Leakage-Diagrams, Importance Sampling, and Composition in the Random Probing Model

Vahid Jahandideh[1], Bart Mennink[2] and Lejla Batina[1]

[1] Radboud University, Nijmegen, The Netherlands
{v.jahandideh,lejla}@cs.ru.nl
[2] Maastricht University, Maastricht, The Netherlands
bart.mennink@maastrichtuniversity.nl

Abstract. Masking security in low-noise regimes remains poorly understood: generic reductions from noisy leakage to the random probing model (RPM) are often too loose to yield meaningful leakage rates, and existing RPM evaluations rely on costly simulation or numerically propagated bounds. This work develops analytic and algorithmic tools for the RPM security of masked gadgets and their composition. For noisy Hamming-weight leakage over $\mathbb{F}_{2^u}$, we instantiate a tighter $\mathbb{F}_2$-linear reduction, yielding a concrete mapping from SNR to probing rate. For $\mathbb{F}_q$-linear circuits, a vector-space representation characterizes RPM leakage as an erasure event. We prove a plug-and-play composition theorem: circuits whose gadgets are separated by linear refreshes are simulatable with error bounded by a sum of per-gadget erasure parameters at a boundary-inflated rate. The inflation is governed by *bridges*—boundary-only relations created by refresh leakage—which we bound analytically via leakage diagrams and which explain why SNI does not ensure RPM composition. We further introduce an importance-sampling estimator for rare leakage events, enabling evaluation in low-rate, high-order regimes, and implement all methods in LAPSE, a tool for linear gadgets.

Keywords: Side-Channel · Masking · Noisy Leakage · Random Probing Model

1 Introduction

Side-channel protection in *low-noise* settings—for instance, single-purpose embedded devices where measurements can be averaged and aligned—remains challenging [BS21]. In such settings, the signal-to-noise ratio can be high enough that classical first- and second-order masking no longer provide a comfortable security margin, motivating higher-order masking. Masking at order t encodes each sensitive variable into $t+1$ shares so that an adversary must jointly exploit all $t+1$ shares to recover the underlying value. For an isolated variable, increasing the order provably improves resistance in standard leakage models [CJRR99, DFS16, JMB25]. For full cryptographic computations, however, many intermediate values interact across time and across gadgets; consequently, the security gained by increasing the masking order depends critically on the leakage model, the circuit structure, and how randomness is consumed and refreshed. The central question of this work is how to obtain non-vacuous, compositional security bounds in such *constant-noise* regimes, where the noise level does not grow with the masking order.

Masked circuits and the role of refresh. A cryptographic computation is typically represented as an acyclic circuit over basic operations (e.g., add and mult). A masking compiler replaces each gate with a *gadget* operating on shared encodings, and often inserts *refresh* gadgets that re-randomize a sharing using fresh randomness. Early circuit-level analyses in the noisy-leakage setting used information-theoretic decompositions under

strong assumptions, including (idealized) leak-free refresh [PR13]. Duc et al. [DDF19] identified key limitations of this view, notably that (i) resulting bounds force the noise to grow with the masking order, (ii) leak-free refresh is hard to justify in practice, and (iii) the metric is misaligned with standard cryptographic security notions. They also proposed a reduction from noisy leakage to the *random probing model* (RPM), which resolves (ii) and (iii) and enables circuit-level reasoning; the noise-versus-order tension (i), however, persists in generic reductions and is one of the gaps this work addresses.

Where the gaps are. Despite substantial progress, obtaining meaningful (i.e., non-vacuous) and compositional bounds for constant-noise regimes remains difficult. Research has since advanced along three main directions:

Tighter reductions to RPM. In RPM, each wire leaks independently with probability p. This abstraction is mathematically convenient, but generic reductions can map even moderate-noise leakage functions to uninformative parameters (often yielding $p \approx 1$). Several works improve the tightness of such reductions under additional structural assumptions [DFS15, PGMP19, ORR+24, JMB25], thereby keeping p in a range where RPM security claims can be meaningful.

Gadgets secure at constant noise. Analyses based on the ISW multiplication and its variants underpin many classic masking proofs [DFŻ19, BFO23, MS23, BCGR24, BDF24], but ISW-style multiplication is known to fail at constant noise in practical settings [BCPZ16]. Only recently have constructions and evaluations targeted multiplication gadgets intended to remain secure in constant-noise/RPM-like regimes [BNR25, JMB24, BC26].

Security estimation and evaluation in RPM. In large protected circuits, the expected number of leaked wires under RPM can easily exceed the masking order, making direct reliance on threshold-style reasoning ineffective unless p decreases with circuit size. Expansion-based compilers improve asymptotic behavior but at significant cost and with primarily asymptotic guarantees [BCP+20, BRT21]. More recent approaches, including cardinal-style analyses and Monte Carlo estimators, offer progress for constant p [BRR25, BNR25, JMB24, BC26], but provable regimes can be narrow (e.g., $p \leq 0.02$ or $p \leq 0.001$) and the resulting bounds rarely admit closed forms amenable to asymptotic analysis.

1.1 Our Contribution

This paper contributes tools and bounds for compositional security in the random probing model, with an emphasis on constant-noise regimes:

- **Concrete noisy-to-RPM parameters for HW leakage (linear case).** We instantiate the tighter reduction of [JMB25] for noisy Hamming-weight leakage and derive the resulting probing rate p as a function of SNR for $\mathbb{F}_2$-linear computations. This complements the general (nonlinear / arbitrary-field) treatment in [BCGR24] and provides interpretable leakage-rate values (Section 3).
- **From local security to global simulation in refreshed block-separated designs.** We show how, under a refreshed block-separated structure (i.e., gadgets connected only through refreshed sharings), circuit-level simulatability can be related to per-gadget leakage parameters at a leakage rate p that does not scale with circuit size or masking order. This links simulation-based metrics used e.g. in [BRR25] to simpler per-gadget metrics such as the adversarial advantage and makes explicit which boundary assumptions are needed for composition (Section 4).

- **Rare-event estimation via importance sampling.** We adapt importance sampling to RPM leakage-pattern sampling, enabling efficient estimation when the success probability is far below the Monte Carlo resolution, e.g., in low-p and higher-order masking regimes (Section 5.3).

- **Leakage-diagram bounds for refresh gadgets.** Using the leakage-diagram technique of [DFŻ19], we give a complete, explicit derivation of provable bounds for a class of refresh gadgets, making explicit technical steps that are left implicit in prior work (Section 5.4).

- **Composition via leakage-diagrams; clarifying the role of SNI in RPM.** We derive analytic bounds on the composition parameters (leakage-rate inflation at refresh boundaries) from [JMB24] using leakage-diagrams, replacing Monte Carlo-only estimation. Our analysis identifies *bridges*—boundary-only relations created by refresh leakage—as the key mechanism governing dependence across refreshes. This "bridge view" explains why *Strong Non-Interference* (SNI) [BBD+16], although central for composition in the threshold probing model (TPM), does not directly yield RPM composition guarantees (Section 6).

- **Multiplication in the RPM.** We revisit RPM-secure multiplication: we recall why the classical ISW multiplication fails in the RPM, outline the idea behind recent patches, and use our leakage-diagram and bridge-style tools to give a conditional formal bound for its linear Expand stage and a clearly labeled heuristic estimate for its Compress stage, clarifying how multiplication can be upper-bounded via associated $\mathbb{F}_2$-linear subcircuits (Section 7). A fully formal treatment of multiplication remains open.

- **Tool support: LAPSE.** To compute the security parameters of linear gadgets, we developed LAPSE (Linear Algebraic Probing Security Estimator). Given a Python-like description of a linear gadget, LAPSE builds the corresponding vector-space representation and computes or estimates the relevant security parameters: exactly when feasible, and otherwise by Monte Carlo estimation enhanced with importance sampling. Unlike more general-purpose verification tools such as IronMask [BMRT22], LAPSE is not intended to analyze arbitrary masked implementations. Its current scope is linear gadgets and the vector-space (span) and bridge parameters used in our local-to-global RPM analysis (Section 8).

Scope and limitations. The composition theorem (Theorem 2) splits the evaluation of a refresh-separated circuit into per-gadget evaluations. The per-gadget erasure parameters that it relies on are established formally for $\mathbb{F}_q$-linear gadgets only. For multiplication, part of the proof is heuristic. Consequently, a formal end-to-end bound for a complete masked circuit containing multiplications, such as AES, is left to future work.

Methodological overview. The paper follows a local-to-global strategy. We first instantiate the connection between noisy Hamming-weight leakage and RPM leakage rates, thereby relating physical leakage parameters to the probing probability p. We then express RPM leakage in linear gadgets as a vector-space erasure event. This yields local security parameters that can be computed exactly for small instances and estimated for larger ones. Since the relevant events become rare in the low-p regime, we introduce an importance-sampling estimator to obtain meaningful estimates beyond the range of plain Monte Carlo. Finally, we use leakage diagrams to analyze refresh-boundary bridge events, which allows these local parameters to be incorporated into a compositional security statement.

Organization. Section 2 recalls the necessary background on masking, noisy leakage, and the random probing model. Section 3 presents the noisy-to-RPM reduction for Hamming-weight leakage; Section 4 develops the local-to-global simulation argument; Section 5 introduces the vector-space view of RPM leakage together with exact, Monte Carlo, and importance-sampling evaluation, and derives leakage-diagram bounds for refresh gadgets; Section 6 treats composition and bridges; Section 7 discusses RPM-secure multiplication; and Section 8 describes LAPSE.

2 Preliminaries

Notation. Lowercase letters denote parameters (e.g., n for the number of shares) or realizations of random variables, and uppercase letters denote random variables or values of interest (e.g., X). We write $\mathbb{E}[\cdot]$ and $\mathrm{Var}[\cdot]$ for expectation and variance. Boldface letters (e.g., $\boldsymbol{X}$) denote vectors or tuples, and calligraphic letters (e.g., $\mathcal{L}$) denote leakage variables, leakage sets, or families of wires. For a finite set S, $|S|$ denotes its cardinality; for a circuit C, $|\mathsf{C}|$ denotes its size, defined below. Masked objects are prefixed by $\mathbb{S}\cdot$ (e.g., $\mathbb{S}\mathsf{C}$). Unless stated otherwise, all variables and operations are over $\mathbb{F}_q$. The notation $R \xleftarrow{\$} \mathbb{F}_q$ means that R is sampled uniformly from $\mathbb{F}_q$. Additional notation is introduced where first used.

Statistical distance. For distributions P, Q on the same finite set, we write

$$\mathsf{SD}(P,Q) \triangleq \frac{1}{2}\sum_{x} |P(x) - Q(x)|.$$

For random variables A, B, we write $\mathsf{SD}(A,B)$ for the statistical distance between their induced distributions. For random variables A, B, C, we define

$$\mathsf{SD}(A, (B \mid C)) \triangleq \mathbb{E}_{c \leftarrow C}\left[\mathsf{SD}\big(A,\ (B \mid C{=}c)\big)\right].$$

Definition 1 (Erasure channel)**.** The *erasure channel* with parameter $p \in [0,1]$ is the randomized map

$$\phi_p : \mathbb{F}_q \to \mathbb{F}_q \cup \{\perp\}, \qquad \phi_p(x) = \begin{cases} x & \text{with probability } p, \\ \perp & \text{with probability } 1-p, \end{cases}$$

where $\perp \notin \mathbb{F}_q$ denotes an erasure.

Wire-enumeration convention (static single assignment). Security metrics in probing-style models depend on how the circuit wires (intermediate variables) are represented and counted. We use a *static single assignment* (SSA) representation: inputs appear first, and each operation produces exactly one new wire whose value is never updated. To model fan-out explicitly, we include a copy primitive.

For example, if a value U is used twice, we represent this explicitly with copy wires: $W_1 \leftarrow U$, $W_2 \leftarrow \mathsf{copy}(W_1)$, and we then use W_1 and W_2 as inputs to later operations.

We enforce bounded fan-out (for example, out-degree at most 2) by inserting copy operations; larger effective fan-out is obtained by chaining copies. This yields a directed acyclic graph (DAG), which is the representation used throughout the paper. For a complete example of SSA enumeration, see Example 2.

2.1 Masked Circuits and Gadgets

Let C be a circuit over $\mathbb{F}_q$. We denote by Σ_{C} the set of wires of C, including its input wires and intermediate wires. We refer to the values carried by these wires as *native variables*, or simply *natives*. A native may be secret (e.g., a key-dependent intermediate) or public (e.g., a nonce-dependent value).

Sharing convention. We use n to denote the number of shares. This is a structural parameter of the masked representation. In the threshold probing model, one often considers the usual threshold $t = n - 1$; in the random probing model, however, n is not itself a security order.

Definition 2 (n-sharing)**.** Let $X \in \mathbb{F}_q$. An *n-sharing* of X is a random vector $\boldsymbol{X} = (X_1, \ldots, X_n) \in \mathbb{F}_q^n$ distributed uniformly over the affine subspace

$$\left\{ (x_1, \ldots, x_n) \in \mathbb{F}_q^n : \sum_{i=1}^{n} x_i = X \right\}.$$

Equivalently, $\boldsymbol{X}$ satisfies $\sum_{i=1}^{n} X_i = X$ and every proper subtuple of shares is uniformly distributed (and hence statistically independent of X).

A masked circuit $\mathbb{S}\mathsf{C}_n$ is obtained by encoding each input into an n-sharing and replacing each operation of C by a masked implementation acting on encoded values.

Let G range over basic gates, such as add and mult.

Definition 3 (Gadget)**.** A gadget $\mathbb{S}\mathsf{G}$ for a gate G is a masked subcircuit that, given encoded inputs, produces encoded outputs representing G applied to the corresponding native values. Unless stated otherwise, we consider gadgets with n-shared inputs and n-shared outputs.

Thus, $\mathbb{S}\mathsf{C}_n$ is obtained by replacing each gate G in C with a corresponding gadget $\mathbb{S}\mathsf{G}$. Gadgets are built from primitives such as $\{+, \times, \mathsf{copy}, \xleftarrow{\$}\}$. To reduce inter-gadget leakage correlations, one often inserts a refresh gadget $\mathbb{S}\mathsf{R}$ that re-randomizes an encoding of a value X using fresh randomness while preserving the underlying value.

Size. Circuit size is measured in the SSA/DAG representation as the number of operations. Typically, $|\mathbb{S}\mathsf{G}| = \mathcal{O}(n)$ for linear operations and $|\mathbb{S}\mathsf{G}| = \mathcal{O}(n^2)$ for masked multiplication; refresh gadgets typically lie between these two regimes. Hence, for a circuit C, one typically has

$$|\mathbb{S}\mathsf{C}_n| = \mathcal{O}\big(|\mathsf{C}|\, n^2\big).$$

2.2 Leakage Models and Security Metrics

The adversary knows the masked circuit $\mathbb{S}\mathsf{C}_n$ and, in each execution, obtains leakage from wires in $\Sigma_{\mathbb{S}\mathsf{C}_n}$.

- **Threshold probing model (TPM) [ISW03].** The adversary adaptively selects up to t wires and learns their exact values.
- **Random probing model (RPM) [ISW03, DDF19, BCP+20].** Each wire $X \in \Sigma_{\mathbb{S}\mathsf{C}_n}$ leaks independently with probability p. Equivalently, the adversary observes $\phi_p(X)$ for every wire, with independent channel randomness. We denote the full leakage by
$$\mathcal{L} = \big(\phi_p(X)\big)_{X \in \Sigma_{\mathbb{S}\mathsf{C}_n}}.$$
Throughout this paper, we refer to the parameter p as the *leakage rate*, also known as the *probing rate*.
- **Noisy leakage [DDF19].** For each wire $X \in \Sigma_{\mathbb{S}\mathsf{C}_n}$, the adversary observes $\mathsf{L}(X) \in \mathbb{R}$ for some randomized leakage function $\mathsf{L} : \mathbb{F}_q \to \mathbb{R}$. Following [DDF19], we say that L is δ-*noisy* if, for uniform $X \xleftarrow{\$} \mathbb{F}_q$,
$$\mathsf{SD}\big(X,\ (X \mid \mathsf{L}(X))\big) \leq \delta.$$

Local security: advantage. Let $X \in \Sigma_{\mathsf{C}}$ be a native variable of interest whose value is uniformly distributed over $\mathbb{F}_q$, and let $\mathcal{L}$ denote the leakage observed in one execution. The associated single-target guessing advantage is

$$\mathsf{Adv}_X(\mathcal{L}) \triangleq \mathbb{E}_{\mathcal{L}}\left[\max_{x\in\mathbb{F}_q} \Pr[X = x \mid \mathcal{L}]\right] - \frac{1}{q},$$

where the expectation is over the leakage (including the channel randomness). Equivalently, the optimal guessing success probability is $\frac{1}{q} + \mathsf{Adv}_X(\mathcal{L})$.

Global security: leakage simulation. For proof-oriented composition results, we also use a simulation-based notion of security.

Definition 4 (ε-close leakage simulation)**.** Let $\mathcal{W}$ be a subset of wire indices in $\Sigma_{\mathbb{S}\mathsf{C}_n}$, and let $\mathcal{L}_{\mathcal{W}}$ denote the leakage restricted to $\mathcal{W}$. We say that $\mathcal{L}_{\mathcal{W}}$ is *ε-simulatable* if there exists a randomized simulator Sim such that, for every secret input S and public input P of C,

$$\mathsf{SD}\big(\mathcal{L}_{\mathcal{W}}(S,P),\ \mathsf{Sim}(P,\mathcal{W})\big) \le \varepsilon.$$

The interpretation of ε depends on the leakage model:

- **TPM.** The masked circuit $\mathbb{S}\mathsf{C}_n$ is secure if, for every $\mathcal{W}$ with $|\mathcal{W}| \le t$, the leakage is simulatable with $\varepsilon = 0$.
- **RPM / Noisy leakage.** Here $\mathcal{W}$ is the random set of leaking wires, and the statistical distance in Definition 4 is taken on average over the choice of $\mathcal{W}$. One typically requires simulatability with $\varepsilon \le 2^{-\kappa}$ for a security parameter κ [BCP+20], or, asymptotically, that $\varepsilon = \varepsilon(n)$ be negligible in n, i.e., for every polynomial $\mathrm{poly}(\cdot)$, there exists n_0 such that for all $n \ge n_0$,

$$\varepsilon(n) \le \frac{1}{\mathrm{poly}(n)}.$$

3 Reduction of Noisy Hamming Weight to Random Probing

Scope of this section. Before turning to our main results in the random probing model, we present a complementary contribution that connects the abstract leakage rate p to measurable physical quantities, namely the channel noise level as captured by the signal-to-noise ratio (SNR). We do this for the widely used noisy Hamming-weight leakage model. Readers interested only in the RPM-specific contributions may skip this section and continue with Section 4.

Context. Side-channel leakage is naturally modeled as *noisy leakage*, but proving security directly in this model is difficult. Duc et al. [DDF19] showed that security against δ-noisy leakage can be reduced to security in the random probing model (RPM) with a suitable probing rate p. Unfortunately, the resulting value of p can be overly pessimistic: a circuit may be secure under the original noisy leakage model, while the reduction yields an RPM parameter for which known guarantees become vacuous. In this section, we instantiate this reduction for the practically important *noisy Hamming-weight* leakage model. The general Duc et al. parameter for this leakage was computed in [BCGR24]; here we compute the tighter binary-inner-product parameter introduced by Jahandideh et al. [JMB25], which gives a more tangible interpretation of concrete leakage rates for $\mathbb{F}_2$-linear subcircuits.

Background: Duc et al. reduction. Let $X \in \mathbb{F}_q$ be uniform and let $L = \mathsf{L}(X)$ be its leakage. Duc et al. [DDF19] show that the joint distribution (X, L) can be simulated from an erasure view $\phi_p(X)$ for any $p \geq p_{\min}$, where

$$p_{\min} \;=\; 1 - q \sum_{\beta} \min_{\alpha \in \mathbb{F}_q} \Pr\big(L = \beta,\; X = \alpha\big) \;\leq\; q\,\delta, \tag{1}$$

where for continuous leakage the sum over β is replaced by an integral. In terms of mutual information (in bits) $\mathsf{I}_{\mathrm{bits}}(X; L)$, one can bound [DFS19, BCGR24]

$$\frac{\mathsf{I}_{\mathrm{bits}}(X; L)}{\log_2 q} \;\leq\; p_{\min} \;\leq\; q\sqrt{\frac{\ln 2}{2}\,\mathsf{I}_{\mathrm{bits}}(X; L)}.$$

This reduction is fully general: it does not rely on any special structure of $\mathbb{F}_q$ or of the underlying circuit. However, it can be loose, especially when q is large.

Tighter reduction for $\mathbb{F}_2$-linear circuits. For $q = 2^u$ and $\mathbb{F}_2$-linear circuits, Jahandideh et al. [JMB25] reduce leakage not to $\phi_p(X)$ itself, but to erasures of binary linear forms $\phi_p(\langle H, X\rangle)$, where $H \in \{0,1\}^u \setminus \{0\}$. Writing $X = (x_1, \ldots, x_u) \in \mathbb{F}_2^u$ and $H = (h_1, \ldots, h_u) \in \mathbb{F}_2^u \setminus \{0\}$ in bit representation, we define

$$\langle H, X\rangle \;=\; \bigoplus_{i=1}^{u} (h_i x_i),$$

where $h_i x_i$ and $\oplus$ are over $\mathbb{F}_2$.

For a fixed H, the optimal binary probing parameter is

$$p_{\min,H} \;=\; 2P^H_{\mathrm{succ}} - 1,$$

where P^H_{succ} denotes the optimal success probability of guessing $\langle H, X\rangle$ from the leakage L. The relevant bottleneck parameter is then

$$p^b_{\min} \;\triangleq\; \max_{H \neq 0} p_{\min,H}.$$

The reduction guarantees that if an $\mathbb{F}_2$-linear circuit is RPM-secure at rate $p^b_{\min}$, then it is secure against the original noisy leakage $\mathsf{L}(X)$. Moreover, one always has $p^b_{\min} \leq p_{\min}$, and the improvement is typically strict; Figure 1 quantifies the gap for noisy Hamming-weight leakage.

3.1 Noisy Hamming-Weight Leakage

Let $X \in \mathbb{F}_{2^u}$ be uniform and consider the noisy Hamming-weight leakage

$$L = \mathsf{L}(X) = \mathsf{HW}(X) + N, \qquad N \sim \mathcal{N}(0, \sigma^2), \qquad \mathsf{HW}(X) \in \{0, \ldots, u\}.$$

For this leakage model, the signal-to-noise ratio is

$$\mathrm{SNR} = \frac{u}{4\sigma^2}$$

as in [BCPZ16]. A direct computation of (1) gives

$$p_{\min} = 1 - 2\,\mathsf{Q}\Big(\frac{u}{2\sigma}\Big)$$

for the general Duc et al. reduction [BCGR24, Prop. 2], where Q denotes the Gaussian tail function, $\mathsf{Q}(x) = \frac{1}{\sqrt{2\pi}} \int_x^\infty e^{-t^2/2}\,dt$.

The following lemma gives an explicit expression for $p^b_{\min}$ in terms of u and σ, and hence as a function of (u, SNR).

Lemma 1. *For $\mathbb{F}_2$-linear circuits under noisy Hamming-weight leakage, the corresponding binary probing parameter is*

$$p^b_{\min} = \max_{1\le s\le u} \frac{1}{2^u} \int_{-\infty}^{\infty} \left| \sum_{a=0}^{u} K_a(s)\, \varphi_\sigma(l-a) \right| dl,$$

where $K_a(s)$ is the Krawtchouk polynomial

$$K_a(s) \;=\; \sum_{k=0}^{a} (-1)^k \binom{s}{k} \binom{u-s}{a-k}, \qquad \textit{and} \qquad \varphi_\sigma(t) = \frac{1}{\sqrt{2\pi}\sigma} e^{-t^2/(2\sigma^2)}.$$

Proof. Fix $H \in \{0,1\}^u \setminus \{0\}$, let $s = \mathsf{HW}(H)$, and define the target bit

$$B = \langle H, X \rangle \in \{0,1\}.$$

Let

$$A = \mathsf{HW}(X).$$

For $x \in \mathbb{F}_2^u$ with $\mathsf{HW}(x) = a$, the bit $\langle H, x\rangle$ is the parity of the overlap $k = |\mathrm{supp}(x) \cap \mathrm{supp}(H)|$ between the supports of x and H. Since exactly $\binom{s}{k}\binom{u-s}{a-k}$ such x have overlap k, summing $(-1)^k$ over all overlap sizes gives

$$\#\{x : \mathsf{HW}(x) = a,\ \langle H, x\rangle = 0\} - \#\{x : \mathsf{HW}(x) = a,\ \langle H, x\rangle = 1\} \;=\; K_a(s),$$

with $K_a(s)$ as in the lemma statement, and hence the identities

$$\Pr(A = a, B = 0) = \frac{1}{2^{u+1}} \Big(\binom{u}{a} + K_a(s) \Big), \qquad \Pr(A = a, B = 1) = \frac{1}{2^{u+1}} \Big(\binom{u}{a} - K_a(s) \Big).$$

Therefore, the conditional density of L given $B = b$ is a Gaussian mixture

$$f_b(l) \;=\; \sum_{a=0}^{u} \alpha_a^{(b)}\, \varphi_\sigma(l-a),$$

with mixture weights

$$\alpha_a^{(0)} = \Pr(A = a \mid B = 0) = \frac{\binom{u}{a} + K_a(s)}{2^u}, \qquad \alpha_a^{(1)} = \Pr(A = a \mid B = 1) = \frac{\binom{u}{a} - K_a(s)}{2^u}.$$

Since $H \neq 0$ and X is uniform, B is uniform on $\{0,1\}$ (equivalently, $\sum_{a=0}^{u} K_a(s) = 0$ for $s \ge 1$). The Bayes-optimal decision rule is therefore the likelihood-ratio test

$$\widehat{B}(l) = \arg\max_{b\in\{0,1\}} f_b(l),$$

and the corresponding optimal success probability is

$$P^H_{\mathrm{succ}} = \frac{1}{2} + \frac{1}{4} \int_{-\infty}^{\infty} |f_0(l) - f_1(l)|\, dl.$$

Substituting the expressions for f_0 and f_1 gives

$$P^H_{\mathrm{succ}} = \frac{1}{2} + \frac{1}{2^{u+1}} \int_{-\infty}^{\infty} \left| \sum_{a=0}^{u} K_a(s)\, \varphi_\sigma(l-a) \right| dl.$$

Hence,

$$p_{\min,H} = 2P^H_{\mathrm{succ}} - 1 = \frac{1}{2^u} \int_{-\infty}^{\infty} \left| \sum_{a=0}^{u} K_a(s)\, \varphi_\sigma(l-a) \right| dl.$$

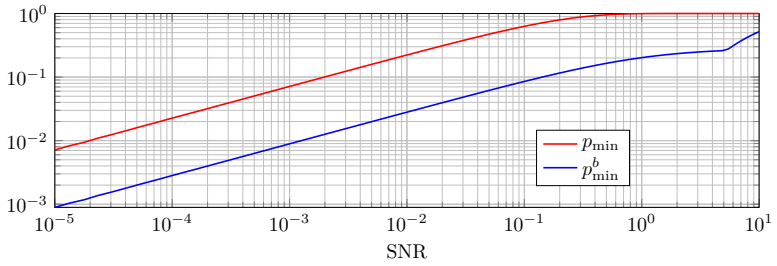


Figure 1: General reduction parameter $p_{\min}$ (red) and binary-inner-product parameter $p^b_{\min}$ (blue) as functions of the SNR, for noisy Hamming-weight leakage with $u = 8$ (log–log scale).

For uniform X, this quantity depends on H only through $s = \mathsf{HW}(H)$; denote its common value by $p_{\min}(s)$. Therefore,

$$p^b_{\min} = \max_{1 \le s \le u} p_{\min}(s),$$

which is the claimed expression. □

Lemma 1 reduces the computation of $p^b_{\min}$ to a one-dimensional integral depending on (u, σ, s), which can be evaluated numerically. For the practically relevant case $u = 8$, Figure 1 plots both $p_{\min}$ and $p^b_{\min}$ as functions of the SNR.

Interpretation of the results. State-of-the-art RPM-secure compilers and analyses typically operate in the regime of small probing rates, roughly $p \lesssim 0.02$ [JMB24, BCP+20, BRT21]. From Figure 1, the general Duc et al. reduction for noisy Hamming-weight leakage maps this regime to SNR $\lesssim 10^{-4}$ (for $u = 8$), whereas the $\mathbb{F}_2$-linear reduction based on binary inner products permits substantially lower noise levels (higher SNR), about SNR $\approx 5 \times 10^{-3}$, because $p^b_{\min} \ll p_{\min}$ in this range. At high SNR, the contrast becomes stark: for instance at SNR $= 1$, the general reduction yields $p_{\min} \approx 1$ (and thus becomes essentially vacuous), while the binary-inner-product reduction still gives $p^b_{\min} \approx 0.2$. Although $p \approx 0.2$ is beyond what current full compilers can handle, it is already within reach for some standalone $\mathbb{F}_2$-linear components such as linear refresh gadgets [JMB24].

Why $\mathbb{F}_2$-linearity matters. $\mathbb{F}_2$-linear subcircuits are ubiquitous in masked implementations. All refresh gadgets considered in this work are $\mathbb{F}_2$-linear, and multiplication by a *public* scalar (applied sharewise) is also $\mathbb{F}_2$-linear. In contrast, multiplying two n-sharings is not $\mathbb{F}_2$-linear; nevertheless, for certain designs (see Section 7), its RPM security can be upper-bounded by the security of associated $\mathbb{F}_2$-linear subcircuits [JMB24].

Remaining gap. While reducing to binary inner products improves the parameter ($p^b_{\min} \le p_{\min}$), the reduction can still be pessimistic at circuit level: even in the binary setting, an adversary with the original noisy leakage can be strictly weaker than an adversary given the corresponding random-probing (erasure) leakage. We illustrate this with a toy example.

Example 1 (Noisy leakage versus its RPM reduction: a strict gap)**.** Let the leakage be a binary symmetric channel (BSC),

$$\mathsf{L}(Z) = Z \oplus e, \qquad \Pr[e = 1] = P_e \le \tfrac{1}{2},$$

with independent noise for each leakage sample. Let X, R be uniform bits and define

$$Y = X \oplus R.$$

The adversary observes the noisy leakage $\mathsf{L}(X) = X \oplus e_1$, $\mathsf{L}(R) = R \oplus e_2$, and $\mathsf{L}(Y) = Y \oplus e_3$, where e_1, e_2, e_3 are independent with $\Pr[e_i = 1] = P_e$. Consider the derived observation

$$\mathsf{L}(R) \oplus \mathsf{L}(Y) = X \oplus e_2 \oplus e_3,$$

which is again a BSC observation of X, now with crossover probability

$$2P_e(1 - P_e) \;\geq\; P_e.$$

Since R is uniform and independent of X, the pair $\{\mathsf{L}(R), \mathsf{L}(Y)\}$ is informative about X only through this XOR. Hence, in the presence of $\mathsf{L}(X)$, it provides only a degraded observation of X and does not change the maximum a posteriori (MAP) decision. In particular,

$$\arg\max_{b\in\{0,1\}} \Pr[X = b \mid \mathsf{L}(X)] = \arg\max_{b\in\{0,1\}} \Pr[X = b \mid \mathsf{L}(X), \mathsf{L}(R), \mathsf{L}(Y)].$$

Therefore, the optimal adversary recovers X with probability

$$P_{\text{noisy}} = 1 - P_e.$$

In the corresponding RPM view, the adversary obtains

$$\{\phi_p(X), \phi_p(R), \phi_p(Y)\}, \qquad p = 2(1 - P_e) - 1 = 1 - 2P_e.$$

The optimal decision rule then succeeds with probability

$$P_{\text{RPM}} = \tfrac{1}{2}(1 + p + p^2 - p^3).$$

Since $P_{\text{noisy}} = 1 - P_e = \frac{1}{2}(1 + p)$, we obtain

$$P_{\text{RPM}} - P_{\text{noisy}} = \tfrac{1}{2}\, p^2(1 - p) > 0 \qquad \text{for all } P_e \in (0, \tfrac{1}{2}),$$

exhibiting a strict gap.

4 From Local Security to Global Security

Scope of this section. This section gives the abstract route from *local* leakage bounds to a *global* simulation guarantee. At this point, we do not yet prove the required gadget-specific assumptions. Instead, we state the composition template used later in the paper and explain how the subsequent sections instantiate it.

Context. The *advantage* Adv is a local metric: it captures an adversary's optimal guessing advantage for a single target secret. Following Prest et al. [PGMP19], we refer to such measures as *direct*. In contrast, simulation-based definitions are *indirect*: security holds if there exists a simulator that can reproduce the leakage distribution (up to statistical distance) without access to the secrets. Because simulation ranges over leakage from many intermediate values, it is a global notion.

Why local security alone does not imply global security. In general, local indistinguishability does not imply global simulatability. Two typical obstacles are: (i) *non-independent leakage*, when one leakage sample depends jointly on several wires, and (ii) *cross-gadget accumulation*, when leakage collected across several gadgets reveals enough information to reconstruct a secret unless refresh/randomness breaks the dependence across gadget boundaries.

Abstract composition setting. We consider a masked circuit $\mathbb{SC}_n$ composed of gadgets $\mathbb{SG}_1, \ldots, \mathbb{SG}_m$, with corresponding target natives $X_1, \ldots, X_m$ (not necessarily independent). We assume:

1. **Independent leakage:** each leakage sample is a randomized function of a single wire and, conditioned on that wire, is independent of all other wires and secrets.
2. **Block separation:** the total leakage can be partitioned as $\mathcal{L} = (\mathcal{L}_1, \ldots, \mathcal{L}_m)$ such that, conditioned on the public inputs P and the corresponding native X_i, each block $\mathcal{L}_i$ is independent of the other blocks and of the other natives.
3. **Erasure-like block view:** for each block, there exists a coupling to an erasure view of its native, i.e., $\mathcal{L}_i$ can be simulated from $\phi_{p_i}(X_i)$ for some parameter p_i.

Why these assumptions are reasonable in our setting. Assumption (1) is standard in masking analyses and is a good approximation for software implementations.[1] Assumption (3) will be justified for linear gadgets in Section 5, where we show that, under RPM leakage, a linear block either reveals its native or is statistically independent of it. Assumption (2) does not come for free: later, in Theorem 1, we explain how leakage from refresh boundaries can be reduced to this block-separated form by replacing each real refresh with an ideal leak-free one at the cost of additional exposure and an inflated effective probing rate. Without such separation, leakage may remain correlated across adjacent gadgets, and the clean composition argument below no longer applies.

Lemma 2 (Erasure blocks imply global simulatability)**.** *Let $X_1, \ldots, X_m$ be arbitrary $\mathbb{F}_q$-valued random variables and let $\mathcal{L}' = (\phi_{p_1}(X_1), \ldots, \phi_{p_m}(X_m))$. There exists a simulator* Sim *(independent of $X_1, \ldots, X_m$) such that*

$$\mathsf{SD}(\mathcal{L}', \mathsf{Sim}) \;\le\; 1 - \prod_{i=1}^{m}(1-p_i) \;\le\; \sum_{i=1}^{m} p_i.$$

Proof. Let Sim output $(\perp, \ldots, \perp)$ deterministically. The statistical distance between any distribution and a point mass at z equals $1 - \Pr[z]$, and, by independence of the erasure channels, $\Pr[\mathcal{L}' = (\perp, \ldots, \perp)] = \prod_{i=1}^{m}(1-p_i)$. The first bound follows (in fact with equality for this simulator); the second is the union bound. □

Interpretation. Lemma 2 is the abstract local-to-global step used throughout the paper: once each leakage block is reduced to an erasure event for its native, global simulation follows immediately—by Assumption (3) and data processing, the same bound carries over to the real leakage $\mathcal{L}$. The purpose of the next sections is therefore to instantiate the blockwise erasure parameters p_i and to justify the block-separation assumption for the concrete gadgets studied here.

5 Random Probing Leakage in Linear Circuits

Scope of this section. This section has three goals. First, we develop a vector-space representation for $\mathbb{F}_q$-linear circuits and use it to formalize the erasure-channel behavior of random probing leakage. Second, we present three ways to evaluate the resulting erasure parameters in practice: exact computation for small circuits, Monte Carlo estimation for moderate values of (n, p), and *importance sampling* for the low-p/high-order regime where

[1] It excludes dependent leakage, e.g., caused by coupling, glitches, or noise correlated across time samples; preprocessing such as principal component analysis (PCA) or noise whitening can reduce such correlations in practice, but quantifying the remaining dependence is outside the scope of this work.

rare events are difficult to observe directly. Third, we derive analytic upper bounds using leakage diagrams. Throughout this section, the exact/MC/IS methods serve as practical evaluation tools, whereas the leakage-diagram bounds are mainly proof-oriented and are typically looser.

5.1 Vector-Space Representation

We begin with *linear* circuits, namely circuits built only from linear gadgets (addition, copy, scalar multiplication, and randomness generation), and excluding multiplications of two n-sharings. As a running example, we consider $\mathbb{S}$R-multi [DFŻ19] (Algorithm 1), which repeats the simple refresh of [RP10] for $k \geq 1$ consecutive rounds.

In such a circuit, every wire is an $\mathbb{F}_q$-linear combination of the input shares and the fresh random variables. Fix an SSA enumeration of the wires, and let the basis consist of the n input-share symbols together with the r randomness symbols introduced by $\xleftarrow{\$}$. Each wire W then admits a representation vector $\vec{w} \in \mathbb{F}_q^{n+r}$ such that

$$W = \langle \vec{w},\, (X_1, \ldots, X_n, R_1, \ldots, R_r) \rangle\,.$$

The set of all wire-vectors spans a subspace $\mathcal{V}_{\mathbb{S}\mathsf{C}_n} \subseteq \mathbb{F}_q^{n+r}$.

We use the tool LAPSE to construct these vectors on the fly: each instruction produces a new vector derived from its operands, and each sampling instruction $R \xleftarrow{\$} \mathbb{F}_q$ introduces a new basis element represented by the corresponding unit vector. In this section, LAPSE is used to evaluate local RPM security for linear gadgets via exact computation, Monte Carlo estimation, and importance sampling.

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

Input: $\boldsymbol{X} = \{X_1, \ldots, X_n\},\ k$
Output: $\boldsymbol{Y} = \{Y_1, \ldots, Y_n\}, \sum X_i = \sum Y_i$

for $j = 1$ **to** k **do**
  $R_n \leftarrow 0$
  for $i = 1$ **to** $n-1$ **do**
    $R_i \xleftarrow{\$} \mathbb{F}_q$
    $R_n \leftarrow R_n - R_i$
  for $i = 1$ **to** n **do**
    $X_i \leftarrow X_i + R_i$
$\boldsymbol{Y} \leftarrow \boldsymbol{X}$ ▷ Copy to outputs
return $\boldsymbol{Y}$

Example 2 (Vector space and SSA wire enumeration)**.** For $(n = 2, k = 1)$, Algorithm 1 introduces one fresh random variable R_1 and sets $R_2 = -R_1$. An SSA enumeration of the intermediates is

$$\Sigma_{\mathbb{S}\text{R-multi},\, n=2} = \{X_1,\ X_2,\ R_1,\ R_2,\ X_1 + R_1,\ X_2 + R_2,\ Y_1,\ Y_2\}.$$

Using the basis (X_1, X_2, R_1), the corresponding vectors (in that order) are

$$(1,0,0),\ (0,1,0),\ (0,0,1),\ (0,0,-1),\ (1,0,1),\ (0,1,-1),\ (1,0,1),\ (0,1,-1).$$

The (unshared) native $X = X_1 + X_2$ corresponds to the vector $(1,1,0)$.

Erasure-channel viewpoint. For linear circuits handling a single secret X, the following property holds.

Lemma 3 (Lemma 1 [JMB24])**.** *Consider a linear circuit over $\mathbb{F}_q$ whose wires are $\mathbb{F}_q$-linear functions of (X, U), where U denotes the collection of fresh random variables sampled independently of X. Under the random probing model, a leakage instance $\mathcal{L}$ satisfies exactly one of the following:*

1. *(*reveal*) X is uniquely determined by $\mathcal{L}$; or*
2. *(*erase*) $\mathcal{L}$ is statistically independent of X (equivalently, $\Pr[\mathcal{L} \in \cdot \mid X = x]$ is the same for all $x \in \mathbb{F}_q$).*

Concretely, let $S \subseteq \Sigma_{\mathbb{SC}_n}$ be the set of leaked wires. If the representation vector of X lies in the $\mathbb{F}_q$-span of the vectors corresponding to S, then X is recovered uniquely; we write $\mathsf{is\text{-}in\text{-}span}(X, S) = 1$. Otherwise $\mathsf{is\text{-}in\text{-}span}(X, S) = 0$, in which case S (and hence $\mathcal{L}$) is independent of X. LAPSE checks membership in $\mathrm{span}(S)$ via Gaussian elimination.

Let $\mathsf{E}_X(n, p)$ be the probability (over the random RPM leakage pattern) that X is revealed, i.e.,

$$\mathsf{E}_X(n, p) \triangleq \Pr_{S \leftarrow \mathrm{RPM}(p)}\big[\mathsf{is\text{-}in\text{-}span}(X, S) = 1\big].$$

For uniform X, when X is not revealed the adversary can do no better than uniform guessing, hence

$$\Pr[\text{success}] = \mathsf{E}_X(n, p) + \big(1 - \mathsf{E}_X(n, p)\big)\tfrac{1}{q}.$$

The resulting advantage over random guessing is therefore

$$\mathsf{Adv}_X(n, p) = \Pr[\text{success}] - \tfrac{1}{q} = \tfrac{q-1}{q}\,\mathsf{E}_X(n, p).$$

In the remainder of this section we thus estimate $\mathsf{E}_X(n, p)$, which fully determines $\mathsf{Adv}_X(n, p)$ for linear circuits.

Exact method. Using the vector-space representation, we can compute $\mathsf{E}_X(n, p)$ *exactly* as

$$\mathsf{E}_X(n, p) = \sum_{S \subseteq \Sigma_{\mathbb{SC}_n}} \mathsf{is\text{-}in\text{-}span}(X, S) \cdot \Pr[S] = \mathbb{E}\big[\mathsf{is\text{-}in\text{-}span}(X, S)\big].$$

Let $T := |\Sigma_{\mathbb{SC}_n}|$ be the number of SSA wires. Under RPM, wires leak independently with rate p, so

$$\Pr[S] = p^{|S|}\,(1 - p)^{T - |S|}.$$

This approach enumerates all 2^T leakage patterns and is feasible only for small T.

5.2 Monte Carlo Estimation

When T is large, exact enumeration is infeasible. A standard alternative is Monte Carlo (MC) sampling as given in Algorithm 2. For a target (n, p), we estimate $\mathsf{E}_X(n, p) = \Pr[X \in \mathrm{span}(S)]$, where $S \subseteq \Sigma_{\mathbb{SC}_n}$ is the random set of leaked wires (each wire included independently with probability p).

The estimator $\widehat{\mathsf{E}}_X$ is *unbiased*, i.e., $\mathbb{E}[\widehat{\mathsf{E}}_X] = \mathsf{E}_X(n, p)$. Its cost is $\mathcal{O}\big(N \cdot C_{\mathrm{span}}\big)$, where C_{span} is the cost of one span-membership test (via Gaussian elimination; with bitset implementations this is near-quadratic in practice), in contrast with the $\mathcal{O}(2^T)$ cost of exact enumeration.

Algorithm 2 Monte Carlo estimation of $\mathsf{E}_X(n, p)$

Input: $\mathbb{SC}$, (n, p), N
Output: $\widehat{\mathsf{E}}_X \approx \mathsf{E}_X(n, p)$

$b \leftarrow 0$
for $i = 1$ **to** N **do**
  $\mathcal{L} \leftarrow \emptyset$
  for each $W \in \Sigma_{\mathbb{SC}_n}$ **do**
    Sample $B \leftarrow \mathrm{Ber}(p)$ ▷ Bernoulli distribution, $B \leftarrow 1$ w.p. p
    if $B = 1$ **then**
      $\mathcal{L} \leftarrow \mathcal{L} \cup \{W\}$
  $b \leftarrow b + \mathsf{is\text{-}in\text{-}span}(X, \mathcal{L})$
return $\widehat{\mathsf{E}}_X \leftarrow \dfrac{b}{N}$

Accuracy of Monte Carlo. Let $Z_i \in \{0,1\}$ be i.i.d. with $\Pr(Z_i = 1) = \mathsf{E}_X(n,p)$ and define $B \coloneqq \sum_{i=1}^{N} Z_i$ and $\widehat{\mathsf{E}}_X \coloneqq B/N$. Then

$$\mathbb{E}[\widehat{\mathsf{E}}_X] = \mathsf{E}_X(n,p), \qquad \mathsf{Var}[\widehat{\mathsf{E}}_X] = \frac{\mathsf{E}_X(n,p)\big(1-\mathsf{E}_X(n,p)\big)}{N}.$$

Hoeffding's inequality [Hoe63] gives, for any $\epsilon > 0$,

$$\Pr\Big(\big|\widehat{\mathsf{E}}_X - \mathsf{E}_X(n,p)\big| \geq \epsilon\Big) \;\leq\; 2e^{-2N\epsilon^2}.$$

Thus, we can guarantee an absolute accuracy of $\epsilon = 10^{-3}$, using $N \approx 2 \times 10^6$ samples, with confidence around 95%.

For small probabilities ($\mathsf{E}_X(n,p) \ll 1$), the *relative* standard error satisfies

$$\frac{\sqrt{\mathsf{Var}(\widehat{\mathsf{E}}_X)}}{\mathsf{E}_X(n,p)} \;\approx\; \frac{1}{\sqrt{\mathsf{E}_X(n,p)\,N}},$$

so achieving 10% relative error requires $N \approx 100/\mathsf{E}_X$ (and 1% requires $N \approx 10{,}000/\mathsf{E}_X$).

Limitation: rare events. MC can estimate an event only if it occurs in the sample. If $\theta \coloneqq \Pr(\mathsf{e}) \ll 1/N$, then with high probability no trial exhibits e and the estimator returns 0. Indeed,

$$\Pr(\text{no occurrences in } N \text{ trials}) = (1-\theta)^N \;\leq\; e^{-\theta N}.$$

To see at least one occurrence with confidence $1-\alpha$, it suffices that

$$N \;\geq\; \frac{\ln(1/\alpha)}{\theta}.$$

For instance, if $p = 10^{-3}$ and $n = 3$, the crude scaling $\theta \approx p^n = 10^{-9}$ suggests that observing even a single hit with 90% confidence requires $N \gtrsim 2.3 \times 10^9$ trials, which is beyond practical budgets. Consequently, naive MC becomes ineffective precisely in the low-p, higher-n regimes of interest.

5.3 Importance Sampling

Importance sampling (IS) estimates probabilities beyond the reach of naive Monte Carlo [KvD78]. To estimate the probability of a rare event e, we sample leakage patterns from a *tilted* distribution under which e is no longer rare, and reweight samples to recover the target probability.

Change of measure for RPM. Let $T \coloneqq |\Sigma_{\mathbb{S}\mathsf{C}_n}|$ and let $S \subseteq \Sigma_{\mathbb{S}\mathsf{C}_n}$ denote the (random) set of leaked wires. Under RPM with leakage rate p, the probability of a leakage pattern S is

$$\Pr_p(S) = p^{|S|}(1-p)^{T-|S|}.$$

Fix a *tilted* rate $\tilde{p} \in (0,1)$ and sample S instead from $\Pr_{\tilde{p}}$. Since $p, \tilde{p} \in (0,1)$, $\Pr_{\tilde{p}}(S) > 0$ whenever $\Pr_p(S) > 0$, and for any function $f(S)$,

$$\mathbb{E}_p[f(S)] = \sum_S f(S) \Pr_p(S) = \sum_S f(S) \underbrace{\frac{\Pr_p(S)}{\Pr_{\tilde{p}}(S)}}_{A(S)} \Pr_{\tilde{p}}(S) = \mathbb{E}_{\tilde{p}}\big[f(S)\,A(S)\big],$$

where the likelihood-ratio weight is

$$A(S) = \left(\frac{p}{\tilde{p}}\right)^{|S|} \left(\frac{1-p}{1-\tilde{p}}\right)^{T-|S|}.$$

Our target event is $f(S) = \mathsf{is\text{-}in\text{-}span}(X,S)$, hence

$$\mathsf{E}_X(n,p) = \mathbb{E}_p[f(S)] = \mathbb{E}_{\tilde{p}}\big[f(S)\,A(S)\big].$$

IS estimator (unbiased). The IS estimation approach is given in Algorithm 3. Its unbiasedness follows immediately from the change-of-measure identity:

$$\mathbb{E}\left[\widehat{\mathsf{E}}_X^{\mathrm{IS}}\right] = \mathsf{E}_X(n,p).$$

Algorithm 3 Importance sampling estimation of $\mathsf{E}_X(n,p)$

Input: $\mathbb{SC}$, (n,p), $\tilde{p}$, N
Output: $\widehat{\mathsf{E}}_X^{\mathrm{IS}} \approx \mathsf{E}_X(n,p)$

$b \leftarrow 0$
for $i = 1$ **to** N **do**
  $\mathcal{L}_i \leftarrow \emptyset$
  for each $W \in \Sigma_{\mathbb{SC}_n}$ **do**
    Sample $B \leftarrow \mathrm{Ber}(\tilde{p})$ ▷ $B \leftarrow 1$ w.p. $\tilde{p}$
    if $B = 1$ **then**
      $\mathcal{L}_i \leftarrow \mathcal{L}_i \cup \{W\}$
  $A_i \leftarrow \left(\frac{p}{\tilde{p}}\right)^{|\mathcal{L}_i|} \left(\frac{1-p}{1-\tilde{p}}\right)^{T-|\mathcal{L}_i|}$
  $b \leftarrow b + \mathsf{is\text{-}in\text{-}span}(X, \mathcal{L}_i) \cdot A_i$
return $\widehat{\mathsf{E}}_X^{\mathrm{IS}} \leftarrow \frac{b}{N}$

Variance and estimation. Let $I_i \coloneqq \mathsf{is\text{-}in\text{-}span}(X, S_i)$ and $Y_i \coloneqq I_i A_i$. A consistent variance estimator is

$$\widehat{\mathsf{Var}}\left(\widehat{\mathsf{E}}_X^{\mathrm{IS}}\right) = \frac{1}{N(N-1)} \sum_{i=1}^{N} \left(Y_i - \widehat{\mathsf{E}}_X^{\mathrm{IS}}\right)^2.$$

If $\tilde{p}$ is chosen too far from p, the weights A_i can become very large, which increases variance. We choose $\tilde{p}$ so that the success probability under the tilted measure, $\Pr_{\tilde{p}}[I = 1]$, lies in a moderate range (e.g., 5%–30%). A short pilot run can tune $\tilde{p}$ by bisection. For simplicity, we apply a *single* tilt parameter $\tilde{p}$ to all wires.

Example 3. We evaluate $\mathsf{E}_X(n,p)$ for $\mathbb{S}$R-multi (Alg. 1) using LAPSE. Given a threshold on the SSA wire count T below which exact enumeration is feasible, LAPSE proceeds as follows for each (n,p): it computes $\mathsf{E}_X(n,p)$ exactly when possible; otherwise it runs Monte Carlo, and switches to importance sampling when Monte Carlo falls into the rare-event (no-hit) regime. Figure 2 reports the resulting estimates.

Limitations of importance sampling. Importance sampling is effective only when the tilted probing rate $\tilde{p}$ is chosen in a reasonable range. If $\tilde{p}$ is too far from the target rate p, the likelihood ratios become highly variable and the estimator variance increases. On the other hand, if $\tilde{p}$ is chosen too close to p, the algorithm may remain in the no-hit regime and fail to generate enough informative rare events. This tradeoff is visible in Figure 2: for moderate choices of the tilt, $\log_{10}(\mathsf{E}_X(n,p))$ exhibits an approximately constant-slope decay as the number of shares increases, whereas in the extreme low-p/high-n corner the finite-sample estimates gradually deviate from this pattern. Thus, importance sampling substantially expands the practical working range of Monte Carlo estimation, and in particular enables the estimation of probabilities far below the naive $1/N$ resolution limit, but it does not remove the rare-event difficulty entirely.

Relation to PERSEUS. PERSEUS [BC26] faces the same difficulty, namely failure probabilities far below the resolution $1/N$ of naive Monte Carlo, but addresses it with a different technique. It combines Monte Carlo sampling with a pruning strategy that reduces

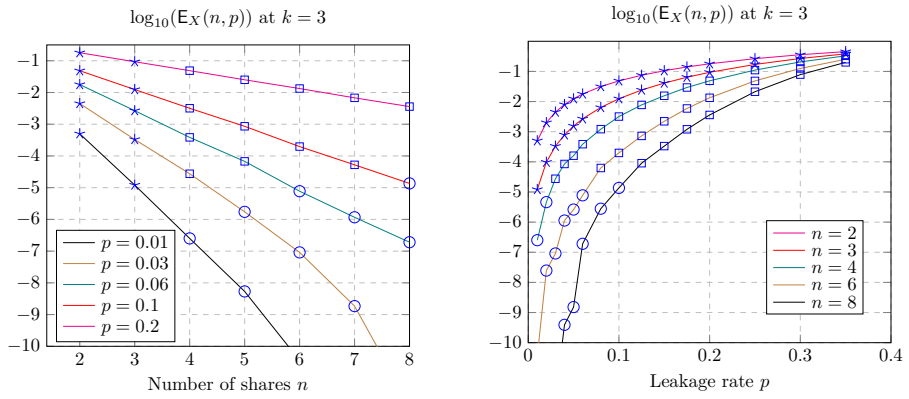


Figure 2: Evaluation of $\mathsf{E}_X(n,p)$ over a grid of (n,p) for $\mathbb{S}$R-multi with $k=3$. Stars: exact enumeration (Section 5.1); squares: Monte Carlo (Section 5.2); circles: importance sampling (Section 5.3). Each estimate uses $N=10^5$ trials. Importance sampling allows reliable estimates well below the naive Monte Carlo resolution $1/N$ (here $\log_{10}(1/N)=-5$). Estimates below 10^{-10} (low p, $n\geq 6$) lie outside the plotted range.

the sampling space, whereas we rely on importance sampling, a rare-event simulation technique that keeps the full space of leakage patterns, samples it at a tilted rate $\tilde{p}$, and reweights the samples so that the estimator remains unbiased at the original rate p.

5.4 Leakage-Diagram View

Dziembowski et al. [DFŻ19] introduced the *leakage-diagram* technique, which turns the event $\mathsf{is\text{-}in\text{-}span}(X,\mathcal{L})=1$ for $\mathbb{S}$R-multi into a graph-connectivity problem. For parameters (n,k), the diagram is an $(n+1)\times(k+1)$ grid of vertices. Each leaked wire induces a local linear constraint, which we represent by adding an edge between two neighboring vertices; see Figure 3. They show that the native $X=\sum_{i=1}^{n}X_i$ is recoverable from $\mathcal{L}$ if and only if the resulting diagram $G(\mathcal{L})$ contains a path connecting the leftmost column to the rightmost column. For completeness, we restate the construction and derive an explicit upper bound on $\mathsf{E}_X(n,p)$.

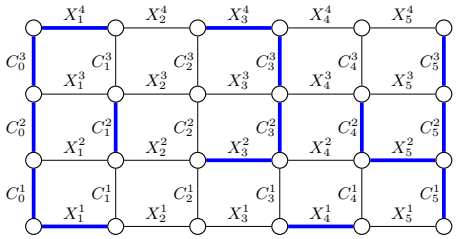


Figure 3: An example of a leakage diagram with $(n=5,k=3)$. A leakage instance is shown in blue. The vertical edges in the leftmost and rightmost columns are always present (also blue), since $C_0^j=C_n^j=0$ are public constants.

Leakage-diagram for $\mathbb{S}$R-multi. Write the input shares to the j-th refresh as $\boldsymbol{X}^j=\{X_1^j,\ldots,X_n^j\}$ for $j\in\{1,\ldots,k\}$, where the refresh maps $\boldsymbol{X}^j\mapsto\boldsymbol{X}^{j+1}$. In refresh round j, sample $R_1^j,\ldots,R_{n-1}^j\xleftarrow{\$}\mathbb{F}_q$ and define

$$C_0^j\leftarrow 0,\qquad C_i^j\leftarrow C_{i-1}^j+R_i^j\quad(1\leq i<n),$$

and

$$R_n^j \leftarrow -C_{n-1}^j, \qquad C_n^j \leftarrow 0.$$

The shares update as

$$X_i^{j+1} \leftarrow X_i^j + R_i^j \qquad (1 \le i \le n).$$

Thus the relevant wire families are $\{X_i^j\}$, $\{C_i^j\}$, and the random wires $\{R_i^j\}$ used to build them.

Lemma 4 (R-wire substitution)**.** *Let $\mathcal{L}$ be a leakage instance on wires $\{X_i^j, C_i^j, R_i^j\}$ of* $\mathbb{S}$R-multi*, and define*

$$\tau(\mathcal{L}) \;=\; (\mathcal{L} \cap \{X_i^j, C_i^j\}) \;\cup \bigcup_{R_i^j \in \mathcal{L}} \{\, C_{i-1}^j,\, C_i^j \,\}.$$

Then $\mathrm{span}(\tau(\mathcal{L})) \supseteq \mathrm{span}(\mathcal{L})$*. In particular, if X is revealed by $\mathcal{L}$, then X is revealed by $\tau(\mathcal{L})$, and hence*

$$\mathsf{E}_X(n,p) \;\le \Pr_{\mathcal{L} \leftarrow \mathrm{RPM}(p)}\big[\textsf{is-in-span}(X, \tau(\mathcal{L})) = 1\big].$$

Proof. It suffices to show that every wire of $\mathcal{L}$ lies in $\mathrm{span}(\tau(\mathcal{L}))$. This is immediate for X- and C-wires, and for each $R_i^j \in \mathcal{L}$ it follows from $R_i^j = C_i^j - C_{i-1}^j$ with $C_{i-1}^j, C_i^j \in \tau(\mathcal{L})$. □

Lemma 5 (Leakage-diagram characterization [DFŻ19])**.** *Consider* $\mathbb{S}$R-multi *with parameters (n,k), and assume leakage is restricted to wires $\{X_i^j, C_i^j\}$. Define a graph $G(\mathcal{L})$ with vertex set $\{v_{i,j} : 0 \le i \le n,\ 0 \le j \le k\}$, see Figure 3. Add a horizontal edge $(v_{i-1,j}, v_{i,j})$ whenever X_i^{j+1} leaks, and a vertical edge $(v_{i,j-1}, v_{i,j})$ whenever C_i^j leaks. (Additionally, since $C_0^j = C_n^j = 0$ are public constants, the vertical edges of the leftmost and rightmost columns are always present.) Then*

$$\textsf{is-in-span}(X, \mathcal{L}) = 1 \quad \Longleftrightarrow \quad G(\mathcal{L}) \textit{ contains a path from the left to the right column.}$$

If no such path exists, then $\mathcal{L}$ is statistically independent of X.

By Lemma 4, it therefore suffices to bound the probability that the diagram $G(\tau(\mathcal{L}))$ connects the two boundary columns. The edges of $G(\tau(\mathcal{L}))$ are determined by the original $\mathrm{RPM}(p)$ leakage as follows. A horizontal edge $h_{i,j}$ is present if and only if $X_i^{j+1} \in \mathcal{L}$: it has a single *cause*, of probability p, and distinct horizontal edges have distinct causes. A vertical edge $v_{i,j}$ $(1 \le i \le n-1)$ is present if and only if $C_i^j \in \mathcal{L}$ or $R_i^j \in \mathcal{L}$ or $R_{i+1}^j \in \mathcal{L}$: it has at most three causes, each of probability p, and each cause covers at most two edges (an R-wire covers the two adjacent vertical edges of its round; C- and X-wires cover one edge each).

Lemma 6 (Cause-selection bound)**.** *Let $F = F_h \cup F_v$ be a fixed set of horizontal and vertical (non-boundary) edges of the leakage diagram. Under RPM(p) leakage on $\{X_i^j, C_i^j, R_i^j\}$,*

$$\Pr_p\big[F \subseteq E\big(G(\tau(\mathcal{L}))\big)\big] \;\le\; p^{|F_h|}\,\big(3\sqrt{p}\big)^{|F_v|}.$$

Proof. Horizontal and vertical edges have disjoint cause sets (X-wires versus C- and R-wires), hence their presences are independent, and since each horizontal edge has a unique cause, $\Pr[F_h \text{ present}] = p^{|F_h|}$. For the vertical part, if all edges of F_v are present then some choice function σ selects one leaked cause per edge. There are at most $3^{|F_v|}$ such functions, and each selects at least $|F_v|/2$ distinct wires, because every wire is a cause of at most two edges. A union bound over choice functions gives $\Pr[F_v \text{ present}] \le 3^{|F_v|}\, p^{|F_v|/2} = (3\sqrt{p})^{|F_v|}$. □

Corollary 1. *For* $\mathbb{S}$R-multi *with parameters* (n, k) *and any* $p < 1/81$,

$$\mathsf{E}_X(n,p) \;\le\; \frac{k+1}{3} \cdot \frac{(3p)^n}{1-9\sqrt{p}}.$$

Proof. By Lemmas 4 and 5, it suffices to bound the probability of the event $\mathsf{Conn}_{n,k}$ that $G(\tau(\mathcal{L}))$ contains a path between the boundary columns. Contract the leftmost column to a source s and the rightmost column to a sink t; this is valid because their vertical edges are always present.

A simple s–t path of length ℓ must advance from column 0 to column n and thus contains $h \ge n$ horizontal edges; in particular $\ell \ge n$. Its edges are distinct, so by Lemma 6 it is present with probability at most

$$p^h\,(3\sqrt{p})^{\ell-h} \;\le\; p^n\,(3\sqrt{p})^{\ell-n},$$

where the inequality replaces each surplus horizontal factor p by the larger factor $3\sqrt{p}$.

Every simple s–t path is a walk with no immediate backtracking. From s there are at most $k+1$ choices for the first step (one per row), and at most 3 choices for each subsequent step, so there are at most $(k+1)\,3^{\ell-1}$ such walks of length ℓ; this upper-bounds the number of simple paths of length ℓ. By a union bound over all lengths $\ell \ge n$,

$$\begin{aligned}\Pr[\mathsf{Conn}_{n,k}] \;&\le\; (k+1)\sum_{\ell\ge n} 3^{\ell-1}\,p^n\,(3\sqrt{p})^{\ell-n} \;=\; (k+1)\,p^n\,3^{n-1}\sum_{m\ge 0}(9\sqrt{p})^m \\ &=\; \frac{k+1}{3}\cdot\frac{(3p)^n}{1-9\sqrt{p}},\end{aligned}$$

where the series converges since $9\sqrt{p} < 1$, i.e., $p < 1/81$. □

Hence, $\mathsf{E}_X(n,p)$ for $\mathbb{S}$R-multi decreases exponentially in n at any constant rate $p \le 0.01$.

Comparison with an empirical fit. As a compact summary of the Monte Carlo data, one may fit an expression of the form

$$\alpha(\beta p)^{\gamma n},$$

as in [JMB24]; such fits, however, are only justified over the tested parameter range. For $\mathbb{S}$R-multi with $k = 3$, the data in Figure 2 is well approximated by

$$\mathsf{E}_X(n,p) \;\approx\; 0.7(2.5p)^n.$$

Corollary 1 gives a formal analytic bound with looser constants and the restriction $p < 1/81$, but it has the same p^n scaling observed empirically.

6 Composition in the Random Probing Model

Scope of this section. In this section, we explain how leakage from a real refresh gadget can be reduced to leakage from an ideal (leak-free) refresh gadget, at the cost of an increased effective probing rate on the boundary shares. This reduction is a key step toward the block-separated setting required by the local-to-global composition framework of Section 4. The main obstacle is the presence of *bridges*, namely boundary-only dependencies created by refresh leakage. We quantify their effect both empirically, via Monte Carlo and importance-sampling estimation, and analytically, via leakage diagrams.

Context. In a masked circuit, many gadgets operate on correlated sharings, and their interaction can degrade the overall side-channel security of the underlying natives. Proving security therefore requires suitable *composition* properties. In the threshold probing model (TPM), *strong non-interference* (SNI), introduced by Barthe et al. [BBD+16], is a standard sufficient condition for secure composition.

Definition 5 (SNI [BBD+16])**.** Let $\mathbb{SG}$ be a gadget that takes an input sharing $\boldsymbol{X} = \{X_1, \ldots, X_n\}$ and produces an output sharing $\boldsymbol{Y} = \{Y_1, \ldots, Y_n\}$, and let $\Sigma_{\mathbb{SG}}$ denote its internal wires. The gadget is t_n-*SNI* if for every choice of sets $S_Y \subseteq \{Y_1, \ldots, Y_n\}$ and $S_I \subseteq \Sigma_{\mathbb{SG}}$ with $|S_Y| + |S_I| \leq t_n$, the joint values of the wires in $S_Y \cup S_I$ are simulatable from the public inputs and at most $|S_I|$ input shares.

If each gadget in a circuit is t_n-SNI, then any adversarial choice of at most t_n probes across the whole circuit can be *propagated back* to the circuit inputs without exceeding the probing threshold (typically $t_n = n - 1$). In the random probing model (RPM), however, leakage is a *random* set of probed wires and may contain far more than t_n wires. As a result, TPM-style probe propagation does not directly yield tight RPM composition bounds, and several alternatives have been proposed.

Propagation approach. Belaïd et al. [BCP+20] proposed a basic RPM composition rule: if each gadget's RPM leakage together with t output probes can be simulated from t input shares except with failure probability ϵ, then a circuit of such gadgets fails with probability at most $\min\{1, |\mathsf{C}|\epsilon\}$. Cassiers et al. [CFOS21] refined this by tracking *which* input-share subsets are needed for which output-share subsets, which becomes costly as n grows, and Belaïd et al. [BRR25] instead track only the *number* of required input shares via per-gadget *envelopes* propagated backward through the circuit. While powerful, the resulting envelopes (with refinements in [BNR25]) are typically obtained numerically, making closed-form and asymptotic analyses challenging.

Localization approach. In contrast to propagating leakage to the circuit inputs, Jahandideh et al. [JMB24] proposed to *localize* leakage by inserting refresh gadgets between computational gadgets. They show how RPM leakage $\mathcal{L} = \{\mathcal{L}_1, \ldots, \mathcal{L}_m\}$ sampled at rate p can be replaced by leakage $\mathcal{L}' = \{\mathcal{L}'_1, \ldots, \mathcal{L}'_m\}$ sampled at a larger rate $p' > p$ such that the blocks $\mathcal{L}'_i$ are conditionally independent given their corresponding natives. Our explicit notion of boundary relations (and, in particular, cross-boundary bridges) below is a reformulation that makes such boundary dependencies easier to isolate analytically. In this section, we review this approach and, using the leakage-diagram technique, derive analytical bounds for the resulting composition parameters that were previously estimated only via Monte Carlo.

6.1 On the Role of Refresh Gadgets

Let $\boldsymbol{X}$ and $\boldsymbol{Y}$ denote the input and output n-sharings of a (linear) refresh gadget $\mathbb{SR}$. If $\mathbb{SR}$ is *leak-free*, i.e., $\mathcal{L}_{\mathbb{SR}} = \emptyset$, then the refresh randomness makes the output sharing fresh: conditioned on the underlying native

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

the input and output sharings are independent. Concretely, for any realizations $\boldsymbol{x}, \boldsymbol{y}$,

$$\Pr(\boldsymbol{Y} = \boldsymbol{y} \mid X, \boldsymbol{X} = \boldsymbol{x}) \;=\; \Pr(\boldsymbol{Y} = \boldsymbol{y} \mid X).$$

Consequently, in a circuit where gadgets are separated by leak-free refreshes, the circuit-wide leakage partitions into independent per-gadget blocks, as exploited in [PR13].

In the presence of leakage $\mathcal{L}_{\mathbb{SR}} \neq \emptyset$, the sharings $\boldsymbol{X}$ and $\boldsymbol{Y}$ may become dependent. Intuitively, leaked internal wires allow the adversary to eliminate internal variables and derive additional linear relations involving only boundary shares. Following [JMB24], for each leakage instance one can describe the induced linear constraints as a full-rank system

$$\begin{cases} \boldsymbol{P}_{\text{Inf}}(\boldsymbol{X}, \boldsymbol{Y}) = \boldsymbol{0}, \\ \boldsymbol{P}_{\text{Non-Inf}}(\boldsymbol{X}, \boldsymbol{Y}, \Sigma_{\mathbb{SR}}) = \boldsymbol{0}, \end{cases}$$

where $\boldsymbol{P}_{\text{Non-Inf}}$ is independent of the native X and does not affect the joint distribution of $(\boldsymbol{X}, \boldsymbol{Y})$ beyond correctness, while $\boldsymbol{P}_{\text{Inf}}$ appears only due to leakage and may introduce additional constraints involving boundary shares. Any such relation supported only on boundary shares will be called a *boundary relation*. When such a relation involves shares from both $\boldsymbol{X}$ and $\boldsymbol{Y}$, it can couple the two sharings across the refresh boundary.

Definition 6 (Boundary relations and bridges)**.** In a refresh gadget $\mathbb{SR}$ with boundary sharings $\boldsymbol{X} = \{X_1, \ldots, X_n\}$ and $\boldsymbol{Y} = \{Y_1, \ldots, Y_n\}$, a *boundary relation* is a nontrivial $\mathbb{F}_q$-linear relation of the form

$$\sum_{i=1}^{n} a_i X_i \; + \; \sum_{i=1}^{n} b_i Y_i \; = \; c,$$

with coefficients $a_i, b_i, c \in \mathbb{F}_q$, that involves only boundary shares and no internal intermediates.

A boundary relation is called:

- *input-only* if $b_1 = \cdots = b_n = 0$,
- *output-only* if $a_1 = \cdots = a_n = 0$,
- a *cross-boundary bridge*, or bridge for short, if it involves at least one input share and at least one output share.

A (cross-boundary) bridge is called *essential* if there is no $\lambda \in \mathbb{F}_q$ such that, after adding

$$\lambda\Big(\sum_{i=1}^{n} X_i - \sum_{i=1}^{n} Y_i\Big) = 0,$$

the resulting relation becomes input-only or output-only.

Thus, input-only and output-only boundary relations capture additional leakage on one side of the refresh boundary, whereas essential bridges are precisely the relations that prevent a clean separation between the input and output sharings. In particular, they are what enables leakage to propagate across refresh boundaries.

Without leakage, we do not expect any nontrivial boundary relation beyond the correctness relation

$$\sum_{i=1}^{n} X_i - \sum_{i=1}^{n} Y_i = 0.$$

Leakage can eliminate intermediates and thereby create additional boundary relations, including essential cross-boundary bridges.

Example 4 (How leakage creates boundary relations)**.** Suppose $Y_i \leftarrow X_i + R_i$ for some internal intermediate R_i. If $R_i \in \mathcal{L}$, then the relation

$$Y_i - X_i = R_i$$

becomes a boundary-only linear relation over (X_i, Y_i) with a known right-hand side, and hence an essential bridge, occurring with probability p.

If instead

$$Y_i \leftarrow (X_i + R_i) + R_i',$$

then the same bridge requires both $R_i, R_i' \in \mathcal{L}$ and therefore occurs with probability p^2.

SNI refresh and bridges. Given the central role of SNI for TPM composition, one might expect SNI to play a similar role in the RPM (cf. the discussion of SNI in [BC26]). The bridge view shows that SNI does not, by itself, control RPM composition across refresh boundaries.

Consider any refresh that operates by adding an encoding of zero, i.e., with boundary assignments of the form

$$Y_i \;=\; X_i + M_i,$$

where $(M_1, \ldots, M_n)$ is an internally computed sharing of 0. This class contains genuinely $(n-1)$-SNI gadgets, such as the quasi-linear refresh of [GJR18, GPRV21]. Yet, regardless of SNI, if M_i leaks then the relation $Y_i - X_i = M_i$ becomes a bridge: for a fixed i this occurs with probability p, and for some $i \in \{1, \ldots, n\}$ with probability $1 - (1-p)^n \approx np$.

As a concrete instance, we use the $\mathbb{S}$R-Random gadget (Alg. 4) [BRR25], where the *randomness parameter* λ denotes the number of random zero-sum injection steps (the loop bound in Alg. 4). For sufficiently large λ, $\mathbb{S}$R-Random is $(n-1)$-SNI except with probability negligible in λ (for any finite λ there is a nonzero probability that some share is never refreshed), yet it admits bridges with probability $\approx np$ as above.

Algorithm 4 $\mathbb{S}$R-Random

Input: $\boldsymbol{X} = \{X_1, \ldots, X_n\}$, λ
Output: $\boldsymbol{Y} = \{Y_1, \ldots, Y_n\}$

for $i = 1$ **to** n **do**
  $M_i = 0$
for $t = 1$ **to** λ **do** ▷ λ denotes the number of injection iterations.
  $R \xleftarrow{\$} \mathbb{F}_q$, $i, j \xleftarrow{\$} \{1, 2, \ldots, n\}$
  $M_i \leftarrow M_i + R$
  $M_j \leftarrow M_j - R$
for $i = 1$ **to** n **do**
  $Y_i = X_i + M_i$
return $\boldsymbol{Y}$

In contrast, $\mathbb{S}$R-multi with $k = n$ is also $(n-1)$-SNI [CGPZ16], but we will prove that the probability of an essential cross-boundary bridge in this refresh gadget is negligible in n for the leakage rates of interest.

6.2 Composition Rule in RPM

Boundary relations, and in particular essential cross-boundary bridges, obstruct the separation of leakage across refresh-separated gadgets. A convenient reduction, introduced in [JMB24], is to *expose* the participating boundary shares: this removes the boundary dependencies at the cost of an inflated leakage rate on the boundary shares.

Affected shares and the exposure rate. Let $\mathcal{L}_{\mathbb{S}\mathsf{R}}$ be the RPM(p) leakage of $\mathbb{S}$R, and write

$$\mathsf{Aff}(\mathcal{L}_{\mathbb{S}\mathsf{R}}) \;\subseteq\; \boldsymbol{X} \cup \boldsymbol{Y}$$

for the set of *affected* boundary shares: those that appear in at least one boundary relation implied by $\mathcal{L}_{\mathbb{S}\mathsf{R}}$. If $\mathbb{S}$R is *symmetric*, every boundary share is affected with the same probability, and we define, for a fixed boundary share $U \in \boldsymbol{X} \cup \boldsymbol{Y}$,

$$\beta(n, p) \;\triangleq\; \Pr\big[\, U \in \mathsf{Aff}(\mathcal{L}_{\mathbb{S}\mathsf{R}}) \;\big|\; U \notin \mathcal{L}_{\mathbb{S}\mathsf{R}} \,\big] :$$

the probability that U appears in a boundary relation although it has not leaked itself. Direct leakage of U is accounted for separately below, so $\beta(n, p)$ measures only *indirect* affection, created by the leakage of the remaining wires.

The following reduction is [JMB24, Theorem 1] restated in the language of boundary relations, which makes the boundary-only dependencies explicit; this explicit form is what the leakage-diagram analysis of Section 6.3 exploits to replace Monte Carlo estimation of $\beta(n,p)$ by analytic bounds.

Theorem 1 (Refresh-boundary exposure reduction, following [JMB24])**.** *Consider a linear refresh gadget* $\mathbb{S}\mathsf{R}$ *carrying the native* X *with boundary sharings* $\boldsymbol{X}$ *and* $\boldsymbol{Y}$*, under RPM(p). There is a reduction to an ideal leak-free refresh such that the original leakage is dominated, for monotone distinguishers, by the following leakage:*

1. *an* erasure view *of the native, namely* $\phi_{\mathsf{E}_X(n,p)}(X)$*, where* $\mathsf{E}_X(n,p)$ *is as in Section 5;*
2. *independent leakage of each boundary share in* $\boldsymbol{X} \cup \boldsymbol{Y}$ *with inflated rate*

$$p' \;=\; 1-(1-p)(1-\beta(n,p)) \;=\; p+\beta(n,p)-p\,\beta(n,p).$$

Proof sketch. If the native X is revealed (probability $\mathsf{E}_X(n,p)$), item 1 accounts for it. Otherwise, revealing every affected and every directly leaked boundary share is only more informative than the boundary relations themselves; afterwards no boundary relation couples the remaining shares, and the residual leakage is simulatable as for a leak-free refresh (see the proof of [JMB24, Theorem 1] for the simulation mechanics). A fixed share U is exposed either by leaking directly or by appearing in a boundary relation without having leaked itself, so $\Pr[U \text{ exposed}] = p+(1-p)\,\beta(n,p) = 1-(1-p)(1-\beta(n,p)) = p'$ exactly. As in [JMB24, Theorem 1], the exposures of distinct shares are treated as independent; this is exact for each share's marginal rate. □

This restores the block-separated setting of Section 4 at the effective boundary rate p'. The parameter $\beta(n,p)$ can be estimated by sampling leakage instances and identifying the affected shares [JMB24]; our LAPSE tool combines this Monte Carlo procedure with importance sampling. For $\mathbb{S}\mathsf{R}$-multi we additionally prove analytic bounds via the leakage diagram, culminating in Corollary 2.

6.3 Leakage-Diagram Bounds for Composition Parameters

We now use leakage diagrams to analyze the refresh gadget $\mathbb{S}\mathsf{R}$-multi under RPM leakage. This viewpoint allows us to upper-bound the probability of boundary relations created by leakage. In the leakage diagram, essential bridges and one-sided boundary relations arise through different connectivity patterns. We begin with the characterization of essential bridges.

Essential bridges. We first show that an essential cross-boundary bridge can arise only if the top row and the bottom row of the leakage diagram become connected through leaked wires, see Figure 4. We then bound the probability of this event.

Lemma 7 (Essential bridges ⇔ top–bottom interior connectivity)**.** *Consider* $\mathbb{S}\mathsf{R}$-multi *with parameters* (n,k) *and its leakage diagram. Let* $v_{i,j}$ *denote the vertex in column* $i \in \{0,\dots,n\}$ *and row* $j \in \{0,\dots,k\}$*, and define the* potential *at this vertex by*

$$V_{i,j} \;\triangleq\; \sum_{t=1}^{i} X_t^{j+1}.$$

Thus $V_{0,j}=0$ *and* $V_{n,j}=X$ *for all* j*.*

Let $G(\mathcal{L})$ *be the subgraph containing exactly the interior edges whose corresponding wires leak (X-horizontal or C-vertical edges), together with the always-present boundary-column edges in columns* 0 *and* n*.*

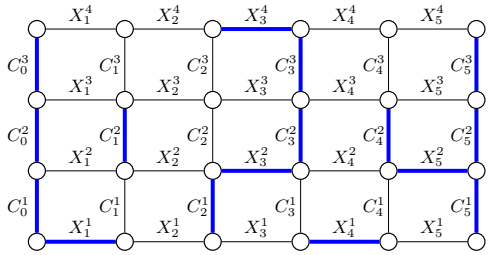


Figure 4: An example of a leakage diagram with $(n = 5, k = 3)$, with leaking and known wires specified in blue. There is a path connecting bottom and top rows, making an essential boundary bridge. The parity equation of this bridge is $(X_2^1) \oplus (X_1^4 \oplus X_2^4) = c$, where c is known by the given leakage.

Then a leakage instance $\mathcal{L}$ induces an essential bridge *between the bottom sharing $\boldsymbol{X}^1$ and the top sharing $\boldsymbol{X}^{k+1}$ if and only if there exist interior columns $a, b \in \{1, \ldots, n-1\}$ such that*

$$v_{a,0} \text{ is connected to } v_{b,k} \text{ by a path in } G(\mathcal{L}) \text{ that never visits columns } 0 \text{ or } n.$$

Proof. Each leaked edge $e = (u, v)$ determines a linear equation of the form

$$V(v) - V(u) = c_e$$

for a known constant c_e (coming from the leaked X- or C-wire label). Hence, within any connected component of $G(\mathcal{L})$, all pairwise differences $V(v) - V(u)$ are determined as signed sums of edge labels along a path.

Suppose first that there is an interior path from some $v_{a,0}$ to some $v_{b,k}$. Then $\mathcal{L}$ determines the difference

$$V_{b,k} - V_{a,0} = \sum_{t=1}^{b} X_t^{k+1} - \sum_{t=1}^{a} X_t^1 = c$$

for some known constant c. This is a boundary-only linear relation involving both the bottom and top sharings, hence a cross-boundary bridge. Because $a, b \in \{1, \ldots, n-1\}$ and the path avoids the boundary columns 0 and n, this relation is not a multiple of the correctness relation

$$\sum_{t=1}^{n} X_t^{k+1} - \sum_{t=1}^{n} X_t^1 = 0,$$

nor can it be reduced to an input-only or output-only relation by adding such a multiple. Therefore, the bridge is essential.

Conversely, suppose that $\mathcal{L}$ induces an essential bridge between $\boldsymbol{X}^1$ and $\boldsymbol{X}^{k+1}$. Any such bridge is a boundary-only relation involving both sharings and, modulo the correctness relation, can be written as

$$V_{b,k} - V_{a,0} = c$$

for some $a, b \in \{1, \ldots, n-1\}$ and some known constant c. Since this difference is determined by the leakage, the two vertices $v_{a,0}$ and $v_{b,k}$ must belong to the same connected component of $G(\mathcal{L})$. Hence there exists a path between them. Because the bridge is essential, this path cannot be routed entirely through column 0 or column n; therefore there exists such a path staying in the interior columns. □

To bound the probability of the connectivity events, we pass to an *idealized diagram* with independent edges; this step is justified by the covering result of [DFŻ19].

Lemma 8 (Idealized diagram via covering [DFŻ19, Lemma 3])**.** *Let $\mathcal{L} \leftarrow \mathrm{RPM}(p)$ on the wires of* $\mathbb{S}\mathsf{R}$-multi*, and let $\mathcal{U}$ denote the random set of interior edges of $G(\tau(\mathcal{L}))$. Let*

$\mathcal{Q}$ *denote the* idealized diagram*: each interior vertical edge is present independently with probability*

$$p_{\mathrm{LD}} \triangleq p + 2\sqrt{3p},$$

each horizontal edge is present independently with probability p*, and the boundary columns are always present. Then the distribution of* $\mathcal{Q}$ *covers the distribution of* $\mathcal{U}$ *in the sense of [DFŻ19, Definitions 1–2]. Consequently, for every monotone (edge-increasing) event* E*,*

$$\Pr[\mathcal{U} \in E] \;\le\; \Pr[\mathcal{Q} \in E].$$

Proof. Each horizontal edge has the single cause X_i^{j+1}, distinct across edges, so the horizontal edges are already independent with probability p and independent of the vertical edges, whose causes are C- and R-wires. Within one round j, a leaked C_i^j contributes the vertical edge at column i and a leaked R_i^j contributes the two adjacent vertical edges at columns $i-1$ and i (the substitution τ of Lemma 4); this is exactly the leakage projection of a refreshing bundle in [DFŻ19], and by [DFŻ19, Lemma 3] the resulting per-round edge process is covered by the standard distribution with parameter $p + 2\sqrt{3p}$. Distinct rounds involve disjoint wires and are therefore independent, so the per-round coverings combine into a covering of the product distribution. Finally, covering in the sense of [DFŻ19, Definition 1] moves probability mass only from edge sets to their supersets; hence it can only increase the probability of a monotone event. □

Every event we bound below is monotone in the edge set, so the bounds computed in the idealized diagram apply to the true leakage diagram $G(\tau(\mathcal{L}))$.

Lemma 9 (Vertical interior connectivity bound)**.** *Consider* $\mathbb{S}$R-multi *with parameters* (n,k) *under RPM(p) with* $3p_{\mathrm{LD}} < 1$*. Let* $\mathsf{VConn}_{n,k}$ *be the event that there exist* $a, b \in \{1,\dots,n-1\}$ *such that* $v_{a,0}$ *is connected to* $v_{b,k}$ *in* $G(\tau(\mathcal{L}))$ *by a path staying entirely within the interior columns* $\{1,\dots,n-1\}$*. Then*

$$\Pr[\mathsf{VConn}_{n,k}] \;\le\; (n-1)\cdot\frac{(3p_{\mathrm{LD}})^k}{1-3p_{\mathrm{LD}}}.$$

Proof. Since $\mathsf{VConn}_{n,k}$ is monotone, we may bound its probability in the idealized diagram. Fix a starting vertex $v_{a,0}$ with $a \in \{1,\dots,n-1\}$. Any path from $v_{a,0}$ to the top row contains at least k vertical edges, since each vertical step changes the row index by one; hence its length satisfies $\ell \ge k$. Every simple path of length ℓ starting at $v_{a,0}$ is in particular a walk without immediate backtracking: there are at most 3 choices for the first step (left, right, or up—row 0 has no downward neighbor) and at most 3 choices for each subsequent step, so at most 3^ℓ such walks. In the idealized diagram the edges are independent and each is present with probability at most p_{LD}, so a fixed path of length ℓ is present with probability at most p_{LD}^ℓ. By the union bound over lengths and over the $n-1$ starting vertices,

$$\Pr[\mathsf{VConn}_{n,k}] \;\le\; (n-1)\sum_{\ell=k}^{\infty} 3^\ell\, p_{\mathrm{LD}}^\ell \;=\; (n-1)\cdot\frac{(3p_{\mathrm{LD}})^k}{1-3p_{\mathrm{LD}}},$$

the series converging since $3p_{\mathrm{LD}} < 1$. □

One-sided boundary relations. We now consider boundary relations supported on only one side of the refresh boundary, namely on the input shares or on the output shares. By the top–bottom symmetry of the diagram, it suffices to analyze input-only relations. A path of $G(\tau(\mathcal{L}))$ between two bottom-row vertices $v_{a,0}$ and $v_{b,0}$ $(0 \le a < b \le n)$ reveals the difference of their potentials,

$$V_{b,0} - V_{a,0} \;=\; X_{a+1}^1 + \cdots + X_b^1 \;=\; c,$$

an input-only relation *enclosing* the interval of shares $X^1_{a+1}, \ldots, X^1_b$ (for $b = n$, use $V_{n,0} = X$ and reduce by the correctness relation). Since the input-only relations implied by an instance are spanned, modulo correctness, by these interval relations, a fixed share is affected exactly when some enclosing path exists:

$$X^1_i \text{ is in an input-only relation } \iff \exists\, 0 \le a < i \le b \le n :\ v_{a,0} \sim v_{b,0} \text{ in } G(\tau(\mathcal{L})).$$

Figure 5 illustrates two such enclosing paths.

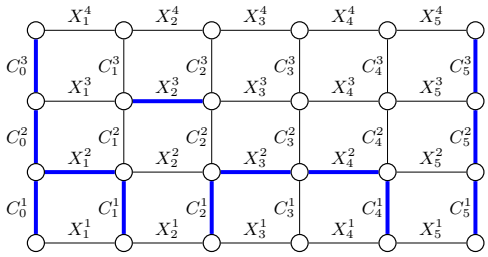


Figure 5: The same diagram layout with a different leakage instance, for $(n = 5, k = 3)$. The leakage instance exhibits a path of length 3 enclosing X^1_1, and a path of length 4 enclosing $\{X^1_3, X^1_4\}$. The parity equation of the latter is $X^1_3 \oplus X^1_4 = c$, where c is known by the given leakage.

Lemma 10 (One-sided affection)**.** *Assume* $3p_{\mathrm{LD}} < 1$*, consider* $\mathbb{S}$R-multi *with parameters* (n, k) *under RPM(p), and fix an input share* X^1_i*,* $1 \le i \le n$*. Conditioned on* X^1_i *not having leaked, the probability that it is contained in an input-only boundary relation while the native is not revealed is at most* $B_i(n, p)$*, where*

$$B_i(n,p) \triangleq \begin{cases} \dfrac{3p_{\mathrm{LD}}\, p}{(1-3p_{\mathrm{LD}})^2}, & i \in \{1, n\}, \\ \dfrac{p}{(1-3p_{\mathrm{LD}})^2}\left(\dfrac{9p^2_{\mathrm{LD}}}{1-9p^2_{\mathrm{LD}}} + \dfrac{6p \cdot 3p_{\mathrm{LD}}}{1-3p_{\mathrm{LD}}}\right), & 2 \le i \le n-1. \end{cases}$$

By the top–bottom symmetry of the diagram, the same bounds hold for the output shares.

Proof. Under the conditioning, the bottom-row edge of X^1_i is absent from the diagram, and all edge configurations counted below are determined by the leakage of the *remaining* wires; their conditional and unconditional probabilities therefore coincide.

By the enclosure criterion above, we must bound the probability that *some* pair $0 \le a < i \le b \le n$ has $v_{a,0}$ connected to $v_{b,0}$, i.e., we count all enclosing paths, over all admissible intervals $[a+1, b] \ni i$ at once. Consider an arbitrary such path. It connects a vertex in a column $\le i-1$ to a vertex in a column $\ge i$, so it crosses the gap between columns $i-1$ and i; since the edge of X^1_i is absent, every crossing is a horizontal edge X^{j+1}_i with $j \in \{1, \ldots, k\}$. Split the path at its *first* crossing: the part before the crossing connects $v_{a,0}$ to $v_{i-1,j}$ within columns $\le i-1$, and the part after connects $v_{i,j}$ to $v_{b,0}$. Thus the left part reaches row 0 (if $a \ge 1$) or column 0 (if $a = 0$), the right part reaches row 0 (if $b \le n-1$) or column n (if $b = n$), and reaching *both* outer columns is excluded, as it would connect the two boundary columns and reveal the native (Lemma 5). The existence of such an anchored configuration is a monotone event, so we bound it in the idealized diagram, where the anchor and the two parts occupy distinct, independent edges.

Case $i \in \{1, n\}$. Here the crossing edge has an endpoint in a boundary column, so the corresponding part is empty, and the opposite part cannot reach the other boundary column (excluded above); it is an interior-column path from height j to its first bottom-row visit, of length $\ell \ge j$, counted by at most 3^ℓ non-backtracking walks, each present with

probability at most p_{LD}^{ℓ}. With the horizontal anchor of probability p,

$$\Pr[\,\cdot\,] \;\le\; \sum_{j=1}^{k} p \sum_{\ell \ge j} 3^{\ell} p_{\mathrm{LD}}^{\ell} \;\le\; \frac{3p_{\mathrm{LD}}\, p}{(1-3p_{\mathrm{LD}})^2}.$$

Case $2 \le i \le n-1$. Counting non-backtracking walks as above, a part reaching row 0 from row j contributes at most $(3p_{\mathrm{LD}})^j/(1-3p_{\mathrm{LD}})$, while a part reaching column 0 (resp. column n) contains at least $i-1$ (resp. $n-i$) horizontal edges and contributes at most $(3p)^{i-1}/(1-3p_{\mathrm{LD}})$ (resp. $(3p)^{n-i}/(1-3p_{\mathrm{LD}})$). Of the four endpoint combinations, (column 0, column n) is excluded, since it reveals the native; summing the remaining three over the anchor height j and bounding $(3p)^{i-1}, (3p)^{n-i} \le 3p$ gives

$$\Pr[\,\cdot\,] \;\le\; \frac{p}{(1-3p_{\mathrm{LD}})^2} \sum_{j=1}^{k} \Big[(3p_{\mathrm{LD}})^{2j} + 6p\,(3p_{\mathrm{LD}})^{j} \Big] \;\le\; B_i(n,p). \qquad \square$$

Putting it all together. The following corollary combines the bounds for essential bridges and one-sided boundary relations to upper-bound the exposure rate $\beta(n,p)$ for $\mathbb{S}$R-multi; it is the single place where numeric thresholds enter. Bridges that are not essential reduce, modulo the correctness relation, to one-sided relations; since boundary relations are counted beyond the correctness relation (the system $\boldsymbol{P}_{\mathrm{Inf}}$ of Section 6.1), their contribution is covered by the one-sided term.

Corollary 2. *For* $\mathbb{S}$R-multi *under RPM(p),*

$$\beta(n,p) \;\le\; \mathsf{E}_X(n,p) \;+\; (n-1)\cdot \frac{(3p_{\mathrm{LD}})^k}{1-3p_{\mathrm{LD}}} \;+\; \max_i B_i(n,p), \qquad p_{\mathrm{LD}} = p + 2\sqrt{3p}.$$

If $p \le 1/450$, *then* $\max_i B_i(n,p) \le 2p$; *if, in addition,* $(n-1)(3p_{\mathrm{LD}})^k \le p\,(1-3p_{\mathrm{LD}})/2$—*which holds for* $k = n$ *as soon as* $(n-1)2^{-n} \le p/4$—*and* $(k+1)(3p)^{n-1} \le 1/4$, *then*

$$\beta(n,p) \;\le\; 3p \qquad \textit{and hence} \qquad p' \;=\; 1-(1-p)\big(1-\beta(n,p)\big) \;\le\; 4p.$$

Proof. Fix a boundary share U and condition on $U \notin \mathcal{L}_{\mathbb{S}\mathsf{R}}$; decompose over the native-reveal event. The reveal event is monotone in the leaked wire set, and conditioning on $U \notin \mathcal{L}$ removes one potential edge, which cannot increase the probability of a monotone event; hence the reveal term is at most $\mathsf{E}_X(n,p)$. Outside the reveal event, either the instance induces an essential bridge—which, by Lemma 7, requires the interior top–bottom connectivity event $\mathsf{VConn}_{n,k}$, again monotone, so Lemma 9 applies under the conditioning—or U is contained in a one-sided relation, an event of conditional probability at most $B_i(n,p)$ by Lemma 10. The union bound gives the first claim.

Now let $p \le 1/450$, so that $3p_{\mathrm{LD}} \le \frac{1}{2}$ and $9\sqrt{p} \le \frac{1}{2}$. With $u = 3p_{\mathrm{LD}}$, the boundary-adjacent bound satisfies $u/(1-u)^2 \le 2$ for $u \le \frac{1}{2}$, so $B_1(n,p) \le 2p$; for the interior bound, $p_{\mathrm{LD}}^2 \le 13p$ gives a bracket of at most $12p_{\mathrm{LD}}^2 + 6p \le 162p$, whence $B_i(n,p) \le 4 \cdot 162\, p^2 = 648\, p^2 \le 2p$, since $p \le 1/450 < 2/648$. Under the two additional hypotheses, the bridge term is at most $p/2$, and by Corollary 1, $\mathsf{E}_X(n,p) \le p \cdot \frac{(k+1)(3p)^{n-1}}{1-9\sqrt{p}} \le p \cdot \frac{1}{4} \cdot 2 = p/2$. Hence $\beta(n,p) \le p/2 + p/2 + 2p = 3p$ and $p' \le p + \beta(n,p) \le 4p$. $\square$

A plug-and-play composition statement. We can now combine the pieces into a single circuit-level bound, making the local-to-global transition explicit: Theorem 1 restores block separation at the boundary rate p', Lemma 3 turns each block's leakage into an erasure of its native, and Lemma 2 converts the erasures into a global simulation distance. Figure 6 illustrates the reduction.

Theorem 2 (Plug-and-play composition)**.** *Let* $\mathbb{SC}_n$ *be a masked circuit consisting of gadget blocks* $\mathbb{SG}_1, \mathbb{SG}_2, \ldots$ *separated by linear refresh gadgets* $\mathbb{SR}$ *satisfying the assumptions of Theorem 1, with target natives* $X_1, \ldots, X_m$ *(counting the natives carried by the refreshes among them). Under RPM(p) with independent leakage, and treating boundary-share exposures as independent per share as in Theorem 1, the full circuit leakage is* ε*-simulatable (Definition 4) with*

$$\varepsilon \;\le\; \sum_{i=1}^{m} \mathsf{E}_{X_i}(n, p'), \qquad p' = p + \beta(n,p) - p\,\beta(n,p).$$

Proof sketch. Apply Theorem 1 to every refresh boundary. This replaces each real refresh by a leak-free one, so the circuit becomes block-separated in the sense of Section 4 (Assumption 2), at the cost of leaking each boundary share independently at the inflated rate p'; interior wires still leak at rate $p \le p'$. The reveal event of a block is monotone in the set of leaked wires, and RPM(p) leakage can be coupled inside RPM(p') leakage wire by wire, so each block's reveal probability is at most $\mathsf{E}_{X_i}(n, p')$. By Lemma 3, each block's leakage is then an erasure view $\phi_{\mathsf{E}_{X_i}(n,p')}(X_i)$ of its native (Assumption 3), and Lemma 2 with data processing (Section 4) yields the claimed bound. □

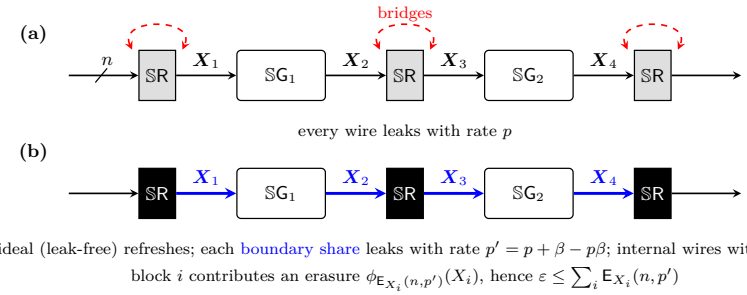


Figure 6: Using Theorem 2. (a) A refresh-separated circuit under RPM(p): every wire leaks with rate p, and refresh leakage can create bridges (dashed) that couple the sharings across a refresh boundary. (b) After applying Theorem 1 to every refresh: refreshes are ideal (black), each boundary share leaks independently at the inflated rate p', and internal wires still leak at rate p. Each block then contributes an erasure of its native, and the erasures compose by Lemma 2.

Scope of Theorem 2. The theorem splits the leakage evaluation of a composed circuit into independent evaluations of its gadgets; the only cost of composition is that each gadget is evaluated at the inflated rate p' instead of p. The main burden of proof lies on the refreshes, which we require to be linear: Theorem 1 and our bounds on $\beta(n,p)$ rely on their linear structure. The gadgets themselves are only required to admit a per-gadget erasure parameter. For linear gadgets, this parameter is the reveal probability $\mathsf{E}_X(n,p)$, established by Lemma 3; for any other gadget, such as a multiplication, any function of (n,p) bounding the adversary's success in the sense of Assumption 3 can be used instead, evaluated at p' rather than p. For multiplication, such a bound is currently available only in part: formally for Expand (Proposition 1) and heuristically for Compress (Section 7). A formal end-to-end bound for a complete circuit such as AES is therefore left to future work.

7 Random-Probing-Secure Multiplication Gadget

An RPM-secure multiplication gadget is essential for masking general circuits. Recent works study multiplication security in the RPM using both the cardinal approach [BNR25]

and Monte Carlo estimation [JMB24, BC26]. In this section we (i) recall why the classical ISW multiplication [ISW03] is insecure in the RPM, and (ii) outline the main idea behind recent patches. Using the tools developed earlier, we then give a *conditional formal bound* for the Expand stage—discharging the ideal-refresh assumption via Theorem 1—together with a clearly labeled heuristic estimate for Compress, and we state explicitly the technical obstacles that currently prevent a fully formal treatment. This discussion also clarifies why, for the purpose of bounding RPM leakage, the multiplication gadget can be upper-bounded via associated $\mathbb{F}_2$-linear subcircuits, as already hinted in Section 3.

Overview of ISW multiplication. A multiplication gadget takes two n-sharings $\boldsymbol{X} = \{X_1, \ldots, X_n\}$ and $\boldsymbol{Y} = \{Y_1, \ldots, Y_n\}$ and outputs an n-sharing $\boldsymbol{Z} = \{Z_1, \ldots, Z_n\}$ of $Z = XY$. The classical ISW construction [ISW03] first forms all cross-products X_iY_j and then combines them (with fresh randomness) into n output shares. While ISW achieves threshold probing security, it is not secure in the random probing model [BCPZ16,BCGR24].

Root cause in the RPM. The issue is *share reuse*: each input share X_i participates in n cross-products $X_iY_1, \ldots, X_iY_n$. In an SSA/wire model this typically induces n (correlated) opportunities for X_i to be revealed through wires that depend on X_i. A simple (and intentionally pessimistic) heuristic is to treat these as n independent chances, yielding

$$\Pr[X_i \text{ effectively leaks}] \approx 1 - (1-p)^n.$$

If this happens for all $i \in \{1, \ldots, n\}$, then the native $X = \sum_i X_i$ is recovered, with probability on the order of

$$\Pr[\text{all } X_i \text{ leak}] \approx \left(1 - (1-p)^n\right)^n.$$

For any fixed $p > 0$, this expression tends to 1 as $n \to \infty$, reflecting that increasing the masking order does *not* repair the leakage amplification caused by reuse:

$$\forall p > 0 : \quad \lim_{n \to \infty} \left(1 - (1-p)^n\right)^n = 1.$$

A more accurate analysis of the vulnerable "multiplication phase" can be carried out via rook-domination polynomials [BCGR24]. For practical demonstrations of this weakness in concrete leakage settings, see [BCPZ16].

7.1 Towards RPM-Secure Multiplication

Battistello et al. [BCPZ16] proposed an RPM-oriented multiplication gadget (with subsequent refinements in [JMB24, BNR25, BC26]) that mitigates the share-reuse amplification underlying the RPM insecurity of ISW. The gadget proceeds in three stages: (i) Expand each input sharing into n^2 refreshed multiplicands, (ii) multiply them entrywise, and (iii) Compress the n^2 products back to an n-sharing using the usual ISW-style compression.

Expand-by-refresh idea. Let $\boldsymbol{X} = \{X_1, \ldots, X_n\}$ and $\boldsymbol{Y} = \{Y_1, \ldots, Y_n\}$ be n-sharings of X and Y. The starting point is the identity

$$\begin{aligned}\sum \left[\{X_1, \ldots, X_n\} \times \{Y_1, \ldots, Y_n\}\right] = \\ \sum \left[\mathbb{S}\mathsf{R}(\boldsymbol{X}_L) \times \mathbb{S}\mathsf{R}(\boldsymbol{Y}_L)\right] + \sum \left[\mathbb{S}\mathsf{R}(\boldsymbol{X}_L) \times \mathbb{S}\mathsf{R}(\boldsymbol{Y}_R)\right] + \\ \sum \left[\mathbb{S}\mathsf{R}(\boldsymbol{X}_R) \times \mathbb{S}\mathsf{R}(\boldsymbol{Y}_L)\right] + \sum \left[\mathbb{S}\mathsf{R}(\boldsymbol{X}_R) \times \mathbb{S}\mathsf{R}(\boldsymbol{Y}_R)\right],\end{aligned}$$

where $(\boldsymbol{X}_L, \boldsymbol{X}_R)$ (resp. $(\boldsymbol{Y}_L, \boldsymbol{Y}_R)$) is the partition of $\boldsymbol{X}$ (resp. $\boldsymbol{Y}$) into left and right halves, and each occurrence of $\mathbb{S}\mathsf{R}(\cdot)$ denotes an *independent* refresh invocation with fresh randomness. The same divide-and-refresh logic applies recursively to half-sized subsets

of shares, and the actual multiplications are carried out only when no further division is possible. Algorithm 5 shows how the X-side operands are recursively derived and stored in an $n \times n$ matrix $\boldsymbol{A}$. The Y-side matrix $\boldsymbol{B}$ is derived analogously, with the block ordering

$$\boldsymbol{B} = \begin{bmatrix} \mathsf{Expand}\left(\mathbb{SR}(\boldsymbol{Y}_L)\right) & \mathsf{Expand}\left(\mathbb{SR}(\boldsymbol{Y}_R)\right) \\ \mathsf{Expand}\left(\mathbb{SR}(\boldsymbol{Y}_L)\right) & \mathsf{Expand}\left(\mathbb{SR}(\boldsymbol{Y}_R)\right) \end{bmatrix}.$$

After expanding $\boldsymbol{X}$ and $\boldsymbol{Y}$ in $\boldsymbol{A}$ and $\boldsymbol{B}$ matrices, we then compute the entrywise product

$$\boldsymbol{C} \;=\; \boldsymbol{A} \odot \boldsymbol{B}, \qquad C_{i,j} = A_{i,j} B_{i,j}.$$

Finally, Algorithm 6 applies the standard ISW compression to map the n^2 products in $\boldsymbol{C}$ to an output n-sharing $\boldsymbol{Z}$ of $Z = XY$.

Algorithm 5 Expand

Initial Input: An n-sharing $\boldsymbol{X}$
Final Output: An $n \times n$ Expansion

$n = \mathsf{Length}(\boldsymbol{X})$
if $n = 1$ **then**
  return $\left[X_1\right]$
$m = \lfloor n/2 \rfloor$
$\boldsymbol{X}_L = \{X_1, \ldots, X_m\}$
$\boldsymbol{X}_R = \{X_{m+1}, \ldots, X_n\}$
$\boldsymbol{A} = \begin{bmatrix} \mathsf{Expand}(\mathbb{SR}(\boldsymbol{X}_L)) & \mathsf{Expand}(\mathbb{SR}(\boldsymbol{X}_L)) \\ \mathsf{Expand}(\mathbb{SR}(\boldsymbol{X}_R)) & \mathsf{Expand}(\mathbb{SR}(\boldsymbol{X}_R)) \end{bmatrix}$
return $\boldsymbol{A}$

Algorithm 6 Compress

Input: The n^2 multiplications in $\boldsymbol{C}_{i,j}$
Output: An n-sharing $\boldsymbol{Z}$ for Z

for $i = 1$ **to** n **do**
  for $j = i + 1$ **to** n **do**
    $R \xleftarrow{\$} \mathbb{F}_q$
    $C_{i,j} = C_{i,j} + R$
    $C_{j,i} = C_{j,i} - R$
for $i = 1$ **to** n **do**
  $Z_i = C_{i,i}$
  for $j = 1$ **to** n, $j \neq i$ **do**
    $Z_i = Z_i + C_{i,j}$
return $\boldsymbol{Z} = \{Z_1, \ldots, Z_n\}$

RPM security estimate. The gadget consists of linear stages (Expand on $\boldsymbol{X}, \boldsymbol{Y}$ and linear Compress) separated by n^2 nonlinear products $C_{i,j} = A_{i,j} B_{i,j}$. We upper-bound RPM($p$) leakage by a local augmentation at each multiplication: if any of $(A_{i,j}, B_{i,j}, C_{i,j})$ leaks, we additionally reveal the other two. This exposes the nonlinear layer as a leakage interface and inflates the effective rate to

$$p_\times \;=\; 1 - (1-p)^3 \;\leq\; 3p.$$

Note that within any one linear stage the augmented exposures of distinct interface wires are *independent* (distinct triples involve disjoint wires), so each linear stage can be analyzed under honest i.i.d. RPM($p_\times$) leakage on its interface; the perfect correlation of exposures *across* stages (through a shared triple) is absorbed by the final union bound. Conservatively, we evaluate the linear stages at the inflated rate $p_\times$ throughout, using Theorem 1 for the internal refreshes.

Expand estimate. Assume $n = 2^\ell$, and let $\mathsf{E}^{\mathsf{Exp}}(n, \rho)$ denote the probability that leakage within Expand at rate ρ reveals

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

Rather than *assuming* ideal internal refreshes, we discharge this assumption by Theorem 1: each internal refresh is replaced by a leak-free one at the cost of the inflated boundary rate

$$p'_\times \;=\; p_\times + \beta(n, p_\times) - p_\times\, \beta(n, p_\times).$$

With leak-free refreshes, the four quadrant expansions carry independent re-sharings of the two half-sums, and each half-sum is expanded twice with fresh randomness. This yields the recursion

$$\mathsf{E}^{\mathsf{Exp}}(n, p'_\times) = \left(1 - (1 - \mathsf{E}^{\mathsf{Exp}}(n/2, p'_\times))^2\right)^2, \qquad \mathsf{E}^{\mathsf{Exp}}(1, p'_\times) = p'_\times.$$

Using $1-(1-u)^2 \le 2u$ as a conservative upper bound yields the following.

Proposition 1 (Conditional Expand bound)**.** *Let $n = 2^\ell$ and let the internal refreshes of Expand satisfy the assumptions of Theorem 1 with boundary-exposure rate β. Then, under the interface augmentation above,*

$$\mathsf{E}^{Exp}(n, p'_\times) \;\le\; \tfrac{1}{4}\,(4p'_\times)^n,$$

which decreases exponentially in n whenever $p'_\times < \frac{1}{4}$.

Compress estimate (heuristic). Fix i. The circuit computes Z_i as a length-n sum and implicitly forms intermediate partial sums. The dominant reconstruction patterns for Z_i are: either Z_i leaks directly, or some checkpoint partial sum $S_{i,t}$ leaks together with all remaining $(n-1-t)$ addends needed to complete the sum (this enumeration is heuristic in that it ignores reconstructions combining leaked randomness wires across rows). For each t, the corresponding pattern requires at least $1+(n-1-t)$ leaked wires, and therefore has probability at most $p_\times^{n-t}$. Summing over the n possible checkpoints,

$$\Pr[Z_i \text{ is known}] \;\lesssim\; \sum_{u=1}^{n} p_\times^u \;=\; \frac{p_\times(1-p_\times^n)}{1-p_\times} \;\le\; 2p_\times \qquad (p_\times \le \tfrac{1}{2}).$$

Heuristically, since the randomness injected into different rows and the terms $C_{i,j}$ are largely disjoint, the events "Z_i is reconstructible" are close to independent in practice. Thus, recovering the native

$$Z = \sum_i Z_i$$

requires reconstructing essentially all n output shares, which suggests

$$\mathsf{E}_Z^{\mathsf{Comp}}(n, p_\times) \lesssim (2p_\times)^n.$$

Combined estimate. Combining the above bounds for the three natives gives

$$\Pr[\text{recover } X \text{ or } Y \text{ or } Z] \;\lesssim\; 2\,\mathsf{E}^{\mathsf{Exp}}(n, p'_\times) \;+\; \mathsf{E}_Z^{\mathsf{Comp}}(n, p_\times) \;\lesssim\; \frac{1}{2}(4p'_\times)^n + (2p_\times)^n.$$

For small p (where $\beta(n, p_\times) = O(p_\times)$ and hence $p'_\times \le c\,p$ for a constant c), this suggests exponential decay in n at any sufficiently small constant rate.

Toward a formal claim, and remaining obstacles. The Expand part of the above is formal conditioned on the stated assumptions (Proposition 1), and it accounts for internal-refresh bridges through Theorem 1—addressing the composition role of refreshes inside the gadget. If the Compress estimate were established formally, the multiplication gadget would slot directly into Theorem 2 as one of the gadget blocks, completing the composition picture. Three technical obstacles currently prevent this:

1. **Randomness reuse in Compress.** Each fresh R is added to $C_{i,j}$ and subtracted from $C_{j,i}$, coupling rows i and j; this is the intra-gadget analogue of the bridges of Section 6, and a leakage-diagram analysis of the resulting dependency graph (rows as columns of a diagram) is open.

2. **Exhaustiveness of reconstruction patterns.** The checkpoint enumeration for Z_i ignores reconstructions that combine leaked randomness wires across rows; bounding these requires the analysis of item 1.

3. **Independence across output shares.** The step from $\Pr[Z_i \text{ known}] \leq 2p_\times$ to $\mathsf{E}_Z^{\mathsf{Comp}} \lesssim (2p_\times)^n$ assumes near-independence of the per-row events; the shared randomness of item 1 is precisely what could violate it, so items 1–3 stand or fall together.

We leave the resolution of these obstacles, and hence a fully formal multiplication bound, to future work; concrete values of $\mathsf{E}_Z^{\mathsf{Comp}}$ can meanwhile be estimated with LAPSE, since Compress is linear in its inputs. As a consequence, Theorem 2 cannot yet be instantiated formally for complete masked circuits containing multiplications, such as AES; such a bound, and its quantitative comparison with propagation-based and cardinal bounds [BRR25, BNR25], are left to future work.

8 LAPSE

We implement our evaluation methods in LAPSE (Linear Algebraic Probing Security Estimator), a tool for analyzing *linear* masked gadgets in the random probing model.[2] It computes or estimates the two local quantities used throughout this work: the recovery probability $\mathsf{E}_X(n,p)$ of a target native and the boundary-share exposure rate $\beta(n,p)$ of Theorem 1.

LAPSE takes a Python-like description of a linear gadget over a finite field, compiles it into static single assignment (SSA) form, and associates to each leakable wire a vector over the source variables (input shares and fresh randomness). RPM leakage questions then reduce to linear algebra on the leaked wire vectors: a target native is recoverable exactly when its vector lies in their span, and a boundary share is affected exactly when it appears in a boundary relation of the reduced leaked system. Accordingly, the tool evaluates

$$\mathsf{E}_X(n,p) = \Pr\big[\text{leakage recovers } X\big] \quad \text{and} \quad \beta(n,p) = \Pr\big[U \in \mathsf{Aff}(\mathcal{L}_{\mathbb{SR}}) \;\big|\; U \notin \mathcal{L}_{\mathbb{SR}}\big]$$

for a fixed boundary share U of a symmetric refresh gadget; the direct leakage of U is excluded, matching the exposure-rate formula of Theorem 1. For small gadgets these quantities are computed exactly by enumerating all leakage patterns; for larger ones, LAPSE uses Monte Carlo sampling and switches to importance sampling when the target event is too rare to observe reliably at the original probing rate, reporting variance/confidence information alongside the estimate. The tool is intended for local linear-algebraic analysis and boundary-relation checks; it is not a general-purpose analyzer for arbitrary masked circuits.

Comparison with IronMask. IronMask [BMRT22] is a general-purpose verifier of probing and random-probing properties (e.g., SNI, RPC, and RPE), complete for linear and quadratic gadgets such as standard multiplications; in the random probing setting, it enumerates probe tuples up to a chosen size and computes the exact number of failure tuples of each size. LAPSE is restricted to linear gadgets, but it targets the two quantities required by our composition theorem, $\mathsf{E}_X(n,p)$ and $\beta(n,p)$, and estimates them by sampling, using importance sampling in the low-p, high-n regime, where the dominant failure tuples have size about n and exhaustive enumeration becomes infeasible. The two tools are thus complementary: exact verification of general gadgets of moderate size on one side, statistical estimation of linear-gadget parameters at larger scale on the other.

9 Conclusion

We studied the random probing model (RPM) as a concrete and composable abstraction for side-channel leakage, with an emphasis on *analytic* security estimates beyond simulation-

[2] https://github.com/vahid-jahandideh/LAPSE.git

heavy regimes. We (i) made noisy Hamming-weight → RPM parameters explicit over $\mathbb{F}_{2^u}$ and showed that an $\mathbb{F}_2$-linear inner-product reduction yields substantially tighter effective probing rates; (ii) characterized RPM leakage in $\mathbb{F}_q$-linear circuits as an erasure event via a vector-space representation, relating $\mathsf{E}_X(n,p)$ to adversarial advantage; (iii) adapted importance sampling to estimate rare RPM recovery events in low-p/high-n regimes; and (iv) analyzed composition across refresh boundaries via *bridges*: we derived analytic leakage-diagram bounds for the exposure rate $\beta(n,p)$ of the localization reduction of [JMB24]—previously accessible only by Monte Carlo—clarified why TPM notions such as SNI do not capture the RPM boundary phenomenon, and combined the pieces into a plug-and-play composition theorem (Theorem 2), with the compact rule $p' \leq 4p$ for $\mathbb{S}$R-multi at practical rates and moderate refresh depth (Corollary 2). We implemented these techniques in LAPSE, supporting exact, Monte Carlo, and importance-sampling evaluation of RPM parameters.

9.1 Future research directions

Sharper composition bounds. Our one-sided bounds scale as $p^{3/2}$ rather than the empirically observed p^2: the covering parameter $p_{\mathrm{LD}} = \Theta(\sqrt{p})$ of [DFŻ19, Lemma 3] must account, uniformly over all edge configurations, for the two diagram edges that a single leaked R-wire creates at once, whereas along the dominant enclosing paths no two edges can share a cause. A finer, path-by-path accounting of shared causes should therefore yield $O(p^2)$ bounds, matching the Monte Carlo estimates. Extending the analytic treatment beyond linear gadgets is a further natural step.

Impact of the field size q. Recent works advocate masking over prime fields, motivated in part by improved concrete security of a *single* encoding as the prime q grows [CMM+23, GMM+24, FMM+24, BGM+25], in line with noisy-leakage analyses suggesting that prime fields amplify noise more effectively than extension fields [DFS16]; the q-dependence appearing in noisy-to-RPM reductions (questioned already in [DFS19]) has instead been attributed to looseness of the reduction [BCG+23, BDF24]. Working directly in the RPM, the relevance of q is less clear: the linear gadgets studied here admit the same vector-space and leakage-diagram structure over any $\mathbb{F}_q$, our LAPSE experiments—including on $\mathbb{S}$R-multi—show no meaningful dependence of $\mathsf{E}_X(n,p)$ or $\beta(n,p)$ on q, and the analytic bounds of this work are q-independent. Whether q has an inherent effect on RPM security at the level of complete *circuits*—beyond reduction-induced changes in the effective probing rate—remains open, and clarifying it would help assess circuit-level benefits of prime-field masking beyond isolated encodings.

Acknowledgments

We thank the anonymous reviewers and our shepherd for their careful reading and constructive feedback, which helped improve the clarity and presentation of this paper.

Lejla Batina and Vahid Jahandideh are supported by the European Union under Grant Agreement No. 101225691 (QARC). Lejla Batina additionally acknowledges support from the European Union under Grant Agreement No. 101217315 (SECURE-NET) and from the Dutch Research Council (NWO) through the projects PROACT (NWA.1215.18.014), TTW PREDATOR (19782), TTW BASES (20858), CiCS within the Gravitation research programme (024.006.037), and ESCAPE within the Open Competition programme (OCENW.XL.23.112).

Bart Mennink is supported by the Netherlands Organisation for Scientific Research (NWO) under grant VI.Vidi.203.099.

References

[BBD+16] Gilles Barthe, Sonia Belaïd, François Dupressoir, Pierre-Alain Fouque, Benjamin Grégoire, Pierre-Yves Strub, and Rébecca Zucchini. Strong Non-Interference and Type-Directed Higher-Order Masking. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 116–129, 2016.

[BC26] Sonia Belaïd and Gaëtan Cassiers. PERSEUS – Probabilistic Evaluation of Random Probing SEcurity Using Efficient Sampling. In Joan Daemen and Emmanuel Thomé, editors, *Advances in Cryptology – EUROCRYPT 2026*, pages 506–536, Cham, 2026. Springer Nature Switzerland.

[BCG+23] Julien Béguinot, Wei Cheng, Sylvain Guilley, Yi Liu, Loïc Masure, Olivier Rioul, and François-Xavier Standaert. Removing the Field Size Loss from Duc et al.'s Conjectured Bound for Masked Encodings. In Elif Bilge Kavun and Michael Pehl, editors, *Constructive Side-Channel Analysis and Secure Design - 14th International Workshop, COSADE 2023, Munich, Germany, April 3-4, 2023, Proceedings*, volume 13979 of *Lecture Notes in Computer Science*, pages 86–104. Springer, 2023.

[BCGR24] Julien Béguinot, Wei Cheng, Sylvain Guilley, and Olivier Rioul. Formal Security Proofs via Doeblin Coefficients: Optimal Side-Channel Factorization from Noisy Leakage to Random Probing. In Leonid Reyzin and Douglas Stebila, editors, *Advances in Cryptology - CRYPTO 2024 - 44th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2024, Proceedings, Part VI*, volume 14925 of *Lecture Notes in Computer Science*, pages 389–426. Springer, 2024.

[BCP+20] Sonia Belaïd, Jean-Sébastien Coron, Emmanuel Prouff, Matthieu Rivain, and Abdul Rahman Taleb. Random Probing Security: Verification, Composition, Expansion and New Constructions. In Daniele Micciancio and Thomas Ristenpart, editors, *Advances in Cryptology – CRYPTO 2020*, pages 339–368, Cham, 2020. Springer International Publishing.

[BCPZ16] Alberto Battistello, Jean-Sébastien Coron, Emmanuel Prouff, and Rina Zeitoun. Horizontal Side-Channel Attacks and Countermeasures on the ISW Masking Scheme. In Benedikt Gierlichs and Axel Y. Poschmann, editors, *Cryptographic Hardware and Embedded Systems – CHES 2016*, pages 23–39, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg.

[BDF24] Gianluca Brian, Stefan Dziembowski, and Sebastian Faust. From Random Probing to Noisy Leakages Without Field-Size Dependence. In Marc Joye and Gregor Leander, editors, *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*, volume 14654 of *Lecture Notes in Computer Science*, pages 345–374. Springer, 2024.

[BFO23] Francesco Berti, Sebastian Faust, and Maximilian Orlt. Provable Secure Parallel Gadgets. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2023(4):420–459, 2023.

[BGM+25] Brieuc Balon, Lorenzo Grassi, Pierrick Méaux, Thorben Moos, François-Xavier Standaert, and Matthias Johann Steiner. mid-pSquare: Leveraging the Strong

Side-Channel Security of Prime-Field Masking in Software. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2025(4):486–519, 2025.

[BMRT22] Sonia Belaïd, Darius Mercadier, Matthieu Rivain, and Abdul Rahman Taleb. IronMask: Versatile Verification of Masking Security. In *43rd IEEE Symposium on Security and Privacy, SP 2022, San Francisco, CA, USA, May 22-26, 2022*, pages 142–160. IEEE, 2022.

[BNR25] Sonia Belaïd, Victor Normand, and Matthieu Rivain. Masked Circuit Compiler in the Cardinal Random Probing Composability Framework. In Goichiro Hanaoka and Bo-Yin Yang, editors, *Advances in Cryptology - ASIACRYPT 2025 - 31st International Conference on the Theory and Application of Cryptology and Information Security, Melbourne, VIC, Australia, December 8-12, 2025, Proceedings, Part II*, volume 16246 of *Lecture Notes in Computer Science*, pages 375–406. Springer, 2025.

[BRR25] Sonia Belaïd, Matthieu Rivain, and Mélissa Rossi. New Techniques for Random Probing Security and Application to Raccoon Signature Scheme. In Serge Fehr and Pierre-Alain Fouque, editors, *Advances in Cryptology - EUROCRYPT 2025 - 44th Annual International Conference on the Theory and Applications of Cryptographic Techniques Madrid, Spain, May 4-8, 2025, Proceedings, Part VIII*, volume 15608 of *Lecture Notes in Computer Science*, pages 94–123. Springer, 2025.

[BRT21] Sonia Belaïd, Matthieu Rivain, and Abdul Rahman Taleb. On the Power of Expansion: More Efficient Constructions in the Random Probing Model. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology - EUROCRYPT 2021 - 40th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, October 17-21, 2021, Proceedings, Part II*, volume 12697 of *Lecture Notes in Computer Science*, pages 313–343. Springer, 2021.

[BS21] Olivier Bronchain and François-Xavier Standaert. 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.

[CFOS21] Gaëtan Cassiers, Sebastian Faust, Maximilian Orlt, and François-Xavier Standaert. Towards Tight Random Probing Security. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology - CRYPTO 2021 - 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16-20, 2021, Proceedings, Part III*, volume 12827 of *Lecture Notes in Computer Science*, pages 185–214. Springer, 2021.

[CGPZ16] Jean-Sébastien Coron, Aurélien Greuet, Emmanuel Prouff, and Rina Zeitoun. Faster Evaluation of SBoxes via Common Shares. In Benedikt Gierlichs and Axel Y. Poschmann, editors, *Cryptographic Hardware and Embedded Systems - CHES 2016 - 18th International Conference, Santa Barbara, CA, USA, August 17-19, 2016, Proceedings*, volume 9813 of *Lecture Notes in Computer Science*, pages 498–514. Springer, 2016.

[CJRR99] Suresh Chari, Charanjit S. Jutla, Josyula R. Rao, and Pankaj Rohatgi. Towards Sound Approaches to Counteract Power-Analysis Attacks. In Michael Wiener, editor, *Advances in Cryptology — CRYPTO' 99*, pages 398–412, Berlin, Heidelberg, 1999. Springer Berlin Heidelberg.

[CMM+23] Gaëtan Cassiers, Loïc Masure, Charles Momin, Thorben Moos, and François-Xavier Standaert. Prime-Field Masking in Hardware and its Soundness against Low-Noise SCA Attacks. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2023(2):482–518, Mar. 2023.

[DDF19] Alexandre Duc, Stefan Dziembowski, and Sebastian Faust. Unifying Leakage Models: From Probing Attacks to Noisy Leakage. *J. Cryptol.*, 32(1):151–177, 2019.

[DFS15] Stefan Dziembowski, Sebastian Faust, and Maciej Skorski. Noisy Leakage Revisited. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology - EUROCRYPT 2015*, pages 159–188, Berlin, Heidelberg, 2015. Springer Berlin Heidelberg.

[DFS16] Stefan Dziembowski, Sebastian Faust, and Maciej Skórski. Optimal Amplification of Noisy Leakages. In Eyal Kushilevitz and Tal Malkin, editors, *Theory of Cryptography - 13th International Conference, TCC 2016-A, Tel Aviv, Israel, January 10-13, 2016, Proceedings, Part II*, volume 9563 of *Lecture Notes in Computer Science*, pages 291–318. Springer, 2016.

[DFS19] Alexandre Duc, Sebastian Faust, and François-Xavier Standaert. Making Masking Security Proofs Concrete (Or How to Evaluate the Security of Any Leaking Device), Extended Version. *J. Cryptol.*, 32(4):1263–1297, 2019.

[DFŻ19] Stefan Dziembowski, Sebastian Faust, and Karol Żebrowski. Simple Refreshing in the Noisy Leakage Model. In Steven D. Galbraith and Shiho Moriai, editors, *Advances in Cryptology – ASIACRYPT 2019*, pages 315–344, Cham, 2019. Springer International Publishing.

[FMM+24] Sebastian Faust, Loïc Masure, Elena Micheli, Maximilian Orlt, and François-Xavier Standaert. Connecting Leakage-Resilient Secret Sharing to Practice: Scaling Trends and Physical Dependencies of Prime Field Masking. In Marc Joye and Gregor Leander, editors, *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*, volume 14654 of *Lecture Notes in Computer Science*, pages 316–344. Springer, 2024.

[GJR18] Dahmun Goudarzi, Antoine Joux, and Matthieu Rivain. How to Securely Compute with Noisy Leakage in Quasilinear Complexity. In Thomas Peyrin and Steven D. Galbraith, editors, *Advances in Cryptology - ASIACRYPT 2018 - 24th International Conference on the Theory and Application of Cryptology and Information Security, Brisbane, QLD, Australia, December 2-6, 2018, Proceedings, Part II*, volume 11273 of *Lecture Notes in Computer Science*, pages 547–574. Springer, 2018.

[GMM+24] Lorenzo Grassi, Loïc Masure, Pierrick Méaux, Thorben Moos, and François-Xavier Standaert. Generalized Feistel Ciphers for Efficient Prime Field Masking. In Marc Joye and Gregor Leander, editors, *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*, volume 14653 of *Lecture Notes in Computer Science*, pages 188–220. Springer, 2024.

[GPRV21] Dahmun Goudarzi, Thomas Prest, Matthieu Rivain, and Damien Vergnaud. Probing Security through Input-Output Separation and Revisited Quasilinear Masking. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2021(3):599–640, 2021.

[Hoe63] Wassily Hoeffding. Probability Inequalities for Sums of Bounded Random Variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.

[ISW03] Yuval Ishai, Amit Sahai, and David Wagner. Private Circuits: Securing Hardware against Probing Attacks. In Dan Boneh, editor, *Advances in Cryptology - CRYPTO 2003*, pages 463–481, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg.

[JMB24] Vahid Jahandideh, Bart Mennink, and Lejla Batina. An Algebraic Approach for Evaluating Random Probing Security With Application to AES. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2024(4):657–689, 2024.

[JMB25] Vahid Jahandideh, Bart Mennink, and Lejla Batina. A Decomposition Approach for Evaluating Security of Masking. In Goichiro Hanaoka and Bo-Yin Yang, editors, *Advances in Cryptology - ASIACRYPT 2025 - 31st International Conference on the Theory and Application of Cryptology and Information Security, Melbourne, VIC, Australia, December 8-12, 2025, Proceedings, Part I*, volume 16245 of *Lecture Notes in Computer Science*, pages 580–612. Springer, 2025.

[KvD78] Teun Kloek and Herman K. van Dijk. Bayesian Estimates of Equation System Parameters: An Application of Integration by Monte Carlo. *Econometrica*, 46(1):1–19, 1978.

[MS23] Loïc Masure and François-Xavier Standaert. Prouff and Rivain's Formal Security Proof of Masking, Revisited - Tight Bounds in the Noisy Leakage Model. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology - CRYPTO 2023 - 43rd Annual International Cryptology Conference, CRYPTO 2023, Santa Barbara, CA, USA, August 20-24, 2023, Proceedings, Part III*, volume 14083 of *Lecture Notes in Computer Science*, pages 343–376. Springer, 2023.

[ORR+24] Maciej Obremski, João Ribeiro, Lawrence Roy, François-Xavier Standaert, and Daniele Venturi. Improved Reductions from Noisy to Bounded and Probing Leakages via Hockey-Stick Divergences. In Leonid Reyzin and Douglas Stebila, editors, *Advances in Cryptology – CRYPTO 2024*, pages 461–491, Cham, 2024. Springer Nature Switzerland.

[PGMP19] Thomas Prest, Dahmun Goudarzi, Ange Martinelli, and Alain Passelègue. Unifying Leakage Models on a Rényi Day. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2019, Proceedings, Part I*, volume 11692 of *Lecture Notes in Computer Science*, pages 683–712. Springer, 2019.

[PR13] Emmanuel Prouff and Matthieu Rivain. Masking against Side-Channel Attacks: A Formal Security Proof. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology - EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Athens, Greece, May 26-30, 2013. Proceedings*, volume 7881 of *Lecture Notes in Computer Science*, pages 142–159. Springer, 2013.

[RP10] Matthieu Rivain and Emmanuel Prouff. Provably Secure Higher-Order Masking of AES. In Stefan Mangard and François-Xavier Standaert, editors, *Cryptographic Hardware and Embedded Systems, CHES 2010*, pages 413–427, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.