Unifying Leakage Models: From Probing Attacks to Noisy Leakage

Alexandre Duc

University of Applied Sciences and Arts Western Switzerland (HES-SO/HEIG-VD),
Yverdon-les-Bains, Switzerland

Stefan Dziembowski

University of Warsaw, Warsaw, Poland
S.Dziembowski@crypto.edu.pl

Sebastian Faust

TU Darmstadt, Darmstadt, Germany

Communicated by Kenneth G. Paterson.

Received 17 January 2015 / Revised 19 February 2018

Online publication 5 June 2018

Abstract. A recent trend in cryptography is to formally show the leakage resilience of cryptographic implementations in a given leakage model. One of the most prominent leakage model—the so-called bounded leakage model—assumes that the amount of leakage that an adversary receives is a-priori bounded. Unfortunately, it has been pointed out by several works that the assumption of bounded leakages is hard to verify in practice. A more realistic assumption is to consider that leakages are sufficiently noisy, following the engineering observation that real-world physical leakages are inherently perturbed by physical noise. While already the seminal work of Chari et al. (in: CRYPTO, pp 398–412, 1999) study security of side-channel countermeasures in the noisy model, only recently Prouff and Rivain (in: Johansson T, Nguyen PQ (eds) EUROCRYPT, volume 7881 of lecture notes in 931 computer science, pp 142–159, Springer, 2013) offer a full formal analysis of the masking countermeasure in a physically motivated noise model. In particular, the authors show that a block-cipher implementation that uses the Boolean masking scheme is secure against a very general class of noisy leakage functions. While this is an important step toward better understanding the security of masking schemes, the analysis of Prouff and Rivain has several shortcomings including in particular requiring leak-free gates. In this work, we provide an alternative security proof in the same noise model that overcomes these challenges. We achieve this goal by a new reduction from noisy leakage to the important model of probing adversaries (Ishai et al. in: CRYPTO, pp 463–481, 2003). This reduction is the main technical contribution of our work that significantly simplifies the formal security analysis of masking schemes against realistic side-channel leakages.

S. Dziembowski: Received founding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013)/ERC Grant Agreement Number 207908. Part of this work was done while Stefan Dziembowski was on leave from Sapienza University of Rome, Italy. S. Faust: Partially funded by the Emmy Noether Program FA 1320/1-1 of the German Research Foundation (DFG).

Keywords. Leakage-resilient cryptography, Noisy leakage, Probing attacks.

1. Introduction

Physical side-channel attacks that exploit leakage emitting from devices are an important threat for cryptographic implementations. Prominent sources of such physical leakages include the running time of an implementation [20], its power consumption [21] or electromagnetic radiation emitting from it [30]. A large body of recent applied and theoretical research attempts to incorporate the information an adversary obtains from the leakage into the security analysis and develops countermeasures to defeat common side-channel attacks [2, 6, 12, 17, 24, 34, 35]. While there is still a large gap between what theoretical models can achieve and what side-channel information is measured in practice, some recent important works propose models that go better in line with the perspective of cryptographic engineering [28, 33, 34]. Our work follows this line of research by analyzing the security of a common countermeasure—the so-called masking countermeasure—in the model of Prouff and Rivain [28]. Our analysis works by showing that security in certain theoretical leakage models implies security in the model of [28] and hence may be seen as a first attempt to *unify* the large class of different leakage models used in recent results.

The Masking Countermeasure. A large body of work on cryptographic engineering has developed countermeasures to defeat side-channel attacks (see, e.g., [22] for an overview). While many countermeasures are specifically tailored to protect particular cryptographic implementations (e.g., key updates or shielded hardware), a method that generically works for most cryptographic schemes is *masking* [4, 16, 27, 35]. The basic idea of a masking scheme is to secretly share all sensitive information, including the secret key and all intermediate values that depend on it, thereby making the leakage independent of the secret data. The most prominent masking scheme is the Boolean masking: a bit b is encoded by a random bit string $(b_1, \dots, b_n)$ such that $b = b_1 \oplus \dots \oplus b_n$. The main difficulty in designing masking schemes is to develop masked operations, which securely compute on encoded data and ensure that *all* intermediate values are protected.

Masking Against Noisy Leakages. Besides the fact that masking can be used to protect arbitrary computation, it has the advantage that it can be analyzed in formal security models. The first work that formally studies the soundness of masking in the presence of leakage is the seminal work of Chari et al. [6]. The authors consider a model where each share b_i of an encoding is perturbed by Gaussian noise and show that the number of noisy samples needed to recover the encoded secret bit b grows exponentially with the number of shares. As stated in [6], this model matches real-world physical leakages that inherently are noisy. Moreover, many practical solutions exist to amplify leakage noise (see for instance the works of [7, 8, 22]).

One limitation of the security analysis given in [6] is the fact that it does not consider leakage emitting from masked computation. This shortcoming has been addressed in the recent important work of Prouff and Rivain [28], who extend at Eurocrypt 2013 the noisy leakage model of Chari et al. [6] to also include leakage from the masked operations. Specifically, they show that a variant of the construction of Ishai et al. [17] is secure even when there is noisy leakage from all the intermediate values that are produced during

the computation. The authors of [28] also generalize the noisy leakage model of Chari et al. [6] to a wider range of leakage functions instead of considering only the Gaussian one. While clearly noisy leakage is closer to physical leakage occurring in real world, the security analysis of [28] has a number of shortcomings which puts strong limitations in which settings the masking countermeasure can be used and achieves the proved security statements. In particular, like earlier works on leakage resilient cryptography [10, 13] the security analysis of Prouff and Rivain relies on so-called *leak-free* gates. Moreover, security is shown in a restricted adversarial model that assumes that plaintexts are chosen uniformly during an attack and the adversary does not exploit joint information from the leakages and, e.g., the ciphertext. We discuss these shortcomings in more detail in the next section.

1.1. The Work of Prouff and Rivain [28]

Prouff and Rivain [28] analyze the security of a block-cipher implementation that is masked with an additive masking scheme working over a finite field $\mathbb{F}$. More precisely, let t be the security parameter then a secret $s \in \mathbb{F}$ is represented by an encoding $(X_1, \dots, X_t)$ such that each $X_i \leftarrow \mathbb{F}$ is uniformly random subject to $s = X_1 \oplus \dots \oplus X_t$. As discussed above the main difficulty in designing secure masking schemes is to devise masked operations that work on masked values. To this end, Prouff and Rivain use the original scheme of Ishai et al. [17] augmented with some techniques from [5, 31] to work over larger fields and to obtain a more efficient implementation. The masked operations are built out of several smaller components. First, a leak-free operation that refreshes encodings, i.e., it takes as input an encoding $(X_1, \dots, X_t)$ of a secret s and outputs a freshly and independently chosen encoding of the same value. Second, a number of leaky elementary operations that work on a constant number of field elements. For each of these elementary operations the adversary is given leakage $f(X)$, where X are the inputs of the operation and f is a noisy function. Clearly, the noise level has to be high enough so that given $f(X)$ the values of X is not completely revealed. To this end, the authors introduce the notion of a *bias*, which informally says that the statistical distance between the distribution of X and the conditional distribution $X|f(X)$ is bounded by some parameter.

While noisy leakages are certainly a step in the right direction to model physical leakage, we detail below some of the limitations of the security analysis of Prouff and Rivain [28]:

1. *Leak-Free Components.* The assumption of leak-free computation has been used in earlier works on leakage resilient computation [10, 13]. It is a strong assumption on the physical hardware and, as stated in [28], an important limitation of the current proof approach. The leak-free component of [28] is a simple operation that takes as input an encoding and refreshes it. While the computation of this operation is supposed to be completely shielded against leakage, the inputs and the outputs of this computation may leak. Notice that the leak-free component of [28] depends on the computation that is carried out in the circuit by taking inputs. In particular, this means that the computation of the leak-free component depends on secret information, which makes it harder to protect in practice and is different from earlier works that use leak-free components [10, 13].

2. *Random Message Attacks.* The security analysis is given only for random (known) message attacks. This is in contrast to most works in cryptography, which usually consider at least a chosen message attack. Hence, the proof does not cover chosen plaintext or chosen ciphertext attacks. We notice, however, that it is not clear whether chosen message attacks can improve the adversary's success probability in the context of DPA attacks [36].
3. *Mutual-Information-Based Security Statement.* The final statement of Theorem 4 in [28] only gives a bound on the mutual information of the key and the leakages from the cipher. While a mutual information analysis is a common method in side-channel analysis to evaluate the security of countermeasures [33], it has important shortcomings such as not including information that an adversary may learn from exploiting joint information from the leakages and plaintext/ciphertext pairs. Notice that such use of mutual information gets particularly problematic under continuous leakage attacks, since multiple plaintext/ciphertext pairs *information theoretically completely reveal* the secret key. The more standard security notion used in cryptography and also for the analysis of masking schemes, e.g., in the work of Ishai et al., uses a simulation-based approach and does not have these drawbacks.

1.2. Our Contribution

We show in this work how to eliminate limitations 1–3 by a simple and elegant simulation-based argument and a reduction to the so-called t -probing adversarial setting [17] (that in this paper we call the *t -threshold-probing model* to emphasize the difference between this model and the *random-probing model* defined later). The t -threshold-probing model considers an adversary that can learn the value of t intermediate values that are produced during the computation and is often considered as a good approximation for modeling higher-order attacks. We notice that limitation 4 from above is what enables our security analysis. The fact that the noise is independent for each elementary operation allows us to formally prove security under an *identical* noise model as [28], but using a simpler and improved analysis. In particular, we are able to show that the original construction of Ishai et al. satisfies the standard simulation-based security notion under noisy leakages without relying on any leak-free components. We emphasize that our techniques are very different (and much simpler) than the recent breakthrough result of Goldwasser and Rothblum [15] who show how to eliminate leak-free gates in the bounded leakage model. We will further discuss related works in Sect. 1.3.

Our proof considers three different leakage models and shows connections between them. One may view our work as a first attempt to “reduce” the number of different leakage models, which is in contrast to many earlier works that introduced new leakage settings. Eventually, we are able to reduce the security in the noisy leakage model to the security in the t -threshold-probing model. This shows that, for the particular choice of parameters given in [28], security in the t -threshold-probing model implies security in the noisy leakage model. This goes in line with the common approach of showing security against t -order attacks, which usually requires to prove security in the t -threshold-probing model. Moreover, it shows that the original construction of Ishai et al. that has been used in many works on masking (including the work of Prouff and Rivain) is indeed

a sound approach for protecting against side-channel leakages when assuming that they are sufficiently noisy. We give some more details on our techniques below.

From Noisy Leakages to Random Probes. As a first step in our security proof we show that we can simulate any adversary in the noisy leakage model of Prouff and Rivain with an adversary in a simpler noise model that we name a *random-probing adversary* and is similar to a model introduced in [17]. In this model, an adversary recovers an intermediate value with probability ϵ and obtains a special symbol $\perp$ with probability $1 - \epsilon$. This reduction shows that this model is worth studying, although from the engineering perspective it may seem unnatural.

From Random Probes to the t -Threshold-Probing Model. We show how to go from the random-probing adversary setting to the more standard t -threshold-probing adversary of Ishai et al. [17]. This step is rather easy as due to the independency of the noise we can apply Chernoff’s bound almost immediately. One technical difficulty is that the work of Prouff and Rivain considers joint noisy leakage from elementary operations, while the standard t -threshold-probing setting only talks about leakage from wires. Notice, however, that the elementary operations of [28] only depend on two inputs and, hence, it is not hard to extend the result of Ishai et al. to consider “gate probing adversary” by tolerating a loss in the parameters. Finally, our analysis enables us to show security of the masking based countermeasure without the limitations 1–3 discussed above.

Leakage Resilient Circuits with Simulation-Based Security. In our security analysis we use the framework of leakage resilient circuits introduced in the seminal work of Ishai et al. [17]. A circuit compiler takes as input the description of a cryptographic scheme C with secret key K , e.g., a circuit that describes a block cipher, and outputs a transformed circuit C' and corresponding key K' . The circuit $C'[K']$ shall implement the same functionality as C running with key K , but additionally is resilient to certain well-defined classes of leakage. Notice that while the framework of [17] talks about circuits the same approach applies to software implementations, and we only follow this notation to abstract our description.

Moreover, our work uses the well-established simulation paradigm to state the security guarantees we achieve. Intuitively, simulation-based security says that whatever attack an adversary can carry out when knowing the leakage; he can also run (with similar success probability) by just having black-box access to C . In contrast to the approach based on Shannon information theory, our analysis includes attacks that exploit joint information from the leakage and plaintext/ciphertext pairs. It seems impossible to us to incorporate the plaintext/ciphertext pairs into an analysis based on Shannon information theory. To see this, consider a block-cipher execution, where, clearly, when given a couple of plaintext/ciphertext pairs, the secret key is information theoretically revealed.¹ The authors of [28] are well aware of this problem and explicitly exclude such joint information. A consequence of the simulation-based security analysis is that we require an additional mild assumption on the noise—namely, that it is efficiently computable

¹ More concretely: imagine an adversary that attacks a block-cipher implementation E_K , where K is the secret key. Then just by launching a known-plaintext attack he can obtain several pairs $V = (M_0, E_K(M_0)), (M_1, E_K(M_1)), \dots$. Clearly a small number of such pairs is usually enough to determine K *information theoretically*. Hence it makes no sense to require that “ K is information-theoretically hidden given V and the side-channel leakage.”

(see Sect. 3.1 for more details). While this is a standard assumption made in most works on leakage resilient cryptography, we emphasize that we can easily drop the assumption of efficiently computable noise (and hence considering the same noise model as [28]), when we only want to achieve the weaker security notion considered in [28]. Notice that in this case we are still able to eliminate the limitations 1 and 2 mentioned above.

1.3. Related Work

Masking & Leakage Resilient Circuits. A large body of work has proposed various masking schemes and studies their security in different security models (see, e.g., [4, 16, 27, 31, 35]). The already mentioned t -threshold-probing model has been considered in the work of Rivain and Prouff [31], who show how to extend the work of Ishai et al. to larger fields and propose efficiency improvements. In [29] it was shown that techniques from multiparty computation can be used to show security in the t -threshold-probing model. The work of Standaert et al. [35] studies masking schemes using the information theoretic framework of [33] by considering the Hamming weight model. Many other works analyze the security of the masking countermeasure, and we refer the reader for further details to [28].

With the emergence of leakage resilient cryptography [2, 12, 24] several works have proposed new security models and alternative masking schemes. The main difference between these new security models and the t -threshold-probing model is that they consider *joint leakages* from large parts of the computation. The work of Faust et al. [13] extends the security analysis of Ishai et al. beyond the t -threshold-probing model by considering leakages that can be described by low-depth circuits (so-called AC^0 leakages). Faust et al. use leak-free component that have been eliminated by Rohtblum in [32] using computational assumptions. The recent work of Miles and Viola [25] proposes a new circuit transformation using alternating groups and shows security with respect to AC^0 and TC^0 leakages.

Another line of work considers circuits that are provably secure in the so-called continuous bounded leakage model [10, 14, 15, 18]. In this model, the adversary is allowed to learn arbitrary information from the computation of the circuit as long as the amount of information is bounded. The proposed schemes rely additionally on the assumption of “only computation leaks information” of Micali and Reyzin [24].

Noisy Leakage Models. The work of Faust et al. [13] also considers circuit compilers for noisy models. Specifically, they propose a construction with security in the binomial noise model, where each value on a wire is flipped independently with probability $p \in (0, 1/2)$. In contrast to the work of [28] and our work the noise model is restricted to binomial noise, but the noise rate is significantly better (constant instead of linear noise). Similar to [28] the work of Faust et al. also uses leak-free components. Besides these works on masking schemes, several works consider noisy leakages for concrete cryptographic schemes [12, 19, 26]. Typically, the noise model considered in these works is significantly stronger than the noise model that is considered for masking schemes. In particular, no strong assumption about the independency of the noise is made.

2. Preliminaries

We start with some standard definitions and lemmas about the statistical distance. If $\mathcal{A}$ is a set, then $U \leftarrow \mathcal{A}$ denotes a random variable sampled uniformly from $\mathcal{A}$. Recall that if A and B are random variables over the same set $\mathcal{A}$ then the *statistical distance between A and B* is denoted as $\Delta(A; B)$, and defined as $\Delta(A; B) = \frac{1}{2} \sum_{a \in \mathcal{A}} |\mathbb{P}(A = a) - \mathbb{P}(B = a)| = \sum_{a \in \mathcal{A}} \max\{0, \mathbb{P}(A = a) - \mathbb{P}(B = a)\}$. If $\mathcal{X}, \mathcal{Y}$ are some events, then by $\Delta((A|\mathcal{X}); (B|\mathcal{Y}))$ we will mean the distance between variables A' and B' , distributed according to the conditional distributions $P_{A|\mathcal{X}}$ and $P_{B|\mathcal{Y}}$. If $\mathcal{X}$ is an event of probability 1, then we also write $\Delta(A; (B|\mathcal{Y}))$ instead of $\Delta((A|\mathcal{X}); (B|\mathcal{Y}))$. If C is a random variable, then by $\Delta(A; (B|C))$ we mean $\sum \mathbb{P}(C = c) \cdot \Delta(A; (B|(C = c)))$.

If A, B , and C are random variables, then $\Delta((B; C) | A)$ denotes $\Delta((BA); (CA))$. It is easy to see that it is equal to $\sum_a \mathbb{P}(A = a) \cdot \Delta((B|A = a); (C|A = a))$. If $\Delta(A; B) \leq \epsilon$, then we say that A and B are ϵ -close. The “ $\stackrel{d}{=}$ ” symbol denotes the equality of distributions, i.e., $A \stackrel{d}{=} B$ if and only if $\Delta(A; B) = 0$.

2.1. Basic Probability-Theoretic Facts

Here, we state some basic lemmas that will be used later in some proofs.

Lemma 1. *Let A, B be two (possibly correlated) random variables. Let B' be a variable distributed identically to B but independent from A . We have*

$$\Delta(A; (A|B)) = \Delta((B; B') | A). \quad (1)$$

Proof. We have

$$\begin{aligned} \Delta(A; (A|B)) &= \sum_b \frac{1}{2} \cdot \mathbb{P}(B = b) \cdot \sum_a |\mathbb{P}(A = a) - \mathbb{P}(A = a|B = b)| \\ &= \frac{1}{2} \sum_{a,b} |\mathbb{P}(B = b) \cdot \mathbb{P}(A = a) - \mathbb{P}(B = b) \cdot \mathbb{P}(A = a|B = b)| \\ &= \frac{1}{2} \sum_{a,b} |\mathbb{P}(B' = b \wedge A = a) - \mathbb{P}(B = b \wedge A = a)| \\ &= \Delta((B; B') | A), \end{aligned} \quad (2)$$

where in (2) we used the fact that B' is a variable distributed identically to B and it is independent from A . $\square$

Lemma 2. *For any random variables A and B and an event $\mathcal{E}$ we have*

$$\Delta((A | \neg \mathcal{E}); B) \leq \Delta(A; B) + \mathbb{P}(\mathcal{E}),$$

where $\neg \mathcal{E}$ denotes the negation of $\mathcal{E}$.

Proof. We have

$$\begin{aligned}
 & \Delta((A \mid \neg\mathcal{E}); B) \\
 &= \frac{1}{2} \sum_a |\mathbb{P}(A = a \mid \neg\mathcal{E}) - \mathbb{P}(B = a)| \\
 &\leq \underbrace{\frac{1}{2} \sum_a |\mathbb{P}(A = a) - \mathbb{P}(A = a \mid \neg\mathcal{E})|}_{=:(*)} + \underbrace{\frac{1}{2} \sum_a |\mathbb{P}(A = a) - \mathbb{P}(B = a)|}_{=\Delta(A; B)}, \quad (3)
 \end{aligned}$$

where (3) comes from the triangle inequality. Now, (*) is equal to

$$\sum_{a \in \mathcal{A}} \mathbb{P}(A = a) - \mathbb{P}(A = a \mid \neg\mathcal{E}), \quad (4)$$

where $\mathcal{A}$ is the set of all a 's such that $\mathbb{P}(A = a) \geq \mathbb{P}(A = a \mid \neg\mathcal{E})$. Clearly we have that $\mathbb{P}(A = a \mid \neg\mathcal{E}) \geq \mathbb{P}(A = a \wedge \neg\mathcal{E})$, and hence (4) is at most equal to $\mathbb{P}(A = a) - \mathbb{P}(A = a \wedge \neg\mathcal{E})$. We therefore have

$$\begin{aligned}
 & \sum_{a \in \mathcal{A}} \mathbb{P}(A = a) - \mathbb{P}(A = a \wedge \neg\mathcal{E}) \\
 &= \mathbb{P}(A \in \mathcal{A}) - \mathbb{P}(A \in \mathcal{A} \wedge \neg\mathcal{E}) \\
 &\leq \mathbb{P}(\mathcal{E})
 \end{aligned}$$

Thus, altogether, (4) is at most equal to $\mathbb{P}(\mathcal{E}) + \Delta(A; B)$. This finishes the proof. $\square$

We will also need the following standard fact.

Lemma 3. (Chernoff bound, see, e.g., [9], Theorem 1.1) *Let $Z = \sum_{i=1}^n Z_i$, where Z_i 's are random variables independently distributed over $[0, 1]$. Then for every $\xi \in [0, 1]$ we have*

$$\mathbb{P}(Z \geq (1 + \xi)\mathbb{E}(Z)) \leq \exp\left(-\frac{\xi^2}{3}\mathbb{E}(Z)\right).$$

3. Noise from Set Elements

We start with describing the basic framework for reasoning about the noise from elements of a finite set $\mathcal{X}$. Later, in Sect. 4, we will consider the leakage from the vectors over $\mathcal{X}$, and then, in Sect. 5, from the entire computation. The reason why we can smoothly use the analysis from Sect. 3.1 in the later sections is that, as in the work of Prouff and Rivain, we require that the noise is independent for all elementary operations. By elementary operations, [28] considers the basic underlying operations over the underlying field $\mathcal{X}$ used in a masked implementation. In this work, we consider the same setting and type of underlying operations (in fact, notice that our construction is identical to theirs—except

that we eliminate the leak-free gates and prove a stronger statement). Notice that instead of talking about elementary operations, we consider the more standard term of “gates” that was used in the work of Ishai et al. [17].

3.1. Modeling Noise

Let us start with a discussion defining what it means that a randomized function $Noise : \mathcal{X} \rightarrow \mathcal{Y}$ is “noisy”. We will assume that $\mathcal{X}$ is finite and rather small: typical choices for $\mathcal{X}$ would be $\text{GF}(2)$ (the “Boolean case”), or $\text{GF}(2^8)$, if we want to deal with the AES circuit. The set $\mathcal{Y}$ corresponds to the set of all possible noise measurements and may be infinite, except when we require the “efficient simulation” (we discuss it further at the end of this section). As already informally described in Sect. 1.1, our basic definition is as follows: we say that the function $Noise$ is δ -noisy if

$$\delta = \Delta(X; (X|Noise(X))). \quad (5)$$

Of course for (5) to be well-defined we need to specify the distribution of X . The idea to define noisy functions by comparing the distributions of X and “ X conditioned on $Noise(X)$ ” comes from [28], where it is argued that the most natural choice for X is a random variable distributed uniformly over $\mathcal{X}$. We also adopt this convention and assume that $X \leftarrow \mathcal{X}$. We would like to stress, however, that in our proofs we will apply $Noise$ to inputs $\hat{X}$ that are not necessarily uniform and in this case the value of $\Delta(\hat{X}; (\hat{X}|Noise(\hat{X})))$ may obviously be some non-trivial function of δ . Of course if $X \leftarrow \mathcal{X}$ and $X' \leftarrow \mathcal{X}$, then $Noise(X')$ is distributed identically to $Noise(X)$, and hence, by Lemma 1, Eq. (5) is equivalent to:

$$\delta = \Delta((Noise(X); Noise(X')) | X), \quad (6)$$

where X and X' are uniform over $\mathcal{X}$. Note that at the beginning this definition may be a bit counter-intuitive, as *smaller* δ means *more* noise: in particular we achieve “full noise” if $\delta = 0$, and “no noise” if $\delta \approx 1$. Let us compare this definition with the definition of [28]. In a nutshell: the definition of [28] is similar to ours, the only difference being that instead of the statistical distance Δ in [28] the authors use a distance based on the Euclidean norm. More precisely, they start with defining d as: $d(X; Y) := \sqrt{\sum_{x \in \mathcal{X}} (\mathbb{P}(X = x) - \mathbb{P}(Y = y))^2}$, and using this notion they define β as:

$$\beta(X|Noise(X)) := \sum_{y \in \mathcal{Y}} \mathbb{P}(Noise(X) = y) \cdot d(X; (X|Noise(X) = y))$$

(where X is uniform). In the terminology of [28] a function $Noise$ is “ δ -noisy” if $\delta = \beta(X|Noise(X))$. Observe that the right hand side of our noise definition in Eq. (5) can be rewritten as:

$$\sum_{y \in \mathcal{Y}} \mathbb{P}(Noise(X) = y) \cdot \Delta(X; (X|Noise(X) = y)),$$

hence the only difference between their approach and ours is that we use Δ where they use the distance d . The authors do not explain why they choose this particular measure. We believe that our choice to use the standard definition of statistical distance Δ is more natural in this setting, since, unlike the “ d ” distance, it has been used in hundreds of cryptographic papers in the past. The popularity of the Δ distance comes from the fact that it corresponds to an intuitive concept of the “indistinguishability of distributions”—it is well-known, and simple to verify, that $\Delta(X; Y) \leq \delta$ if and only if no adversary can distinguish between X and Y with advantage better than δ .² Hence, e.g., (6) can be interpreted as:

δ is the maximum probability, over all adversaries $\mathcal{A}$, that $\mathcal{A}$ distinguishes between the noise from a uniform X that is *known to him*, and a uniform X' that is *unknown to him*.

It is unclear to us if a d distance has a similar interpretation. We emphasize, however, that the choice whether to use Δ or β is not too important, as the following inequalities between these measures hold for every X and Y distributed over $\mathcal{X}$ (cf. [28]):

$$\frac{1}{2} \cdot d(X; Y) \leq \Delta(X; Y) \leq \frac{\sqrt{|\mathcal{X}|}}{2} \cdot d(X; Y),$$

and consequently

$$\frac{1}{2} \cdot \beta(X|Noise(X)) \leq \Delta(X; (X|Noise(X))) \leq \frac{\sqrt{|\mathcal{X}|}}{2} \cdot \beta(X|Noise(X)). \quad (7)$$

Hence, we decide to stick to the “ Δ distance” in this paper. However, to allow for comparison between our work and the one of [28] we will at the end of the paper present our results also in terms of the β measure. (This translation will be straightforward, thanks to the inequalities in (7).) In [28] (cf. Theorem 4) the result is stated in form of Shannon information theory. While such an information theoretic approach may be useful in certain settings [33], we follow the more “traditional” approach and provide an efficient simulation argument. As discussed in the introduction, this also covers a setting where the adversary exploits joint information of the leakage and, e.g., the plaintext/ciphertext pairs. We emphasize, however, that our results can easily be expressed in the information theoretic language, thanks to the following bound: $I(A; B) \leq (2N / \ln 2) \cdot \Delta(A; (A|B))$, where A is uniformly distributed over a set of cardinality N . This result comes from Proposition 1 in [28] combined with the inequality $\beta(A|B) \leq 2\Delta(A; (A|B))$, cf. (7).

3.1.1. The Issue of “Efficient Simulation”

To achieve the strong simulation-based security notion, we need an additional requirement on the leakage, namely, that the leakage can efficiently be “simulated”—which typically requires that the noise function is efficiently computable. In fact, for our proofs to go through we actually need something slightly stronger, namely that *Noise is efficiently decidable* by which we mean that (a) there exists a randomized poly-time algorithm that

² This formally means that for every $\mathcal{A}$ we have $|\mathbb{P}(\mathcal{A}(X) = 1) - \mathbb{P}(\mathcal{A}(Y) = 1)| \leq \delta$.

computes it, and (b) the set $\mathcal{Y}$ is finite and for every x and y the value of $\mathbb{P}(\text{Noise}(x) = y)$ is computable in polynomial time. While (b) may look like a strong assumption we note that in practice for most “natural” noise functions (like the Gaussian noise with a known parameter, measured with a very good, but finite, precision) it is easily satisfiable.

Recall that the results of [28] are stated without taking into consideration the issue of the “efficient simulation”. Hence, if one wants to compare our results with [28], then one can simply drop the efficient decidability assumption on the noise. To keep our presentation concise and clean, also in this case the results will be presented in a form “for every adversary $\mathcal{A}$ there exists an (inefficient) simulator $\mathcal{S}$ ”. Here the “inefficient simulator” can be an arbitrary machine, capable, e.g., of sampling elements from *any* probability distributions.

3.2. Simulating Noise by ϵ -Identity Functions

Lemma 4 below is our main technical tool. Informally, it states that every δ -noisy function $\text{Noise} : \mathcal{X} \rightarrow \mathcal{Y}$ can be represented as a composition $\text{Noise}' \circ \varphi$ of efficiently computable randomized functions Noise' and φ , where φ is a “ $\delta \cdot |\mathcal{X}|$ -identity function”, defined in Definition 1 below.

Definition 1. A randomized function $\varphi : \mathcal{X} \rightarrow \mathcal{X} \cup \{\perp\}$ is an ϵ -identity if for every x we have that either $\varphi(x) = x$ or $\varphi(x) = \perp$ and $\mathbb{P}(\varphi(x) \neq \perp) = \epsilon$.

This will allow us to reduce the “noisy attacks” to the “random-probing attacks”, where the adversary learns each wire (or a gate, see Sect. 5.5) of the circuit with probability ϵ . Observe also, that thanks to the assumed independence of noise, the events that the adversary learns each element are independent, which, in turn, will allow us to use the Chernoff bound to prove that with a good probability the number of wires that the adversary learns is small.

Lemma 4. *Let $\text{Noise} : \mathcal{X} \rightarrow \mathcal{Y}$ be a δ -noisy function. Then there exist $\epsilon \leq \delta \cdot |\mathcal{X}|$ and a randomized function $\text{Noise}' : \mathcal{X} \cup \{\perp\} \rightarrow \mathcal{X}$ such that for every $x \in \mathcal{X}$ we have*

$$\text{Noise}(x) \stackrel{d}{=} \text{Noise}'(\varphi(x)), \quad (8)$$

where $\varphi : \mathcal{X} \rightarrow \mathcal{X} \cup \{\perp\}$ is the ϵ -identity function. Moreover, if Noise is efficiently decidable then $\text{Noise}'(\varphi(x))$ is computable in time that is expected polynomial in $|\mathcal{X}|$.

Before we proceed to the proof let us remark that the “ $|\mathcal{X}|$ factor” loss in this lemma (when going from δ to ϵ) is in general unavoidable. More concretely, in the subsequent work [11] (Sect. 5) it is shown that there exist δ -noisy functions that can be reduced (in the sense of Lemma 4) to an ϵ -identity function only if ϵ is (at least) approximately equal to $\delta \cdot |\mathcal{X}|/2$.

Proof of Lemma 4. We consider only the case when Noise is efficiently decidable, and hence the Noise' function that we construct will be efficiently computable. The case when Noise is not efficiently decidable is handled in an analogous way (the proof is actually

simpler as the only difference is that we do not need to argue about the efficiency of the sampling algorithms).

Very informally speaking, our proof is based on an extension of the standard observation that for any two random variables A and B one can find two events $\mathcal{A}$ and $\mathcal{B}$ such that the distributions $P_{A|\mathcal{A}}$ and $P_{B|\mathcal{B}}$ are equal and $\mathbb{P}(\mathcal{A}), \mathbb{P}(\mathcal{B}) = \Delta(A; B)$ (see, e.g., [23, Section 1.3]).

Let X and X' be uniform over $\mathcal{X}$. For every $y \in \mathcal{Y}$ define

$$\pi(y) = \min_{x \in \mathcal{X}} (\mathbb{P}(\text{Noise}(x) = y)).$$

Clearly π is computable in time polynomial in $|\mathcal{X}|$. Obviously π is usually not a probability distribution since it does not sum up to 1. The good news is that it sums up “almost” to 1 provided δ is sufficiently small. This is shown below. Let $\epsilon := 1 - \sum_{y \in \mathcal{Y}} \pi(y)$. We now have

$$\begin{aligned} \epsilon &= \overbrace{\sum_{y \in \mathcal{Y}} \mathbb{P}(\text{Noise}(X') = y)}^{=1} - \sum_{y \in \mathcal{Y}} \pi(y) \\ &= \sum_{y \in \mathcal{Y}} \mathbb{P}(\text{Noise}(X') = y) - \min_{x \in \mathcal{X}} (\mathbb{P}(\text{Noise}(x) = y)) \\ &= \sum_{y \in \mathcal{Y}} \max_{x \in \mathcal{X}} (\mathbb{P}(\text{Noise}(X') = y) - \mathbb{P}(\text{Noise}(x) = y)) \\ &\leq \sum_{y \in \mathcal{Y}} \sum_{x \in \mathcal{X}} \max(0, \mathbb{P}(\text{Noise}(X') = y) - \mathbb{P}(\text{Noise}(x) = y)) \end{aligned} \quad (9)$$

$$\begin{aligned} &= \sum_{x \in \mathcal{X}} \Delta(\text{Noise}(x); \text{Noise}(X')) \\ &= |\mathcal{X}| \cdot \Delta((\text{Noise}(X); \text{Noise}(X')) \mid X) \\ &= \delta \cdot |\mathcal{X}|, \end{aligned} \quad (10)$$

where (9) comes from the fact that the maximum of positive values cannot be larger than their sum³, and (10) follows from the assumption that the *Noise* function is δ -noisy. Our construction of the *Noise'* function is based on the standard technique of *rejection sampling*. Let *Noise'*(x) be a distribution defined as follows: for every $y \in \mathcal{Y}$ and every $x \neq \perp$ let:

$$\mathbb{P}(\text{Noise}'(x) = y) = (\mathbb{P}(\text{Noise}(x) = y) - \pi(y)) / \epsilon, \quad (11)$$

³ More precisely, for every $\{Z_x\}_{x \in \mathcal{X}}$ we have:

$$\begin{aligned} \max_{x \in \mathcal{X}} (Z_x) &\leq \sum_{x: Z_x \geq 0} Z_x \\ &= \sum_x \max(0, Z_x), \end{aligned}$$

where in our case $Z_x := \mathbb{P}(\text{Noise}(x) = y)$.

and otherwise:

$$\mathbb{P}(\text{Noise}'(\perp) = y) = \pi(y)/(1 - \epsilon). \quad (12)$$

We will later show how to sample Noise' efficiently. Obviously this will automatically imply that (11) and (12) define probability distributions over $\mathcal{Y}$ (which may not be obvious at the first sight). First, however, let us show (8). To this end take any $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ and observe that

$$\begin{aligned} \mathbb{P}(\text{Noise}'(\varphi(x)) = y) &= \mathbb{P}(\varphi(x) = x) \cdot \mathbb{P}(\text{Noise}'(x) = y) + \mathbb{P}(\varphi(x) = \perp) \cdot \mathbb{P}(\text{Noise}'(\perp) = y) \\ &= \epsilon \cdot (\mathbb{P}(\text{Noise}(x) = y) - \pi(y))/\epsilon + (1 - \epsilon) \cdot \pi(y)/(1 - \epsilon) \\ &= \mathbb{P}(\text{Noise}(x) = y) - \pi(y) + \pi(y) \\ &= \mathbb{P}(\text{Noise}(x) = y). \end{aligned}$$

Which implies (8). What remains is to show how to sample Noise' efficiently. Let us first show an efficient algorithm $\text{Alg}_1(x)$ for computing $\text{Noise}'(x)$ for $x \neq \perp$:

$\text{Alg}_1(x)$:

1. Sample y from $\text{Noise}(x)$.
2. With probability $\pi(y)/\mathbb{P}(\text{Noise}(x) = y)$ resample y , i.e.: go back to Step 1.
3. Output y .

We now argue that $\text{Alg}_1(x)$ indeed computes $\text{Noise}'(x)$ efficiently. Let $R_1 \in \{1, 2, \dots\}$ be a random variable denoting the number of times the algorithm $\text{Alg}_1(x)$ performed Step 1. First observe that the probability of jumping back to Step 1 in Step 2 is equal to

$$\sum_y \mathbb{P}(\text{Noise}(x) = y) \cdot \pi(y) / \mathbb{P}(\text{Noise}(x) = y) = \sum_y \pi(y) \quad (13)$$

$$= 1 - \epsilon \quad (14)$$

Therefore the probability of *not* jumping back to Step 1 in Step 2 is ϵ , and hence the expected number $\mathbb{E}(R_1)$ of the executions of Step 1 in $\text{Alg}_1(x)$ is equal to $\sum_{i=1}^n i \cdot (1 - \epsilon)^{i-1} \cdot \epsilon = 1/\epsilon$. Moreover for every $i = 0, 1, \dots$ we have:

$$\begin{aligned} \mathbb{P}(\text{Alg}_1(x) = y \wedge R_1 = i \mid R_1 \geq i) &= \mathbb{P}(\text{Noise}(x) = y) \cdot (1 - (\pi(y)/\mathbb{P}(\text{Noise}(x) = y))) \\ &= \mathbb{P}(\text{Noise}(x) = y) - \pi(y) \end{aligned}$$

Hence

$$\begin{aligned} \mathbb{P}(\text{Alg}_1(x) = y) &= \sum_{i=0}^{\infty} \mathbb{P}(\text{Alg}_1(x) = y \wedge R_1 = i) \\ &= \sum_{i=0}^{\infty} \mathbb{P}(\text{Alg}_1(x) = y \wedge R_1 = i \mid R_1 \geq i) \cdot \mathbb{P}(R_1 \geq i) \end{aligned}$$

$$\begin{aligned}
&= (\mathbb{P}(\text{Noise}(x) = y) - \pi(y)) \cdot \sum_{i=1}^{\infty} \mathbb{P}(R_1 \geq i) \\
&= (\mathbb{P}(\text{Noise}(x) = y) - \pi(y)) \cdot \mathbb{E}(R_1) \\
&= (\mathbb{P}(\text{Noise}(x) = y) - \pi(y))/\epsilon,
\end{aligned}$$

as required in (11). We now present an efficient algorithm Alg_2 for computing $\text{Noise}'(\perp)$. Fix an arbitrary element $x_0 \in \mathcal{X}$, and execute the following.

Alg₂ :

1. Sample y from $\text{Noise}(x_0)$.
2. With probability $1 - (\pi(y)/\mathbb{P}(\text{Noise}(x_0) = y))$ resample y , i.e.: go back to Step 1.
3. Output y .

By a similar argument as in the case of Alg_1 we obtain that the expected number R_2 of times the algorithm Alg_2 performs Step 1 is equal to $\mathbb{E}(R_2) = 1/(1 - \epsilon)$. Moreover for every $i = 1, 2, \dots$ we have:

$$\mathbb{P}(\text{Alg}_2 = b \wedge R_2 = i \mid R_2 \geq i) = \pi(y),$$

which, in turn, implies that $\mathbb{P}(\text{Alg}_2(x) = y) = \pi(y)/(1 - \epsilon)$, and hence the output of Alg_2 satisfies (12). Clearly, the expected running time of both algorithms is polynomial in $|\mathcal{X}|$ and $\mathbb{E}(R)$, where R is the number of execution of Step 1 in Alg_1 or Alg_2 . We obviously have

$$\begin{aligned}
\mathbb{E}(R) &= \mathbb{E}(R_1 \mid \varphi(x) \neq \perp) \cdot \mathbb{P}(\varphi(x) \neq \perp) + \mathbb{E}(R_2 \mid \varphi(x) = \perp) \cdot \mathbb{P}(\varphi(x) = \perp) \\
&= (1/\epsilon) \cdot \epsilon + (1/(1 - \epsilon)) \cdot (1 - \epsilon) \\
&= 2.
\end{aligned}$$

Hence, the expected running time of $\text{Noise}'(\varphi(x))$ is polynomial in $|X|$. $\square$

4. Leakage from Vectors

In this section we describe the leakage models relevant to this paper. We start with describing the models abstractly, by considering leakage from an arbitrary sequence $(x_1, \dots, x_\ell) \in \mathcal{X}^\ell$, where $\mathcal{X}$ is some finite set and ℓ is a parameter. The adversary $\mathcal{A}$ will be able to obtain some partial information about $(x_1, \dots, x_\ell)$ via the games described below. Note that we do not specify the computational power of $\mathcal{A}$, as the definitions below make sense for both computationally bounded and infinitely powerful $\mathcal{A}$.

Noisy Model. For $\delta \geq 0$ a δ -noisy adversary on $\mathcal{X}^\ell$ is a machine $\mathcal{A}$ that plays the following game against an oracle that knows $(x_1, \dots, x_\ell) \in \mathcal{X}^\ell$:

1. $\mathcal{A}$ specifies a sequence $\{\text{Noise}_i : \mathcal{X} \rightarrow \mathcal{Y}\}_{i=1}^\ell$ of noisy functions such that every Noise_i is δ'_i -noisy, for some $\delta'_i \leq \delta$ and mutually independent noises.

2. $\mathcal{A}$ receives $Noise_1(x_1), \dots, Noise_\ell(x_\ell)$ and outputs some value $out_{\mathcal{A}}(x_1, \dots, x_\ell)$.

If $\mathcal{A}$ works in polynomial time and the noise functions specified by $\mathcal{A}$ are efficiently decidable, then we say that $\mathcal{A}$ is *poly-time-noisy*.

Random-Probing Model. For $\epsilon \leq 0$ a ϵ -random-probing adversary on $\mathcal{X}^\ell$ is a machine $\mathcal{A}$ that plays the following game against an oracle that knows $(x_1, \dots, x_\ell) \in \mathcal{X}^\ell$:

1. $\mathcal{A}$ specifies a sequence $(\epsilon_1, \dots, \epsilon_\ell)$ such that each $\epsilon_i \leq \epsilon$.
2. $\mathcal{A}$ receives $\varphi_1(x_1), \dots, \varphi_\ell(x_\ell)$ and outputs some value $out_{\mathcal{A}}(x_1, \dots, x_\ell)$, where each φ_i is the ϵ_i -identity function with mutually independent randomness.

A similar model was introduced in the work of Ishai, Sahai and Wagner [17] to obtain a circuit compiler that blows up the size of the circuit linearly in the security parameter d . Also, the work of Ajtai [1] considers the random-probing model and constructs a compiler that for sufficiently large security parameter d achieves security in the random-probing model for a small (but constant) probability ϵ . [1], however, does not give concrete parameters for ϵ and d , and circuits produced by the compiler of [1] result into a huge circuit size blow-up ($O(d^4)$ with large hidden constants).

Threshold-Probing Model. For $t = 0, \dots, \ell$ a t -threshold-probing adversary on $\mathcal{X}^\ell$ is a machine $\mathcal{A}$ that plays the following game against an oracle that knows $(x_1, \dots, x_\ell) \in \mathcal{X}^\ell$:

1. $\mathcal{A}$ specifies a set $\mathcal{I} = \{i_1, \dots, i_{|\mathcal{I}|}\} \subseteq \{1, \dots, \ell\}$ of cardinality at most t ,
2. $\mathcal{A}$ receives $(x_{i_1}, \dots, x_{i_{|\mathcal{I}|}})$ and outputs some value $out_{\mathcal{A}}(x_1, \dots, x_\ell)$.

4.1. Simulating the Noisy Adversary by a Random-Probing Adversary

The following lemma shows that every δ -noisy adversary can be simulated by a $\delta \cdot |\mathcal{X}|$ -random-probing adversary.

Lemma 5. *Let $\mathcal{A}$ be a δ -noisy adversary on $\mathcal{X}^\ell$. Then there exists a $\delta \cdot |\mathcal{X}|$ -random-probing adversary $\mathcal{S}$ on $\mathcal{X}^\ell$ such that for every $(x_1, \dots, x_\ell)$ we have*

$$out_{\mathcal{A}}(x_1, \dots, x_\ell) \stackrel{d}{=} out_{\mathcal{S}}(x_1, \dots, x_\ell). \quad (15)$$

Moreover, if $\mathcal{A}$ is poly-time-noisy, then $\mathcal{S}$ works in time polynomial in $|\mathcal{X}|$.

Proof. Without loss of generality assume that $\mathcal{A}$ simply outputs all the information that he gets. Thus (15) can be rewritten as:

$$(Noise_1(x_1), \dots, Noise_\ell(x_\ell)) \stackrel{d}{=} out_{\mathcal{S}}(x_1, \dots, x_\ell), \quad (16)$$

where $Noise_i$'s are the δ_i -noisy functions chosen by $\mathcal{A}$. By Lemma 4 for each i there exists $\epsilon_i \leq \delta_i \cdot |\mathcal{X}| \leq \delta \cdot |\mathcal{X}|$ and a randomized function $Noise'_i : \mathcal{X} \cup \{\perp\} \rightarrow \mathcal{X}$, such that for every $x \in \mathcal{X}$ we have

$$Noise_i(x) \stackrel{d}{=} Noise'_i(\varphi_i(x)), \quad (17)$$

where $\varphi_i : \mathcal{X}_i \rightarrow \mathcal{X}_i \cup \{\perp\}$ is the ϵ_i -identity function and $Noise'_i(\varphi_i(x))$ is computable in time polynomial in $|\mathcal{X}|$. We now describe the actions of $\mathcal{S}$. The sequence that he specifies is $(\epsilon_1, \dots, \epsilon_\ell)$. After receiving $(y_1, \dots, y_\ell)$ (equal to $(\varphi_1(x_1), \dots, \varphi_\ell(x_\ell))$) he outputs

$$out(x_1, \dots, x_\ell) := (Noise'_1(y_1), \dots, Noise'_\ell(y_\ell))$$

(this clearly takes time that is expected polynomial in $\ell \cdot |\mathcal{X}|$). We now have

$$\begin{aligned} & (Noise'_1(y_1), \dots, Noise'_\ell(y_\ell)) \\ & \stackrel{d}{=} (Noise'_1(\varphi_1(x_1)), \dots, Noise'_\ell(\varphi_\ell(x_\ell))) \\ & \stackrel{d}{=} (Noise_1(x_1), \dots, Noise_\ell(x_\ell)) \end{aligned} \tag{18}$$

where (18) comes from (17). This implies (16) and hence it finishes the proof. $\square$

Intuitively, this lemma easily follows from Lemma 4 applied independently to each element of $(x_1, \dots, x_\ell)$.

4.2. Simulating the Random-Probing Adversary by a Threshold-Probing Adversary

In this section we show how to simulate every δ -random-probing adversary by a threshold adversary. This simulation, unlike the one in Sect. 4, will not be perfect in the sense that the distribution output by the simulator will be identical to the distribution of the original adversary only when conditioned on some event that happens with a large probability. We start with the following lemma, whose proof is a straightforward application of the Chernoff bound.

Lemma 6. *Let $\mathcal{A}$ be an ϵ -random-probing adversary on $\mathcal{X}^\ell$. Then there exists a $(2\epsilon\ell - 1)$ -threshold-probing adversary $\mathcal{S}$ on $\mathcal{X}^\ell$ operating in time linear in the working time of $\mathcal{A}$ such that for every $(x_1, \dots, x_\ell)$ we have*

$$\Delta(out_{\mathcal{A}}(x_1, \dots, x_\ell) ; out_{\mathcal{S}}(x_1, \dots, x_\ell) \mid out_{\mathcal{S}}(x_1, \dots, x_\ell) \neq \perp) = 0, \tag{19}$$

where

$$\mathbb{P}(out_{\mathcal{S}}(x_1, \dots, x_\ell) = \perp) \leq \exp\left(-\frac{\epsilon\ell}{3}\right). \tag{20}$$

Proof. As in the proof of Lemma 5 we assume that the simulated adversary $\mathcal{A}$ outputs all the information that he received. Moreover, since for $\epsilon' \leq \epsilon$ every ϵ' -identity function φ' can be simulated by the ϵ -identity function φ ,⁴ hence we can assume that each ϵ_i

⁴ Just set $\varphi'(x) := \varphi(x)$ with probability ϵ'/ϵ , and $\varphi'(x) = \perp$ otherwise. Then clearly $\mathbb{P}(\varphi'(x) = x) = \epsilon \cdot \epsilon'/\epsilon = \epsilon'$.

specified by $\mathcal{A}$ is equal to ϵ . Thus, we need to show a $2\epsilon\ell$ -threshold-probing simulator $\mathcal{S}$ such that for every $(x_1, \dots, x_\ell) \in \mathcal{X}^\ell$ we have

$$\Delta(\varphi_1(x_1), \dots, \varphi_\ell(x_\ell) ; \text{out}_{\mathcal{S}}(x_1, \dots, x_\ell) \mid \text{out}_{\mathcal{S}}(x_1, \dots, x_\ell) \neq \perp) = 0, \quad (21)$$

(where each φ_i is the ϵ -identity function) and (20) holds. The simulator $\mathcal{S}$ proceeds as follows. First he chooses a sequence $(Z_1, \dots, Z_\ell)$ of independent random variables in the by setting, for each i :

$$Z_i := \begin{cases} 1 & \text{with probability } \epsilon_i \\ 0 & \text{otherwise.} \end{cases}$$

Let Z denote the number of Z_i 's equal to 1, i.e., $Z := \sum_{i=1}^\ell Z_i$. If $Z \geq 2\ell\epsilon$, then $\mathcal{S}$ outputs $\perp$. Otherwise, he specifies the set $\mathcal{I}$ as $\mathcal{I} := \{i : Z_i = 1\}$. He receives $(x_{i_1}, \dots, x_{i_{|\mathcal{I}|}})$. For all the remaining i 's (i.e., those not in the set $\mathcal{I}$) the simulator sets $x_i := \perp$. He outputs $(x_1, \dots, x_\ell)$. It is straightforward to see that $\mathcal{S}$ is $(2\ell\epsilon - 1)$ -threshold-probing and that (21) holds. What remains is to show (20). Since $\mathbb{E}(Z) = \epsilon\ell$,

$$\begin{aligned} \mathbb{P}(Z \geq 2\ell\epsilon) &= \mathbb{P}(Z \geq 2\mathbb{E}(Z)) \\ &\leq \exp\left(-\frac{\epsilon\ell}{3}\right), \end{aligned} \quad (22)$$

where (22) comes from the Chernoff bound with $\xi = 1$ (cf. Lemma 3). This finishes the proof. $\square$

The following corollary combines Lemmas 5 and 6 together, and will be useful in the sequel.

Corollary 1. *Let $d, \ell \in \mathbb{N}$ with $\ell > d$ and let $\mathcal{A}$ be a $d/(4\ell \cdot |\mathcal{X}|)$ -noisy adversary on $\mathcal{X}^\ell$. Then there exists an $(d/2 - 1)$ -threshold-probing adversary $\mathcal{S}$ such that*

$$\Delta(\text{out}_{\mathcal{A}}(x_1, \dots, x_\ell) ; \text{out}_{\mathcal{S}}(x_1, \dots, x_\ell) \mid \text{out}_{\mathcal{S}}(x_1, \dots, x_\ell) \neq \perp) = 0 \quad (23)$$

and $\mathbb{P}(\text{out}_{\mathcal{S}}(x_1, \dots, x_\ell) = \perp) \leq \exp(-d/12)$. Moreover, if $\mathcal{A}$ is poly-time-noisy then $\mathcal{S}$ works in time polynomial $\ell \cdot |\mathcal{X}|$.

Proof. By Lemma 5 there exists a $d/(4\ell)$ -random-probing adversary $\mathcal{A}'$ whose output is distributed identically to the output of $\mathcal{A}$. In turn, by Lemma 6 for $t = 2 \cdot (d/(4\ell)) \cdot \ell = d/2$ there exists a $(t - 1)$ -threshold-probing adversary $\mathcal{S}$ whose output, conditioned on not being equal to $\perp$, is distributed identically to the output of $\mathcal{A}'$, and such that $\mathbb{P}(\text{out}_{\mathcal{S}}(x_1, \dots, x_\ell) = \perp) \leq \exp(-d/12)$.

If $\mathcal{A}$ is poly-time noisy then clearly the expected working time of $\mathcal{A}'$ is polynomial in $\ell \cdot |\mathcal{X}|$. Since the working time of $\mathcal{S}$ is linear in the working time of $\mathcal{A}$ hence this finishes the proof. $\square$

5. Leakage from Computation

In this section we address the main topic of this paper, which is the noise-resilience of cryptographic computations. Our main model will be the model of arithmetic circuits over a finite field. First, in Sect. 5.1 we present our security definitions, and then, in Sect. 5.2 we describe a secure “compiler” that transforms any cryptographic scheme secure in the “black-box” model into one secure against the noisy leakage (it is essentially identical to the transformation of [17] later extended in [31]). Finally, in the last section we present our security results.

5.1. Definitions

A (*stateful arithmetic*) circuit Γ over a field $\mathbb{F}$ is a directed graph whose nodes are called *gates*. Each gate γ can be of one of the following types: an *input gate* γ^{inp} of fan-in zero, an *output gate* γ^{out} of fan-out zero, a *random gate* γ^{rand} of fan-in zero, a *multiplication gate* $\gamma^\times$ of fan-in 2, an *addition gate* γ^+ of fan-in 2, a *subtraction gate* γ^- of fan-in 2, a *constant gate* γ^{const} , and a *memory gate* γ^{mem} of fan-in 1. Following [17] we assume that the fan-out of every gate is at most 3. The only cycles that are allowed in Γ must contain exactly 1 memory gate. The *size* $|\Gamma|$ of the circuit Γ is defined to be the total number of its gates. The numbers of input gates, output gates and memory gates will be denoted $|\Gamma.inp|$, $|\Gamma.out|$, and $|\Gamma.mem|$, respectively.

The computation of Γ is performed in several “rounds” numbered $1, 2, \dots$. In each of them the circuit will take some input, produce an output and update the memory state. Initially, the memory gates of Γ are preloaded with some initial “state” $k_0 \in \mathbb{F}^{|\Gamma.mem|}$. At the beginning of the i th round the input gates are loaded with elements of some vector $a_i \in \mathbb{F}^{|\Gamma.inp|}$ called the *input for the i th round*. The computation of Γ in the i th round depends on a_i and on the memory state k_{i-1} . It proceeds in a straightforward way: if all the input wires of a given gate are known then the value on its output wire can be computed naturally: if γ is a multiplication gate with input wires carrying values a and b , then its output wire will carry the value $a \cdot b$ (where “ $\cdot$ ” is the multiplication operation in $\mathbb{F}$), and the addition and the subtraction gates are handled analogously. We assume that the random gates produce a fresh random field element in each round. The *output of the i th round* is read-off from the output gates and denoted $b_i \in \mathbb{F}^{|\Gamma.out|}$. The *state after the i th round* is contained in the memory gates and denoted k_i . For $k \in \mathbb{F}^{|\Gamma.mem|}$ and a sequence of inputs $(a_1, \dots, a_m)$ (where each $a_i \in \mathbb{F}^{|\Gamma.inp|}$) let $\Gamma(k, a_1, \dots, a_m)$ denote the sequence $(B_1, \dots, B_m)$ where each B_i is the output of Γ with $k_0 = k$ and inputs $a_1, \dots, a_m$ in rounds $1, 2, \dots$. Observe that, since Γ is randomized, hence $\Gamma(k, a_1, \dots, a_m)$ is a random variable.

A *black-box circuit adversary* $\mathcal{A}$ is a machine that adaptively interacts with a circuit Γ via the input and output interface. Then $out \left(\mathcal{A} \stackrel{bb}{\rightleftharpoons} \Gamma(k) \right)$ denotes the output of $\mathcal{A}$ after interacting with Γ whose initial memory state is $k_0 = k$. A δ -noisy circuit adversary $\mathcal{A}$ is an adversary that has the following additional ability: after each i th round $\mathcal{A}$ gets some partial information about the internal state of the computation via the noisy leakage functions. More precisely: let $(X_1, \dots, X_\ell)$ be the random variable denoting the values on the wires of $\Gamma(k)$ in the i th round. Then $\mathcal{A}$ plays the role of a δ -noisy

adversary in a game against $(X_1, \dots, X_\ell)$ (c.f. Sect. 4), namely: he chooses a sequence $\{Noise_i : \mathbb{F} \rightarrow \mathcal{Y}\}_{i=1}^\ell$ of functions such that every $Noise_i$ is δ_i -noisy for some $\delta_i \leq \delta$ and he receives $Noise_1(X_1), \dots, Noise_\ell(X_\ell)$. Let $out \left(\mathcal{A} \stackrel{noisy}{\rightleftharpoons} \Gamma(k) \right)$ denote the output of such an $\mathcal{A}$ after interacting with Γ whose initial memory state is $k_0 = k$.

We can also replace, in the above definition, the “ δ -noisy adversary” with the “ ϵ -random probing adversary”. In this case, after each i th round $\mathcal{A}$ chooses a sequence $(\epsilon_1, \dots, \epsilon_\ell)$ such that each $\epsilon_i \leq \epsilon$ and he learns $\varphi_1(X_1), \dots, \varphi_\ell(X_\ell)$, where each φ_i is the ϵ_i -identity function. Let $out \left(\mathcal{A} \stackrel{rnd}{\rightleftharpoons} \Gamma(k) \right)$ denote the output of such $\mathcal{A}$ after interacting with Γ whose initial memory state is $k_0 = k$.

Analogously we can replace the “ δ -noisy adversary” with the “ t -threshold probing adversary” obtaining an adversary that after each i th round $\mathcal{A}$ learns t elements of $X_1, \dots, X_\ell$. Let $out \left(\mathcal{A} \stackrel{thr}{\rightleftharpoons} \Gamma(k) \right)$ denote the output of such $\mathcal{A}$ after interacting with Γ whose initial memory state is $k_0 = k$.

Definition 2. Consider two stateful circuits Γ and Γ' (over some field $\mathbb{F}$) and a randomized encoding function Enc . We say that Γ' is a (δ, ξ) -noise-resilient implementation of a circuit Γ w.r.t. Enc if the following holds for every $k \in \mathbb{F}^{| \Gamma.inp |}$:

1. the input-output behavior of $\Gamma(k)$ and $\Gamma'(Enc(k))$ is identical, i.e.: for every sequence of inputs $a_1, \dots, a_m$ and outputs $b_1, \dots, b_m$ we have

$$\begin{aligned} \mathbb{P}(\Gamma(k, a_1, \dots, a_m) = (b_1, \dots, b_m)) \\ = \mathbb{P}(\Gamma'(Enc(k), a_1, \dots, a_m) = (b_1, \dots, b_m)) \end{aligned}$$

and

2. for every δ -noisy circuit adversary $\mathcal{A}$ there exists a black-box circuit adversary $\mathcal{S}$ such that

$$\Delta \left(out \left(\mathcal{S} \stackrel{bb}{\rightleftharpoons} \Gamma(k) \right) ; out \left(\mathcal{A} \stackrel{noisy}{\rightleftharpoons} \Gamma'(Enc(k)) \right) \right) \leq \xi. \quad (24)$$

The definition of Γ' being a (ϵ, ξ) -random-probing resilient implementation of a circuit Γ is identical to the one above, except that Point 2 is replaced with:

- 2'. for every ϵ -random-probing circuit adversary $\mathcal{A}$ there exists a black-box circuit adversary $\mathcal{S}$ such that

$$\Delta \left(out \left(\mathcal{S} \stackrel{bb}{\rightleftharpoons} \Gamma(k) \right) ; out \left(\mathcal{A} \stackrel{rnd}{\rightleftharpoons} \Gamma'(Enc(k)) \right) \right) \leq \xi.$$

The definition of Γ' being a (t, ξ) -threshold-probing resilient implementation of a circuit Γ is identical to the one above, except that Point 2 is replaced with:

- 2". for every t -threshold-probing circuit adversary $\mathcal{A}$ there exists a black-box circuit adversary $\mathcal{S}$ such that

$$\Delta \left(\text{out} \left(\mathcal{S} \stackrel{bb}{\hookrightarrow} \Gamma(k) \right) ; \text{out} \left(\mathcal{A} \stackrel{thr}{\hookrightarrow} \Gamma'(\text{Enc}(k)) \right) \right) \leq \xi.$$

In all cases above we will say that Γ' is a an implementation Γ with *efficient simulation* if the simulator $\mathcal{S}$ works in time polynomial in $\Gamma' \cdot |\mathbb{F}|$ as long as $\mathcal{A}$ is poly-time and the noise functions specified by $\mathcal{A}$ are efficiently decidable.

5.2. The Implementation

In this section we describe the circuit compiler of [17], generalized to larger fields in [31]. Let Γ be a stateful arithmetic circuit and let $d \in \mathbb{N}$ be a parameter. The encoding function Enc_+ that we use is also standard and is often called the “additive masking”. It is defined as: $\text{Enc}_+(x) := (X_1, \dots, X_d)$, where $X_1, \dots, X_d$ are uniform such that $X_1 + \dots + X_d = x$.

At a high level, each wire w in the original circuit Γ is represented by a *wire bundle* in Γ' , consisting of d wires $\vec{w} = (w_1, \dots, w_d)$, that carry an *encoding* of w . The gates in C are replaced gate-by-gate with so-called *gadgets*, computing on encoded values. The main difficulty is to construct gadgets that remain “secure” even if their internals may leak.

Because the transformed gadgets in Γ' operate on encodings, Γ' needs to have a subcircuit at the beginning that encodes the inputs and another subcircuit at the end that decodes the outputs. We will deal with the output decoding later. The input encoding is easy to implement for our encoding function Enc_+ : to encode an input x one simply uses the random gates to generate $d - 1$ field elements $x_1, \dots, x_{d-1}$ and then computes x_d as $x_1 + \dots + x_{d-1} - x$. Clearly this can be done using d addition and subtraction gates. Recall that the memory gates of Γ are assumed to be preloaded with field elements that already encode k using the encoding Enc_+ [cf. (24)]; hence, there is no need to encode k .

Each constant gate γ_c^{const} in Γ can be transformed into d constant gates in Γ' , the first of them being γ_c^{const} and the remaining ones being γ_0^{const} . This is trivially correct as $c = c + 0 + \dots + 0$. Every random gate γ^{rand} in Γ is transformed into d random gates in Γ' . This works since, clearly, a uniformly random encoding $(X_1, \dots, X_d)$ encodes a uniformly random element of $\mathbb{F}$.

What remains to show is how the operation (addition, subtraction, and multiplication) gates are handled. Consider a gate γ in Γ . Let a and b be its input wires and let $\vec{a} = (a_1, \dots, a_d)$ and $\vec{b} = (b_1, \dots, b_d)$ be their corresponding wire bundles in Γ' . Let the output wire bundle in Γ' be $(c_1, \dots, c_d)$. The cases when γ is an addition or subtraction gate are actually easy to deal with, thanks to the linearity of the encoding function. For example, if γ is an addition gate γ^+ then each c_i can be computed using an addition gate γ^+ in Γ' with input wires a_i and b_i (this is obviously correct as $(a_1 + b_1) + \dots + (a_d + b_d) = (a_1 + \dots + a_d) + (b_1 + \dots + b_d)$). The subtraction is handled analogously. The only tricky case is when γ is the multiplication gate. In this case the circuit Γ' generates, for

every $1 \leq i < j \leq d$, a random field element $z_{i,j}$ (this is done using the random gates in Γ'). Then, for every $1 \leq j < i \leq d$ it computes $z_{i,j} := a_i b_j + a_j b_i - z_{j,i}$, and finally he computes each c_i (for $i = 1, \dots, d$) as $c_i := a_i b_i + \sum_{i \neq j} z_{i,j}$. To see why this computation is correct consider the sum $c = c_1 + \dots + c_d$ and observe that every $z_{i,j}$ in it appears exactly once with plus sign and once with a minus sign, and hence it cancels out. Moreover each term $a_i b_j$ appears in the formula for c exactly once. Hence c is equal to $\sum_{i,j \in \{1, \dots, n\}} a_i b_j = \left(\sum_{i=1}^d a_i \right) \left(\sum_{j=1}^d b_j \right) = ab$. It is straightforward to verify that the total number of gates in this gadget is $3.5 \cdot d^2$. This finishes the description of the compiler.

The multiplication gadget above turns out to be useful as a building block for “refreshing” of the encoding. More concretely, suppose we have a wire bundle $\vec{a} = (a_1, \dots, a_d)$ and we wish to obtain another bundle $\vec{b} = (b_1, \dots, b_d)$ such that $\vec{b}$ is a fresh encoding of $\text{Dec}_+(\vec{a})$. This can be achieved by a *Refresh* sub-gadget constructed as follows. First, create an encoding $(1, 0, \dots, 0)$ of 1 (using d constant gates), and multiply $(1, 0, \dots, 0)$ and $\vec{a}$ together using the multiplication protocol above. Since $(1, 0, \dots, 0)$ is an encoding of 1, hence the result will be an encoding of $1 \cdot a = a$. The multiplication can be done with $3.5 \cdot d^2$ gates, and hence altogether this gadget uses $3.5 \cdot d^2 + d$ gates.

We can now use the *Refresh* sub-gadget to construct the output gadgets in Γ' . Let γ^{out} be an output gate in Γ with an input wire a . Then in Γ' it is transformed into the following: let $\vec{a}$ be the wire bundle corresponding to a . First apply the *Refresh* sub-gadget, and then calculate the sum $b_1 + \dots + b_d$ (where $(b_1, \dots, b_d)$ is the output of *Refresh*) and output the result.

The refreshing gadget is also useful to provide security of the memory encoding in the multi-round scenario. More precisely, we assume that every memory state gets refreshed at the end of each round by the *Refresh* procedure. It is easy to see that without this “refreshing” the contents of the memory would eventually leak completely to the adversary even if he probes a very limited number (say: 1) of wires in each round. For more details see [17].

5.3. Security in the Probing Model [17]

In [17] it is shown that the compiler from the previous section is secure against probing attacks in which the adversary can probe at most $\lfloor (d-1)/2 \rfloor$ wires in each round.⁵ This parameter may be a bit disappointing as the number of probes that the adversary needs to break the security does not grow with the size of the circuit. This assumption may seem particularly unrealistic for large circuits Γ . Fortunately, [17] also shows a small modification of the construction from Sect. 5.2 that is resilient to a larger number of probes, provided that the number of probes from each gadget is bounded. Before we

⁵ Strictly speaking the proof of [17] considers only the case when $\mathbb{F} = \text{GF}(2)$. It was observed in [31] that it can be extended to any finite field, as the only properties of $\text{GF}(2)$ that are used in the proof are the field axioms. Moreover, additions (XOR) gates are not considered, as AND and NOT gates are sufficient to describe any type of Boolean circuit. A full proof of security for all linear operations (including the field addition) and their composability has recently been provided by Andrychowicz et al. [3]. Moreover, we emphasize that when we use a composable refreshing scheme that is placed between each gadget, then the proof for the addition gate is trivial and just follows from counting the intermediate operations.

present it let us argue why the original construction is not secure against such attacks. To this end, assume that our circuit Γ has a long sequence of wires $a_1, \dots, a_m$, where each a_i (for $i > 1$) is the result of adding to a_{i-1} (using an addition gate) a 0 constant (that was generated using a γ_0^{const} gate). It is easy to see that in the circuit Γ' all the wire bundles $\vec{a}_1, \dots, \vec{a}_m$ (where each $\vec{a}_i$ corresponds to a_i) will be identical. Hence, the adversary that probes even a single wire in each addition gadget in Γ' will learn the encoding of a_1 completely as long as $m \geq d$. Fortunately one can deal with this problem by “refreshing” the encoding after each subtraction and addition gate exactly in the same way as done before, i.e., by using the *Refresh* sub-gadget.

Lemma 7. ([17]) *Let Γ be an arbitrary stateful arithmetic circuit over some field $\mathbb{F}$. Let Γ' be the circuit that results from the procedure described above. Then Γ' is a $(\lfloor (d-1)/2 \rfloor \cdot |\Gamma|, 0)$ -threshold-probing resilient implementation of a circuit Γ (with efficient simulation), provided that the adversary does not probe each gadget more than $\lfloor (d-1)/2 \rfloor$ times in each round.*

We notice that [17] also contains a second transformation with blow-up $\tilde{O}(d |\Gamma|)$. It may be possible that this transformation can provide better noise parameters as is achieved by Theorem 2. However, due to the hidden parameters in the $\tilde{O}$ -notation we do not get a straightforward improvement of our result. In particular, using this transformation the size of the transformed circuit depends also on an additional statistical security parameter, which will affect the tolerated noise level.

5.4. Resilience to Noisy Leakage from the Wires

We now show that the construction from Sect. 5.3 is secure against the noisy leakage. More precisely, we show the following.

Theorem 1. *Let Γ be an arbitrary stateful arithmetic circuit over some field $\mathbb{F}$. Let Γ' be the circuit that results from the procedure described in Sect. 5.3. Then Γ' is a $(\delta, |\Gamma| \cdot \exp(-d/12))$ -noise-resilient implementation of Γ (with efficient simulation), where*

$$\delta := ((28d + 8) |\mathbb{F}|)^{-1} = O(1/(d \cdot |\mathbb{F}|)).$$

Proof. Let $\mathcal{A}$ be a δ -noisy circuit adversary attacking Γ' . We construct an efficient black-box simulator $\mathcal{S}$ such that for every k it holds that

$$\Delta \left(\text{out} \left(\mathcal{S} \stackrel{\text{bb}}{\leftarrow} \Gamma(k) \right) ; \text{out} \left(\mathcal{A} \stackrel{\text{noisy}}{\leftarrow} \Gamma'(\text{Enc}(k)) \right) \right) \leq |\Gamma| \cdot \exp(-d/12). \quad (25)$$

Observe that in our construction every gates gets transformed into a gadget of at most $3.5 \cdot d^2 + d$ gates. Since each gate can have at most 2 inputs, hence the total number of wires in a gadget is $\ell := 7 \cdot d^2 + 2 \cdot d$. Let $\gamma^1, \dots, \gamma^{|\mathcal{G}|}$ be the gates of Γ . For each $i = 1, \dots, \ell$ let the wires in the gadget in Γ' that corresponds to γ^i be denoted $(x_1^i, \dots, x_\ell^i)$. Since $\delta = d/(4\ell |\mathbb{F}|)$, we can use Corollary 1 and simulate the noise from each $(x_1^i, \dots, x_\ell^i)$

by a $(d/2 - 1)$ -threshold-probing adversary $\mathcal{S}^i$ working in time polynomial in $\ell \cdot |\mathcal{X}|$. The simulation is perfect, unless $\mathcal{S}^i$ outputs $\perp$, which, by Corollary 1 happens with probability at most $\exp(-d/12)$. Hence, by the union-bound the probability that *some* $\mathcal{S}^i$ outputs $\perp$ is at most $|\Gamma| \cdot \exp(-d/12)$. Denote this event $\mathcal{E}$.

From Lemma 7 we know that every probing adversary that attacks Γ' by probing at most $\lfloor (d-1)/2 \rfloor \geq d/2 - 1$ wires from each gadget can be perfectly simulated in polynomial time by an adversary $\mathcal{S}$ with a black-box access to Γ . Hence, $\mathcal{A}$ can also be simulated perfectly by a black-box access to Γ conditioned on the fact that $\mathcal{E}$ did not occur. Hence we get

$$\Delta \left(\text{out} \left(\mathcal{S} \stackrel{bb}{\hookrightarrow} \Gamma(k) \right) \mid \neg \mathcal{E} ; \text{out} \left(\mathcal{A} \stackrel{noisy}{\hookrightarrow} \Gamma'(Enc(k)) \right) \right) = 0.$$

This, by Lemma 2 (Sect. 2.1), implies (25). Obviously $\mathcal{S}$ works in time polynomial in $|\Gamma| \cdot d^2 \cdot |\mathbb{F}|$, which is polynomial in $\Gamma' \cdot |\mathbb{F}|$. This finishes the proof. $\square$

In short, this theorem is proven by combining Corollary 1 that reduces the noisy adversary to the probing adversary, with Lemma 7 that shows that the construction from Sect. 5.3 is secure against probing.

5.5. Resilience to Noisy Leakage from the Gates

The model of Prouff and Rivain is actually slightly different than the one considered in the previous section. The difference is that they assume that the noise is generated by the *gates*, not by the *wires*. This can be formalized by assuming that each noise function *Noise* is applied to the “contents of a gate”. We do not need to specify exactly what we mean by this. It is enough to observe that the contents of each gate γ can be described by at most 2 field elements: obviously if γ is a random gate, output gate, or memory gate then its entire state in a given round can be described by one field element, and if γ is an operation gate then it can be described by two field elements that correspond to γ ’s input. Hence, without loss of generality we can assume that the noise function is defined over the domain $\mathbb{F} \times \mathbb{F}$.

Formally, we define a δ -gate-noisy circuit adversary $\mathcal{A}$ as a machine that, besides of having black-box access to a circuit $\Gamma(k)$, can, after each i th round, get some partial information about the internal state of the computation via the δ -noisy leakage functions applied to the gates (in a model described above). Let $\text{out} \left(\mathcal{A} \stackrel{g\text{-noisy}}{\hookrightarrow} \Gamma(k) \right)$ denote the output of such $\mathcal{A}$ after interacting with Γ whose initial memory state is $k_0 = k$.

We can accordingly modify the definition of noise-resilient circuit implementations (cf. Definition 2). We say that Γ' is a (δ, ξ) -input-gate-noise resilient implementation of a circuit Γ w.r.t. *Enc* if for every k and every δ -noisy circuit adversary $\mathcal{A}$ described above there exists a black-box circuit adversary $\mathcal{S}$ working in time polynomial in $\Gamma' \cdot |\mathbb{F}|$ such that

$$\Delta \left(\text{out} \left(\mathcal{S} \stackrel{bb}{\hookrightarrow} \Gamma(k) \right) ; \text{out} \left(\mathcal{A} \stackrel{g\text{-noisy}}{\hookrightarrow} \Gamma'(Enc(k)) \right) \right) \leq \xi. \quad (26)$$

It turns out that the transformation from Sect. 5.3 also works in this model, although with different parameters. More precisely we have the following theorem.⁶

Theorem 2. *Let Γ be an arbitrary stateful arithmetic circuit over some field $\mathbb{F}$. Let Γ' be the circuit that results from the procedure described in Sect. 5.3. Then Γ' is a $(\delta, |\Gamma| \cdot \exp(-d/24))$ -noise-resilient implementation of Γ (with efficient simulation), where*

$$\delta := \left((28d + 8) \cdot |\mathbb{F}|^2 \right)^{-1} = O(1/(d \cdot |\mathbb{F}|^2)). \quad (27)$$

Proof. The proof is similar to the one of Theorem 1 so we only describe the key differences. Let $\mathcal{A}$ be a δ -noisy adversary. The number ℓ corresponds now to the number of gates in each gadget, and hence it is equal to $3.5 \cdot d^2 + d$. It is therefore straightforward to calculate that δ defined in (27) is equal to $(d/2)/(4\ell \cdot |\mathbb{F}|^2)$. Since the *Noise* function has domain of size $|\mathbb{F}|^2$, we can use Corollary 1 obtaining that $\mathcal{A}$ can be simulated by an adversary $\mathcal{S}$ that probes each gadget in less than $d/2$ positions. Since now each “position” corresponds to a gate in the circuit, hence the adversary needs to probe up to two wires to determine its value. Therefore $\mathcal{S}$ probes less than d wires in each gadget. Since d is now $1/2$ of what it was in the proof of Corollary 1, hence the error probability becomes $\exp(-d/12) = \exp(-d/24)$. $\square$

Comparison with [28]. As described in the introduction, our main advantage over [28] is the removal of the assumption about the existence of the leak-free gates, a stronger security model—chosen message attack, instead of a random message attack, and a more meaningful security statement. Still, it is interesting to compare our noise parameters with the parameters of [28]. Let us analyze how much noise is needed by [28] to ensure that the adversary obtains exponentially small information from leakage. The reader should keep in mind that both in our paper, and in [28] “more noise” means that a certain quantity, δ , in our case, is *smaller*. Hence, the larger δ is, the stronger the result becomes (as it means that *less* noise is required for the security to hold).

The main result of [28] is Theorem 4 on page 154. Unfortunately, the statement of this theorem is asymptotic treating $|\mathbb{F}|$ as constant, and hence to get a precise bound on how much noise is required one needs to inspect the proof. The bound on the noise can be deduced from the part of the proof entitled “Security of Type 3 Subsequences”, where the required noise is inversely proportional to “ $\lambda(d)$ ”, and this last value is linear in $d \cdot |\mathbb{F}|^3$ for a general d and linear in $d \cdot |\mathbb{F}|^{3/2}$ for large d ’s (note that $|\mathbb{F}|$ is denoted by N in [28], and d is a security parameter identical to ours). Hence, for a general d , their δ is $O(1/(d \cdot |\mathbb{F}|^3))$.

However, as explained in Sect. 3.1, the notion of distance in [28] is slightly different than the standard “statistical distance” that we use. Fortunately, one can use (7) to translate our bound into their language. It turns out that in this case our and their bounds are asymptotically identical for general d ’s, i.e., $(O(1/(d \cdot |\mathbb{F}|^3)))$. This is shown in

⁶ Note that our result holds only when the number of shares is large. For small values of d (e.g., $d = 2, 3, 4$) like those considered in [35], our result does not give meaningful bounds. This is similar to the work of Prouff and Rivain [28] and it is an interesting open research question to develop security models that work for small security parameters.

Corollary 2 below. Note that this translation is unidirectional, in the sense that their “ $O(1/(d \cdot |\mathbb{F}|^3))$ ” bound does *not* imply a bound “ $O(1/(d \cdot |\mathbb{F}|^2))$ ” in our sense.

Corollary 2. *Let Γ be an arbitrary stateful arithmetic circuit over some field $\mathbb{F}$. Let Γ' be the circuit that results from the procedure described in Sect. 5.3. Then Γ' is a $(\delta', |\Gamma| \cdot \exp(-d/24))$ -noise-resilient implementation of Γ (with efficient simulation) when the noise is defined using the β distance, where*

$$\delta' = \left((14d + 4) \cdot |\mathbb{F}|^3 \right)^{-1} = O(1/(d \cdot |\mathbb{F}|^3)).$$

Proof. From (7) with $\mathcal{X} = \mathbb{F} \times \mathbb{F}$, it follows that if *Noise* is δ' -noisy with respect to the β distance, then it is $(|\mathbb{F}| \cdot \delta'/2)$ -noisy in the standard sense. Since this last value is equal to δ defined in (27), hence we can use Theorem 2 obtaining that Γ' is a $(\delta', |\Gamma| \cdot \exp(-d/24))$ -noise-resilient implementation of Γ when the noise is defined using the β distance. $\square$

Acknowledgements

We would like to thank the anonymous *Eurocrypt* and *Journal of Cryptology* reviewers for their careful reading of our manuscript and their many insightful comments.

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

References

- [1] M. Ajtai, Secure computation with information leaking to an adversary. In *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pages 715–724 (2011)
- [2] A. Akavia, S. Goldwasser, V. Vaikuntanathan, Simultaneous Hardcore Bits and Cryptography against Memory Attacks. In *TCC*, pages 474–495 (2009)
- [3] M. Andrychowicz, S. Dziembowski, S. Faust, Circuit compilers with $o(1/\log(n))$ leakage rate. In *Advances in Cryptology - EUROCRYPT 2016 - 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Proceedings, Part II*, pages 586–615 (2016)
- [4] J. Blömer, J. Guajardo, V. Krummel, Provably Secure Masking of AES. In *Selected Areas in Cryptography*, pages 69–83 (2004)
- [5] C. Carlet, L. Goubin, E. Prouff, M. Quisquater, M. Rivain, Higher-Order Masking Schemes for S-Boxes. In *FSE*, pages 366–384 (2012)
- [6] S. Chari, C.S. Jutla, J.R. Rao, P. Rohatgi, Towards Sound Approaches to Counteract Power-Analysis Attacks. In *CRYPTO*, pages 398–412 (1999)
- [7] C. Clavier, J. Coron, N. Dabbous, Differential Power Analysis in the Presence of Hardware Countermeasures. In *CHES*, pages 252–263 (2000)
- [8] J. Coron, I. Kizhvatov, Analysis and Improvement of the Random Delay Countermeasure of CHES 2009. In *CHES*, pages 95–109 (2010)

- [9] D.P. Dubhashi, A. Panconesi. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press (2009)
- [10] S. Dziembowski, S. Faust. Leakage-Resilient Circuits without Computational Assumptions. In *TCC*, pages 230–247 (2012)
- [11] S. Dziembowski, S. Faust, M. Skorski, Noisy leakage revisited. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part II*, volume 9057, pages 159–188. Springer (2015)
- [12] S. Dziembowski, K. Pietrzak, Leakage-Resilient Cryptography. In *FOCS*, pages 293–302 (2008)
- [13] S. Faust, T. Rabin, L. Reyzin, E. Tromer, V. Vaikuntanathan, Protecting Circuits from Leakage: the Computationally-Bounded and Noisy Cases. In *EUROCRYPT*, pages 135–156 (2010)
- [14] S. Goldwasser, G.N. Rothblum. Securing computation against continuous leakage. In *CRYPTO*, pages 59–79 (2010)
- [15] S. Goldwasser, G.N. Rothblum. How to Compute in the Presence of Leakage. In *FOCS*, pages 31–40 (2012)
- [16] L. Goubin, J. Patarin, DES and Differential Power Analysis (The “Duplication” Method). In *CHES*, pages 158–172 (1999)
- [17] Y. Ishai, A. Sahai, D. Wagner, Private Circuits: Securing Hardware against Probing Attacks. In *CRYPTO*, pages 463–481 (2003)
- [18] A. Juma, Y. Vahlis, Protecting Cryptographic Keys against Continual Leakage. In *CRYPTO*, pages 41–58 (2010)
- [19] J. Katz, V. Vaikuntanathan, Signature Schemes with Bounded Leakage Resilience. In *ASIACRYPT*, pages 703–720 (2009)
- [20] P.C. Kocher, Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In *CRYPTO’96*, pages 104–113 (1996)
- [21] P.C. Kocher, J. Jaffe, B. Jun, Differential Power Analysis. In *CRYPTO’99*, pages 388–397 (1999)
- [22] Stefan Mangard, Elisabeth Oswald, Thomas Popp, *Power Analysis Attacks: Revealing the Secrets of Smart Cards (Advances in Information Security)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA (2007)
- [23] U.M. Maurer, S. Tessaro, A hardcore lemma for computational indistinguishability: Security amplification for arbitrarily weak prgs with optimal stretch. In Daniele Micciancio, editor, *TCC*, volume 5978 of *Lecture Notes in Computer Science*, pages 237–254. Springer (2010)
- [24] S. Micali, L. Reyzin, Physically Observable Cryptography (Extended Abstract). In *TCC*, pages 278–296 (2004)
- [25] E. Miles, E. Viola, Shielding circuits with groups. In *STOC*, pages 251–260 (2013)
- [26] M. Naor, G. Segev, Public-key cryptosystems resilient to key leakage. In *CRYPTO*, pages 18–35 (2009)
- [27] E. Oswald, S. Mangard, N. Pramstaller, V. Rijmen, A Side-Channel Analysis Resistant Description of the AES S-Box. In *FSE*, pages 413–423 (2005)
- [28] E. Prouff, M. Rivain, Masking against Side-Channel Attacks: A Formal Security Proof. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT*, volume 7881 of *Lecture Notes in Computer Science*, pages 142–159. Springer (2013)
- [29] E. Prouff, T. Roche, Higher-Order Glitches Free Implementation of the AES Using Secure Multi-party Computation Protocols. In *CHES*, pages 63–78 (2011)
- [30] J.-J. Quisquater, D. Samyde, ElectroMagnetic Analysis (EMA): Measures and Counter-Measures for Smart Cards. In *E-smart*, pages 200–210 (2001)
- [31] M. Rivain, E. Prouff, Provably Secure Higher-Order Masking of AES. In *CHES*, pages 413–427 (2010)
- [32] G.N. Rothblum, How to Compute under AC0 Leakage without Secure Hardware. In *CRYPTO*, pages 552–569 (2012)
- [33] F.-X. Standaert, T. Malkin, M. Yung, A Unified Framework for the Analysis of Side-Channel Key Recovery Attacks. In *EUROCRYPT*, pages 443–461 (2009)
- [34] François-Xavier Standaert, Olivier Pereira, Yu Yu. Leakage-Resilient Symmetric Cryptography under Empirically Verifiable Assumptions. In *CRYPTO (I)*, pages 335–352 (2013)
- [35] F.-X. Standaert, N.V.-Charvillat, E. Oswald, B. Gierlichs, M. Medwed, M. Kasper, S. Mangard, The World Is Not Enough: Another Look on Second-Order DPA. In *ASIACRYPT*, pages 112–129 (2010)

- [36] N. Veyrat-Charvillon, F.-X. Standaert, Adaptive Chosen-Message Side-Channel Attacks. In Jianying Zhou and Moti Yung, editors, *ACNS*, volume 6123 of *Lecture Notes in Computer Science*, pages 186–199 (2010)