A Decomposition Approach for Evaluating Security of Masking

Vahid Jahandideh, Bart Mennink, and Lejla Batina

Radboud University, Nijmegen, The Netherlands
{v.jahandideh, b.mennink, lejla}@cs.ru.nl

Abstract. Masking is a widely used countermeasure against side-channel attacks, encoding secrets into multiple shares—each potentially subject to leakage. A central question is under what leakage conditions, and to what extent, increasing the number of shares improves security. While this has been studied extensively in low-SNR regimes, scenarios where the adversary gains significant information—such as on low-noise hardware or via static power analysis—remain less understood.
We address this gap by deriving *necessary and sufficient* noise conditions for the security of masked encodings and linear gadgets. Our approach introduces a decomposition technique that reduces leakage analysis over extended fields to binary subproblems involving bit-level projections. This enables the derivation of tight bounds in binary subfields, which are then lifted back to the full field.
Beyond binary settings, we present a general framework for analyzing masking in other structures, including prime fields. As an application, we prove a conjecture by Dziembowski et al. (TCC 2016), showing that for an additive group $\mathbb{G}$ with largest subgroup $\mathbb{H}$, any δ-noisy leakage satisfying $\delta < 1 - \frac{|\mathbb{H}|}{|\mathbb{G}|}$ ensures that masking enhances security.

1 Introduction

Masking to Mitigate Side-Channel Threats. Side-channel information refers to unintended leakages that an adversary can extract from the physical implementation of a cryptographic algorithm. A *leakage model* provides an abstraction for characterizing such leakages. One widely studied model is the *noisy leakage model*, introduced by Prouff and Rivain [29] and subsequently explored in several works [9, 10, 12, 13, 28]. In this model, for each intermediate value $X \in \mathbb{F}_q$ in a computation, the adversary observes a function $\mathsf{L}(X)$, such as a noisy Hamming weight.

A primary countermeasure against side-channel leakage is *masking*. In this approach, a secret value X is split into random *shares* $X_1, \ldots, X_n \in \mathbb{F}_q$ such that $X = X_1 + \cdots + X_n$. Instead of manipulating X directly, the implementation operates on the individual shares, and the adversary can only observe the leakage $\mathsf{L}(X_i)$ from each one. In a standalone encoding, a single secret is encoded, and the only intermediates are the shares themselves. In contrast, a protected circuit may involve multiple secrets and many intermediates. The effectiveness of masking

is typically evaluated by how a chosen *security metric*—such as the adversary's success rate—degrades as the number of shares n increases. Since the seminal work of Chari et al. [7], this methodology has been central to the evaluation of side-channel countermeasures, in both standalone encodings and full protected circuits.

Open Challenge. If $\mathsf{L}(X)$ reveals X entirely, masking provides no benefit. Therefore, the leakage must inherently include some *noise*. Determining the minimal noise level required for masking to be effective—and understanding how the security scales with the number of shares n in borderline cases—remains an open problem. This paper addresses these challenges.

Practical Relevance. Low-noise (high-SNR) conditions arise when $\mathsf{L}(X)$ reveals a substantial amount of information about X. Such scenarios have been reported in various contexts. For instance, *low-noise processors*—particularly small embedded devices such as the ARM Cortex-M0—exhibit inherently lower noise levels in their power consumption [6]. Likewise, *static power analysis*, unlike dynamic power analysis, measures a stable leakage signal over an extended period, resulting in highly precise side-channel observations [26]. Lastly, *averaging* or *horizontal attacks* can combine multiple leakage samples corresponding to the same or related intermediates to produce a clearer, aggregated leakage trace [3].

Security Metrics. Several metrics have been proposed to evaluate the effectiveness of masking:

- *Success rate (SR)*: the probability that an adversary correctly identifies X given $\mathsf{L}(X)$ [32];
- *Statistical distance (SD)*: the distance δ between the prior and posterior distributions of X given $\mathsf{L}(X)$;
- *Mutual information (MI)*: the information shared between X and $\mathsf{L}(X)$, i.e., $\mathsf{MI}(X\,;\,\mathsf{L}(X)) \in [0, \log q]$.

These metrics are interrelated—for instance, Pinsker's inequality links MI and SD [15], and SR is also related to SD (see Lemma 1)—but their relationships are often not tight in practice.

In this work, we focus on the success rate metric. It is intuitive and directly reflects the number of traces needed for an attack to succeed. In some practical divide-and-conquer attacks where secrets are split into many chunks, the adversary must guess each chunk correctly in one attempt, making SR the most relevant metric.

1.1 Evaluating the Security of Single Encodings

Building on the reduction proposed by Duc et al. [9], further work [10] established that $q\delta < 1$ suffices for masking to be effective. Dziembowski et al. [13] improved this bound, proving that for binary extended fields ($q = 2^u$), the threshold $\delta < \frac{1}{2}$

is optimal. However, their analysis does not address the borderline case where $\delta \geq \frac{1}{2}$.

More recently, using mutual information as a metric, Ito et al. [21] suggested a threshold of $\mathsf{MI}(X\,;\,\mathsf{L}(X)) < 0.72$ for all shares. Béguinot et al. [4] relaxed this to ≤ 1 for some shares, assuming that different shares may leak differently. Despite these efforts, the precise characterization of when masking is effective for a given leakage function remains unresolved.

Our Contribution to the Problem. We address this gap by relaxing noise requirements and establishing that masking improves security *if and only if* $\mathsf{L}(X)$ does not fully determine any nontrivial bitwise combination of X. Specifically, for a u-bit variable X, masking is effective only if

$$\mathsf{MI}\big(\langle X, h\rangle\ ;\ \mathsf{L}(X)\big) < 1,$$

for every nonzero $h \in \{1, \ldots, 2^u - 1\}$, where

$$\langle X, h\rangle = \bigoplus_{j=0}^{u-1} x_j h_j$$

denotes the binary inner product of X and h.

Our first result provides a tight security bound for binary fields. We then extend the analysis to larger values of u and arbitrary sharing order n, deriving tight lower and upper bounds on the normalized success rate (SR) metric in terms of $\langle X, h\rangle$ projections.

To demonstrate the practical relevance of this approach, we conduct an experimental evaluation of $\mathsf{MI}(\langle X, h\rangle\ ;\ \mathsf{L}(X))$ across all h, and show how this refined metric:

- Improves accuracy in *leakage certification*;
- Identifies the *required masking order* needed to meet a target security level.

When $\mathsf{MI}(\langle X, h\rangle\ ;\ \mathsf{L}(X)) = 1$ for some h, masking over characteristic-two fields offers no security benefit. To address this limitation, we introduce a general framework based on the SR metric for broader algebraic settings, including prime fields and additive groups. Within this framework, we:

- Show that inner product-based metric is effective for estimating adversary's success rate even in prime fields;
- Prove a conjecture of Dziembowski et al. [13], establishing that for an additive group $\mathbb{G}$ with largest subgroup $\mathbb{H}$, masking is effective if

$$\delta < 1 - \frac{|\mathbb{H}|}{|\mathbb{G}|}.$$

1.2 Evaluating the Security of Protected Circuits

Security guarantees for individual shares in a standalone encoding do not directly extend to protected circuits. Prouff and Rivain [29] initiated the study of noisy leakage in masked circuits, which was later refined by Masure and Standaert [25]. Still, circuit-level analysis remains challenging and often depends on strong assumptions—such as the presence of leak-free refresh gadgets.

To address these challenges, Duc et al. [9] proposed a reduction from the noisy leakage model to the *random probing model* (RPM), enabling the transfer of security guarantees between the two frameworks. While subsequent works [12,27, 28] improved the tightness of the reduction parameters, its applicability degrades in low-noise regimes, where $\mathsf{L}(X)$ may nearly reveal X entirely.

Our Contributions for Circuit-Level Security. We show that our bitwise decomposition strategy enables a tighter reduction from the noisy leakage model to the RPM—even in low-noise conditions. We apply this new approach to both mask encodings and linear masked gadgets.

This direction is further supported by recent work of Jahandideh et al. [22], which demonstrates that analyses based on linear circuits can yield meaningful side-channel security margins—even when the circuit includes certain non-linear components.

1.3 Related Work

A key technical perspective in our work is the analysis of information contained in bitwise combinations of a random variable X, given access to a leakage function $\mathsf{L}(X)$. This recalls the classical result of Goldreich and Levin [17], who introduced the notion of a *hardcore predicate* for one-way functions. Their result states that if it is computationally hard to recover a u-bit string X from an input length-preserving function $\mathsf{f}(X)$, then there exists a bitmask $r \in \{0,1\}^u$ such that the inner product $\langle X, r\rangle$ cannot be predicted significantly better than random guessing—even given both $\mathsf{f}(X)$ and r.

We adopt a similar viewpoint in the context of side-channel leakage: when $\mathsf{L}(X)$ reveals only partial information about X, we study which bitwise projections $\langle X, r\rangle$ remain hidden from the adversary. This characterization serves as a foundation for determining when masking continues to offer meaningful protection.

1.4 Outline

Section 2 introduces the security metrics and their relationships for single variables. Section 3 develops our masking analysis using a decomposition-based approach. Section 4 extends this analysis to masking over additive groups. Finally, Section 5 applies the decomposition framework to linear circuits.

2 Side-Channel Security of a Single Variable

This section formalizes the noisy leakage model and introduces the adversary's *advantage*, a normalized version of the success rate. We relate this metric to the statistical distance δ in Lemma 1. Subsection 2.2 presents the reduction to the random probing model, and Lemma 2 proves its tightness in binary fields. Finally, Lemma 3 derives a bound on the advantage using this reduction.

2.1 Preliminaries

Let X be uniformly distributed over $\mathbb{F}_q$, representing an intermediate value in a cryptographic implementation. Side-channel leakage is modeled by a probabilistic function $\mathsf{L}(X) \in \mathbb{R}^m$, and the adversary's goal is to guess X from $\mathsf{L}(X)$. The optimal strategy is *maximum a posteriori* (MAP) estimation [20]:

$$\hat{X} \xleftarrow{\$} \left\{ \operatorname*{argmax}_{\alpha \in \mathbb{F}_q} \Pr(X = \alpha \mid l) \right\}, \quad \text{where } l \leftarrow \mathsf{L}(X).$$

The MAP success probability varies with the leakage value. For instance, if $\mathsf{L}(X) = \mathsf{HW}(X)$, then observing $l = 0$ reveals X completely.

Define the average success probability:

$$P_c \triangleq \mathop{\mathbb{E}}_{l}[\Pr(\hat{X} = X \mid l)] = \sum_{l} \Pr(\mathsf{L}(X) = l) \cdot \Pr(\hat{X} = X \mid l),$$

and the *advantage* over random guessing as:

$$\mathsf{Adv}_X \triangleq P_c - \frac{1}{q}.$$

Statistical Distance and δ-Noisy Leakage. The *statistical distance* between X and its posterior $X \mid \mathsf{L}(X)$ measures leakage informativeness [29]. It is defined as:

$$\mathsf{SD}(X\,;\, X \mid \mathsf{L}(X)) \triangleq \sum_{l} \Pr(\mathsf{L}(X) = l) \cdot \mathsf{TV}(X\,;\, X \mid l),$$

where the *total variation distance* is:

$$\mathsf{TV}(X\,;\, X \mid l) \triangleq \frac{1}{2} \sum_{\alpha \in \mathbb{F}_q} \left| \Pr(X = \alpha \mid l) - \frac{1}{q} \right| = \sum_{\Pr(X=\alpha|l) > \frac{1}{q}} \left(\Pr(X = \alpha \mid l) - \frac{1}{q} \right).$$

We say $\mathsf{L}(X)$ is δ-noisy if $\mathsf{SD}(X\,;\, X \mid \mathsf{L}(X)) = \delta$. By definition, $\delta \in [0, 1 - \frac{1}{q}]$.

Relation Between δ and Adv_X. A lower δ corresponds to higher noise and thus lower advantage. The following lemma formalizes this relationship.

Lemma 1. *Let $X \in \mathbb{F}_q$ and let $\mathsf{L}(X)$ be δ-noisy. Then:*

$$\frac{\delta}{q-1} \leq \mathsf{Adv}_X \leq \delta.$$

In the binary case ($q = 2$), we have $\mathsf{Adv}_X = \delta$.

Proof. For each l, we have $\max_\alpha \Pr(X = \alpha \mid l) - \frac{1}{q} \leq \mathsf{TV}(X\,;\, X \mid l)$. Hence, averaging over l proves the upper bound. The lower bound follows by noting that at most $q-1$ values can exceed $\frac{1}{q}$. □

2.2 Leakage Simulation

The *erasure channel* [8, 19] is a probabilistic mapping $\phi^\epsilon \colon \mathbb{F}_q \to \{\perp, \mathbb{F}_q\}$ defined as:

$$\phi^\epsilon(X) = \begin{cases} X & \text{with probability } \epsilon, \\ \perp & \text{otherwise.} \end{cases}$$

Duc et al. [9] showed that for sufficiently noisy leakage $\mathsf{L}(X)$, one can construct a function L' such that for any X, the leakages $\mathsf{L}'(\phi^\epsilon(X))$ and $\mathsf{L}(X)$ are statistically indistinguishable:

$$\forall \alpha \in \mathbb{F}_q, \quad \mathsf{TV}\left(\mathsf{L}'(\phi^\epsilon(\alpha))\,;\, \mathsf{L}(\alpha)\right) = 0.$$

This holds if $\epsilon \geq \epsilon_{\min}$, where

$$\epsilon_{\min} \triangleq 1 - \sum_l \min_{\alpha \in \mathbb{F}_q} \Pr(l \mid X = \alpha) \;\leq_{(\mathtt{I})}\; q\delta. \tag{1}$$

This value, known as the *Doeblin coefficient* [5], lies in $[0, 1]$: $\epsilon_{\min} = 0$ means $\mathsf{L}(X)$ is independent of X, while $\epsilon_{\min} = 1$ implies that the noise of leakage is too low, and the technique is not applicable.

The right-hand side of inequality ($\mathtt{I}$) was proved in [9]. We now show that in the binary case, equality holds.

Lemma 2. *Let $(X, \mathsf{L}(X))$ be a joint distribution with $\mathsf{SD}(X\,;\, X \mid \mathsf{L}(X)) = \delta$ and X uniform over $\mathbb{F}_2$. Then:*

$$\epsilon_{\min} = 2\delta.$$

Proof.

$$\begin{aligned}
\epsilon_{\min} &= 1 - \sum_l \min_{\alpha \in \{0,1\}} \Pr(l \mid X = \alpha) \\
&= \sum_l \Pr(\mathsf{L}(X) = l) \left[1 - 2 \min_\alpha \Pr(X = \alpha \mid l)\right] \\
&= \sum_l \Pr(l) \left[\max_\alpha \Pr(X = \alpha \mid l) - \min_\alpha \Pr(X = \alpha \mid l)\right] \\
&= \sum_l \Pr(l) \left[\left|\Pr(X = 1 \mid l) - \tfrac{1}{2}\right| + \left|\Pr(X = 0 \mid l) - \tfrac{1}{2}\right|\right] = 2\delta. \quad \square
\end{aligned}$$

Since $\delta \leq \frac{1}{2}$ for binary X, we have $\epsilon_{\min} < 1$ unless $\mathsf{L}(X)$ fully reveals X. Thus, in the binary case, the residual uncertainty about X is precisely reflected in $\epsilon_{\min}$.

Example 1. Let $X \in \mathbb{F}_{2^u}$ and consider $\mathsf{L}(X) = x_0 \oplus e$, where x_0 is the least significant bit of X and $\Pr(e = 1) = \mathsf{e} \leq \frac{1}{2}$. Then $\mathsf{L}(X)$ leaks noisy information about the LSB only. The posterior of X satisfies:

$$\Pr(X = \alpha \mid l) = \begin{cases} \frac{1-\mathsf{e}}{2^{u-1}} & \text{if } \alpha \in \{0,1\}^{u-1}||l, \\ \frac{\mathsf{e}}{2^{u-1}} & \text{otherwise.} \end{cases}$$

Hence:

$$\mathsf{Adv}_X = \frac{1-2\mathsf{e}}{2^u}, \quad \mathsf{SD}(X\,;\, X \mid \mathsf{L}(X)) = \frac{1}{2} - \mathsf{e}, \quad \epsilon_{\min} = 1 - 2\mathsf{e}. \ \square$$

A Security Reduction. Since $(X, \mathsf{L}(X))$ and $(X, \mathsf{L}'(\phi^{\epsilon}(X)))$ are identically distributed, the adversary's advantage remains unchanged:

$$\mathsf{Adv}_X[\mathsf{L}(X)] = \mathsf{Adv}_X[\mathsf{L}'(\phi^{\epsilon}(X))].$$

We use $\mathsf{Adv}_X[\cdot]$ to specify the leakage source explicitly.

The informativeness of the random variables (RVs) $\phi^{\epsilon}(X)$, $\mathsf{L}'(\phi^{\epsilon_{\min}}(X))$, and $\mathsf{L}(X)$ about X can be expressed with the following chain:

$$X \to \phi^{\epsilon}(X) \to \phi^{\epsilon_{\min}}(X) \to \mathsf{L}'(\phi^{\epsilon_{\min}}(X)) \to \mathsf{L}(X),$$

with $\epsilon \geq \epsilon_{\min}$. As we move right along the chain, we obtain progressively degraded views of X, so any leakage evaluation metric (e.g., success probability) must decrease:

$$\mathsf{Adv}_X[\mathsf{L}(X)] \leq \mathsf{Adv}_X[\phi^{\epsilon}(X)], \qquad \mathsf{SD}(X\,;\, X \mid \mathsf{L}(X)) \leq \mathsf{SD}(X\,;\, X \mid \phi^{\epsilon}(X)). \quad (2)$$

Thus, to prove security under leakage $\mathsf{L}(X)$, it suffices to prove it under $\phi^{\epsilon_{\min}}(X)$. This reduction from δ-noisy to ϵ-random probing leakage was introduced in [9]. The following lemma is a concrete application.

Lemma 3 ([5], Proposition 1). *Let $X \in \mathbb{F}_q$, and let $\mathsf{L}(X)$ be a leakage with Doeblin coefficient $\epsilon_{\min}$. Then:*

$$\mathsf{Adv}_X \leq \frac{q-1}{q} \cdot \epsilon_{\min}.$$

$\epsilon_{\min}$ Is Not Always Tight. While the reduction form δ-noisy to the ϵ-random probing is powerful, it is not always tight. For instance, when $\mathsf{L}(X) = \mathsf{HW}(X)$, equation (1) gives $\epsilon_{\min} = 1$, even though the Hamming weight does not fully determine X. In this case, the reduction becomes ineffective.

One might attribute this to $\mathsf{HW}(X)$ being highly informative. However, consider a simpler leakage function $\mathsf{ZV}(X)$, inspired by the zero-value model [23], defined as:

$$\mathsf{ZV}(X) = \begin{cases} \nu_a & \text{if } X = 0, \\ \nu_b \neq \nu_a & \text{otherwise.} \end{cases} \tag{3}$$

This function merely indicates whether $X = 0$, revealing only one bit of information. Yet, it still results in $\epsilon_{\min} = 1$.

Our refined reduction approach, introduced in the next section, addresses this limitation by enabling a more fine-grained analysis for such cases.

3 Security of Mask Encoding

3.1 Mask Encoding

A standard method for protecting a sensitive variable X is *masking*, where X is split into a random tuple of n shares $\boldsymbol{X} = (X_1, \ldots, X_n)$ such that

$$X = \sum_{i=1}^{n} X_i.$$

The adversary then observes the leakage vector

$$\boldsymbol{L}(\boldsymbol{X}) = \big[\mathsf{L}_1(X_1),\, \ldots,\, \mathsf{L}_n(X_n)\big].$$

We assume for simplicity that all leakage functions are identical, i.e., $\mathsf{L}_i = \mathsf{L}$, and independent. A key question is how the adversary's advantage $\mathsf{Adv}_X[\boldsymbol{l} \leftarrow \boldsymbol{L}(\boldsymbol{X})]$ depends on the number of shares n and the structure of the field $\mathbb{F}_q$.

Intractability of Exact Metrics. The space of possible leakage vectors $\boldsymbol{L}(\boldsymbol{X})$ grows exponentially with the sharing order n, rendering the exact evaluation of security metrics computationally infeasible. A common workaround—adopted also in this work—is to approximate these metrics based on the distribution of $\mathsf{L}(X)$, albeit at the cost of potential inaccuracies.

A Loose Bound. Duc et al. [10] applied the δ-noisy to ϵ-random probing reduction, in which each share is revealed with probability $\epsilon \leq q\delta$. In the random probing model, the adversary learns each share independently with probability ϵ, so the probability of learning all n shares is at most $(q\delta)^n$. The corresponding leakage vector in this model is:

$$\boldsymbol{\phi^{\epsilon}}(\boldsymbol{X}) = \big[\phi^{\epsilon}(X_1), \ldots, \phi^{\epsilon}(X_n)\big].$$

Generalizing from (2), the statistical distance satisfies:

$$\Delta = \mathsf{SD}(X\,;\, X \mid \boldsymbol{L}(\boldsymbol{X})) \leq \mathsf{SD}(X\,;\, X \mid \boldsymbol{\phi^{\epsilon}}(\boldsymbol{X})).$$

If at least one share is not revealed, the posterior distribution of X is uniform, yielding zero statistical distance. Only when all shares are leaked does the distance reach its maximum value of $1 - \frac{1}{q}$. Therefore, we obtain the bound:

$$\Delta \leq \left(1 - \tfrac{1}{q}\right) q^n \delta^n. \tag{4}$$

A Tighter Bound. The bound in (4) increases rapidly with q, yet empirical results in [10] showed no such dependency. This led to the conjecture that the q-factor may be a proof artifact. A tighter bound, removing this dependency, was later proven by Masure et al. [24]:

Lemma 4 ([24], Proposition 4). *Let $X = (X_1, \ldots, X_n)$ be a masking of $X \in \mathbb{F}_q$, and suppose $\mathsf{SD}(X\,;\, X \mid \mathsf{L}(X)) = \delta$. Then*

$$\Delta = \mathsf{SD}(X\,;\, X \mid \boldsymbol{L}(\boldsymbol{X})) \leq 2^{n-1}\delta^n.$$

This result confirms an observation of Dziembowski et al. [13]: when $\delta < \frac{1}{2}$, the posterior distribution of X becomes increasingly uniform as n increases. In this setting, the adversary's best guess converges to $\frac{1}{q}$, and applying Lemma 1 yields:

$$\mathsf{Adv}_X[\boldsymbol{l} \leftarrow \boldsymbol{L}(\boldsymbol{X})] \leq 2^{n-1}\delta^n.$$

Case of $q = 2$. We prove that for binary fields the given bound is tight.

Lemma 5. *In the setting of Lemma 4, if the underlying field is $\mathbb{F}_2$, then*

$$\mathsf{Adv}_X[\boldsymbol{l} \leftarrow \boldsymbol{L}(\boldsymbol{X})] = 2^{n-1}\delta^n.$$

Proof. We extend a technique by Wyner [33], originally developed for *wiretap channels*.

To recover $X = X_1 \oplus \cdots \oplus X_n$ from the leakage vector $\boldsymbol{L}(\boldsymbol{X})$, the adversary estimates each X_i individually. Let $\hat{X}_i$ denote the estimate of X_i based on $\mathsf{L}(X_i)$, and let $\mathsf{e}_i = \Pr(\hat{X}_i \neq X_i)$ be the average error probability, taken over both the uniform choice of $X_i \in \{0, 1\}$ and the leakage randomness:

$$\mathsf{e}_i = \mathbb{E}_{X_i,\, l \leftarrow \mathsf{L}(X_i)} \left[\Pr(\hat{X}_i \neq X_i \mid l) \right].$$

Assume $\mathsf{e}_i \leq \frac{1}{2}$ (the case $\mathsf{e}_i > \frac{1}{2}$ is similar). The corresponding advantage is then

$$\mathsf{Adv}_{X_i} = (1 - \mathsf{e}_i) - \tfrac{1}{2} = \tfrac{1}{2} - \mathsf{e}_i.$$

By Lemma 1, we have $\mathsf{Adv}_{X_i} = \delta$, implying $\mathsf{e}_i = \frac{1}{2} - \delta$. Let $\mathsf{e} = \frac{1}{2} - \delta$ for simplicity.

The adversary computes $\hat{X} = \hat{X}_1 \oplus \cdots \oplus \hat{X}_n$. This estimate equals the true value X if an even number of errors occur. Thus, the success probability is:

$$\begin{aligned}
\Pr(\hat{X} = X) &= \sum_{j=0}^{\lfloor n/2 \rfloor} \binom{n}{2j} \mathsf{e}^{2j}(1-\mathsf{e})^{n-2j} \\
&= \frac{1}{2}\left[\sum_{i=0}^{n} \binom{n}{i} \mathsf{e}^i (1-\mathsf{e})^{n-i} + \sum_{i=0}^{n} \binom{n}{i} (-\mathsf{e})^i (1-\mathsf{e})^{n-i} \right] \\
&=_{(\mathrm{I})} \frac{1}{2}\left[(\mathsf{e} + 1 - \mathsf{e})^n + (-\mathsf{e} + 1 - \mathsf{e})^n \right] \\
&= \frac{1}{2}\left[1^n + (1 - 2\mathsf{e})^n \right] = \frac{1}{2} + 2^{n-1}\delta^n,
\end{aligned}$$

where step (I) uses the binomial expansion:

$$(\pm\mathsf{e} + (1-\mathsf{e}))^n = \sum_{i=0}^{n} \binom{n}{i} (\pm\mathsf{e})^i (1-\mathsf{e})^{n-i}.$$

Subtracting the baseline guessing probability $\frac{1}{2}$, we get:

$$\mathsf{Adv}_X[\boldsymbol{l} \leftarrow \boldsymbol{L}(\boldsymbol{X})] = \Pr(\hat{X} = X) - \tfrac{1}{2} = 2^{n-1}\delta^n. \qquad \square$$

The binary case forms the basis for our reasoning over extended fields. Before proceeding, we highlight an important observation.

Optimality of the Reduction at $q = 2$. Lemma 2 establishes that the δ-noisy to ϵ-random reduction is tight for $q = 2$, with $\epsilon_{\min} = 2\delta$. Substituting this into (4), we obtain:

$$\Delta \leq \mathsf{SD}(X\,;\, X \mid \boldsymbol{\phi}^{\boldsymbol{\epsilon_{\min}}}(\boldsymbol{X})) = \left(1 - \tfrac{1}{2}\right) \cdot 2^n \delta^n = 2^{n-1}\delta^n.$$

On the other hand, Lemma 5 gives $\Delta = 2^{n-1}\delta^n$, implying that

$$\Delta = \mathsf{SD}(X\,;\, X \mid \boldsymbol{\phi}^{\boldsymbol{\epsilon_{\min}}}(\boldsymbol{X})).$$

This equality shows that the reduction is tight not only for a single variable, but also in the context of mask encoding.[1]

Furthermore, when the leakage functions differ across shares—each with corresponding parameter δ_i—the bound generalizes to:

$$\Delta = 2^{n-1} \prod_{i=1}^{n} \delta_i.$$

For $q > 2$, the Bound in Lemma 4 is Loose We now present a concrete example for $q = 4$, illustrating that the bound in Lemma 4 is not tight in general.

Example 2. Let $X \in \mathbb{F}_{2^2}$, and define the leakage function as

$$\mathsf{L}(X) = (x_1 \oplus e_1) \,\|\, (x_0 \oplus e_0),$$

where x_0, x_1 are the bits of X, and e_0, e_1 are independent Bernoulli variables with $\Pr(e_i = 1) = \mathsf{e} < \frac{1}{2}$. A direct computation yields:

$$\delta = \mathsf{SD}(X\,;\, X \mid \mathsf{L}(X)) = \left(\tfrac{1}{2} - \mathsf{e}\right)\left(\tfrac{3}{2} - \mathsf{e}\right). \tag{5}$$

This leakage is effectively the concatenation of two independent binary leakages. From previous results, we know that under such a leakage model, a masked bit is recovered with error probability:

$$\mathsf{e}_n = \tfrac{1}{2}\left[1 - (1-2\mathsf{e})^n\right].$$

[1] That is, for the metrics Δ and Adv_X. The tightness may not extend to other metrics, such as the mutual information between $\boldsymbol{L}(\boldsymbol{X})$ and X.

Using this, we define an equivalent leakage function:

$$\mathsf{L}'(X) = (x_1 \oplus e'_1) \parallel (x_0 \oplus e'_0),$$

where $e'_i \sim \mathrm{Ber}(\mathsf{e}_n)$. Then, by applying the same structure as in (5), we get:

$$\begin{aligned}\Delta &= \mathsf{SD}(X\,;\, X \mid \boldsymbol{L}(\boldsymbol{X})) = \mathsf{SD}(X\,;\, X \mid \mathsf{L}'(X)) \\ &= \left(\tfrac{1}{2} - \mathsf{e}_n\right)\left(\tfrac{3}{2} - \mathsf{e}_n\right) \\ &= 2^{n-1}\left(\tfrac{1}{2} - \mathsf{e}\right)^n \left(1 + \tfrac{1}{2}(1-2\mathsf{e})^n\right).\end{aligned}$$

In contrast, Lemma 4 gives the bound:

$$\Delta \leq 2^{n-1}\delta^n,$$

which is looser than our exact computation:

$$\begin{aligned}\Delta &= 2^{n-1}\left(\tfrac{1}{2} - \mathsf{e}\right)^n \left(1 + \tfrac{1}{2}(1-2\mathsf{e})^n\right) \\ &<_{(\mathsf{I})} 2^{n-1}\left(\tfrac{1}{2} - \mathsf{e}\right)^n \left(1 + \tfrac{1}{2}(1-2\mathsf{e})\right)^n = 2^{n-1}\delta^{n-1},\end{aligned}$$

where step (I) uses the inequality:

$$(1 + \tfrac{1}{2}t^n) < (1 + \tfrac{1}{2}t)^n \quad \text{for } 0 < t < 1,\ n > 1. \quad \square$$

Need for More Fine-Tuned Analysis. Our findings thus far indicate that for $q = 2^u$ with $u > 1$, the standard δ-noisy to ϵ-random probing reduction—such as in the case of $\mathsf{L}(X) = \mathsf{ZV}(X)$—and indirect metric estimates (as illustrated in Example 2) introduce a non-negligible gap. The central contribution of this paper is to close this gap via a new *decomposition-based* approach.

3.2 Decomposition into Binary Subfields

The observation that exact metrics are tractable and the reduction is tight in binary fields motivates a decomposition strategy: we reduce computations in $\mathbb{F}_{2^u}$ to binary relations, where metrics can be efficiently analyzed, and then lift the results back. We begin by outlining the foundational concepts.

Consider two u-bit integers, A and B, and define their bitwise *inner product* as:

$$\langle A, B\rangle = \bigoplus_{i=0}^{u-1} a_i b_i,$$

where a_i and b_i are the ith bits of A and B, respectively.

Let $X \in \mathbb{F}_{2^u}$ be a random variable with a masked encoding $\boldsymbol{X} = (X_1, \ldots, X_n)$. For any integer $h \in \{1, \ldots, 2^u - 1\}$—interpreted via its u-bit binary representation—we can project the equality $X = X_1 \oplus \cdots \oplus X_n$ into the binary domain as:

$$\langle X, h\rangle = \langle X_1, h\rangle \oplus \cdots \oplus \langle X_n, h\rangle.$$

We will later formalize the validity of this mapping in the context of Boolean systems. For now, to illustrate its practical utility, Lemma 6 will show that if the adversary fails to recover any of the $2^u - 1$ binary projections $\langle X, h\rangle$ from the leakage $\mathsf{L}(X)$, then they cannot infer X with meaningful advantage.

Lemma 6. *Given the leakage* $\mathsf{L}(X)$ *for* $X \in \mathbb{F}_{2^u}$*, the adversary's advantage in recovering* X *satisfies:*

$$\frac{1}{2^{u-1}} \max_h \mathsf{Adv}_{\langle X,h\rangle} \;\le\; \mathsf{Adv}_X \;\le\; \frac{1}{2^{u-1}} \sum_{h=1}^{2^u-1} \mathsf{Adv}_{\langle X,h\rangle} \;<\; 2 \max_h \mathsf{Adv}_{\langle X,h\rangle}.$$

Here, $\mathsf{Adv}_{\langle X,h\rangle}$ *denotes the adversary's advantage in recovering the binary inner product* $\langle X, h\rangle$*.*

Proof. Let us denote $\mu_h = \mathsf{Adv}_{\langle X,h\rangle}$. Given the set $\{\mu_1, \ldots, \mu_{2^u-1}\}$, we aim to bound Adv_X from above and below.

Let $\{p_0, p_1, \ldots, p_{2^u-1}\}$ denote the posterior distribution of X given the leakage realization l. For any h, the advantage μ_h can be written as:

$$\begin{aligned}
\mu_h &= \mathop{\mathbb{E}}_{l} \left[\max_{\alpha\in\{0,1\}} \Pr(\langle X, h\rangle = \alpha \mid l) \right] - \frac{1}{2} \\
&= \mathop{\mathbb{E}}_{l} \left[\max_{\alpha\in\{0,1\}} \sum_{i\in[0,2^u-1],\langle i,h\rangle=\alpha} p_i \right] - \frac{1}{2} \\
&= \mathop{\mathbb{E}}_{l} \left[\max_{\alpha\in\{0,1\}} \left(\sum_{\langle i,h\rangle=\alpha} p_i - \frac{1}{2} \right) \right] \\
&=_{(\mathsf{I})} \frac{1}{2}\, \mathbb{E}_l \left[\left| \sum_{\langle i,h\rangle=0} p_i - \sum_{\langle i,h\rangle=1} p_i \right| \right] \\
&= \frac{1}{2}\, \mathbb{E}_l \left[\left| \sum_{i=0}^{2^u-1} p_i (-1)^{\langle i,h\rangle} \right| \right].
\end{aligned}$$

In step (I), we used the fact that $a + b = 1 \Rightarrow |a - b| = 2(\max\{a, b\} - \frac{1}{2})$. Let us now define the following quantity:

$$\theta_h \triangleq \sum_{i=0}^{2^u-1} p_i (-1)^{\langle i,h\rangle},$$

so that $\mathbb{E}_l[|\theta_h|] = 2\mu_h$.
Now, by the inverse Walsh-Hadamard transform, we can express p_i as:

$$p_i = \frac{1}{2^u} \sum_{h=0}^{2^u-1} \theta_h (-1)^{\langle i,h\rangle} = \frac{1}{2^u} \left(\theta_0 + \sum_{h=1}^{2^u-1} \theta_h (-1)^{\langle i,h\rangle} \right),$$

and since $\theta_0 = \sum_i p_i = 1$, we have:

$$p_i - \frac{1}{2^u} = \frac{1}{2^u} \sum_{h=1}^{2^u-1} \theta_h (-1)^{\langle i,h\rangle}. \tag{6}$$

For derivation of upper bounds in the lemma, we can write:

$$\begin{aligned}\mathsf{Adv}_X &= \mathop{\mathbb{E}}_{l}\left[\max_i p_i - \frac{1}{2^u}\right] = \frac{1}{2^u}\mathop{\mathbb{E}}_{l}\left[\max_i[\sum_{h=1}^{2^u-1}\theta_h(-1)^{\langle i,h\rangle}]\right] \\ &\le \frac{1}{2^u}\mathop{\mathbb{E}}_{l}\left[\max_i\left|\sum_{h=1}^{2^u-1}\theta_h(-1)^{\langle i,h\rangle}\right|\right] \le \frac{1}{2^u}\mathop{\mathbb{E}}_{l}\left[\sum_{h=1}^{2^u-1}|\theta_h|\right] \\ &= \frac{1}{2^u}\sum_{h=1}^{2^u-1}2\mu_h = \frac{1}{2^{u-1}}\sum_{h=1}^{2^u-1}\mu_h < 2\max_h\ \mu_h.\end{aligned}$$

To derive the lower bound, let $h^* = \mathrm{argmax}_h\,\mu_h$, and define the random index $J \in \{0,\ldots,2^u-1\}$ uniformly sampled from:

$$\{J \in [0, 2^u-1] \mid (-1)^{\langle J,h^*\rangle} = \mathrm{sign}(\theta_{h^*})\}.$$

We have:

$$\begin{aligned}\mathsf{Adv}_X &= \frac{1}{2^u}\ \mathop{\mathbb{E}}_{l}\left[\max_i \sum_{h=1}^{2^u-1}\theta_h(-1)^{\langle i,h\rangle}\right] \\ &\ge \frac{1}{2^u}\ \mathop{\mathbb{E}}_{l}\left[\mathop{\mathbb{E}}_{J}\left[\sum_{h=1}^{2^u-1}\theta_h(-1)^{\langle J,h\rangle}\right]\right] \\ &= \frac{1}{2^u}\ \mathop{\mathbb{E}}_{l}\left[\sum_{h=1}^{2^u-1}\theta_h\ \mathop{\mathbb{E}}_{J}[(-1)^{\langle J,h\rangle}]\right].\end{aligned}$$

The inner expectation simplifies because: - For $h = h^*$, we have $\mathbb{E}_J[(-1)^{\langle J,h^*\rangle}] = \mathrm{sign}(\theta_{h^*})$; - For $h \neq h^*$, since $\langle J,h\rangle$ is independent of $\langle J,h^*\rangle$, we get $\mathbb{E}_J[(-1)^{\langle J,h\rangle}] = 0$.
Therefore:

$$\mathsf{Adv}_X \ge \frac{1}{2^u}\ \mathop{\mathbb{E}}_{l}[\theta_{h^*}\cdot\mathrm{sign}(\theta_{h^*})] = \frac{1}{2^u}\ \mathop{\mathbb{E}}_{l}[|\theta_{h^*}|] = \frac{1}{2^u}\cdot 2\mu_{h^*} = \frac{\mu_{h^*}}{2^{u-1}}.$$

This completes the proof. □

3.3 Tightness of the Decomposition

We continue with the convention from the previous proof and denote $\mathsf{Adv}_{\langle X,h\rangle}$ by μ_h. The decomposition-based approach bounds the adversary's advantage by:

$$\mathsf{Adv}_X \le \frac{1}{2^{u-1}}\sum_{h\neq 0}\mu_h.$$

Our goal in this subsection is twofold: first, we show that this upper bound is tight (i.e., achievable); second, we demonstrate that for any δ-noisy leakage function, the quantity $\frac{1}{2^{u-1}}\sum_{h\neq 0}\mu_h$ remains strictly below δ.

Achievability. Consider the trivial case where the leakage function fully reveals the variable, i.e., $\mathsf{L}(X) = X$. Then, for every nonzero h, the adversary perfectly learns the binary variable $\langle X, h\rangle$, and thus $\mu_h = \frac{1}{2}$.

Substituting into the decomposition bound gives:

$$\mathsf{Adv}_X \leq \frac{1}{2^{u-1}} \cdot \left((2^u - 1) \cdot \frac{1}{2}\right) = 1 - \frac{1}{2^u}.$$

On the other hand, since the adversary recovers X exactly, we have:

$$\mathsf{Adv}_X = \Pr(\hat{X} = X) - \frac{1}{2^u} = 1 - \frac{1}{2^u}.$$

Hence, the bound is tight in this case.

Lemma 7. *Let $X \in \mathbb{F}_{2^u}$ be a uniform variable with leakage function $\mathsf{L}(X)$ such that $\mathsf{SD}(X\,;\, X \mid \mathsf{L}(X)) = \delta$. Then:*

$$\frac{1}{2^{u-1}} \sum_{h=1}^{2^u-1} \mu_h < \delta.$$

Proof. Let $\{p_0, p_1, \ldots, p_{2^u-1}\}$ denote the posterior distribution of X given a leakage instance l. Since $\langle X, h\rangle$ is a binary random variable, we can express μ_h as:

$$\mu_h = \frac{1}{2} \mathop{\mathbb{E}}_{l} \left[\left| \sum_{\langle i,h\rangle=0} p_i - \frac{1}{2} \right| + \left| \sum_{\langle i,h\rangle=1} p_i - \frac{1}{2} \right| \right]. \tag{7}$$

Let us define:

$$q_i \triangleq p_i - \frac{1}{2^u}, \quad \text{so that} \quad \sum_i q_i = 0.$$

Then, the definitions of δ and μ_h become:

$$\delta = \frac{1}{2} \mathop{\mathbb{E}}_{l} \left[\sum_i |q_i| \right], \quad \mu_h = \frac{1}{2} \mathop{\mathbb{E}}_{l} \left[\left| \sum_{\langle i,h\rangle=0} q_i \right| + \left| \sum_{\langle i,h\rangle=1} q_i \right| \right]. \tag{8}$$

While the triangle inequality immediately gives $\mu_h \leq \delta$, we aim to show that the average over all nonzero h is smaller than $\frac{1}{2}\delta$. This requires a refined argument.

Partitioning Strategy. We begin by partitioning each $q_\beta > 0$ into non-negative components:

$$q_\beta = \sum_{\alpha=0}^{2^u-1} a_{\beta,\alpha}, \quad \text{where } a_{\beta,\alpha} \geq 0.$$

These components are then redistributed to offset the negative entries, by defining:

$$q_\alpha = -\sum_\beta a_{\beta,\alpha}, \quad \text{for all } q_\alpha \leq 0.$$

	q_3	q_4	q_5	q_6	q_7
q_0	$a_{0,3}$	$a_{0,4}$	$a_{0,5}$	$a_{0,6}$	$a_{0,7}$
q_1	$a_{1,3}$	$a_{1,4}$	$a_{1,5}$	$a_{1,6}$	$a_{1,7}$
q_2	$a_{2,3}$	$a_{2,4}$	$a_{2,5}$	$a_{2,6}$	$a_{2,7}$

We illustrate this partitioning strategy with a simple example. Let $u = 3$, and suppose that q_0, q_1, q_2 are positive while q_3 through q_7 are negative. The partitioning can be visualized as the following table:

For brevity, we omit the details of how this partitioning is constructed.

Now consider the effect on absolute value terms such as $|q_\beta + q_\alpha + c|$, where $q_\beta > 0$, $q_\alpha \leq 0$, and c is the sum of the remaining terms. We can show that $|q_\beta + q_\alpha| \leq |q_\beta| + |q_\alpha| - 2a_{\beta,\alpha}$, and we derive:

$$|q_\beta + q_\alpha + c| \leq |q_\beta + q_\alpha| + |c| \leq |q_\beta| + |q_\alpha| - 2a_{\beta,\alpha} + |c|.$$

Loss Factor from Pairwise Cancellation. Applying this to all sums in μ_h, we observe that each pair $(q_\beta > 0, q_\alpha < 0)$ appears in the same half of the partition (either 0 or 1 side) for half of the nonzero h's. Hence, each such pair contributes a "loss" of at least:

$$\frac{2^u - 1}{2} \cdot 2a_{\beta,\alpha}.$$

Summing over all such pairs yields a total loss of:

$$(2^u - 1) \sum_{\beta,\alpha} a_{\beta,\alpha} = (2^u - 1) \sum_{\alpha, q_\alpha < 0} (-q_\alpha),$$

where we used $\sum_\beta a_{\beta,\alpha} = -q_\alpha$. Since $\sum_i q_i = 0$, the total positive mass equals the total negative mass:

$$\sum_{\alpha, q_\alpha < 0} (-q_\alpha) = \sum_{\beta, q_\beta > 0} q_\beta = \frac{1}{2} \sum_i |q_i|.$$

Final Bound. Combining everything:

$$\begin{aligned} \sum_{h=1}^{2^u - 1} \mu_h &\leq \frac{2^u - 1}{2} \mathop{\mathbb{E}}_{l} \left[\sum_i |q_i| \right] - \frac{2^u - 1}{2} \cdot \frac{1}{2} \mathop{\mathbb{E}}_{l} \left[\sum_i |q_i| \right] \\ &= \frac{2^u - 1}{4} \mathop{\mathbb{E}}_{l} \left[\sum_i |q_i| \right] = \frac{2^u - 1}{2} \delta. \end{aligned}$$

Dividing both sides by 2^{u-1}, we get:

$$\frac{1}{2^{u-1}} \sum_{h=1}^{2^u - 1} \mu_h \leq \frac{2^u - 1}{2^u} \delta < \delta,$$

which concludes the proof. □

3.4 Application of Decomposition to Mask Encodings

Lemma 6 expresses the side-channel security of a variable $X \in \mathbb{F}_{2^u}$ in terms of the security of its binary projections $\langle X, h\rangle$. This decomposition reduces the complex task of estimating $\mathsf{Adv}_X[l \leftarrow \mathsf{L}(X)]$ to the simpler computation of the binary advantages $\mu_h = \mathsf{Adv}_{\langle X,h\rangle}[l \leftarrow \mathsf{L}(X)]$. The strength of this approach becomes especially apparent when applied to masked encodings.

Consider the setting of masked encoding, where a secret X is shared as $\boldsymbol{X} = (X_1, \ldots, X_n)$, and the adversary observes $\boldsymbol{L}(\boldsymbol{X}) = (\mathsf{L}(X_1), \ldots, \mathsf{L}(X_n))$. For each nonzero $h \in \{1, \ldots, 2^u - 1\}$, the inner product satisfies the binary relation:

$$\langle X, h\rangle = \langle X_1, h\rangle \oplus \cdots \oplus \langle X_n, h\rangle.$$

That is, $\langle X, h\rangle$ is itself masked by the shares $\langle X_i, h\rangle$, making it a binary secret with a standard masked encoding.

Applying Lemma 5, the advantage of the adversary in recovering $\langle X, h\rangle$ from the masked leakage is given by:

$$\mathsf{Adv}_{\langle X,h\rangle}[\boldsymbol{l} \leftarrow \boldsymbol{L}(\boldsymbol{X})] = 2^{n-1}\left(\mathsf{Adv}_{\langle X,h\rangle}[l \leftarrow \mathsf{L}(X)]\right)^n = 2^{n-1}\left(\mu_h\right)^n,$$

where $\mu_h = \mathsf{Adv}_{\langle X_i,h\rangle}[l \leftarrow \mathsf{L}(X_i)]$ is the advantage from a single share.

Combining this with the upper bound from Lemma 6, we obtain:

$$\mathsf{Adv}_X[\boldsymbol{l} \leftarrow \boldsymbol{L}(\boldsymbol{X})] \le \frac{1}{2^{u-1}} \sum_{h=1}^{2^u-1} 2^{n-1}\mu_h^n = \frac{1}{2^u} \sum_{h=1}^{2^u-1} (2\mu_h)^n <_{(\mathsf{I})} \Delta \le_{(\mathsf{II})} 2^{n-1}\delta^n,$$

where:

- $\Delta = \mathsf{SD}(X\,;\, X \mid \boldsymbol{L}(\boldsymbol{X}))$;
- (I) follows from Lemma 7, applied to the masked case;
- (II) follows from Lemma 4.

For clarity, we restate the result as a theorem.

Theorem 1. *Let μ_h denote the adversary's advantage in predicting $\langle X, h\rangle$ from the leakage $\mathsf{L}(X)$, for $h \in \{1, \ldots, 2^u - 1\}$. Then, in a masked encoding with share leakage $\boldsymbol{L}(\boldsymbol{X})$, the adversary's advantage satisfies:*

$$\frac{1}{2^u} \max_h (2\mu_h)^n \;\le\; \mathsf{Adv}_X[\boldsymbol{l} \leftarrow \boldsymbol{L}(\boldsymbol{X})] \;\le\; \frac{1}{2^u} \sum_{h=1}^{2^u-1} (2\mu_h)^n \;<\; \max_h (2\mu_h)^n.$$

Interpretation in Terms of Mutual Information Theorem 1 shows that if $2\mu_h < 1$ for all h, then the adversary's advantage Adv_X decays exponentially with the number of shares n. In this subsection, we connect the condition $\mu_h < \frac{1}{2}$ to the mutual information between the binary projection $\langle X, h\rangle$ and the leakage $\mathsf{L}(X)$. Specifically, we show that:

$$\mathsf{MI}\big(\langle X, h\rangle\,;\, \mathsf{L}(X)\big) < 1 \quad \Longrightarrow \quad \mu_h < \tfrac{1}{2}.$$

By definition of mutual information for a binary variable [8], we have:

$$\mathsf{MI}\big(\langle X, h\rangle;\, \mathsf{L}(X)\big) \;=\; 1 \;-\; \mathsf{H}\left(\tfrac{1}{2} \pm \mu_h\right), \tag{9}$$

where $\mathsf{H}(\cdot)$ denotes the binary entropy function.[2]

The mutual information reaches its maximum value of 1 only when $\mathsf{H}(\frac{1}{2} \pm \mu_h) = 0$, i.e., when $\mu_h = \frac{1}{2}$. This implies that the leakage $\mathsf{L}(X)$ fully determines $\langle X, h\rangle$. Therefore, whenever

$$\mathsf{MI}\big(\langle X, h\rangle;\, \mathsf{L}(X)\big) < 1,$$

the corresponding advantage μ_h must satisfy $\mu_h < \frac{1}{2}$, ensuring that masking offers meaningful security for the binary component $\langle X, h\rangle$.

Revisiting Previous Examples

Example 3. We revisit Example 2 to highlight the strength of the decomposition approach. For $X \in \mathbb{F}_{2^2}$ with leakage $\mathsf{L}(X) = (x_1 \oplus e_1) \,\|\, (x_0 \oplus e_0)$, where $\Pr(e_0 = 1) = \Pr(e_1 = 1) = \mathsf{e}$, we had:

$$\delta = \mathsf{SD}(X\,;\, X \mid \mathsf{L}(X)) = \left(\tfrac{1}{2} - \mathsf{e}\right)\left(\tfrac{3}{2} - \mathsf{e}\right).$$

Under mask encoding, we previously derived:

$$\Delta = \mathsf{SD}(X\,;\, X \mid \boldsymbol{L}(\boldsymbol{X})) = 2^{n-1}\left(\tfrac{1}{2} - \mathsf{e}\right)^n \left(1 + \tfrac{1}{2}(1 - 2\mathsf{e})^n\right).$$

For this leakage, Example 1 (with $u = 1$) yields:

$$\mu_1 = \mu_2 = \tfrac{1}{2} - \mathsf{e},$$

corresponding to the adversary's advantage in predicting x_0 and x_1, respectively. Similarly, for μ_3, which corresponds to predicting $x_0 \oplus x_1$, we compute:

$$\mu_3 = \tfrac{1}{2} - 2\mathsf{e}(1 - \mathsf{e}).$$

By Lemma 6, the decomposition-based bound for unmasked X becomes:

$$\mathsf{Adv}_X[l \leftarrow \mathsf{L}(X)] \le \tfrac{1}{2}(\mu_1 + \mu_2 + \mu_3) = \left(\tfrac{1}{2} - \mathsf{e}\right)\left(\tfrac{3}{2} - \mathsf{e}\right) = \delta.$$

For masked encoding, Theorem 1 gives:

$$\mathsf{Adv}_X[\boldsymbol{l} \leftarrow \boldsymbol{L}(\boldsymbol{X})] \le 2^{n-2}\left(\mu_1^n + \mu_2^n + \mu_3^n\right) = \Delta.$$

[2] The mutual information of a Binary Symmetric Channel (BSC) between A and B is

$$\mathsf{MI}(A; B) = \mathsf{H}(A) - \mathsf{H}(A \mid B) = 1 - \mathsf{H}(P_e),$$

where P_e is the probability of incorrectly estimating A given B. In our setting, $\mu_h = |P_e - \frac{1}{2}|$, so $P_e = \frac{1}{2} \pm \mu_h$.

This confirms the tightness of our bound. For instance, setting $\mathsf{e} = 0.1$, we get $\delta = 0.56$. Since $\delta > \frac{1}{2}$, Lemma 4 cannot guarantee masking security. Meanwhile, mutual information each share is $\mathsf{MI}(X_i; \mathsf{L}(X_i)) = 1.06$, which exceeds the 0.72 threshold required by Ito et al. [21] for masking to be secure (see Subsection 1.1). Thus, neither criterion provides a conclusive answer—while our method confirms that masking is indeed secure in this case. □

Example 4. Recall the leakage function $\mathsf{ZV}(X)$, previously introduced as:

$$\mathsf{ZV}(X) = \begin{cases} \nu_a & \text{if } X = 0, \\ \nu_b \neq \nu_a & \text{otherwise.} \end{cases}$$

For this model, (1) yields $\epsilon_{\min} = 1$, rendering the δ-noisy to ϵ-random probing reduction inapplicable for analyzing masking security.

In contrast, applying our decomposition approach, we compute each $\mu_h = \frac{1}{2^u}$ using relation (7). Then, by Theorem 1, the adversary's advantage under masking satisfies:

$$\mathsf{Adv}_X[\boldsymbol{l} \leftarrow \boldsymbol{L}(\boldsymbol{X})] \leq 2^{n-u} \sum_{h=1}^{2^u-1} \mu_h^n = (2^u - 1)\, 2^{n-u-nu},$$

which decreases exponentially with n, provided $u > 1$. Thus, even though the standard reduction fails, our approach confirms that masking remains effective in this setting. □

Application to Leakage Certification. Leakage certification laboratories evaluate cryptographic implementations on physical devices to assess their resistance against side-channel attacks (see [11, 30]). These evaluations often rely on estimating metrics such as $\mathsf{MI}(X\,;\,\mathsf{L}(X))$ and $\mathsf{SD}(X\,;\,X \mid \mathsf{L}(X))$. Estimating these quantities requires knowledge of the joint distribution $(X, \mathsf{L}(X))$, which can be derived either through parametric models (e.g., Gaussian with estimated parameters) or non-parametric methods (e.g., histogram-based) [2, 16].

Our work introduces an additional criterion for leakage assessment. Specifically, for a u-bit variable X, we propose verifying the condition:

$$\mathsf{MI}\big(\langle X, h\rangle\,;\,\mathsf{L}(X)\big) < 1 \quad \text{for all } h \in [1, 2^u - 1].$$

Masking provides meaningful side-channel protection if and only if this condition holds for every nontrivial bitwise projection $\langle X, h\rangle$.

Using Equation (9), we can express the bounds from Theorem 1 in terms of mutual information:

$$\frac{1}{2^u}\left(2 \max_h \left|\tfrac{1}{2} - \mathsf{H}^{-1}(I_h)\right|\right)^n \;\leq\; \mathsf{Adv}_X[\boldsymbol{l} \leftarrow \boldsymbol{L}(\boldsymbol{X})] \;<\; \left(2 \max_h \left|\tfrac{1}{2} - \mathsf{H}^{-1}(I_h)\right|\right)^n,$$

where $I_h = 1 - \mathsf{MI}\big(\langle X, h\rangle\,;\,\mathsf{L}(X)\big)$, and H^{-1} is the inverse of the binary entropy function.

3.5 Experimental Results and Determining Masking Order

We perform our experiments on a ChipWhisperer CW308 board hosting an STM32F303 UFO processor,[3] running an unprotected 8-bit software implementation of AES (included in the ChipWhisperer package). Power traces are captured using a ChipWhisperer-Lite, synchronized at a sampling rate four times the target's clock frequency.

The targeted intermediate value is the output of the first S-box in the first round:

$$X = \mathsf{S\text{-}box}(P[0] \oplus K[0]).$$

A correlation power analysis (CPA) using only 20 traces successfully recovers the secret byte, indicating a low noise level and motivating the need to assess the feasibility and necessary order of masking.

To this end, we estimate the bias parameters μ_h through profiling. Specifically, we train a deep neural network to predict X from power traces. The network processes one-dimensional input traces and has the following architecture:

- A 1D convolutional layer with 32 filters of size 3, followed by ReLU activation.
- A max-pooling layer with pool size 2.
- A second 1D convolutional layer with 64 filters of size 3, followed by ReLU activation.
- Another max-pooling layer with pool size 2.
- A flattening layer, followed by a dense layer with 128 ReLU-activated neurons.
- A final dense output layer with 256 neurons (corresponding to all possible S-box output values), followed by softmax activation.

We trained this model on 10,000 labeled traces with ground truth $X = \mathsf{S\text{-}box}(P[0] \oplus K[0])$. In the attack phase, the model's output corresponds to the posterior distribution of X given the leakage in a trace.

Using this posterior, we estimate the binary error probability P_e for computing the bitwise projection $\langle X, h \rangle = \bigoplus_{i=0}^{7} x_i h_i$ for each $h \in \{1, \ldots, 255\}$. The neural network demonstrates high accuracy in estimating these projections.

Figure 1 shows the estimated P_e for all values of h, based on 1,000 attack traces. The minimum error occurs at $h = 255$, which corresponds to $\mathsf{mod}(\mathsf{HW}(X), 2)$. The corresponding (maximum) bias is:

$$\mu_h = (1 - P_e) - \tfrac{1}{2} = 0.32.$$

Substituting this value into the bound from Theorem 1, we obtain the following security bounds for masking in this setting:

$$\frac{1}{256} \cdot (0.64)^n \;\leq\; \mathsf{Adv}_X[\boldsymbol{l} \leftarrow \boldsymbol{L}(\boldsymbol{X})] \;<\; (0.64)^n.$$

These bounds illustrate that while the adversary retains a noticeable advantage for small n, increasing the masking order rapidly improves security.

[3] `https://github.com/newaetech/chipwhisperer`

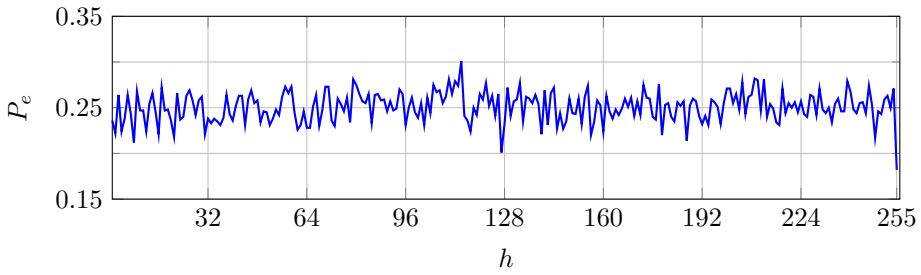

Fig. 1: Binary error probability P_e in estimating $\langle X, h \rangle$.

4 Masking in Odd Prime Fields

Grassi et al. [18], building on earlier results by Dziembowski et al. [13], showed that masking becomes ineffective in the presence of highly informative leakage functions—such as $\mathsf{L}(X) = \mathsf{HW}(X)$—unless the masking is performed over an odd prime field. This insight has sparked further investigation into the security properties of masking in prime fields [14].

In this section, we contribute to this line of work with the following findings:

- In Subsection 4.1, we observe that the adversary's advantage Adv_X decays more rapidly in prime fields.
- In Subsection 4.2, we establish a general condition under which masking guarantees a strict reduction in Adv_X.
- In Subsection 4.3, we show that in prime fields, Adv_X decreases exponentially with n as long as $\min_h \mathsf{MI}(\langle X, h \rangle; \mathsf{L}(X)) < 1$. This condition implies that there exists at least one projection not fully leaked to the adversary—a substantially more relaxed requirement compared to the case of binary extension fields.
- In Subsection 4.4, we prove a conjecture posed by Dziembowski et al. [13] concerning masking in additive groups such as $\mathbb{Z}_m$, where m is neither a prime nor a power of two.

4.1 Adv_X Decays Faster in Prime Fields

We define a leakage class as *symmetric* if, for a uniform $X \in \mathbb{F}_q$, observing an instance l of the leakage transforms the posterior distribution $X \mid l$ into $(p_{e_0}, p_{e_1}, \ldots, p_{e_{q-1}})$, where $\sum_{i=0}^{q-1} p_{e_i} = 1$ and p_{e_i} denotes the probability mass on the value $X + i$.[4] In this notation, p_{e_0} represents the probability of a correct estimation, and the adversary's advantage is:

$$\mathsf{Adv}_X = p_{e_0} - \frac{1}{q}.$$

[4] We refer to this as a symmetric leakage class because it generalizes the binary symmetric channel.

Consider a masked encoding with $n = 2$, where X_1 and X_2 are the shares of X such that $X = X_1 + X_2$. Upon observing the leakages, the adversary forms estimates $\hat{X}_1$ and $\hat{X}_2$, and reconstructs X via $\hat{X}_1 + \hat{X}_2$. The reconstruction is correct if either both estimates are correct or the errors in $\hat{X}_1$ and $\hat{X}_2$ cancel out (i.e., are i and $q - i$). Hence, the updated success probability becomes:

$$p'_{e_0} = (p_{e_0})^2 + \sum_{i=1}^{q-1} p_{e_i} p_{e_{q-i}}. \tag{10}$$

Lemma 8. *For the defined symmetric leakage class, when the field order q is prime, Adv_X decays faster with increasing n.*

Proof. In a prime field, the elements i and $q - i$ are distinct for all $1 \leq i \leq \frac{q-1}{2}$, so we can rewrite Equation (10) as:

$$p'_{e_0} = (p_{e_0})^2 + 2 \sum_{i=1}^{\frac{q-1}{2}} p_{e_i} p_{e_{q-i}}.$$

In contrast, when $q = 2^u$, we have $q - i = i$, so the equation simplifies to:

$$p'_{e_0} = (p_{e_0})^2 + \sum_{i=1}^{q-1} p_{e_i}^2.$$

By applying the inequality $2ab \leq a^2 + b^2$, we obtain:

$$(p_{e_0})^2 + 2 \sum_{i=1}^{\frac{q-1}{2}} p_{e_i} p_{e_{q-i}} \leq (p_{e_0})^2 + \sum_{i=1}^{q-1} p_{e_i}^2.$$

This shows that p'_{e_0}, and hence the adversary's success probability and advantage, are lower in prime fields than in binary extension fields. □

4.2 When Does Adv_X Decrease with Masking?

For a masked encoding of the secret X, we show that if there is no *hole* in the posterior distribution of the shares after observing the leakage vector, then Adv_X strictly decreases.

Definitions. For later reference, we define a leakage instance l as *dummy* if it induces no change in the distribution of $X \mid l$.[5] Under dummy leakage, the peak point of the posterior distribution is $\frac{1}{q}$. A *hole* in a distribution is an element of its domain with zero probability mass. The *support* of a random variable refers to the number of domain elements with non-zero probability mass, and we denote the support size of X by $|X|$.

[5] We refer to it as dummy because any random variable that is independent of X has a similar effect.

Problem Statement. Let X_1 and X_2 be shares of X in $\mathbb{F}_q$, and suppose the adversary observes leakage instances $l_1 \leftarrow \mathsf{L}(X_1)$ and $l_2 \leftarrow \mathsf{L}(X_2)$. The resulting posterior distributions are denoted by $\mathcal{P}^1 = (p_0^1, \ldots, p_{q-1}^1)$ and $\mathcal{P}^2 = (p_0^2, \ldots, p_{q-1}^2)$, respectively. Let $p_{i^*}^1$ and $p_{j^*}^2$ denote the peak probabilities in these distributions.

To estimate the value of X, a maximum a posteriori (MAP) adversary computes the distribution of the sum $(X_1 \mid l_1) + (X_2 \mid l_2)$, denoted by $\mathcal{P} = (p_0, \ldots, p_{q-1})$. The adversary outputs the index of the peak point of $\mathcal{P}$ as $\hat{X}$, and the corresponding advantage is denoted $\mathsf{Adv}_X[l_1, l_2 \leftarrow \mathsf{L}(X_1), \mathsf{L}(X_2)]$. Our goal is to identify conditions ensuring:

$$\mathsf{Adv}_X[l_1, l_2 \leftarrow \mathsf{L}(X_1), \mathsf{L}(X_2)] < \mathsf{Adv}_{X_i}[l_i \leftarrow \mathsf{L}(X_i)],$$

which implies that masking strictly improves the side-channel security of X.

Lemma 9. *If, for at least one non-dummy leakage instance, the posterior distributions $\mathcal{P}^1$ and $\mathcal{P}^2$ have no holes—that is, $\min \mathcal{P}^1 > 0$ and $\min \mathcal{P}^2 > 0$—then the adversary's advantage strictly decreases.*

Proof. For $0 \le i < q$, define $\zeta_i = \frac{p_i^1}{p_{i^*}^1}$ and $\xi_i = \frac{p_i^2}{p_{j^*}^2}$. Since $p_{i^*}^1$ and $p_{j^*}^2$ are the maximum values, we have $\zeta_i \le 1$ and $\xi_i \le 1$. From the normalization $\sum_i p_i^1 = 1$, we obtain:

$$p_{i^*}^1 = \frac{1}{\sum_i \zeta_i}, \quad p_{j^*}^2 = \frac{1}{\sum_i \xi_i}.$$

Let k^* be the index of the peak of $\mathcal{P}$. We show that $p_{k^*} \le \min\{p_{i^*}^1, p_{j^*}^2\}$. Without loss of generality, assume $\min = p_{i^*}^1$. Then:

$$p_{k^*} = \sum_i p_{k^*-i}^1 \, p_i^2 = p_{i^*}^1 p_{j^*}^2 \sum_i \zeta_{k^*-i} \, \xi_i \le p_{i^*}^1,$$

because $p_{i^*}^1 > 0$ and each $\zeta_{k^*-i} \le 1$. The inequality holds since:

$$\sum_i \zeta_{k^*-i} \, \xi_i \le \sum_i \xi_i.$$

Equality would require:

$$\forall i, \quad \zeta_{k^*-i} = 1, \tag{11}$$

which implies $p_i^1 = p_{i^*}^1$ for all i, hence a uniform distribution: $p_{i^*}^1 = \frac{1}{q}$. This contradicts the assumption that the leakage is non-dummy. Thus, the strict inequality $p_{k^*} < \min\{p_{i^*}^1, p_{j^*}^2\}$ must hold for at least one instance.

Therefore, taking expectations over all leakage instances, we obtain:

$$\mathbb{E}_l[p_{k^*}] < \mathbb{E}_l[p_{i^*}^1] = \mathbb{E}_l[p_{j^*}^2] \quad \Rightarrow \quad \mathsf{Adv}_X[\boldsymbol{l} \leftarrow \boldsymbol{L}(\boldsymbol{X})] < \mathsf{Adv}_X[l \leftarrow \mathsf{L}(X)]. \quad \square$$

Reaching a Hole-Free Posterior Distribution. When the leakage function is less noisy—e.g., $\mathsf{L}(X) = \mathsf{HW}(X)$—the posterior distribution of $X_i \mid \mathsf{L}(X_i)$ typically contains holes, preventing the application of Lemma 9. However, increasing the number of shares can eliminate such holes, as we now explain.

Although Lemma 9 applies directly to the case $n = 2$, the result generalizes naturally to higher n. Let $X_1, \ldots, X_{2n}$ be shares of X, with corresponding leakages $\mathsf{L}(X_1), \ldots, \mathsf{L}(X_{2n})$. The task of estimating X from its leakage vector can be decomposed into two stages: first, estimating $S_1 = X_1 + \cdots + X_n$ and $S_2 = X_{n+1} + \cdots + X_{2n}$, and then estimating $X = S_1 + S_2$.

If, beyond a certain threshold n_0, the distribution of $S_1 \mid \mathsf{L}(X_1), \ldots, \mathsf{L}(X_n)$ has no holes, then Lemma 9 implies that Adv_X will decrease with n, indicating security improvement through masking.

Because the shares and leakages are independent, the distribution of $S_1 \mid (\mathsf{L}(X_1), \ldots, \mathsf{L}(X_n))$ equals the convolution $X_1 \mid l_1 + \cdots + X_n \mid l_n$. At any leakage instance, each $X_i \mid l_i$ is a distribution, and our goal is to determine when the support of the sum reaches the full domain $\mathbb{F}_q$, which ensures a hole-free distribution.

Lemma 10 (Generalized Cauchy-Davenport Theorem). *Let $Z_1, \ldots, Z_t$ be independent random variables with supports $|Z_1|, \ldots, |Z_t|$ over a prime field $\mathbb{F}_q$. Then:*

$$|Z_1 + \cdots + Z_t| \geq \min\left\{|Z_1| + \cdots + |Z_t| - t + 1,\ q\right\}.$$

Proof. This is a direct generalization of the classical Cauchy-Davenport theorem, which states the result for $t = 2$. The case $t = 2$ was previously used in the side-channel literature by Dziembowski et al. [13]. □

Using this lemma, we conclude that if each $X_i \mid l_i$ has support size greater than 1 (i.e., $|X_i \mid l_i| > 1$), then there exists some $n \geq n_0$ for which the convolution $X_1 \mid l_1 + \cdots + X_n \mid l_n$ is hole-free. Consequently, Adv_X decreases with increasing n.

4.3 Advantage in Terms of $\langle \boldsymbol{X}, \boldsymbol{h} \rangle$ Projections

In the case of fields with characteristic two ($q = 2^u$), we previously reduced the security of masking to the security of binary projections $\langle X, h \rangle$. In this subsection, we explore the implications of these projections for the security of masking in prime fields. To this end, we rely on the following result derived from a bound by Dziembowski et al. [13].

Lemma 11 ([13], Theorem 1). *Let $\mathbb{F}_q$ be a prime field, and suppose the leakage is δ-noisy. That is, $\mathsf{SD}(X\,;\, X \mid \mathsf{L}(X)) = \delta < 1 - \frac{1}{q}$. Then:*

$$\Delta \leq \frac{1}{q} \cdot 2^{-\frac{n\theta^4}{2}},$$

where $\Delta = \mathsf{SD}(X\,;\, X \mid \boldsymbol{L}(\boldsymbol{X}))$ and $\theta = (1 - \frac{1}{q}) - \delta$.

A straightforward corollary (using Lemma 1) yields:

$$\mathsf{Adv}_X[\boldsymbol{l} \leftarrow \boldsymbol{L}(\boldsymbol{X})] \leq \frac{1}{q} \cdot 2^{-\frac{n\theta^4}{2}}. \tag{12}$$

Now, suppose there exists h such that the leakage $\mathsf{L}(X)$ does not fully determine $\langle X, h\rangle$. That is, the bias μ_h of this bit is bounded by $0 \leq \mu_h < \frac{1}{2}$. In the borderline case where $\mu_h \to \frac{1}{2}$, we can approximate the statistical distance as follows:[6]

$$\mathsf{SD}(X\,;\, X \mid \mathsf{L}(X)) \;\lessapprox\; \frac{1}{2} + \mu_h - \frac{1}{q}.$$

Substituting into the bound (12), we obtain:

$$\mathsf{Adv}_X[\boldsymbol{l} \leftarrow \boldsymbol{L}(\boldsymbol{X})] \leq \frac{1}{q} \cdot 2^{-\frac{n(1/2-\mu_h)^4}{2}}, \tag{13}$$

which indicates that the adversary's advantage decreases exponentially in n as long as $\mu_h < \frac{1}{2}$.

We emphasize that the restriction to a single projection is made solely for illustrative purposes; the bound provided by Lemma 11 remains tighter and more general.

4.4 General Requirement for Secure Masking

For a general additive group $\mathbb{G}$, Dziembowski et al. [13] showed that there exists a leakage function with noise parameter $\delta_0 = 1 - \frac{|\mathbb{H}|}{|\mathbb{G}|}$ that renders masking entirely ineffective. Here, $\mathbb{H}$ denotes the largest proper (non-trivial) subgroup of $\mathbb{G}$. They further conjectured that for any leakage with higher noise—i.e., any δ-noisy leakage with $\delta < \delta_0$—masking should become effective. In this subsection, we prove their conjecture for the case where $\mathbb{G}$ is an abelian group and derive an upper bound on the adversary's advantage.

Lemma 12. *Let X be uniformly distributed over an abelian group $\mathbb{G}$ with largest proper subgroup $\mathbb{H}$. If the leakage satisfies*

$$\delta = \mathsf{SD}(X\,;\, X \mid \mathsf{L}(X)) < 1 - \tfrac{|\mathbb{H}|}{|\mathbb{G}|},$$

then:

$$\mathsf{Adv}_X[\boldsymbol{l} \leftarrow \boldsymbol{L}(\boldsymbol{X})] \leq \frac{1}{\alpha} \cdot 2^{-\frac{n\theta^4}{2}},$$

where $\alpha = [\mathbb{G} : \mathbb{H}] = |\mathbb{G}|/|\mathbb{H}|$ is the index of $\mathbb{H}$ in $\mathbb{G}$, and $\theta = 1 - \frac{|\mathbb{H}|}{|\mathbb{G}|} - \delta$.

[6] This approximation assumes a worst-case scenario where the only uncertainty remaining about X after observing $\mathsf{L}(X)$ lies in the value of $\langle X, h\rangle$, which takes values in $\{0, 1\}$ with probabilities $\frac{1}{2} \pm \mu_h$.

Proof. We base our proof on the structure of the quotient group $\mathbb{K} := \mathbb{G}/\mathbb{H}$. Since $\mathbb{H}$ is the largest proper subgroup of $\mathbb{G}$, the quotient group $\mathbb{K}$ has *prime order* [31], and is therefore a cyclic group isomorphic to $\mathbb{Z}_p$ for some prime $p = |\mathbb{K}|$. In this setting, instead of analyzing the security of the full secret $X \in \mathbb{G}$, we study its image under the natural projection:

$$[X] := X + \mathbb{H} \in \mathbb{K},$$

which identifies the coset of X modulo $\mathbb{H}$.

We now consider the adversary's advantage in guessing the correct coset $[X]$ given the leakage. Since the leakage on X induces leakage on $[X]$, we can apply Lemma 11 to the group $\mathbb{K}$. Rewriting the noise threshold in terms of $|\mathbb{K}| = \alpha = |\mathbb{G}|/|\mathbb{H}|$, the condition becomes:

$$\delta < 1 - \frac{1}{|\mathbb{K}|}.$$

Thus, Lemma 11 yields the following bound on the adversary's advantage for the coset:

$$\mathsf{Adv}_{[X]}[\boldsymbol{l} \leftarrow \boldsymbol{L}(\boldsymbol{X})] \leq \frac{1}{|\mathbb{K}|} \cdot 2^{-\frac{n\theta^4}{2}},$$

where $\theta = 1 - \frac{1}{|\mathbb{K}|} - \delta = 1 - \frac{|\mathbb{H}|}{|\mathbb{G}|} - \delta$, as claimed.

It remains to show that:

$$\mathsf{Adv}_{X}[\boldsymbol{l} \leftarrow \boldsymbol{L}(\boldsymbol{X})] \leq \mathsf{Adv}_{[X]}[\boldsymbol{l} \leftarrow \boldsymbol{L}(\boldsymbol{X})].$$

This inequality holds by the *data-processing inequality*: since $[X]$ is a deterministic function of X, the adversary's advantage in guessing the full secret X cannot exceed the advantage in guessing the coarser value $[X]$.

Hence, we conclude:

$$\mathsf{Adv}_{X}[\boldsymbol{l} \leftarrow \boldsymbol{L}(\boldsymbol{X})] \leq \frac{1}{\alpha} \cdot 2^{-\frac{n\theta^4}{2}}. \quad \square$$

5 Security of Linear Gadgets

Our analysis so far has focused on standalone secrets and their masked encodings. We now extend this to more complex structures, specifically linear *gadgets*. Our goal is to demonstrate the applicability of the proposed decomposition technique for evaluating security in such settings.

Gates and Gadgets. A gadget is a *family* of circuits (one per order n) that securely implements a masked version of a gate. Let $\mathsf{G}\colon (\mathbb{F}_{2^u})^t \to \mathbb{F}_{2^u}$ be a gate with fan-in t. For example, XOR and AND have $t = 2$. The corresponding gadget, denoted $\mathbb{S}\mathsf{G}\colon (\mathbb{F}_{2^u}^n)^t \to \mathbb{F}_{2^u}^n$, takes masked inputs and returns masked outputs.

A *refresh gadget* has fan-in and fan-out 1. It re-randomizes the encoding of a value $\boldsymbol{X}$ while preserving its underlying secret, i.e., $\bigoplus_i X_i = \bigoplus_i X'_i$. A well-known example is the $\mathbb{S}\mathsf{R}\text{-}\mathsf{SNI}$ gadget [1], shown in Algorithm 1.

Algorithm 1 $\mathbb{S}$R-SNI

Input $\boldsymbol{X} = (X_1, \ldots, X_n)$
Output $\boldsymbol{X}' = (X'_1, \ldots, X'_n)$

for $i = 1$ **to** n **do**
  for $j = i + 1$ **to** n **do**
    $r \xleftarrow{\$} \mathbb{F}_{2^u}$
    $X_i = X_i \oplus r$
    $X_j = X_j \oplus r$
return $\boldsymbol{X}' = \boldsymbol{X}$

5.1 Problem Statement

Let $\boldsymbol{\Sigma}_n = \{V_1, \ldots, V_{\mathsf{T}(n)}\}$ be the set of intermediate variables in an $\mathbb{F}_2$-linear gadget processing secret $X \in \mathbb{F}_{2^u}$. Besides the shares of X, $\boldsymbol{\Sigma}_n$ may include random values whose leakage could help an adversary estimate X.

The linearity assumption implies the existence of a matrix $\boldsymbol{P}_n \in \mathbb{F}_2^{\mathsf{P}(n)\times(\mathsf{T}(n)+1)}$ such that:

$$\boldsymbol{P}_n \cdot [X, V_1, \ldots, V_{\mathsf{T}(n)}]^\dagger = \boldsymbol{0}. \tag{14}$$

This system describes all parity constraints among the variables, and any additional relations are linear combinations of these rows. The adversary is assumed to know $\boldsymbol{P}_n$ (or any equivalent form).

Side-channel measurements yield:

$$\boldsymbol{L}_n = [\mathsf{L}(V_1), \ldots, \mathsf{L}(V_{\mathsf{T}(n)})],$$

and the adversary's goal is to estimate X from $\boldsymbol{L}_n$ and $\boldsymbol{P}_n$. We denote their advantage as $\mathsf{Adv}_X[\boldsymbol{l}_n \leftarrow \boldsymbol{L}_n]$.

MAP Adversary and Exact Advantage. Let $\mathcal{S}_n$ be the set of all solutions to (14). Each $\boldsymbol{S} \in \mathcal{S}_n$ is a vector of length $\mathsf{T}(n)+1$, with $\boldsymbol{S}(0)$ denoting the value of X. The MAP adversary outputs:

$$\hat{X} = \underset{\alpha \in \mathbb{F}_{2^u}}{\operatorname{argmax}} \sum_{\substack{\boldsymbol{S} \in \mathcal{S}_n \\ \boldsymbol{S}(0) = \alpha}} \Pr(\boldsymbol{S} \mid \boldsymbol{l}_n).$$

The corresponding advantage is:

$$\mathsf{Adv}_X[\boldsymbol{l}_n \leftarrow \boldsymbol{L}_n] = \mathbb{E}_{\boldsymbol{l}_n}\left[\Pr(\hat{X} = X \mid \boldsymbol{l}_n)\right] - \frac{1}{2^u}.$$

A Non-Tight Upper Bound. To obtain a computable upper bound, we apply the reduction from δ-noisy leakage to ϵ-random probing. We replace each leakage $\mathsf{L}(V_i)$ with its erasure version $\phi^\epsilon(V_i)$, where $\epsilon \geq \epsilon_{\min}$ and $\epsilon_{\min}$ is derived from L (cf. (1)). This yields:

$$\mathsf{Adv}_X[\boldsymbol{l}_n \leftarrow \boldsymbol{L}_n] \leq \mathsf{Adv}_X[l_1, \ldots, l_{\mathsf{T}(n)} \leftarrow \phi^\epsilon(V_1), \ldots, \phi^\epsilon(V_{\mathsf{T}(n)})].$$

We denote this bound by $\mathsf{Adv}(q, n, \epsilon)$, where $q = 2^u$. Since smaller ϵ yields less information, we have:

$$\mathsf{Adv}_X[\boldsymbol{l}_n \leftarrow \boldsymbol{L}_n] \leq \mathsf{Adv}(q, n, \epsilon_{\min}) \leq \mathsf{Adv}(q, n, \epsilon). \tag{15}$$

Jahandideh et al. [22] estimated $\mathsf{Adv}(q, n, \epsilon)$ for various gadgets. For the $\mathbb{S}$R-SNI gadget with $n < 30$ and $\epsilon < 0.15$, they showed:

$$\mathsf{Adv}(q, n, \epsilon) \leq \frac{q-1}{q} \cdot \epsilon^{0.6n}. \tag{16}$$

However, for some leakage functions—e.g., $\mathsf{L}(X) = \mathsf{ZV}(X)$—the reduction gives $\epsilon_{\min} = 1$, leading to a trivial bound:

$$\mathsf{Adv}_X[\boldsymbol{l}_n \leftarrow \boldsymbol{L}_n] \leq 1 - \frac{1}{q}.$$

In the next subsection, we show how our decomposition approach avoids this issue and yields tighter bounds in $\mathbb{F}_{2^u}$ settings.

5.2 Decomposition into Binary Systems

In an $\mathbb{F}_2$-linear system such as

$$\boldsymbol{P}_n \cdot [X, V_1, \ldots, V_{\mathsf{T}(n)}]^\dagger = \mathbf{0},$$

the projections $\langle X, h\rangle$ and $\langle V_i, h\rangle$ also satisfy the same linear structure. More precisely:

Lemma 13. *Let $h \in \{1, \ldots, 2^u - 1\}$. Then the system $\boldsymbol{P}_n \cdot [X, V_1, \ldots, V_{\mathsf{T}(n)}]^\dagger = \mathbf{0}$ implies:*

$$\boldsymbol{P}_n \cdot [\langle X, h\rangle, \langle V_1, h\rangle, \ldots, \langle V_{\mathsf{T}(n)}, h\rangle]^\dagger = \mathbf{0}.$$

Proof. Using the linearity of the binary inner product:

$$\langle V_1 \oplus V_2, h\rangle = \langle V_1, h\rangle \oplus \langle V_2, h\rangle, \quad \langle bV, h\rangle = b \cdot \langle V, h\rangle$$

for $b \in \mathbb{F}_2$, the result follows by applying these identities to each row of $\boldsymbol{P}_n$. □

A Tighter Upper Bound for $\mathsf{Adv}_X[\boldsymbol{l}_n \leftarrow \boldsymbol{L}_n]$. The adversary also obtains side-channel leakage from each intermediate variable. From Lemma 6, we derive the following bound for the secret X:

$$\mathsf{Adv}_X[\boldsymbol{l}_n \leftarrow \boldsymbol{L}_n] \leq \frac{1}{2^{u-1}} \sum_{h=1}^{2^u-1} \mathsf{Adv}_{\langle X,h\rangle}[\boldsymbol{l}_n \leftarrow \boldsymbol{L}_n], \tag{17}$$

where each term $\mathsf{Adv}_{\langle X,h\rangle}[\boldsymbol{l}_n \leftarrow \boldsymbol{L}_n]$ corresponds to the adversary's advantage in estimating a single binary projection $\langle X, h\rangle$ under δ-noisy leakage.

Applying the reduction from δ-noisy leakage to ϵ-random probing, we obtain:

$$\mathsf{Adv}_{\langle X,h\rangle}[\boldsymbol{l}_n \leftarrow \boldsymbol{L}_n] \leq \mathsf{Adv}(2, n, \epsilon^h_{\min}),$$

where $\epsilon^h_{\min}$ depends on both the projection index h and the leakage function $\mathsf{L}(V)$. Importantly, while the systems for X and $\langle X, h\rangle$ are structurally the same, the field size is reduced to $\mathbb{F}_2$.

For binary variables, we previously established that:

$$\epsilon^h_{\min} = 2\mu_h, \tag{18}$$

where μ_h denotes the bias in estimating $\langle V, h\rangle$. One way to compute μ_h is via mutual information. As discussed in Section 3.4, it can be expressed as:

$$\mu_h = \left| \frac{1}{2} - \mathsf{H}^{-1}\left[1 - \mathsf{MI}\big(\langle V, h\rangle\,;\,\mathsf{L}(V)\big)\right] \right|,$$

where H^{-1} is the inverse of the binary entropy function.

Putting all of this together, we obtain:

$$\mathsf{Adv}_X[\boldsymbol{l}_n \leftarrow \boldsymbol{L}_n] \leq \frac{1}{2^{u-1}} \sum_{h=1}^{2^u-1} \mathsf{Adv}(2, n, \epsilon^h_{\min}) \leq 2 \cdot \max_h \mathsf{Adv}(2, n, \epsilon^h_{\min}). \tag{19}$$

This shows that the decomposition approach enables a tighter reduction from the δ-noisy to the ϵ-random probing model for $\mathbb{F}_{2^u}$-valued linear circuits.

Example 5. To illustrate the usefulness of the upper bound in (19), consider the leakage function $\mathsf{ZV}(V)$ and the $\mathbb{S}$R-SNI gadget. As discussed in Example 4, for $\mathsf{ZV}(V)$, we have $\mu_h = \frac{1}{2^u}$, implying:

$$\epsilon^h_{\min} = 2\mu_h = \frac{1}{2^{u-1}}.$$

Applying the bound from (16), we obtain:

$$\mathsf{Adv}_X[\boldsymbol{l}_n \leftarrow \boldsymbol{L}_n] \leq \mathsf{Adv}(2, n, 2\mu_h) \leq \frac{1}{2}\left(\frac{1}{2^{u-2}}\right)^{0.6n},$$

which holds as long as $4\mu_h < 0.15$, i.e., $u \geq 5$.

Without the decomposition approach introduced in this work, the minimum erasure rate would be $\epsilon_{\min} = 1$, yielding a trivial upper bound of $\mathsf{Adv}_X[\boldsymbol{l}_n \leftarrow \boldsymbol{L}_n] \leq 1$. □

6 Conclusion

This work establishes necessary and sufficient noise conditions for the security of masked encodings over binary extension fields. We prove that security requires the leakage to conceal *all* linear combinations of an intermediate's bits—resolving

a longstanding open question about the minimal noise needed for secure masking. This insight is particularly important in high-SNR regimes, where existing noise assumptions are overly conservative.

We also demonstrated how the proposed decomposition approach applies to the analysis of masked circuits, with a focus on linear gadgets. By reducing the problem to binary subfields, we enabled efficient estimation of security metrics and more accurate determination of the noise thresholds required for protection.

While our analysis focused on linear structures, extending this framework to non-linear gadgets and complete circuits remains an open challenge. We leave this as promising future work toward advancing noise-based countermeasures in side-channel security.

References

1. Barthe, G., Belaïd, S., Dupressoir, F., Fouque, P.A., Grégoire, B., Strub, P.Y., Zucchini, R.: Strong Non-Interference and Type-Directed Higher-Order Masking. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security. pp. 116–129 (2016)
2. Batina, L., Gierlichs, B., Prouff, E., Rivain, M., Standaert, F., Veyrat-Charvillon, N.: Mutual information analysis: a comprehensive study. J. Cryptol. **24**(2), 269–291 (2011). https://doi.org/10.1007/S00145-010-9084-8, https://doi.org/10.1007/s00145-010-9084-8
3. Battistello, A., Coron, J.S., Prouff, E., Zeitoun, R.: Horizontal side-channel attacks and countermeasures on the isw masking scheme. In: Gierlichs, B., Poschmann, A.Y. (eds.) Cryptographic Hardware and Embedded Systems – CHES 2016. pp. 23–39. Springer Berlin Heidelberg, Berlin, Heidelberg (2016)
4. Béguinot, J., Cheng, W., Guilley, S., Liu, Y., Masure, L., Rioul, O., Standaert, F.: Removing the Field Size Loss from Duc et al.'s Conjectured Bound for Masked Encodings. In: Kavun, E.B., Pehl, M. (eds.) Constructive Side-Channel Analysis and Secure Design - 14th International Workshop, COSADE 2023, Munich, Germany, April 3-4, 2023, Proceedings. Lecture Notes in Computer Science, vol. 13979, pp. 86–104. Springer (2023). https://doi.org/10.1007/978-3-031-29497-6_5, https://doi.org/10.1007/978-3-031-29497-6_5
5. Béguinot, J., Cheng, W., Guilley, S., Rioul, O.: Formal Security Proofs via Doeblin Coefficients: - Optimal Side-Channel Factorization from Noisy Leakage to Random Probing. In: Reyzin, L., Stebila, D. (eds.) Advances in Cryptology - CRYPTO 2024 - 44th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2024, Proceedings, Part VI. Lecture Notes in Computer Science, vol. 14925, pp. 389–426. Springer (2024). https://doi.org/10.1007/978-3-031-68391-6_12, https://doi.org/10.1007/978-3-031-68391-6_12
6. Bronchain, O., Standaert, F.: Breaking Masked Implementations with Many Shares on 32-bit Software Platforms or When the Security Order Does Not Matter. IACR Trans. Cryptogr. Hardw. Embed. Syst. **2021**(3), 202–234 (2021). https://doi.org/10.46586/TCHES.V2021.I3.202-234
7. Chari, S., Jutla, C.S., Rao, J.R., Rohatgi, P.: Towards sound approaches to counteract power-analysis attacks. In: Wiener, M. (ed.) Advances in Cryptology — CRYPTO' 99. pp. 398–412. Springer Berlin Heidelberg, Berlin, Heidelberg (1999)

8. Cover, T.M., Thomas, J.A.: Elements of Information Theory 2nd Edition (Wiley Series in Telecommunications and Signal Processing). Wiley-Interscience (July 2006)
9. Duc, A., Dziembowski, S., Faust, S.: Unifying leakage models: From probing attacks to noisy leakage. J. Cryptol. **32**(1), 151–177 (2019). https://doi.org/10.1007/S00145-018-9284-1, https://doi.org/10.1007/s00145-018-9284-1
10. Duc, A., Faust, S., Standaert, F.: Making Masking Security Proofs Concrete (Or How to Evaluate the Security of Any Leaking Device), Extended Version. J. Cryptol. **32**(4), 1263–1297 (2019). https://doi.org/10.1007/S00145-018-9277-0, https://doi.org/10.1007/s00145-018-9277-0
11. Durvaux, F., Standaert, F., Veyrat-Charvillon, N.: How to certify the leakage of a chip? In: Nguyen, P.Q., Oswald, E. (eds.) Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings. Lecture Notes in Computer Science, vol. 8441, pp. 459–476. Springer (2014). https://doi.org/10.1007/978-3-642-55220-5_26, https://doi.org/10.1007/978-3-642-55220-5_26
12. Dziembowski, S., Faust, S., Skorski, M.: Noisy leakage revisited. In: Oswald, E., Fischlin, M. (eds.) Advances in Cryptology - EUROCRYPT 2015. pp. 159–188. Springer Berlin Heidelberg, Berlin, Heidelberg (2015)
13. Dziembowski, S., Faust, S., Skórski, M.: Optimal Amplification of Noisy Leakages. In: Kushilevitz, E., Malkin, T. (eds.) Theory of Cryptography - 13th International Conference, TCC 2016-A, Tel Aviv, Israel, January 10-13, 2016, Proceedings, Part II. Lecture Notes in Computer Science, vol. 9563, pp. 291–318. Springer (2016). https://doi.org/10.1007/978-3-662-49099-0_11, https://doi.org/10.1007/978-3-662-49099-0_11
14. Faust, S., Masure, L., Micheli, E., Orlt, M., Standaert, F.: Connecting Leakage-Resilient Secret Sharing to Practice: Scaling Trends and Physical Dependencies of Prime Field Masking. In: Joye, M., Leander, G. (eds.) Advances in Cryptology - EUROCRYPT 2024 - 43rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zurich, Switzerland, May 26-30, 2024, Proceedings, Part IV. Lecture Notes in Computer Science, vol. 14654, pp. 316–344. Springer (2024). https://doi.org/10.1007/978-3-031-58737-5_12, https://doi.org/10.1007/978-3-031-58737-5_12
15. Fedotov, A., Harremoes, P., Topsoe, F.: Refinements of pinsker's inequality. IEEE Transactions on Information Theory **49**(6), 1491–1498 (2003). https://doi.org/10.1109/TIT.2003.811927
16. Gierlichs, B., Batina, L., Tuyls, P., Preneel, B.: Mutual information analysis. In: Oswald, E., Rohatgi, P. (eds.) Cryptographic Hardware and Embedded Systems - CHES 2008, 10th International Workshop, Washington, D.C., USA, August 10-13, 2008. Proceedings. Lecture Notes in Computer Science, vol. 5154, pp. 426–442. Springer (2008). https://doi.org/10.1007/978-3-540-85053-3_27, https://doi.org/10.1007/978-3-540-85053-3_27
17. Goldreich, O., Levin, L.A.: A Hard-Core Predicate for all One-Way Functions. In: Johnson, D.S. (ed.) Proceedings of the 21st Annual ACM Symposium on Theory of Computing, May 14-17, 1989, Seattle, Washington, USA. pp. 25–32. ACM (1989). https://doi.org/10.1145/73007.73010, https://doi.org/10.1145/73007.73010
18. Grassi, L., Masure, L., Méaux, P., Moos, T., Standaert, F.: Generalized Feistel Ciphers for Efficient Prime Field Masking. In: Joye, M., Leander, G. (eds.) Advances

in Cryptology - EUROCRYPT 2024 - 43rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zurich, Switzerland, May 26-30, 2024, Proceedings, Part III. Lecture Notes in Computer Science, vol. 14653, pp. 188–220. Springer (2024). https://doi.org/10.1007/978-3-031-58734-4_7, https://doi.org/10.1007/978-3-031-58734-4_7

19. Guo, Q., Grosso, V., Standaert, F., Bronchain, O.: Modeling soft analytical side-channel attacks from a coding theory viewpoint. IACR Trans. Cryptogr. Hardw. Embed. Syst. **2020**(4), 209–238 (2020). https://doi.org/10.13154/TCHES.V2020.I4.209-238, https://doi.org/10.13154/tches.v2020.i4.209-238
20. Heuser, A., Rioul, O., Guilley, S.: Good is not good enough. In: Batina, L., Robshaw, M. (eds.) Cryptographic Hardware and Embedded Systems – CHES 2014. pp. 55–74. Springer Berlin Heidelberg, Berlin, Heidelberg (2014)
21. Ito, A., Ueno, R., Homma, N.: On the Success Rate of Side-Channel Attacks on Masked Implementations: Information-Theoretical Bounds and Their Practical Usage. In: Yin, H., Stavrou, A., Cremers, C., Shi, E. (eds.) Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security, CCS 2022, Los Angeles, CA, USA, November 7-11, 2022. pp. 1521–1535. ACM (2022). https://doi.org/10.1145/3548606.3560579, https://doi.org/10.1145/3548606.3560579
22. Jahandideh, V., Mennink, B., Batina, L.: An Algebraic Approach for Evaluating Random Probing Security With Application to AES. IACR Trans. Cryptogr. Hardw. Embed. Syst. **2024**(4), 657–689 (2024). https://doi.org/10.46586/TCHES.V2024.I4.657-689
23. Mangard, S., Oswald, E., Popp, T.: Power Analysis Attacks: Revealing the Secrets of Smart Cards (Advances in Information Security). Springer-Verlag, Berlin, Heidelberg (2007)
24. Masure, L., Rioul, O., Standaert, F.: A Nearly Tight Proof of Duc et al.'s Conjectured Security Bound for Masked Implementations. In: Buhan, I., Schneider, T. (eds.) Smart Card Research and Advanced Applications - 21st International Conference, CARDIS 2022, Birmingham, UK, November 7-9, 2022, Revised Selected Papers. Lecture Notes in Computer Science, vol. 13820, pp. 69–81. Springer (2022). https://doi.org/10.1007/978-3-031-25319-5_4, https://doi.org/10.1007/978-3-031-25319-5_4
25. Masure, L., Standaert, F.: Prouff and Rivain's Formal Security Proof of Masking, Revisited - Tight Bounds in the Noisy Leakage Model. In: Handschuh, H., Lysyanskaya, A. (eds.) Advances in Cryptology - CRYPTO 2023 - 43rd Annual International Cryptology Conference, CRYPTO 2023, Santa Barbara, CA, USA, August 20-24, 2023, Proceedings, Part III. Lecture Notes in Computer Science, vol. 14083, pp. 343–376. Springer (2023). https://doi.org/10.1007/978-3-031-38548-3_12, https://doi.org/10.1007/978-3-031-38548-3_12
26. Moos, T.: Static power SCA of sub-100 nm CMOS asics and the insecurity of masking schemes in low-noise environments. IACR Trans. Cryptogr. Hardw. Embed. Syst. **2019**(3), 202–232 (2019). https://doi.org/10.13154/TCHES.V2019.I3.202-232, https://doi.org/10.13154/tches.v2019.i3.202-232
27. Obresmki, M., Ribeiro, J., Roy, L., Standaert, F.X., Venturi, D.: Improved Reductions from Noisy to Bounded and Probing Leakages via Hockey-Stick Divergences. In: Reyzin, L., Stebila, D. (eds.) Advances in Cryptology – CRYPTO 2024. pp. 461–491. Springer Nature Switzerland, Cham (2024)

28. Prest, T., Goudarzi, D., Martinelli, A., Passelègue, A.: Unifying Leakage Models on a Rényi Day. In: Boldyreva, A., Micciancio, D. (eds.) Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2019, Proceedings, Part I. Lecture Notes in Computer Science, vol. 11692, pp. 683–712. Springer (2019). https://doi.org/10.1007/978-3-030-26948-7_24, https://doi.org/10.1007/978-3-030-26948-7_24
29. Prouff, E., Rivain, M.: Masking against side-channel attacks: A formal security proof. In: Johansson, T., Nguyen, P.Q. (eds.) Advances in Cryptology - EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Athens, Greece, May 26-30, 2013. Proceedings. Lecture Notes in Computer Science, vol. 7881, pp. 142–159. Springer (2013). https://doi.org/10.1007/978-3-642-38348-9_9, https://doi.org/10.1007/978-3-642-38348-9_9
30. Renauld, M., Standaert, F., Veyrat-Charvillon, N., Kamel, D., Flandre, D.: A formal study of power variability issues and side-channel attacks for nanoscale devices. In: Paterson, K.G. (ed.) Advances in Cryptology - EUROCRYPT 2011 - 30th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tallinn, Estonia, May 15-19, 2011. Proceedings. Lecture Notes in Computer Science, vol. 6632, pp. 109–128. Springer (2011). https://doi.org/10.1007/978-3-642-20465-4_8, https://doi.org/10.1007/978-3-642-20465-4_8
31. Rotman, J.J.: An Introduction to the Theory of Groups, Graduate Texts in Mathematics, vol. 148. Springer, fourth edn. (1995)
32. Standaert, F.X., Malkin, T.G., Yung, M.: A unified framework for the analysis of side-channel key recovery attacks. In: Joux, A. (ed.) Advances in Cryptology - EUROCRYPT 2009. pp. 443–461. Springer Berlin Heidelberg, Berlin, Heidelberg (2009)
33. Wyner, A.D.: The wire-tap channel. The Bell System Technical Journal **54**(8), 1355–1387 (1975). https://doi.org/10.1002/j.1538-7305.1975.tb02040.x