pietrzak Counterexample. Jain and Pietrzak [JP11] give a negative result for signature schemes. They take any secure signature scheme and change it so that if the message to be signed belongs to a set H , then the signature algorithm simply outputs the entire secret key. The set H is computationally hard to hit (given only the public key), and thus the scheme remains secure. It follows that the scheme remains secure also given leakage of length $O(\log k)$, where k is

the security parameter (more generally, if the underlying problem is 2^λ hard, then the scheme is resilient to $\sim \lambda$ bits of leakage).

They prove that parallel repetition fails, by proving that if the scheme is repeated t times, for some large enough t , then the leakage can in fact give enough information to find a message m that belongs to all the sets H_i , and thus break security completely. They start with a result that relies on common public parameters: a common (seeded) hash function. Then, they suggest to remove this public parameter by replacing the seeded hash function with an explicit hash function, such as SHA256. However, this explicit hash function is also, in some sense, a joint *non-uniform* public parameter.

This counterexample heavily relies on the “help” of the signing oracle when breaking the repeated scheme. The paper also presents a construction of a CCA encryption scheme, where they use the decryption oracle to break the parallel repetition system.

In general, signature schemes are not covered by our amplification theorems. Our theorems (and proofs) only cover public key primitives where the challenger in the security game does not need to know the secret key (beyond providing the adversary with the leakage value). Our results do extend to schemes such as signature schemes or CCA encryption schemes, if they have the property that the challenger (i.e., the signing oracle or the decryption oracle) can be efficiently simulated given only the public key (or given very little information about the secret key), in a way that is computationally indistinguishable even given the leakage. For example, the signature scheme of Katz and Vaikuntanathan [KV09] has this property, and thus its leakage resilience is amplified by parallel repetition. Whether our techniques can be applied to other leakage resilient signature schemes (e.g. [BSW11, MTVY11, GJS11]) is an interesting question that we leave for further research.

1.3 Overview of Our Techniques

In what follows we give a high-level overview of our proofs. For the sake of simplicity, we focus on the case of two-fold parallel repetition. Let \mathcal{E} be any λ -leakage resilient encryption scheme. Our goal is to prove that the scheme \mathcal{E}^2 is 2λ -leakage resilient. For technical reasons, in our actual proof, we manage to show that \mathcal{E}^2 is $(2\lambda - 1)$ -leakage resilient (in both our leakage models).

Our proof is by reduction: Suppose there exists an adversary \mathcal{B} for the parallel repetition scheme \mathcal{E}^2 that leaks $L(pk_1, pk_2, sk_1, sk_2)$, where L reduces the min-entropy of (sk_1, sk_2) by at most $2\lambda - 1$. We construct an adversary \mathcal{A} , that uses \mathcal{B} to break security of \mathcal{E} , and uses a leakage function L' that reduces the min-entropy of the secret key by at most λ .

Intuitively (and, as we will show, falsely), it does not seem too hard to show such a reduction. It only makes sense that when the pair (sk_1, sk_2) loses 2λ bits of entropy, then at least one of the secret keys sk_1, sk_2 “loses” at most λ bits (otherwise the total loss should be more than 2λ). Therefore the adversary \mathcal{A} can sample a key pair by itself and “plant” it either as (pk_1, sk_1) or as (pk_2, sk_2) (at random). Namely, \mathcal{A} will sample a random $i \in \{1, 2\}$, and uniformly sample (pk_i, sk_i) , the key pair of the scheme we actually wish to attack will play the role of (pk_{3-i}, sk_{3-i}) . Upon receiving a leakage function $L(\cdot)$ from \mathcal{B} , the adversary \mathcal{A} will plug the known (pk_i, sk_i) into the function and thus obtain L' to be sent to the challenger. Upon receiving a response from the challenger, it is forwarded back to \mathcal{B} , which can then break security with noticeable probability. Notice that \mathcal{B} 's view in the game is identical to its view in the repeated game against \mathcal{E}^2 , and thus it still breaks the security with the same probability. The only worry is whether the function L' only reduces the key entropy by the allowed amount, which is unfortunately not the case. Assume

that L leaks some 2λ bits on the bit-wise XOR $sk_1 \oplus sk_2$. Then when plugging in a known sk_i , the resulting L' still leaks 2λ bits on sk_{3-i} .

To solve this problem, we must prevent \mathcal{A} from knowing sk_i . This is achieved by having the key pair (pk_i, sk_i) sampled by the leakage function L' , rather than by \mathcal{A} . Namely, $L'(pk, sk)$ is now defined as follows: First, sample (pk_i, sk_i) and set $(pk_{3-i}, sk_{3-i}) = (pk, sk)$. Then run $y \leftarrow L(pk_1, pk_2, sk_1, sk_2)$ to obtain the leakage value. Lastly, output (y, pk_1, pk_2) . Given the output of L' , the adversary \mathcal{A} can forward the value y to \mathcal{B} , that uses it to break the scheme, all without ever being exposed to the value of sk_i .

This seems to give \mathcal{A} the least amount of information possible, so we should hope that now we can prove that the entropy of sk is reduced by at most λ . However, again, this is not true. Suppose that with probability $1/2$, the leakage function L outputs 2λ bits about sk_1 and with probability $1/2$ it outputs 2λ bits about sk_2 . In this case, L indeed reduces the min-entropy of (sk_1, sk_2) by 2λ , and yet for every $i \in \{1, 2\}$ the leakage function $L'(pk, sk)$ reduces the min-entropy of sk by essentially 2λ as well, and thus is not a valid leakage function for the one shot game.

This abnormality results, to some extent, from using min-entropy (as opposed to Shannon entropy) as our entropy measure: If $L'(pk, sk)$ outputs both $y = L(pk_1, pk_2, sk_1, sk_2)$ and sk_{3-i} , then it would indeed leak at most λ bits on sk (with probability $1/2$). The fact that we have *less* information, namely sk_i is not known, might actually *decrease* the min-entropy of the key.

We arrive at a conflict: On one hand, knowing sk_i is a problem, but on the other, not knowing it seems to also be a problem. We show that revealing sk_i only in some cases, enables to prove parallel repetition. We use a simple lemma (Lemma A.1), which essentially shows how to “split-up” the joint min-entropy of two random variables. More precisely, it says that there is a subset S of all possible secret keys sk_1 , such that for every $sk_1 \in S$, the the random variable $sk_2|sk_1$ has high min-entropy. Moreover, given the additional bit of information that $sk_1 \notin S$, causes sk_1 to have high min-entropy (which decreases as the size of S shrinks).

We proceed by a specific analysis for each of our two relaxed models. For explanatory reasons, we first discuss leakage with advice (our second relaxation) and then go back to the almost leakage resilience model (our first relaxation).

Point-Wise λ -Leakage with Advice. In this model, the adversary \mathcal{A} will leak $L'(pk, sk)$, which is a randomized leakage function, defined by choosing a random $\tau \in \{1, 2\}$, setting $(sk_\tau, pk_\tau) = (sk, pk)$, choosing a new fresh key pair (sk_i, pk_i) , where $i = 3 - \tau$, and outputting $L(pk_1, pk_2, sk_1, sk_2)$. In addition, it will use one bit of advice which is whether $sk_i \in S$. If so, the leakage function $L'(pk, sk)$ outputs sk_i in addition to $L(pk_1, pk_2, sk_1, sk_2)$, and otherwise it outputs only $L(pk_1, pk_2, sk_1, sk_2)$. Now we can prove that indeed, for many pairs (pk, sk) , the leakage $L'(pk, sk)$ leaks at most λ bits about sk (and \mathcal{B} breaks \mathcal{E}^2 on the corresponding keys).

Note that the leakage function L' sometimes leaks more than it should. Namely, in some cases the value $y \leftarrow L'(pk, sk)$ reduces the min-entropy of sk by λ ; but in other cases it reduces the min-entropy of sk by more than λ ,³ and in these cases it is an invalid leakage function. For this reason, we need to consider the *point-wise* λ -leakage definition. In addition, note that L' used only one bit of additional advice. Therefore when going from \mathcal{E} to \mathcal{E}^t the reduction uses $\log t$ bits of advice.

³This happens when the set S is very small, yet $sk_\tau \in S$.

Point-Wise Almost λ -Leakage. In this model, the idea of the reduction is the following: The adversary \mathcal{A} will leak $L'(sk, pk)$, which is a randomized leakage function, defined by choosing a random $\tau \in \{1, 2\}$, setting $(sk_\tau, pk_\tau) = (sk, pk)$, choosing a new fresh key pair (sk_i, pk_i) , where $i = 3 - \tau$, and outputting $L(pk_1, pk_2, sk_1, sk_2)$, and in addition with probability $1/2$ outputting sk_i .

As in the model with advice, the leakage function L' might leak more than λ bits about sk , and thus we use the point-wise definition. In the analysis, we distinguish between the case that the set S is noticeable and the case that it is negligible. In the former, with non-negligible probability the leakage function L' will sample $sk_i \in S$ and will output it. In this case the leakage function is legal. If the set S is negligible, we claim the distribution of the secret key sk_τ is statistically close to the distribution of sk_τ conditioned on the event that $sk_i \notin S$ (as this event happens only with negligible probability). Therefore, if L' did not output the secret key sk_i , the secret key sk_τ is statistically close to a distribution with high enough min-entropy. Due to this analysis, we need to relax our leakage model *almost* λ -leakage resilient.

Since the analysis in this model is asymptotic, we are not able to extend it beyond a constant number of repetitions. See discussion in Section 5.

1.4 Paper Organization

We define our generalized notion of public-key primitives in Section 2, where we also define parallel repetition and leakage attacks on such primitives. Our model of point-wise leakage resilience is presented in Section 3. In Section 4 we explain why all known leakage resilient schemes are also point-wise leakage-resilient.

Our parallel repetition theorems for a constant number of repetitions and for a polynomial number of repetitions are presented in Sections 5 and 6, respectively. The proofs are deferred to Appendices C and D, respectively. In Section 7 we discuss what our theorems imply for schemes that are only computationally indistinguishable from being secure in our model.

Appendix A contains the min-entropy splitting lemma that is used for all our proofs.

2 Public-Key Primitives, Parallel-Repetition, Leakage Attacks

In this section we give a definition of a *public key primitive* which generalizes one-way relations and public-key encryption under chosen plaintext attack (CPA). We then show how to define parallel repetition with respect to public-key primitives in a way that, again, generalizes the intuitive notions of parallel repetition for either one-way relations or public-key encryption.

2.1 A Unified Framework for Public-Key Primitives

We use the following formalization that generalizes both one-way relations and public-key encryption.

Definition 2.1 (public-key primitive). *A public-key primitive $\mathcal{E} = (G, V)$ is a pair of PPT algorithms such that*

- *The key generator G generates a pair of secret and public keys: $(sk, pk) \leftarrow G(1^k)$.*
- *The verifier V is an oracle machine such that $V^{\mathcal{O}(pk)}(pk)$ either accepts or rejects.*

Definition 2.2 (secure public-key primitive). *A public-key primitive $\mathcal{E} = (G, V)$ is secure if for any PPT oracle break, it holds that*

$$\Pr_{(sk, pk) \leftarrow G(1^k)} [V^{\text{break}(pk)}(pk)] = \text{negl}(k) .$$

To be concrete, for one-way relations, the breaker needs to send a candidate secret key sk (= inversion of the public key), and the verifier runs the relation’s verification procedure. To see why public key encryption can be stated in these terms, requires some work. The reason it is not immediate is that typically, we would consider the interaction between the verifier and the breaker, to be the following: The verifier gives the breaker a challenge ciphertext $\text{Enc}_{pk}(b)$, and he accepts if the breaker succeeds in guessing b . However, the breaker can clearly cause the verifier to accept with probability $1/2$, where we need to ensure that the breaker succeeds only with negligible probability. This technical annoyance can be fixed by considering the game where the verifier sends $\text{poly}(k)$ challenge ciphertexts to the breaker, each encrypting a random bit. The breaker succeeds if it succeeded in guessing significantly more than $1/2$ of the bits encrypted. The formal definition and precise analysis are much more cumbersome. In order not to break the flow of the presentation, we defer it to Appendix E.

Note that our verifier (which corresponds to the challenger in “security game” based definitions) only gets the public key as input and not the secret key. If the secret key was also given, then all public-key encryption schemes, signature schemes, and one-way relations, would trivially fit into this framework. However, in this work, we only consider primitives where the verifier V does not use the secret key sk to verify, but uses only the public key pk . An example of such a primitive is public-key encryption (under CPA). However, signature schemes or CCA secure encryption schemes do not fall into this category, since for these primitives the verifier in the definition above does need to know the secret key sk in order to simulate the signing oracle, in the case of signature schemes, and to simulate the decryption oracle, in the case of CCA encryption schemes.

2.2 Parallel Repetition

Definition 2.3 (t -parallel repetition). *For any public-key primitive $\mathcal{E} = (G, V)$ and any $t \in \mathbb{N}$, its t -parallel repetition, denoted $\mathcal{E}^t = (G^t, V^t)$, is in itself a public-key primitive defined as follows*

- *The key generator $(sk^t, pk^t) \leftarrow G^t(1^k)$ generates $(sk_i, pk_i) \leftarrow G(1^k)$ for all $i \in [t]$ and outputs $sk^t \triangleq (sk_1, \dots, sk_t)$, $pk^t \triangleq (pk_1, \dots, pk_t)$.*
- *The verifier $(V^t)^{\mathcal{O}(pk^t)}(pk^t)$, runs $V^{\mathcal{O}(pk^t, i)}(pk_i)$ for all $i \in [t]$, and accepts if and only if they all accept.*

A direct product of t schemes $\mathcal{E}_1 \times \dots \times \mathcal{E}_t$ is defined similarly.

While it is straightforward that our definition captures the notion of parallel repetition for one-way relations (where the goal is to find legal pre-images for all input public-keys), let us be a little more explicit about how the above captures parallel repetition for public-key encryption.

Lemma 2.1. *Let $\mathcal{E} = (G, V)$ be a public-key primitive that represents a public-key encryption scheme and let $t \in \mathbb{N}$. Then there exists a public key encryption scheme that is represented by \mathcal{E}^t .*

Moreover, this scheme is obtained by secret sharing the message into t shares and encrypting share i with pk_i . To decrypt, decrypt all shares and restore the message.

The proof is straightforward and is omitted.

2.3 Leakage Attacks

In this section, we generalize the notion of leakage attacks to our public-key primitive framework. Note that we do not define what it means for a scheme to be secure, only present a model for an attack.

Definition 2.4 (leakage attack). *We consider adversaries of the form $\mathcal{A} = (\text{leak}_{\mathcal{A}}, \text{break}_{\mathcal{A}})$, where $\text{leak}_{\mathcal{A}}, \text{break}_{\mathcal{A}}$ are (possibly randomized) functions. We refer to $\text{leak}_{\mathcal{A}}$ as the leakage function and to $\text{break}_{\mathcal{A}}$ as the breaker.*

A leakage attack (with security parameter k) of an adversary $\mathcal{A} = (\text{leak}_{\mathcal{A}}, \text{break}_{\mathcal{A}})$ on a public-key primitive $\mathcal{E} = (G, V)$ is the following process.

- **Initialize:** Generate a key pair $(sk, pk) \xleftarrow{\$} G(1^k)$.
- **Leak:** Apply the leakage function on the key pair to obtain the leakage value $y \leftarrow \text{leak}_{\mathcal{A}}(pk, sk)$.
- **Break:** \mathcal{A} succeeds if $V^{\text{break}(pk, y)}(pk)$ accepts.

3 Point-Wise Leakage Resilience

In this work, we consider “noisy leakage” functions, which are only allowed to reduce the (average) min-entropy of the secret key by a bounded amount. However, we relax the min-entropy restriction, and consider a *point-wise* definition, where we require that the specific leakage value is legal (as opposed to requiring that the leakage function is *always* legal).

We define our new model below. Then, in Sections 3.1, 3.2, we present two relaxed versions of point-wise leakage resilience that we need in order to prove our parallel repetition theorems. Finally, in Section 3.3 we show that all of these notions are strictly stronger than the old bounded-leakage model of [AGV09]. Namely, security w.r.t. to our definitions imply, as a special case, security w.r.t. bounded leakage.

Definition 3.1 (point-wise λ -leakage). *Let $\mathcal{E} = (G, V)$ be a public key primitive. A possibly randomized leakage function leak is λ -leaky at point (pk, y) , where pk is a public key and y is a leakage value (in the image of leak), if*

$$\mathbf{H}_{\infty}(S_{pk, y}) \geq \mathbf{H}_{\infty}(S_{pk}) - \lambda ,$$

where S_{pk} is the distribution of secret keys conditioned on the public key being pk , and $S_{pk, y}$ is the distribution of secret keys conditioned on both the public key being pk and on $\text{leak}(pk, sk) = y$.

Definition 3.2 (point-wise λ -leakage resilience). *A public-key primitive $\mathcal{E} = (G, V)$ is point-wise λ -leakage-resilient if for any PPT adversary $\mathcal{A} = (\text{leak}_{\mathcal{A}}, \text{break}_{\mathcal{A}})$ it holds that*

$$\text{Adv}_{\mathcal{E}, \lambda}[\mathcal{A}] \triangleq \Pr [(\text{leak}_{\mathcal{A}} \text{ is } \lambda\text{-leaky at } (pk, y)) \wedge (\mathcal{A}(pk, y) \text{ succeeds})] = \text{negl}(k) ,$$

where the probability is taken over $(sk, pk) \leftarrow G(1^k)$, over the random coin tosses of $\mathcal{A} = (\text{leak}_{\mathcal{A}}, \text{break}_{\mathcal{A}})$, and over the random coin tosses of the verifier in the verification game.

In order to obtain our direct product theorems for leakage resilience, we relax the point-wise leakage resilience definition in two (incomparable) ways.

3.1 First Relaxation: Almost Leakage Resilience

In this relaxation, instead of requiring that sk has high min-entropy conditioned on pk and $y = \text{leak}(pk, sk)$, we require that the distribution of sk (conditioned on pk, y) is statistically close to one that has high min-entropy.

Definition 3.3 (close to λ -leaky). *A leakage function leak is μ -close to λ -leaky at point (pk, y) if there exists a distribution $\tilde{S}_{pk,y}$ that is μ -close to $S_{pk,y}$ and*

$$\mathbf{H}_\infty(\tilde{S}_{pk,y}) \geq \mathbf{H}_\infty(S_{pk}) - \lambda .$$

Definition 3.4 (resilience to almost λ -leakage). $\mathcal{E} = (G, V)$ is point-wise almost λ -leakage-resilient if for any PPT adversary $\mathcal{A} = (\text{leak}_\mathcal{A}, \text{break}_\mathcal{A})$ and for any negligible function μ , it holds that

$$\text{Adv}_{\mathcal{E}, \lambda, \mu}[\mathcal{A}] \triangleq \Pr [(\text{leak}_\mathcal{A} \text{ is } \mu\text{-close to } \lambda\text{-leaky at } (pk, y)) \wedge (\mathcal{A}(pk, y) \text{ succeeds})] = \text{negl}(k) .$$

where the probability is taken over $(sk, pk) \leftarrow G(1^k)$, over the random coin tosses of $\mathcal{A} = (\text{leak}_\mathcal{A}, \text{break}_\mathcal{A})$, and over the random coin tosses of the verifier in the verification game.

Under this definition we obtain a direct-product theorem for constant number of repetitions.

3.2 Second Relaxation: Leakage Resilience with Advice

To obtain a direct-product theorem for a super-constant number of repetitions, we use a slightly different (and incomparable) model, where we do not allow statistical closeness, but rather allow the attacker to get a logarithmic number of bits of (possibly inefficient) advice.

Definition 3.5 (PPT_a). *We say that a function f is PPT_a computable if the function f_a , defined below, is PPT computable. The function f_a is identical to f , except that the last a bits of its output are truncated.*

We say that an adversary $\mathcal{A} = (\text{leak}_\mathcal{A}, \text{break}_\mathcal{A})$ is a PPT_a adversary if $\text{leak}_\mathcal{A}$ is PPT_a computable and $\text{break}_\mathcal{A}$ is PPT computable.

Definition 3.6 (point-wise λ -leakage with advice). *A public-key primitive $\mathcal{E} = (G, V)$ is resilient to point-wise λ -leakage and logarithmic advice if for any $\text{PPT}_{O(\log k)}$ adversary $\mathcal{A} = (\text{leak}_\mathcal{A}, \text{break}_\mathcal{A})$ it holds that*

$$\text{Adv}_{\mathcal{E}, \lambda}[\mathcal{A}] \triangleq \Pr [(\text{leak}_\mathcal{A} \text{ is } \lambda\text{-leaky at } (pk, y)) \wedge (\mathcal{A}(pk, y) \text{ succeeds})] = \text{negl}(k) ,$$

where the probability is taken over $(sk, pk) \leftarrow G(1^k)$, over the random coin tosses of $\mathcal{A} = (\text{leak}_\mathcal{A}, \text{break}_\mathcal{A})$, and over the random coin tosses of the verifier in the verification game.

3.3 Relation to Bounded Leakage

To conclude, we prove that point-wise λ -leakage resilience implies the basic form of λ -bounded leakage.

Definition 3.7 ([AGV09]). *A public-key primitive $\mathcal{E} = (G, V)$ is λ -bounded leakage resilient if any PPT adversary $\mathcal{A} = (\text{leak}_\mathcal{A}, \text{break}_\mathcal{A})$ for which the output of $\text{leak}_\mathcal{A}$ is at most λ bits, succeeds with negligible probability.*

Lemma 3.1. *If $\mathcal{E} = (G, V)$ is point-wise λ -leakage resilient then it is also λ -bounded leakage resilient.*

We note that point-wise almost λ -leakage resilience, and λ -leakage resilience with logarithmic advice, are stronger notions of security (they give the adversary more power) and thus the above immediately applies to these notions as well.

Since this is not the main focus of this work, we only sketch the proof in Appendix B.

4 Why Known Schemes are Point-Wise Leakage Resilient

In this section, we show that leakage resilience is amplified by parallel repetition for, essentially, all known schemes that are resilient to non-trivial (i.e. super-logarithmic) leakage. To show this, we sketch a proof template that is shared among all (non trivial) leakage resilient results, and we show that this proof template proves security also w.r.t. our leakage models (the point-wise almost λ -leakage model, and the point-wise λ -leakage with logarithmic advice model).

The Proof Template. The proof template for proving leakage resilience is very simple, and works in two hybrid steps. Recall that the adversary first gets a pair $(pk, y = L(pk, sk))$, where L is a poly-size leakage function chosen by \mathcal{A} . Then it chooses messages m_0, m_1 and gets a challenge ciphertext $c_b \leftarrow \text{Enc}_{pk}(m_b)$. The adversary wins if it guesses the bit b correctly.

The first step in the template is to replace the challenge c_b with an “illegally” generated ciphertext c_b^* , such that $(sk, pk, c_b) \stackrel{c}{\approx} (sk, pk, c_b^*)$ (and it is efficient to generate c_b^* given sk, pk, b). Due to computational indistinguishability, the adversary’s success probability should remain unchanged. We note that there is no entropy involved in this part, only a requirement that L is efficiently computable.

The second step is completely information theoretic: It is proven that if the distribution of the secret key conditioned on pk, y , which we denote by $S_{pk, y}$, has sufficient min-entropy, then c_b^* carries no information on b (or, more precisely, that conditioned on the view of the adversary, b is statistically close to uniform). Therefore, no adversary can guess its value with non-negligible advantage.

Point-Wise Leakage Resilience. The above proof template also proves point-wise leakage resilience. The second step of the hybrid works in a point-wise manner and therefore we only need to worry about the first step. In the first step, clearly computational indistinguishability still holds, but proving that the point-wise advantage remains unchanged is a bit harder, since we cannot efficiently check the point-wise advantage. Nevertheless, we argue that if the advantage of \mathcal{A} is non-negligible, then it drops by a factor of at most two. Such a claim is sufficient for the next level of the template.

To see why this is the case, consider an adversary \mathcal{A} that has non-negligible point-wise advantage ϵ when given (pk, y, c_b) , but less than $\epsilon/2$ when given (pk, y, c_b^*) . Recall that the advantage measures the probability of both \mathcal{A} succeeding (in the verification game) and pk, y being point-wise λ -leaky. It follows that with non-negligible probability over pk, y , the conditional success probability of \mathcal{A} , conditioned on pk, y , drops by at least $\epsilon/4$ (otherwise the advantage, which measures over a subset of the pk, y , couldn’t have dropped).

A distinguisher $\mathcal{B}(sk, pk, c_b/c_b^*)$ is defined as follows: First, compute the leakage $y := L(sk, pk)$. Then generate many samples of c_b/c_b^* and use them to evaluate the success probability of \mathcal{A} conditioned on pk, y in the two cases. If indeed pk, y are such that the success probability drops, use \mathcal{A} to distinguish between the two cases. If no noticeable change in the success probability was noticed, then output a random guess. Putting it all together, we get a polynomial distinguisher between (pk, y, c_b) and (pk, y, c_b^*) , in contradiction to the hardness assumption.

We note that this is true even if y is not fully known to the distinguisher: say $O(\log k)$ bits of y are not known, the distinguisher can still try all options and check if for either one the success probability changes by $\epsilon/4$.

Our Relaxed Models of Point-Wise Leakage Resilience. Our first relaxation, of allowing the secret key to be statistically close to λ -leakage resilient, only effects the second step of the template. We can still argue that b is statistically close to uniform by adding another hybrid where the conditional distribution $S_{pk,y}$ is replaced with a statistically indistinguishable $\tilde{S}_{pk,y}$ that has high min-entropy.

Our second relaxation, of allowing logarithmic advice, goes into the first step (this is the only step where we care about the complexity of L). As we explained above, our argument works even if a logarithmic part of the leakage value is not known. Therefore we will use only the efficient part of the leakage function and computational indistinguishability will still hold.

Computationally Indistinguishable Schemes. For some schemes, such as [AGV09, LW10], leakage resilient is proven by showing that they are computationally indistinguishable from another scheme which, in turn, is proven leakage resilient using the template. We show in Section 7 that this still implies that parallel repetition amplifies leakage.

5 Direct-Product Theorem for a Constant Number of Repetitions

In this section, we prove a direct-product theorem for a constant number of repetitions, w.r.t. point-wise almost leakage-resilience as defined in Section 3.1.

Theorem 5.1. *Let $c \in \mathbb{N}$ be a constant, and for every $i \in [c]$, let $\mathcal{E}_i = (G_i, V_i)$ be a point-wise almost λ_i -leakage-resilient public-key primitive. Then, $\mathcal{E}_1 \times \dots \times \mathcal{E}_c$ is point-wise almost λ -leakage-resilient, where $\lambda = \sum_{i=1}^c (\lambda_i - 1)$.*

It suffices to prove this theorem for $c = 2$, and apply it successively. In order to simplify notation, we prove it for the case of parallel repetition, where $\mathcal{E}_1 = \mathcal{E}_2$, the proof extends readily to the case of direct product.

Lemma 5.2. *Let $\mathcal{E} = (G, V)$ be a point-wise almost λ -leakage-resilient public-key primitive. Then \mathcal{E}^2 is point-wise almost $(2\lambda - 1)$ -leakage-resilient.*

Before we present the outline of the proof, let us make a few remarks.

1. Note that there is a loss of one bit in the amplification. Namely, we go from λ to $(2\lambda - 1)$ instead of just 2λ . While some loss in the parameters is implied by our techniques, more detailed analysis can show that the composed scheme is in fact $(2\lambda - \delta)$ -leakage resilient for any $\delta(k) = 1/\text{poly}(k)$. Thus the loss incurred is less than a single bit. As our result is qualitative in nature, we chose not to overload with the additional complication.

2. While at first glance one could imagine that Theorem 5.1 should extend beyond constant c , we were unable to prove such an argument. The reason is that super-constant repetition gives a different scheme for each value of the security parameter. This means that we cannot use Theorem 5.1 as black-box. More importantly, our proof techniques rely on the asymptotic behavior of the scheme so we were not able to even change the proof to apply for a super-constant number of repetitions.

A result for the more general case of any polynomial number of repetitions is presented, in the slightly different and incomparable “advice” model, in Section 6.

Finally, we remark that known negative results for security of parallel repetition are already effective for a constant number of repetitions. Thus our result contrasts them even for this case.

Proof overview of Lemma 5.2. We consider an adversary \mathcal{B} that succeeds in the parallel repetition game, and construct an adversary \mathcal{A} that succeeds in the single instance game. The straightforward proof strategy would be to “plant” the “real” key pair, that is given as input to \mathcal{A} , as one of the key pairs that are input to \mathcal{B} , and sample the other pair uniformly.⁴ In such case, the input to \mathcal{B} is distributed identically to the parallel repetition case and indeed \mathcal{B} will succeed with noticeable probability. However, we may no longer be able to claim that our leakage leaves sufficient entropy in the secret key. We are guaranteed by the functionality of \mathcal{B} that the key pair (sk_1, sk_2) is left with sufficient min-entropy but it is still possible that neither sk_1 nor sk_2 have *any* min-entropy by themselves.

To solve the above we use Lemma A.1, which essentially says how to split-up the joint entropy of two random variables. Specifically it says that either sk_1 or $sk_2|sk_1$ will have sufficient min-entropy, depending on whether sk_1 belongs to a hard-to-recognize set R , and conditioned on the knowledge of whether $sk_1 \in R$. Namely, either $sk_1|\mathbb{1}_{sk_1 \in R}$ or $sk_2|(sk_1, \mathbb{1}_{sk_1 \in R})$ have high min-entropy. If we could compute the bit $\mathbb{1}_{sk_1 \in R}$, we would have been done (and indeed if we are allowed one bit of inefficient leakage, an easier proof follows, see Section 6). Since this is impossible, we turn to case analysis:

Obviously, if $\Pr[sk_1 \in R] = \text{negl}(k)$, then we can always guess that $\mathbb{1}_{sk_1 \in R} = 0$ and be right almost always. This implies that in such case $sk_2|sk_1$ is statistically indistinguishable from having high min-entropy, as we wanted.

For the second case, if $\Pr[sk_1 \in R] \geq 1/\text{poly}(k)$, then $sk_2|(sk_1, \mathbb{1}_{sk_1 \in R})$ will have high min-entropy for a noticeable part of the time. To complete the analysis here, we notice that

$$\mathbf{H}_\infty(sk_2|(sk_1, \mathbb{1}_{sk_1 \in R})) = \mathbf{H}_\infty(sk_2|sk_1).$$

This is because R is a well defined set and thus $\mathbb{1}_{sk_1 \in R}$ is a deterministic (though hard to compute) function of sk_1 . It follows that $sk_2|sk_1$ will have high min-entropy for a noticeable fraction of the time, which completes the proof.

For the formal proof, see Appendix C.

⁴We note that even this step is impossible when relying on “secretly generated” public parameters as in the scheme presented in [LW10] (or rather, the scheme that is computationally indistinguishable to theirs and actually has entropic leakage resilient).

6 Direct-Product Theorem for Polynomially Many Repetitions

In this section we present a direct product theorem that applies to any polynomial number of repetitions. This theorem is relative to the advice model defined in Section 3.2. For the sake of simplicity, we will assume that the number of repetitions is a power of 2, although the same techniques can be used for any number.

Theorem 6.1. *Let $\mathcal{E} = (G, V)$ be a public-key primitive that is resilient to point-wise λ -leakage and logarithmic advice. Let $t = t(k)$ be a polynomially bounded function of the security parameter such that $t(k)$ is always a power of 2. Then \mathcal{E}^t is resilient to point-wise $t(\lambda - 1)$ -leakage and logarithmic advice.*

Towards proving the theorem, we present the following lemma, which is a parameterized special case of the above theorem, and will imply the theorem by successive applications.

Lemma 6.2. *For any public-key primitive $\mathcal{E} = (G, V)$ and any PPT_{-a} adversary $\mathcal{B} = (\text{leak}_{\mathcal{B}}, \text{break}_{\mathcal{B}})$ for \mathcal{E}^2 , there exists a $\text{PPT}_{-(a+1)}$ adversary $\mathcal{A} = (\text{leak}_{\mathcal{A}}, \text{break}_{\mathcal{A}})$ for \mathcal{E} , such that for all k ,*

$$\text{Adv}_{\mathcal{E}, \lambda}[\mathcal{A}] \geq (1/4) \cdot \text{Adv}_{\mathcal{E}^2, (2\lambda-1)}[\mathcal{B}] . \quad (1)$$

In what follows, we prove the theorem given the lemma. The proof of the lemma appears in Appendix D.

Proof of Theorem 6.1. Let \mathcal{B}_0 be a PPT_{-a} adversary for \mathcal{E}^t . Then by Lemma 6.2, there exists a $\text{PPT}_{-(a+1)}$ adversary \mathcal{A} for $\mathcal{E}^{t/2}$, that satisfies Equation (1). We denote this adversary by

$$\mathcal{B}_1 \triangleq \mathcal{A}(\mathcal{B}_0, \mathcal{E}^{t/2}).$$

More generally, we consider a series of adversaries $\mathcal{B}_1, \dots, \mathcal{B}_{\log t}$, defined by

$$\mathcal{B}_i \triangleq \mathcal{A}(\mathcal{B}_{i-1}, \mathcal{E}^{t \cdot 2^{-i}}) .$$

Lemma 6.2 implies that $\mathcal{B}_{\log t}$ is a $\text{PPT}_{-(a+\log t)}$ adversary and that

$$\text{Adv}_{\mathcal{E}, \lambda}[\mathcal{B}_{\log t}] \geq (1/t^2) \cdot \text{Adv}_{\mathcal{E}^t, t(\lambda-1)}[\mathcal{B}_0] ,$$

and the theorem follows. □

7 Leakage from Computationally Indistinguishable Schemes

Our definition of point-wise leakage resilience is based on the residual min-entropy of the secret key, conditioned on the leakage value. In the literature, starting with [NS09], this is referred to as “resilience to *noisy* leakage”. It is self evident that schemes where the public key is an injective function of the secret key cannot be proven leakage resilient in this respect. This is because even leaking the secret key in its entirety, which obviously breaks security, does not reduce its min entropy conditioned on the public key (the conditional min-entropy is 0 to begin with, and it stays 0 after the leakage). We do know, however, of such injective public-key encryption schemes that are proven to be leakage resilient with respect to the weaker notion of “length bounded leakage”. There, the

restriction on the leakage function is that it has bounded length. Notable examples are the scheme of [AGV09] and the scheme of [LW10] (which was introduced as a counterexample for parallel repetition of length-bounded leakage resilience, see Section 1.2). While at first glance it may seem that our result is completely powerless with regards to such schemes, we show in this section that for all known schemes, and specifically for the schemes of [AGV09, LW10], our theorem in fact *does* imply parallel repetition.

The key observation upon revisiting the proofs of security of [AGV09, LW10], is that in both cases, the proof is by presenting a second scheme in which the key distribution is computationally indistinguishable from the original scheme (but may have undesired features such as worse efficiency of key generation), and proving that this second scheme is resilient to leakage of bounded length. This implies that the original scheme is resilient to bounded leakage as well (since otherwise one can distinguish the key generation processes). The second scheme, in these two cases, is in fact resilient to noisy leakage. Furthermore, the second scheme in the two cases adheres to our notion of point-wise leakage resilience.

In light of the above, we put forth the following corollary of Theorems 5.1 and 6.1.

Corollary 7.1. *Let $\mathcal{E} = (G, V)$ be a public-key primitive, and let G' be such that $G(1^k) \stackrel{c}{\approx} G'(1^k)$. Then:*

1. *If $\mathcal{E}' = (G', V)$ is point-wise almost λ -leakage resilient, then \mathcal{E}^t is $t \cdot (\lambda - 1)$ -bounded leakage resilient for any constant $t \in \mathbb{N}$.*
2. *If $\mathcal{E}' = (G', V)$ is point-wise λ -leakage resilient with logarithmic advice, then \mathcal{E}^t is $t \cdot (\lambda - 1)$ -bounded leakage resilient for any polynomial $t = t(k)$.*

Proof. The proof of the two parts is almost identical: We use either Theorem 5.1 or Theorem 6.1 to show that $(\mathcal{E}')^t$ is point-wise almost $t \cdot (\lambda - 1)$ -leakage resilient, or, respectively, leakage resilient with logarithmic advice. By Lemma 3.1, this means that $(\mathcal{E}')^t$ is $t \cdot (\lambda - 1)$ -bounded leakage resilient.

By a hybrid argument $(G')^t(1^k) \stackrel{c}{\approx} G^t(1^k)$.⁵ Therefore, it must be that \mathcal{E}^t is also $t \cdot (\lambda - 1)$ -bounded leakage resilient (otherwise there is a distinguisher between the key generators). This completes the proof. \square

Using Corollary 7.1, we can show that t -parallel repetition of the schemes of [AGV09, LW10] indeed amplifies their leakage resilience.

References

- [ADN⁺10] Joël Alwen, Yevgeniy Dodis, Moni Naor, Gil Segev, Shabsi Walfish, and Daniel Wichs. Public-key encryption in the bounded-retrieval model. In Henri Gilbert, editor, *EUROCRYPT*, volume 6110 of *Lecture Notes in Computer Science*, pages 113–134. Springer, 2010.
- [ADW09] Joël Alwen, Yevgeniy Dodis, and Daniel Wichs. Leakage-resilient public-key cryptography in the bounded-retrieval model. In Halevi [Hal09], pages 36–54.

⁵It is important to note that this will not be necessarily true if either G or G' relies on secret parameters. This relates to our discussion of the result of [LW10] in Section 1.2.

- [AGV09] Adi Akavia, Shafi Goldwasser, and Vinod Vaikuntanathan. Simultaneous hardcore bits and cryptography against memory attacks. In Omer Reingold, editor, *TCC*, volume 5444 of *Lecture Notes in Computer Science*, pages 474–495. Springer, 2009.
- [BG10] Zvika Brakerski and Shafi Goldwasser. Circular and leakage resilient public-key encryption under subgroup indistinguishability - (or: Quadratic residuosity strikes back). In Tal Rabin, editor, *CRYPTO*, volume 6223 of *Lecture Notes in Computer Science*, pages 1–20. Springer, 2010.
- [BS09] Zvika Brakerski and Gil Segev. Personal communication, 2009.
- [BSW11] Elette Boyle, Gil Segev, and Daniel Wichs. Fully leakage-resilient signatures. In Kenneth G. Paterson, editor, *EUROCRYPT*, volume 6632 of *Lecture Notes in Computer Science*, pages 89–108. Springer, 2011.
- [DFR⁺07] Ivan B. Damgård, Serge Fehr, Renato Renner, Louis Salvail, and Christian Schaffner. A tight high-order entropic quantum uncertainty relation with applications. In *Proceedings of the 27th annual international cryptology conference on Advances in cryptology, CRYPTO’07*, pages 360–378, Berlin, Heidelberg, 2007. Springer-Verlag.
- [DGK⁺10] Yevgeniy Dodis, Shafi Goldwasser, Yael Tauman Kalai, Chris Peikert, and Vinod Vaikuntanathan. Public-key encryption schemes with auxiliary inputs. In Micciancio [Mic10], pages 361–381.
- [DKL09] Yevgeniy Dodis, Yael Tauman Kalai, and Shachar Lovett. On cryptography with auxiliary input. In Michael Mitzenmacher, editor, *STOC*, pages 621–630. ACM, 2009.
- [DORS08] Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. *SIAM J. Comput.*, 38(1):97–139, 2008.
- [DP08] Stefan Dziembowski and Krzysztof Pietrzak. Leakage-resilient cryptography. In *FOCS*, pages 293–302. IEEE Computer Society, 2008.
- [FKPR10] Sebastian Faust, Eike Kiltz, Krzysztof Pietrzak, and Guy N. Rothblum. Leakage-resilient signatures. In Micciancio [Mic10], pages 343–360.
- [GJS11] Sanjam Garg, Abhishek Jain, and Amit Sahai. Leakage-resilient zero knowledge. In Phillip Rogaway, editor, *CRYPTO*, volume 6841 of *Lecture Notes in Computer Science*, pages 297–315. Springer, 2011.
- [Hal09] Shai Halevi, editor. *Advances in Cryptology - CRYPTO 2009, 29th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2009. Proceedings*, volume 5677 of *Lecture Notes in Computer Science*. Springer, 2009.
- [JP11] Abhishek Jain and Krzysztof Pietrzak. Parallel repetition for leakage resilience amplification revisited. In Yuval Ishai, editor, *TCC*, volume 6597 of *Lecture Notes in Computer Science*, pages 58–69. Springer, 2011.

- [KV09] Jonathan Katz and Vinod Vaikuntanathan. Signature schemes with bounded leakage resilience. In Mitsuru Matsui, editor, *ASIACRYPT*, volume 5912 of *Lecture Notes in Computer Science*, pages 703–720. Springer, 2009.
- [LW10] Allison B. Lewko and Brent Waters. On the insecurity of parallel repetition for leakage resilience. In Trevisan [Tre10], pages 521–530.
- [Mic10] Daniele Micciancio, editor. *Theory of Cryptography, 7th Theory of Cryptography Conference, TCC 2010, Zurich, Switzerland, February 9-11, 2010. Proceedings*, volume 5978 of *Lecture Notes in Computer Science*. Springer, 2010.
- [MTVY11] Tal Malkin, Isamu Teranishi, Yevgeniy Vahlis, and Moti Yung. Signatures resilient to continual leakage on memory and computation. In Yuval Ishai, editor, *TCC*, volume 6597 of *Lecture Notes in Computer Science*, pages 89–106. Springer, 2011.
- [NS09] Moni Naor and Gil Segev. Public-key cryptosystems resilient to key leakage. In Halevi [Hal09], pages 18–35.
- [Pie09] Krzysztof Pietrzak. A leakage-resilient mode of operation. In Antoine Joux, editor, *EUROCRYPT*, volume 5479 of *Lecture Notes in Computer Science*, pages 462–482. Springer, 2009.
- [Tre10] Luca Trevisan, editor. *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*. IEEE Computer Society, 2010.
- [Wul07] Jürg Wullschlegler. Oblivious-transfer amplification. In Moni Naor, editor, *EUROCRYPT*, volume 4515 of *Lecture Notes in Computer Science*, pages 555–572. Springer, 2007.

A How to Split Min-Entropy

We present a lemma that shows that the joint min-entropy of two random variables can be split between them under some condition. Variants of this lemma appeared in previous works (e.g. [DFR⁺07, Wul07]), this formulation is from [BS09].

Lemma A.1 (min-entropy split). *Let X, Y be such that $\mathbf{H}_\infty(X, Y) \geq a + b$, for $a, b > 0$. Then there exists a set R_X , which is a subset of the support of X such that both:*

1. *For all $x \in R_X$, it holds that $\mathbf{H}_\infty(Y|X = x) \geq b$.*
2. *$\mathbf{H}_\infty(X|X \notin R_X) \geq a - \log(1/\epsilon)$, where $\epsilon \triangleq \Pr[X \notin R_X]$.*

Proof. Define

$$R_X \triangleq \{x : \Pr[X = x] \geq 2^{-a}\} .$$

Then for all $x \in R_X$ and for all y , it holds that $\Pr[Y = y|X = x] \leq 2^{-b}$, and thus $\mathbf{H}_\infty(Y|X = x) \geq b$. In addition, $\mathbf{H}_\infty(X|X \notin R_X) \geq a + \log \Pr[X \notin R_X]$, i.e. $\mathbf{H}_\infty(X|X \notin R_X) \geq a - \log(1/\epsilon)$. \square

B Proof Sketch of Lemma 3.1

Let \mathcal{E} be as above, and let \mathcal{B} be a λ -bounded adversary for \mathcal{E} that succeeds with probability $1/q(k)$ for some polynomial $q(k)$.

We first show that for any polynomial $r(k)$, there exists an adversary \mathcal{A} that is $(\lambda - \log r(k))$ -bounded and succeeds with probability $1/(2q(k))$. The adversary \mathcal{A} will have $\text{leak}_{\mathcal{A}}$ be identical to $\text{leak}_{\mathcal{B}}$, except the last $\log r(k)$ bits be truncated. The breaker $\text{break}_{\mathcal{A}}(pk, y)$ will simulate $\text{break}_{\mathcal{B}}(pk, y||a)$ with all possible suffixes $a \in \{0, 1\}^{\log r(k)}$. For each of the $r(k)$ possible values of a , it will estimate $\eta_a \triangleq \Pr[V^{\text{break}_{\mathcal{B}}(pk, y||a)}(pk)]$ by simulating the verifier with fresh randomness a polynomial number of times. It will choose a^* to be the value with the highest η_{a^*} . It will then interface with the real verifier by simulating $\text{break}_{\mathcal{B}}(pk, y||a^*)$. It is clear that \mathcal{A} would achieve at least as good performance as \mathcal{B} if it always estimated η_a correctly. One can choose the parameters so that the Chernoff bound implies that with probability $1/(4q)$ the estimation error for either value of a is at most $1/(8q)$. This is sufficient to achieve the required success probability.

Once we have \mathcal{A} , we show that its point-wise λ -leakage advantage is at least $1/2q - 1/r$. To do this, we will prove that $\text{leak}_{\mathcal{A}}$ is λ -leaky with all but $1/r$ probability and use the union bound. Choosing $r = 4q$ will then finish the proof of the lemma.

We prove an even stronger version of this last claim. We will show that the above holds true for every value of pk . From now on let us fix some value for pk . Then by conditional min-entropy arguments [DORS08]

$$-\log \mathbb{E}[2^{-\mathbf{H}_{\infty}(S_{pk, y})}] \geq \mathbf{H}_{\infty}(S_{pk}) - (\lambda - \log r) ,$$

and therefore

$$\mathbb{E}[2^{\mathbf{H}_{\infty}(S_{pk}) - \mathbf{H}_{\infty}(S_{pk, y})}] \leq 2^{\lambda/r} .$$

It follows from Markov that

$$\Pr[2^{\mathbf{H}_{\infty}(S_{pk}) - \mathbf{H}_{\infty}(S_{pk, y})} \geq 2^{\lambda}] \leq 1/r ,$$

or stated differently

$$\Pr[\mathbf{H}_{\infty}(S_{pk}) - \mathbf{H}_{\infty}(S_{pk, y}) \geq \lambda] \leq 1/r ,$$

which shows that with all but $1/r$ probability, \mathcal{A} is λ -leaky as required. \square

C Proof of Lemma 5.2

Suppose for the sake of contradiction that \mathcal{E}^2 is not point-wise $(2\lambda - 1)$ -leakage-resilient. Namely, suppose that there exists a PPT adversary \mathcal{B} , a negligible function μ , and a non-negligible function ϵ , such that

$$\begin{aligned} \text{Adv}_{\mathcal{E}^2, 2\lambda-1, \mu}[\mathcal{B}] &= \Pr[(\text{leak}_{\mathcal{B}} \text{ is } \mu\text{-close to } (2\lambda - 1)\text{-leaky at } (pk_1, pk_2, y_{\mathcal{B}})) \wedge (\text{break}_{\mathcal{B}} \text{ succeeds})] \\ &= \Pr \left[\left(\exists \tilde{S}_{pk_1, pk_2, y_{\mathcal{B}}}^2 \text{ s.t. } \mathbf{H}_{\infty}(\tilde{S}_{pk_1, pk_2, y_{\mathcal{B}}}^2) \geq \mathbf{H}_{\infty}(S_{pk_1, pk_2}^2) - (2\lambda - 1) \right) \wedge \right. \\ &\quad \left. (\text{break}_{\mathcal{B}} \text{ succeeds}) \right] \\ &\geq \epsilon(k) , \end{aligned}$$

where $\tilde{S}_{pk_1, pk_2, y_{\mathcal{B}}}^2$ can be any distribution that is μ -close to $S_{pk_1, pk_2, y_{\mathcal{B}}}^2$, and where the probabilities are over $(sk_1, pk_1), (sk_2, pk_2) \leftarrow G(1^k)$, and over the random coin tosses of $\mathcal{B} = (\text{leak}_{\mathcal{B}}, \text{break}_{\mathcal{B}})$.

The fact that ϵ is non-negligible implies that there exists an infinite set $K \subseteq \mathbb{N}$ and a polynomial q such that $\epsilon(k) \geq 1/q(k)$. Throughout the proof, we focus our attention on k 's in K .

Define an adversary $\mathcal{A} = (\text{leak}_{\mathcal{A}}, \text{break}_{\mathcal{A}})$ for \mathcal{E} as follows.

- The function $\text{leak}_{\mathcal{A}}(sk, pk)$ is defined as follows:

1. Sample $\tau \xleftarrow{\$} [2]$ and define $(sk_{\tau}, pk_{\tau}) \triangleq (sk, pk)$.
2. For $i = 3 - \tau$, sample $(sk_i, pk_i) \leftarrow G(1^k)$.
3. Compute $y_{\mathcal{B}} \leftarrow \text{leak}_{\mathcal{B}}((sk_1, sk_2), (pk_1, pk_2))$.
4. Output $y_{\mathcal{A}} = (pk_1, pk_2, y_{\mathcal{B}}, \mathbb{1}_{\tau=2} \cdot sk_1)$.

Where the value $\mathbb{1}_{\tau=2} \cdot sk_1$ means that if $\tau = 2$ then we output sk_1 as part of our leakage, and otherwise this part is null.⁶

- The function $\text{break}_{\mathcal{A}}(pk, y_{\mathcal{A}})$ emulates $\text{break}_{\mathcal{B}}(pk_1, pk_2, y_{\mathcal{B}})$ by playing the role of the verification oracle for $i \neq \tau$ and using the actual verification oracle for $i = \tau$.

It immediately follows that \mathcal{A} is a PPT adversary. It remains to analyze its success probability. By definition, for any negligible function ν ,

$$\text{Adv}_{\mathcal{E}, \lambda, \nu}[\mathcal{A}] = \Pr [(\text{leak}_{\mathcal{A}} \text{ is } \nu\text{-close to } \lambda\text{-leaky at } (pk, y_{\mathcal{A}})) \wedge (\mathcal{A} \text{ succeeds})] .$$

Let ρ denote the random tape of $\text{break}_{\mathcal{A}}$ and the random tape of the verifier in the verification game (we emphasize that ρ does not include the random tape of $\text{leak}_{\mathcal{A}}$). Then, it holds that for any negligible function ν ,

$$\begin{aligned} \text{Adv}_{\mathcal{E}, \lambda, \nu}[\mathcal{A}] &= \Pr [(\text{leak}_{\mathcal{A}} \text{ is } \nu\text{-close to } \lambda\text{-leaky at } (pk, y_{\mathcal{A}})) \wedge (\mathcal{A} \text{ succeeds})] = \\ &= \mathbb{E}_{pk_1, pk_2, y_{\mathcal{B}}, \rho} \left[\Pr [(\text{leak}_{\mathcal{A}} \text{ is } \nu\text{-close to } \lambda\text{-leaky at } (pk, y_{\mathcal{A}})) \wedge (\mathcal{A} \text{ succeeds}) \mid pk_1, pk_2, y_{\mathcal{B}}, \rho] \right] . \end{aligned}$$

Note that conditioning on $(pk_1, pk_2, y_{\mathcal{B}}, \rho)$, the latter probability is only over $\tau \leftarrow [2]$ and over $sk_i \mid pk_1, pk_2, y_{\mathcal{B}}$.

We prove the following claim.

Claim C.0.1. *There exists a polynomial p and a negligible function ν such that for infinitely many k 's in K , and for all $(pk_1, pk_2, y_{\mathcal{B}}, \rho)$ (corresponding to the security parameter k), it holds that*

$$\begin{aligned} &\Pr [(\text{leak}_{\mathcal{A}} \text{ is } \nu\text{-close to } \lambda\text{-leaky at } (pk, y_{\mathcal{A}})) \wedge (\mathcal{A} \text{ succeeds}) \mid pk_1, pk_2, y_{\mathcal{B}}, \rho] \geq \\ &\frac{1}{p(k)} \cdot \Pr [(\text{leak}_{\mathcal{B}} \text{ is } \mu\text{-close to } (2\lambda - 1)\text{-leaky at } (pk_1, pk_2, y_{\mathcal{B}})) \wedge (\mathcal{B} \text{ succeeds}) \mid pk_1, pk_2, y_{\mathcal{B}}, \rho] , \end{aligned}$$

where we envision an invocation of \mathcal{B} that runs in the same probability space as \mathcal{A} : Namely, its key pair is $((sk_1, sk_2), (pk_1, pk_2))$, its leakage value is $y_{\mathcal{B}}$ and its verification game uses the same random tape ρ .

⁶Assume w.l.o.g that this null value cannot be a legal secret key and thus this serves as an indication that $\tau = 1$.

Note that Claim C.0.1 completes the proof of the theorem as it implies that for infinitely many k 's in K ,

$$\text{Adv}_{\mathcal{E},\lambda,\nu}[\mathcal{A}] \geq \frac{1}{p} \cdot \text{Adv}_{\mathcal{E}^2,2\lambda-1,\mu}[\mathcal{B}] .$$

Proof of Claim C.0.1. Let us fix values for (pk_1, pk_2, y_B, ρ) and keep them fixed throughout the proof (thus omitting the conditioning in the probability brackets). Notice that conditioned on these values,

$$\Pr[(\text{leak}_B \text{ is } \mu\text{-close to } (2\lambda - 1)\text{-leaky at } (pk_1, pk_2, y_B)) \wedge (\mathcal{B} \text{ succeeds})] \in \{0, 1\}.$$

Namely, the values (pk_1, pk_2, y_B, ρ) completely determine whether leak_B is μ -close to $(2\lambda - 1)$ -leaky and whether \mathcal{B} succeeds. Moreover, notice that if the above probability is 0 then the claim holds trivially (since the right-hand side is zero). Thus, we assume that for our fixed values (pk_1, pk_2, y_B, ρ) the above probability is 1; i.e., \mathcal{B} succeeds and leak_B is μ -close to $(2\lambda - 1)$ -leaky at (pk_1, pk_2, y_B) .

Since the verification game of \mathcal{A} simulates that of \mathcal{B} using the same randomness, it follows that if \mathcal{B} succeeds in the game then so does \mathcal{A} . Therefore all that is left to prove is that there exists a polynomial p and a negligible function ν such that

$$\Pr[\text{leak}_A \text{ is } \nu\text{-close to } \lambda\text{-leaky at } (pk, y_A)] \geq \frac{1}{p}, \quad (2)$$

for some polynomial p and negligible function ν .

The fact that leak_B is μ -close to $(2\lambda - 1)$ -leaky at (pk_1, pk_2, y_B) implies that there exists a distribution $\tilde{S}_{pk_1, pk_2, y_B}^2$ that is μ -close to S_{pk_1, pk_2, y_B}^2 such that

$$\mathbf{H}_\infty(\tilde{S}_{pk_1, pk_2, y_B}^2) \geq \mathbf{H}_\infty(S_{pk_1, pk_2, y_B}^2) - (2\lambda - 1).$$

The fact that

$$\mathbf{H}_\infty(S_{pk_1, pk_2, y_B}^2) = \mathbf{H}_\infty(S_{pk_1}) + \mathbf{H}_\infty(S_{pk_2}),$$

together with the equation above, implies that

$$\mathbf{H}_\infty(\tilde{S}_{pk_1, pk_2, y_B}^2) \geq \mathbf{H}_\infty(S_{pk_1}) + \mathbf{H}_\infty(S_{pk_2}) - (2\lambda - 1) .$$

Consider the distribution

$$(\tilde{sk}_1, \tilde{sk}_2) \leftarrow \tilde{S}_{pk_1, pk_2, y_B}^2.$$

By Lemma A.1 it holds that there exists a set R such that for all $s_1 \in R$,

$$\mathbf{H}_\infty(\tilde{sk}_2 | pk_1, pk_2, y_B, \tilde{sk}_1 = s_1) \geq \mathbf{H}_\infty(S_{pk_2}) - \lambda , \quad (3)$$

and also that if $\Pr[\tilde{sk}_1 \in R] \leq 1/2$ then

$$\mathbf{H}_\infty(\tilde{sk}_1 | pk_1, pk_2, y_B, \tilde{sk}_1 \notin R) \geq \mathbf{H}_\infty(S_{pk_1}) - \lambda . \quad (4)$$

We proceed with a case analysis, where we distinguish between the case that $\Pr[\tilde{sk}_1 \in R]$ is noticeable and the case that it is negligible.

- **Case 1.** There exists a polynomial p such that for infinitely many k 's in K ,

$$\Pr[\widetilde{sk}_1 \in R] \geq \frac{1}{p(k)}.$$

In this case, we prove that with probability $\frac{1}{6p}$ (over the random coin tosses of $\text{leak}_{\mathcal{A}}$), $\text{leak}_{\mathcal{A}}$ is $4p\mu$ -close to λ -leaky on $(pk, y_{\mathcal{A}})$.

To this end, assume that both $\tau = 2$ and $sk_1 \in R$. Note that this happens with probability at least $\frac{1}{2} \left(\frac{1}{p(k)} - \mu \right) \geq \frac{1}{3p(k)}$. This is the case, since these two events are independent, where one which occurs with probability $1/2$ and the other which occurs with probability at least $\frac{1}{p(k)} - \mu$. (The latter follows from the fact that sk_1 and \widetilde{sk}_1 are μ -close.) Under this assumption,

$$y_{\mathcal{A}} = (pk_1, pk_2, y_{\mathcal{B}}, \mathbb{1}_{\tau=2} \cdot sk_1) = (pk_1, pk_2, y_{\mathcal{B}}, sk_1) .$$

Thus, Equation (3) implies that to prove that \mathcal{A} is statistically close to λ -leaky at $(pk, y_{\mathcal{A}})$ it suffices to prove that the distribution of $sk_2|(pk_1, pk_2, y_{\mathcal{B}}, sk_1)$ is statistically close to the distribution $\widetilde{sk}_2|(pk_1, pk_2, y_{\mathcal{B}}, \widetilde{sk}_1 = sk_1)$. We cannot prove this for every sk_1 , however we prove that with probability at least $\frac{1}{2}$ over $s_1 \stackrel{\$}{\leftarrow} sk_1|(pk_1, pk_2, y_{\mathcal{B}})$, the distribution

$$\widetilde{sk}_2|(pk_1, pk_2, y_{\mathcal{B}}, \widetilde{sk}_1 = s_1) ,$$

which we take to be the distribution $\widetilde{sk}|pk, y_{\mathcal{A}}$, is $4p\mu$ -close to the distribution

$$sk_2|(pk_1, pk_2, y_{\mathcal{B}}, sk_1 = s_1) ,$$

which is identical to the distribution $sk|pk, y_{\mathcal{A}}$.

To this end, let

$$B \triangleq \left\{ s_1 \in R : \widetilde{sk}_2|(pk_1, pk_2, y_{\mathcal{B}}, \widetilde{sk}_1 = s_1) \text{ is not } 4p\mu\text{-close to } sk_2|(pk_1, pk_2, y_{\mathcal{B}}, sk_1 = s_1) \right\}$$

We next prove that

$$\Pr[sk_1 \in B] \leq \frac{1}{2p}. \tag{5}$$

This follows from the following calculation.

$$\begin{aligned}
\mu &\geq \frac{1}{2} \sum_{s_1, s_2} \left| \Pr[(sk_1, sk_2) = (s_1, s_2)] - \Pr[(\widetilde{sk}_1, \widetilde{sk}_2) = (s_1, s_2)] \right| \geq \\
&\frac{1}{2} \sum_{s_1 \in B} \sum_{s_2} \left| \Pr[(sk_1, sk_2) = (s_1, s_2)] - \Pr[(\widetilde{sk}_1, \widetilde{sk}_2) = (s_1, s_2)] \right| = \\
&\frac{1}{2} \sum_{s_1 \in B} \sum_{s_2} \left| \Pr[sk_1 = s_1] \cdot \Pr[sk_2 = s_2 \mid sk_1 = s_1] - \Pr[\widetilde{sk}_1 = s_1] \cdot \Pr[\widetilde{sk}_2 = s_2 \mid \widetilde{sk}_1 = s_1] \right| \geq \\
&\frac{1}{2} \sum_{s_1 \in B} \sum_{s_2} \left| \Pr[sk_1 = s_1] \cdot \Pr[sk_2 = s_2 \mid sk_1 = s_1] - \Pr[sk_1 = s_1] \cdot \Pr[\widetilde{sk}_2 = s_2 \mid \widetilde{sk}_1 = s_1] \right| - \\
&\frac{1}{2} \sum_{s_1 \in B} \sum_{s_2} \left| \Pr[\widetilde{sk}_1 = s_1] \cdot \Pr[\widetilde{sk}_2 = s_2 \mid \widetilde{sk}_1 = s_1] - \Pr[sk_1 = s_1] \cdot \Pr[\widetilde{sk}_2 = s_2 \mid \widetilde{sk}_1 = s_1] \right| \geq \\
&\frac{1}{2} \sum_{s_1 \in B} \Pr[sk_1 = s_1] \cdot \sum_{s_2} \left| \Pr[sk_2 = s_2 \mid sk_1 = s_1] - \Pr[\widetilde{sk}_2 = s_2 \mid \widetilde{sk}_1 = s_1] \right| - \\
&\frac{1}{2} \sum_{s_1 \in B} \left| \Pr[\widetilde{sk}_1 = s_1] - \Pr[sk_1 = s_1] \right| \cdot \sum_{s_2} \Pr[\widetilde{sk}_2 = s_2 \mid \widetilde{sk}_1 = s_1] \geq \\
&\frac{1}{2} \sum_{s_1 \in B} \Pr[sk_1 = s_1] \cdot \sum_{s_2} \left| \Pr[sk_2 = s_2 \mid sk_1 = s_1] - \Pr[\widetilde{sk}_2 = s_2 \mid \widetilde{sk}_1 = s_1] \right| - \mu \geq \\
&\sum_{s_1 \in B} \Pr[sk_1 = s_1] \cdot 4p\mu - \mu.
\end{aligned}$$

This implies that

$$\Pr[sk_1 \in B] = \sum_{s_1 \in B} \Pr[sk_1 = s_1] \leq \frac{2\mu}{4p\mu} = \frac{1}{2p},$$

as desired. Equation (5) implies that

$$\Pr[sk_1 \in B \mid sk_1 \in R] = \frac{\Pr[(sk_1 \in B) \wedge (sk_1 \in R)]}{\Pr[sk_1 \in R]} \leq \frac{p}{2p} = \frac{1}{2}.$$

All in all, we get that with probability at least $\frac{1}{6p(k)}$ (which is the probability that $\tau = 2$ and $sk_1 \in R$, and $sk_1 \notin B$), $\text{leak}_{\mathcal{A}}$ is $4p\mu$ -close to λ -leaky. Namely,

$$\Pr[\text{leak}_{\mathcal{A}} \text{ is } 4p\mu\text{-close to } \lambda\text{-leaky at } (pk, y_{\mathcal{A}})] \geq \frac{1}{6p(k)}.$$

This completes Case 1.

- **Case 2.** $\Pr[\widetilde{sk}_1 \in R] = \text{negl}(k)$. This, together with the fact that sk_1 and \widetilde{sk}_1 are statistically indistinguishable, implies that $\Pr[sk_1 \in R] = \text{negl}(k)$. Assume that $\tau = 1$, which occurs with probability $1/2$. In this case,

$$y_{\mathcal{A}} = (pk_1, pk_2, y_{\mathcal{B}}, \mathbb{1}_{\tau=2} \cdot sk_1) = (pk_1, pk_2, y_{\mathcal{B}}, 0).$$

The fact that $\Pr[\widetilde{sk}_1 \in R] = \text{negl}(k)$, together with Equation (4), implies that

$$\mathbf{H}_\infty(\widetilde{sk}_1 | pk_1, pk_2, y_B, \widetilde{sk}_1 \notin R) \geq \mathbf{H}_\infty(S_{pk_1}) - \lambda .$$

Thus, it remains to prove that $\widetilde{sk}_1 | (pk_1, pk_2, y_B, \widetilde{sk}_1 \notin R)$ is statistically close to $sk_1 | (pk_1, pk_2, y_B)$. This follows from the fact that $\widetilde{sk}_1 | (pk_1, pk_2, y_B, \widetilde{sk}_1 \notin R)$ is statistically close to $\widetilde{sk}_1 | (pk_1, pk_2, y_B)$ (which follows from the fact that $\Pr[\widetilde{sk}_1 \in R] = \text{negl}(k)$), and from the fact that $\widetilde{sk}_1 | (pk_1, pk_2, y_B)$ is statistically close to $sk_1 | (pk_1, pk_2, y_B)$.

Therefore, in this case, with probability $1/2$ (which is the probability that $\tau = 1$), leak_A is statistically close to being λ -leaky.

We conclude that in both cases, for infinitely many k 's in K , with probability $1/\text{poly}$, leak_A is statistically close to being λ -leaky at (pk, y_A) , as desired. \blacksquare

The proof of the claim concludes the proof of the lemma. \square

D Proof of Lemma 6.2

Let $\mathcal{B} = (\text{leak}_B, \text{break}_B)$ be a PPT_{-a} adversary for \mathcal{E}^2 . We define a $\text{PPT}_{-(a+1)}$ adversary $\mathcal{A} = (\text{leak}_A, \text{break}_A)$ for \mathcal{E} .

- The function $\text{leak}_A(sk, pk)$ is defined as follows:

1. Sample $\tau \xleftarrow{\$} [2]$ and define $(sk_\tau, pk_\tau) \triangleq (sk, pk)$.
2. For $i = 3 - \tau$, sample $(sk_i, pk_i) \leftarrow G(1^k)$.
3. Compute $y_B \leftarrow \text{leak}_B((sk_1, sk_2), (pk_1, pk_2))$.
4. Compute an additional inefficient ‘‘advice’’ bit, as follows. Consider the random variables $(X, Y) \triangleq S_{pk_1, pk_2, y_B} = (sk_1, sk_2) | (pk_1, pk_2, y_B)$, and consider the set R as defined in Lemma A.1. Our final bit is the indicator of whether $sk_1 \in R$.
5. Output $y_A = (pk_1, pk_2, y_B, \mathbb{1}_{\tau=2} \cdot sk_1, \mathbb{1}_{sk_1 \in R})$.
Where the value $\mathbb{1}_{\tau=2} \cdot sk_1$ means that if $\tau = 2$ then we output sk_1 as part of our leakage, and otherwise this part is null.⁷

- The function $\text{break}_A(pk, y_A)$ will emulate $\text{break}_B(pk_1, pk_2, y_B)$ by playing the role of the verification oracle for $i \neq \tau$ and using the actual verification oracle for $i = \tau$.

It immediately follows that \mathcal{A} is a $\text{PPT}_{-(a+1)}$ adversary. It remains to analyze its success probability. By definition

$$\text{Adv}_{\mathcal{E}, \lambda}[\mathcal{A}] = \Pr[(\text{leak}_A \text{ is } \lambda\text{-leaky at } (pk, y_A)) \wedge (\mathcal{A} \text{ succeeds})] .$$

Let ρ denote the random tape of break_A and the random tape of the verifier in the verification game (we emphasize that ρ does not include the random tape of leak_A). Then, it holds that

⁷Assume w.l.o.g that this null value cannot be a legal secret key and thus this serves as an indication that $\tau = 1$.

$$\begin{aligned} \text{Adv}_{\mathcal{E},\lambda}[\mathcal{A}] &= \Pr [(\text{leak}_{\mathcal{A}} \text{ is } \lambda\text{-leaky at } (pk, y_{\mathcal{A}})) \wedge (\mathcal{A} \text{ succeeds})] = \\ & \mathbb{E}_{pk_1, pk_2, y_{\mathcal{B}}, \rho} \left[\Pr [(\text{leak}_{\mathcal{A}} \text{ is } \lambda\text{-leaky at } (pk, y_{\mathcal{A}})) \wedge (\mathcal{A} \text{ succeeds}) \mid pk_1, pk_2, y_{\mathcal{B}}, \rho] \right] . \end{aligned}$$

Note that conditioning on $(pk_1, pk_2, y_{\mathcal{B}}, \rho)$, the latter probability is only over $\tau \leftarrow [2]$ and over $sk_i \mid pk_1, pk_2, y_{\mathcal{B}}$.

We will prove the following claim.

Claim D.0.2. *For all $(pk_1, pk_2, y_{\mathcal{B}}, \rho)$, it holds that*

$$\begin{aligned} \Pr [(\text{leak}_{\mathcal{A}} \text{ is } \lambda\text{-leaky at } (pk, y_{\mathcal{A}})) \wedge (\mathcal{A} \text{ succeeds}) \mid pk_1, pk_2, y_{\mathcal{B}}, \rho] &\geq \\ (1/4) \cdot \Pr [(\text{leak}_{\mathcal{B}} \text{ is } (2\lambda - 1)\text{-leaky at } (pk_1, pk_2, y_{\mathcal{B}})) \wedge (\mathcal{B} \text{ succeeds}) \mid pk_1, pk_2, y_{\mathcal{B}}, \rho] , \end{aligned}$$

where we envision an invocation of \mathcal{B} that runs in the same probability space as \mathcal{A} : Namely, its key pair is $((sk_1, sk_2), (pk_1, pk_2))$, its leakage value is $y_{\mathcal{B}}$ and its verification game uses the same random tape ρ .

Note that Claim D.0.2 completes the proof of the theorem as it implies that

$$\text{Adv}_{\mathcal{E},\lambda}[\mathcal{A}] \geq (1/4) \cdot \text{Adv}_{\mathcal{E}^2, (2\lambda-1)}[\mathcal{B}] .$$

Proof of Claim D.0.2. Let us fix values for $(pk_1, pk_2, y_{\mathcal{B}}, \rho)$ and keep them fixed throughout the proof (thus omitting the conditioning in the probability brackets). Notice that conditioned on these values,

$$\Pr [(\text{leak}_{\mathcal{B}} \text{ is } (2\lambda - 1)\text{-leaky at } (pk_1, pk_2, y_{\mathcal{B}})) \wedge (\mathcal{B} \text{ succeeds})] \in \{0, 1\} .$$

Namely, the values $(pk_1, pk_2, y_{\mathcal{B}}, \rho)$ completely determine whether $\text{leak}_{\mathcal{B}}$ is $(2\lambda - 1)$ -leaky and whether \mathcal{B} succeeds. Moreover, notice that if the above probability is 0 then the claim holds trivially (since the right-hand side is zero).

We thus focus our attention on values $(pk_1, pk_2, y_{\mathcal{B}}, \rho)$ for which the above probability is 1; i.e., \mathcal{B} succeeds and $\text{leak}_{\mathcal{B}}$ is $(2\lambda - 1)$ -leaky at $(pk_1, pk_2, y_{\mathcal{B}})$.

Since the verification game of \mathcal{A} simulates that of \mathcal{B} using the same randomness, it follows that if \mathcal{B} succeeds in the game then so does \mathcal{A} . Therefore all that is left to prove is that

$$\Pr [\text{leak}_{\mathcal{A}} \text{ is } \lambda\text{-leaky at } (pk, y_{\mathcal{A}})] \geq 1/4 . \quad (6)$$

We now let (X, Y) denote the distribution $S_{pk_1, pk_2, y_{\mathcal{B}}}^2$. Since \mathcal{B} is $(2\lambda - 1)$ -leaky at $(pk_1, pk_2, y_{\mathcal{B}})$, it holds that

$$\mathbf{H}_{\infty}(X, Y) \geq \mathbf{H}_{\infty}(S_{pk_1, pk_2}^2) - (2\lambda - 1) = (\mathbf{H}_{\infty}(S_{pk_1}) - \lambda + 1) + (\mathbf{H}_{\infty}(S_{pk_2}) - \lambda) .$$

Using Lemma A.1, we conclude that there exists a set R such that

$$\mathbf{H}_{\infty}(X \mid X \notin R) \geq \mathbf{H}_{\infty}(S_{pk_1}) - \lambda + 1 + \log \Pr[X \notin R] , \quad (7)$$

and such that for all $x \in R$

$$\mathbf{H}_{\infty}(Y \mid X = x) \geq \mathbf{H}_{\infty}(S_{pk_2}) - \lambda . \quad (8)$$

We proceed with case analysis as to whether R is such that $\Pr[X \in R] \geq 1/2$.

- Assume $\Pr[X \in R] \geq 1/2$. If $\tau = 2$ and $sk_1 \in R$, which are independent events, each of probability at least $1/2$, then,

$$y_{\mathcal{A}} = (pk_1, pk_2, y_{\mathcal{B}}, \mathbb{1}_{\tau=2} \cdot sk_1, \mathbb{1}_{sk_1 \in R}) = (pk_1, pk_2, y_{\mathcal{B}}, sk_1, 1) ,$$

and therefore, by Eq. (8), \mathcal{A} is λ -leaky at $(pk, y_{\mathcal{A}})$.

- Assume $\Pr[X \in R] < 1/2$, then $\log \Pr[X \notin R] \geq -1$. In this case, if $\tau = 1$ and $sk_1 \notin R$, which are again independent with probability $1/2$, it holds that

$$y_{\mathcal{A}} = (pk_1, pk_2, y_{\mathcal{B}}, \mathbb{1}_{\tau=2} \cdot sk_1, \mathbb{1}_{sk_1 \in R}) = (pk_1, pk_2, y_{\mathcal{B}}, 0, 0) ,$$

and therefore, by Eq. (7), \mathcal{A} is λ -leaky at $(pk, y_{\mathcal{A}})$.

We conclude that in both cases, \mathcal{A} is λ -leaky at $(pk, y_{\mathcal{A}})$ with probability at least $1/4$. ■

The proof of the claim concludes the proof of the lemma. □

E Public-Key Encryption as Public-Key Primitive

We show that public-key encryption can be stated in terms of public-key primitives, defined in Section 2.1. The basic idea is that an adversary that succeeds with $1/2 + 1/\text{poly}(k)$ probability in the CPA game, will succeed, given a large enough number of challenges q , on at least $1/2 + 1/\text{poly}(k)$ fraction of them almost always. However, if the scheme is CPA secure then no adversary can succeed on a long enough sequence of inputs beyond beyond the standard deviation. Translating this intuition to an actual game turned out to be more complicated than apparent at first glance.

E.1 Definition and CPA-Security

A public-key encryption scheme is a tuple of PPT algorithms $(\text{Gen}, \text{Enc}, \text{Dec})$ such that

- $\text{Gen}(1^k)$ is the *key generation* algorithm that takes the security parameter as input, and outputs a key pair (sk, pk) .
- $\text{Enc}(pk, m)$ is the *encryption algorithm* that takes as input the public key pk and a message m taken from a set M which is the *message space* of the scheme. It outputs a ciphertext c .
- $\text{Dec}(sk, c)$ is the *decryption algorithm* that takes the secret key sk and a ciphertext c , and outputs a message $m' \in M$.

The correctness requirement is that for all $m \in M$ it holds that

$$\Pr[\text{Dec}(sk, \text{Enc}(pk, m)) \neq m] = \text{negl}(k),$$

where the probability is over the key generation and the randomness of Enc, Dec .

Definition E.1 (CPA-security). *Let $(\text{Gen}, \text{Enc}, \text{Dec})$ be a public-key encryption scheme and let \mathcal{A} be some adversary. We define the CPA-experiment conducted between a challenger and the adversary \mathcal{A} as follows:*

1. The challenger draws $(sk, pk) \leftarrow \text{Gen}(1^k)$ and sends pk to \mathcal{A} .
2. \mathcal{A} chooses two messages of equal length $m_0, m_1 \in M$ and sends them to the challenger.
3. The challenger samples a random bit $b \xleftarrow{\$} \{0, 1\}$ and sends $c_b := \text{Enc}(pk, m_b)$ to \mathcal{A} .
4. \mathcal{A} decides on a “guess” $b^* \in \{0, 1\}$ and sends it to the challenger.

The scheme is CPA-secure if for any polynomial time \mathcal{A} it holds that

$$\Pr[b = b^*] \leq 1/2 + \text{negl}(k) ,$$

where the probability is over all the randomness in the experiment.

We note that the CPA experiment can be conducted even when the challenger has no knowledge of sk (one can think of pk as a common input to the challenger and to \mathcal{A}).

E.2 Public-Key Primitive

Let $(\text{Gen}, \text{Enc}, \text{Dec})$ be a public key encryption scheme, and consider the following public-key primitive $\mathcal{E} = (G, V)$:

- The key generator G is identical to Gen . Namely, $G(1^k)$ runs $(sk, pk) \leftarrow \text{Gen}(1^k)$ and outputs (sk, pk) .
- The verifier V , given a public key pk runs as follows. Recall that V is conducting a verification game with its oracle $\mathcal{O}(pk)$, who knows the public key pk .
 1. \mathcal{O} sends two messages $m_0, m_1 \in M$, and a unary value 1^q to V .
 2. V asks \mathcal{O} to answer $128 \cdot q^3 \cdot k$ challenges as follows. For all $i \in [\ell]$, where $\ell = 128 \cdot q^2 \cdot k$, and for all $j \in [q]$, V samples a bit $b_{i,j} \xleftarrow{\$} \{0, 1\}$ and computes $c_{i,j} \leftarrow \text{Enc}(pk, m_{b_{i,j}})$. It sends $\{c_{i,j}\}_{(i,j) \in [\ell] \times [q]}$ to \mathcal{O} .
 3. The oracle \mathcal{O} returns a sequence of bits $\{b_{i,j}^*\}_{(i,j) \in [\ell] \times [q]}$.
 4. For all $i \in [\ell]$, we let W_i be an indicator variable for the event where for strictly more than $q/2$ values of j , it holds that $b_{i,j}^* = b_{i,j}$. Note that W_i can be efficiently computed by V for all i .
 5. After computing all values W_1, \dots, W_ℓ , the verifier outputs 1 if and only if all of them are true. Namely if $\bigwedge_{i \in [\ell]} W_i$.

We note that for any polynomial time oracle, the above experiment runs in polynomial time, simply because q cannot be more than \mathcal{O} 's running time.

We now argue that \mathcal{E} is a secure public-key primitive if and only if $(\text{Gen}, \text{Enc}, \text{Dec})$ is CPA-secure.

Lemma E.1. *If $\mathcal{E} = (G, V)$ is a secure public-key primitive then $(\text{Gen}, \text{Enc}, \text{Dec})$ is CPA-secure.*

Proof. Assume towards contradiction that there exists a (non-uniform) probabilistic polynomial time adversary \mathcal{A} such that $\Pr[b^* = b] \geq 1/2 + \epsilon$, where $\epsilon = \epsilon(k)$ is non-negligible. We assume without loss of generality that \mathcal{A} is deterministic. We can assume this by simply fixing the “best” randomness for \mathcal{A} .

We define an adversary \mathcal{B} , that uses \mathcal{A} to break the security of the game of \mathcal{E} , as follows.

1. Upon receiving pk , \mathcal{B} sends pk to \mathcal{A} and gets back messages m_0, m_1 , which are forwarded to the verifier V . \mathcal{B} sets $q \leftarrow O(k/\epsilon^2)$ (see the analysis for the exact value of q), and sends 1^q to V .
2. For any $c_{i,j}$ throughout the different phases of the attack, \mathcal{B} does the same: It computes $b_{i,j}^*$ as \mathcal{A} 's response to the ciphertext $c_{i,j}$. Note that \mathcal{B} rewinds \mathcal{A} to the point in its execution that is right before receiving the challenge ciphertext.

To analyze the performance of \mathcal{B} we first note that with probability at least $\epsilon/2$ over pk , it holds that

$$\Pr[b^* = b|pk] \geq 1/2 + \epsilon/2.$$

Conditioned on this event, it holds by the Chernoff bound that taking $O(k/\epsilon^2)$ samples, the probability that the number of successes deviates by more than $\epsilon/2$ from the expectation is at most 2^{-k} . We set q to be this value. [**Yael's note: I would add in a footnote the precise Chernoff bound that we are using.**]

Applying the union bound over all ℓ executions of the basic procedure, we have that \mathcal{B} wins in the breaking game with probability at least $\frac{\epsilon}{2} \cdot (1 - \ell \cdot 2^{-k})$ which is non-negligible. \square

Lemma E.2. *If $(\text{Gen}, \text{Enc}, \text{Dec})$ is CPA-secure then $\mathcal{E} = (G, V)$ is a secure public-key primitive.*

Proof. Assume towards contradiction that there exists a (non-uniform) probabilistic polynomial time adversary \mathcal{A} that breaks \mathcal{E} with probability ϵ , where $\epsilon = \epsilon(k)$ is non-negligible. As in the proof of Lemma E.1 we assume without loss of generality that \mathcal{A} is deterministic. We define an adversary \mathcal{B} for the CPA game as follows. Upon receiving pk , \mathcal{B} simulates \mathcal{A} to obtain m_0, m_1 and 1^q and forwards m_0, m_1 to the challenger.

In addition, \mathcal{B} tests the conditional success probability of \mathcal{A} given pk and the random tape ρ of \mathcal{A} . Namely, \mathcal{B} computes a binary variable S (that corresponds to a probabilistic event) such that both $\Pr[S] \geq \epsilon/4$ and $\Pr[\mathcal{A} \text{ succeeds} | S] \geq \epsilon/2$. Essentially, S is computed by simulating the verification game many times and estimating the success probability up to (say) $\epsilon/10$ accuracy with exponentially low estimation error. The variable S denotes the event where the estimated success probability is more than (say) $2\epsilon/3$. The first property follows since \mathcal{A} succeeds with probability $2\epsilon/3$ with probability at least $\epsilon/3$ over the values of pk, ρ . The second property follows from the estimation accuracy. We omit further details but emphasize that S depends only on pk, ρ .

If the event S does not happen, then \mathcal{B} outputs a random guess $b^* \xleftarrow{\$} \{0, 1\}$ regardless of the challenge it receives. From this point and on, we describe the case where S indeed occurs.

Upon receiving a challenge ciphertext c from the challenger, \mathcal{B} randomly samples $i^* \xleftarrow{\$} [\ell]$ and $j^* \xleftarrow{\$} [q]$. It then simulates the entire breaking game of \mathcal{A} , playing the role of the verifier V , and setting $c_{i^*, j^*} := c$ (all other $c_{i,j}$ are generated appropriately according to the breaking game).

After the breaking game is complete, \mathcal{B} can compute W_i for all $i \neq i^*$. It proceeds by computing $W^* \triangleq \bigwedge_{i < i^*} W_i$, and outputting $b^* := b_{i^*, j^*}^*$ if W^* holds, and a random value $b^* \xleftarrow{\$} \{0, 1\}$ otherwise.

To analyze the performance of \mathcal{B} , we first notice that with probability at least $\epsilon/2$ over pk , the success probability of \mathcal{A} is at least $\epsilon/2$. Let us condition on this case, bearing in mind that it is correctly identified by \mathcal{B} with all but 2^{-k} probability.

A successful break of \mathcal{A} means that the event $\bigwedge_{i \in [\ell]} W_i$ indeed occurs. Therefore in our case

$$\epsilon/2 \leq \Pr \left[\bigwedge_{i \in [\ell]} W_i \right] = \prod_{i \in [\ell]} \Pr \left[W_i \mid \bigwedge_{h < i} W_h \right],$$

where the probability is only over the generation of the ciphertexts $c_{i,j}$ (all other randomness has been fixed).

Denoting $\delta_i \triangleq \Pr[W_i \mid \bigwedge_{h < i} W_h]$, we have that $\epsilon/2 \leq \prod_{i \in [\ell]} \delta_i$. We claim that this implies that most δ_i 's are in fact very close to 1.

Claim E.2.1. *Let $\epsilon/2 \leq \prod_{i \in [\ell]} \delta_i$, where $\epsilon, \delta_1, \dots, \delta_\ell \in [0, 1]$. Then for all $t \geq 2$, it holds that*

$$\Pr_{i \stackrel{\$}{\leftarrow} [\ell]} [\delta_i < 1 - 1/t] < 2t \cdot \ln(4/\epsilon) / \ell.$$

Proof. Let I be the set of i 's such that δ_i is larger than the threshold and assume towards contradiction that $|I| \geq 2t \ln(4/\epsilon)$. Then it holds that

$$\prod_{i \in [\ell]} \delta_i \leq \prod_{i \in I} \delta_i < (1 - 1/t)^{|I|} \leq e^{-|I|/(2t)} \leq \epsilon/4,$$

in contradiction to $\epsilon/2 \leq \prod_{i \in [\ell]} \delta_i$. ■

We proceed in the analysis of the success probability of \mathcal{B} (recall that we are still conditioned on pk):

$$\begin{aligned} \Pr[b^* = b] &= \Pr \left[\neg \bigwedge_{i < i^*} W_i \right] \cdot \frac{1}{2} + \Pr \left[b_{i^*,j^*} = b_{i^*,j^*}^* \mid \bigwedge_{i < i^*} W_i \right] \cdot \Pr \left[\bigwedge_{i < i^*} W_i \right] \\ &= \frac{1}{2} + \Pr \left[\bigwedge_{i < i^*} W_i \right] \cdot \left(\Pr \left[b_{i^*,j^*} = b_{i^*,j^*}^* \mid \bigwedge_{i < i^*} W_i \right] - \frac{1}{2} \right). \end{aligned}$$

In addition

$$\begin{aligned} \Pr \left[b_{i^*,j^*} = b_{i^*,j^*}^* \mid \bigwedge_{i < i^*} W_i \right] &\geq \Pr \left[b_{i^*,j^*} = b_{i^*,j^*}^* \mid W_{i^*}, \bigwedge_{i < i^*} W_i \right] - \Pr[\neg W_{i^*} \mid \bigwedge_{i < i^*} W_i] \\ &\geq \frac{1}{2} + 1/(2q) - \mathbb{E}_{i^* \stackrel{\$}{\leftarrow} [\ell]} [1 - \delta_{i^*}] \\ &\geq \frac{1}{2} + 1/(2q) - \frac{2t \ln(4/\epsilon)}{\ell} - 1/t, \end{aligned}$$

for all $t \geq 2$. Where the last equation follows from Claim E.2.1. Taking $t = 8q$, we get that $\frac{2t \ln(4/\epsilon)}{\ell} = \frac{\ln(4/\epsilon)}{8q \cdot k} \leq 1/(8q)$ as well (since $\ln(4/\epsilon) = O(\log k) \leq k$), and therefore

$$\Pr \left[b_{i^*,j^*} = b_{i^*,j^*}^* \mid \bigwedge_{i < i^*} W_i \right] \geq \frac{1}{2} + 1/(2q) - 1/(8q) - 1/(8q) = \frac{1}{2} + 1/(4q).$$

Plugging into the above equation, this implies that

$$\Pr[b^* = b] = \frac{1}{2} + \Pr \left[\bigwedge_{i < i^*} W_i \right] \cdot \frac{1}{4q} \geq \frac{1}{2} + \frac{\epsilon}{8q},$$

since $\Pr \left[\bigwedge_{i < i^*} W_i \right] \geq \Pr \left[\bigwedge_{i \in [\ell]} W_i \right] \geq \epsilon/2$.

Finally, we compute the total success probability of \mathcal{B} in the CPA experiment. Recalling the $\Pr[S] \geq \epsilon/4$, we have that

$$\begin{aligned} \Pr[b = b^*] &= \Pr[\neg S] \cdot \frac{1}{2} + \Pr[S] \cdot \Pr[b = b^* | S] \\ &= \frac{1}{2} + \Pr[S] \cdot \left(\Pr[b = b^* | S] - \frac{1}{2} \right) \\ &\geq \frac{1}{2} + \Pr[S] \cdot \frac{\epsilon}{8q} \\ &\geq \frac{1}{2} + \frac{\epsilon^2}{32q}. \end{aligned}$$

Since q is bounded by the absolute running time of \mathcal{A} , it holds that $\frac{\epsilon^2}{32q}$ is non-negligible and the result follows. \square