Tight Estimate of the Local Leakage Resilience of the Additive Secret-sharing Scheme & its Consequences

Hemanta K. Maji
Department of Computer Science, Purdue University, USA

Hai H. Nguyen
Department of Computer Science, Purdue University, USA

Anat Paskin-Cherniavsky
Department of Computer Science, Ariel University, Israel

Tom Suad
Department of Computer Science, Ariel University, Israel

Mingyuan Wang
Department of EECS, University of California Berkeley, USA

Xiuyu Ye
Department of Computer Science, Purdue University, USA

Albert Yu
Department of Computer Science, Purdue University, USA

Abstract

Innovative side-channel attacks have repeatedly exposed the secrets of cryptosystems. Benhamouda, Degwekar, Ishai, and Rabin (CRYPTO–2018) introduced local leakage resilience of secret-sharing schemes to study some of these vulnerabilities. In this framework, the objective is to characterize the unintended information revelation about the secret by obtaining independent leakage from each secret share. This work accurately quantifies the vulnerability of the additive secret-sharing scheme to local leakage attacks and its consequences for other secret-sharing schemes.

Consider the additive secret-sharing scheme over a prime field among k parties, where the secret shares are stored in their natural binary representation, requiring λ bits – the security parameter. We prove that the reconstruction threshold $k = \omega(\log \lambda)$ is necessary to protect against local physical-bit probing attacks, improving the previous $\omega(\log \lambda / \log\log \lambda)$ lower bound. This result is a consequence of accurately determining the distinguishing advantage of the "parity-of-parity" physical-bit local leakage attack proposed by Maji, Nguyen, Paskin-Cherniavsky, Suad, and Wang (EUROCRYPT–2021). Our lower bound is optimal because the additive secret-sharing scheme is perfectly secure against any $(k-1)$-bit (global) leakage and (statistically) secure against (arbitrary) one-bit local leakage attacks when $k = \omega(\log \lambda)$.

Any physical-bit local leakage attack extends to (1) physical-bit local leakage attacks on the Shamir secret-sharing scheme with adversarially-chosen evaluation places, and (2) local leakage attacks on the Massey secret-sharing scheme corresponding to any linear code. In particular, for Shamir's secret-sharing scheme, the reconstruction threshold $k = \omega(\log \lambda)$ is necessary when the number of parties is $n = \mathcal{O}(\lambda \log \lambda)$. Our analysis of the "parity-of-parity" attack's distinguishing advantage establishes it as the best-known local leakage attack in these scenarios.

Our work employs Fourier-analytic techniques to analyze the "parity-of-parity" attack on the additive secret-sharing scheme. We accurately estimate an exponential sum that captures the vulnerability of this secret-sharing scheme to the parity-of-parity attack, a quantity that is also closely related to the "discrepancy" of the Irwin-Hall probability distribution.

2012 ACM Subject Classification Theory of computation → Cryptographic primitives; Security and privacy → Cryptanalysis and other attacks

Keywords and phrases leakage resilience, additive secret-sharing, Shamir's secret-sharing, physical-

3rd Conference on Information-Theoretic Cryptography (ITC 2022).
Editor: Dana Dachman-Soled; Article No. 16; pp. 16:1–16:19
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

bit probing leakage attacks, Fourier analysis

Digital Object Identifier 10.4230/LIPIcs.ITC.2022.16

Funding Hemanta K. Maji, Hai H. Nguyen, Mingyuan Wang, Xiuyu Ye, and Albert Yu are supported in part by an NSF CRII Award CNS–1566499, NSF SMALL Awards CNS–1618822 and CNS–2055605, the IARPA HECTOR project, MITRE Innovation Program Academic Cybersecurity Research Awards (2019–2020, 2020–2021), a Ross-Lynn Research Scholars Grant (2021–2022), a Purdue Research Foundation (PRF) Award (2017–2018), and The Center for Science of Information, an NSF Science and Technology Center, Cooperative Agreement CCF–0939370. Anat Paskin-Cherniavsky and Tom Suad are supported by the Ariel Cyber Innovation Center in conjunction with the Israel National Cyber directorate in the Prime Minister's Office. Mingyuan Wang is also supported in part by DARPA under Agreement No. HR00112020026, AFOSR Award FA9550-19-1-0200, NSF CNS Award 1936826, and research grants by the Sloan Foundation, and Visa Inc. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Government or DARPA.

1 Introduction

Innovative and sophisticated side-channel attacks, beginning with [13, 14], have repetitively exposed the secrets of cryptosystems. Over the last few decades, there have been extensive studies on the security and efficiency of cryptosystems against various models of potential attacks (refer to the excellent survey [12]).

Benhamouda, Degwekar, Ishai, and Rabin [2] recently introduced local leakage resilience of secret-sharing schemes to investigate some of these vulnerabilities (this primitive is also implicitly studied by Goyal and Kumar [6]). Leakage-resilient cryptography aims to provide provable security in the presence of known attacks and even unforeseen attacks. Secret-sharing schemes are crucial building blocks for nearly all threshold cryptography. In leakage-resilient secret-sharing, the objective is to characterize the unintended information revelation about the secret by obtaining independent leakage from each secret share. The secret-sharing scheme is *locally leakage-resilient* if the joint distribution of the leakage from every secret share is (statistically) independent of the secret.

Interestingly, the local leakage resilience of Shamir's secret-sharing scheme is closely related to the problem of repairing Reed-Solomon codes [8, 9, 21, 7, 3, 17]. To break the leakage-resilience of a secret-sharing scheme, the adversary does not need to reconstruct the whole secret; obtaining partial information to distinguish any two secrets is sufficient. For example, in a linear secret-sharing scheme over characteristic-two fields, a suitable one-bit leakage from each share determines the "least significant bit" of the secret. The adversary's objective is to leak as small and simple a leakage as possible to achieve as significant a distinguishing advantage as possible.

The *physical-bit leakage model* is a realistic (and analytically-tractable) leakage model where an adversary probes physical bits in the memory hardware [11, 10, 4]. In the context of local leakage resilience of secret-sharing schemes, parties store their secret shares in their natural *binary representation*. The adversary chooses a bounded number of positions to probe the memory hardware storing these secret shares. The adversary's objective is to use this leakage to obtain some partial information about the secret. If the adversary's view is statistically independent of the secret, the secret-sharing scheme is secure against the local leakage; an *indistinguishability-based definition* captures this intuition [2].

This work characterizes the vulnerability of the additive secret-sharing scheme to the "parity-of-parity" physical-bit local leakage attack proposed by [15]. Next, we explore the consequences of this result to the leakage resilience of other linear secret-sharing schemes (in particular, Shamir's secret-sharing scheme).

Summary of known attacks

Consider the *additive secret-sharing scheme* among k parties over a prime field. Benhamouda et al. [2] proposed a one-bit local leakage attack with a distinguishing advantage of $\geqslant 1/k^k$.[1] Recently, Maji et al. [15] proposed the "parity-of-parity" attack, where the secret shares are stored in their natural binary representation, and the attacker leaks the least significant bit from every secret share. Adams et al. [1] proved that the "parity-of-parity" attack has a *distinguishing advantage* $\geqslant (1/2^k \cdot k!) \approx (\mathrm{e}/2)^k/k^k$. Therefore, the threshold k must be $\omega(\log\lambda/\log\log\lambda)$ for the additive secret-sharing scheme to be secure, where λ is the security parameter. Since the physical-bit probing attack is a significantly weak leakage attack, their result poses a pressing threat to the secret-sharing scheme's security. Furthermore, a local leakage attack on the additive secret-sharing scheme extends to Shamir's secret-sharing schemes for adversarially-chosen evaluation places [15].

Using a probabilistic argument, Nielsen and Simkin [19] presented a leakage attack on Shamir's secret-sharing scheme. They showed the existence of a leakage function and a secret such that the leakage is consistent with the secret with a probability of at least 1/2. Their attack requires $m \geqslant \frac{k\log p}{n-k}$ bits of leakage from *each* secret share, where n is the number of parties and k is the reconstruction threshold. This result is not applicable when, for example, the number of parties $n = k$, the reconstruction threshold.

Summary of our results

This work presents a tight analysis of the parity-of-parity attack (Figure 1). We prove that this attack has a distinguishing advantage of $\geqslant \frac{1}{2}\cdot(2/\pi)^k$, which, in turn, implies that the threshold k must be $\omega(\log\lambda)$ for the additive secret-sharing scheme to be secure. Observe that our result *qualitatively improves* the lower bounds of [2] and [1] while relying only on *physical-bit* local leakage.

Our result shows that the simplistic parity-of-parity physical-bit probing attack is asymptotically optimal. The distinguishing advantage of any local leakage attack (possibly performing more sophisticated leakages) cannot be significantly higher because Benhamouda et al. [2] proved that the distinguishing advantage of *any* one-bit local leakage attack on the additive secret-sharing scheme is $\leqslant 2.47\cdot(2/\pi)^k$. Furthermore, due to the $(k-1)$ independence of the additive secret-sharing scheme, any (global) $(k-1)$ bits of leakage has *no advantage* in distinguishing any two secrets.

Maji et al. [15] and Adams et al. [1] showed that any physical-bit local leakage attack extends to a physical-bit leakage attack on Shamir's secret-sharing scheme with adversarially-chosen evaluation places. Previously, the best-known distinguishing advantage was $\geqslant 1/(2^k\cdot k!)$ [1], where k is the reconstruction threshold of Shamir's secret-sharing scheme. Our work improves this lower bound to $\geqslant \frac{1}{2}\cdot(2/\pi)^k$, which implies that $k = \omega(\log\lambda)$ is necessary for security against physical-bit local leakage attacks.

This attack also translates into a local leakage attack on the Massey secret-sharing scheme corresponding to any linear code (refer to Appendix C for a definition); for example, Shamir's

[1] This attack performs a computation on the entire secret share and leaks one bit of information from it. We emphasize that this attack is *not* a physical-bit attack.

secret-sharing scheme with arbitrary evaluation places. Before our work, to ensure local leakage-resilience, the lower bound on the reconstruction threshold of Shamir's secret-sharing scheme was (1) $k = \omega(\log\lambda/\log\log\lambda)$, if $n = \mathcal{O}(\lambda\log\lambda/\log\log\lambda)$ [1], and (2) $k \geqslant n/(\lambda+1)$, if $n = \omega(\lambda\log\lambda/\log\log\lambda)$ [19]. Our results improve the lower bound to $k = \omega(\log\lambda)$ when the number of parties $n = \mathcal{O}(\lambda\log\lambda)$.

Technically, we obtain our lower bound through a Fourier-analytic approach and an accurate estimation of an appropriate exponential sum. As a consequence of our result, we also improve the bound on the "discrepancy" of the Irwin-Hall probability distribution, a fundamental property of any real-valued probability distribution proposed in [15].

2 Our Contribution

We begin with some notation to facilitate an overview of our results.

Secret-sharing schemes and local leakage resilience

Fix a prime field F of order p. The elements of F are naturally represented as λ-bit binary strings corresponding to the elements $\{0, 1, \ldots, p-1\}$, where $2^{\lambda-1} \leqslant p < 2^{\lambda}$. Fix a linear secret-sharing scheme over F among n parties with a reconstruction threshold k. Note that the secret and the secret shares are all elements of F. The number of bits in the representation of the secret and the secret shares is the security parameter λ.

Our work considers a (static) adversary who obtains $m = 1$ physical-bit leakage from each secret share. A one-bit physical-bit leakage function $\tau = (\tau_1, \tau_2, \ldots, \tau_n)$ is a collection of functions $\tau_i : F \to \{0, 1\}$ such that, on input $x \in F$, function τ_i outputs the ℓ_i-th physical-bit of x for some $1 \leqslant \ell_i \leqslant \lambda$, for all $1 \leqslant i \leqslant n$. For instance, $\ell_i = 1$ refers to the least significant bit and $\ell_i = \lambda$ refers to the most significant bit. Let $\tau(s)$ be the joint distribution of the leakage function τ over the sample space $\{0, 1\}^n$ defined by the experiment: (1) sample random secret shares $(s_1, s_2, \ldots, s_n) \in F^n$ for the secret $s \in F$ and (2) output the leakage $(\tau_1(s_1), \tau_2(s_2), \ldots, \tau_n(s_n))$.

A secret-sharing scheme is *ε-locally leakage-resilient* against one physical-bit probing attacks, if, for any pair of secrets $s^{(0)}, s^{(1)} \in F$, the leakage distributions $\tau(s^{(0)})$ and $\tau(s^{(1)})$ have statistical distance $\leqslant \varepsilon$. As per convention, we want to ensure that the parameter ε decays faster than any inverse-polynomial in the security parameter λ, represented as $\varepsilon = \mathsf{negl}(\lambda)$.

Additive secret-sharing scheme

Consider the *additive secret-sharing scheme* with k parties over a finite field F (possibly of composite order). For a secret $s \in F$, this secret-sharing scheme chooses random secret shares $s_1, \ldots, s_k \in F$ such that $s_1 + \cdots + s_k = s$. We assume that if F is a prime field, parties store the secret shares $s_1, \ldots, s_k$ in their natural binary representation. However, if F is a composite order field of characteristic p, then the secret shares are stored as a vector of F_p elements, where every F_p element is represented in its natural binary representation.[2]

[2] The degree-a extension of the field F_p, i.e. the finite field F_{p^a}, is isomorphic to $F_p[X]/\pi(X)$, where $\pi(X)$ is a degree-a irreducible polynomial. Therefore, every element $s \in F_{p^a}$ has a natural $(s_1, \ldots, s_a) \in F_p^a$ representation, each element in turn has a λ-bit binary representation.

Parity-of-parity attack

Maji et al. [15] introduced the *parity-of-parity* local physical-bit leakage attack on the additive secret sharing scheme over fields of arbitrary characteristic. If F is a prime field (of an odd order), then the attacker leaks the least significant bit of each secret share, i.e., the leaked bit indicates whether the secret share $s_i \in \{0, 2, \dots, |F|-1\}$ or $s_i \in \{1, 3, \dots, |F|-2\}$. Finally, the attack predicts the parity of the secret using the parity of these leaked parities. If F is a degree-a extension of the prime field F_p, then every secret share $s_i \in F$ has equivalent representation $(s_{i,1}, \dots, s_{i,a}) \in F_p^a$. For some fixed index $j \in \{1, 2, \dots, a\}$, the attacker leaks the parity of the element $s_{i,j}$ from the i-th secret share. Over extension fields, this attack predicts the parity of s_j, where the secret $s = (s_1, \dots, s_a) \in F^a$.

For example, if F has characteristic 2, observe that the parity of the j-th coordinate of all the secret shares (as vectors in F_2) yields the j-th coordinate of the secret, which completely breaks the leakage-resilience of the additive secret-sharing scheme.

Adams et al. [1] proved that the advantage of this attacker is maximized when the secrets are $s^{(0)} = 0$ and $s^{(1)} = (p-1)/2$. Furthermore, they proved that the advantage of this attack is $\geqslant 1/(2^k \cdot k!)$.

Our results

Given ε and k, our objective is to identify whether there are two distinct secrets $s^{(0)}, s^{(1)} \in F$ such that the parity-of-parity attack has (at least) ε-advantage in distinguishing the secret shares that these secrets generate. Without loss of generality, assume that F is a prime field of order $\geqslant 2$, because the characteristic of the field determines the vulnerability of the additive secret-sharing scheme. We prove the following result.

▶ **Theorem 1.** *Consider the additive secret sharing scheme with k parties over the prime field F. There exist two secrets $s^{(0)}, s^{(1)} \in F$ such that the parity-of-parity attack has ε-advantage in distinguishing the secret shares of $s^{(0)}$ from the secret shares of $s^{(1)}$, where*

$$\varepsilon \geqslant \frac{1}{2} \cdot \left(\frac{2}{\pi}\right)^k .$$

▶ Remark 2. Our bound captures the intuition that, for a fixed k, with increasing p, the insecurity of the additive secret-sharing scheme reduces. As $p \to \infty$, the insecurity tends (from above) to the limit $\frac{1}{2}\left(\frac{2}{\pi}\right)^k$, a constant. Intuitively, the "most-secure additive secret-sharing scheme" corresponds to the case when the order of the finite field is an "infinitely large prime p." This phenomenon and an exponential lower bound in k were originally conjectured in [15] based on empirical evidence (refer to Figure 1). Recently, [1] made partial progress towards non-trivially lower-bounding the advantage of the parity-of-parity attack by proving $\varepsilon \geqslant 1/(2^k \cdot (k-1)!) - (3(k-1)^2 + 1)/p$.[3] However, this insecurity bound is increasing in p; thus, their work could not substantiate this conjecture. Our result substantiates the empirical evidence of [15] and positively resolves their conjecture.

Our lower bound is (asymptotically) *optimal* and also proves the optimality of the parity-of-parity attack in the following sense. Over prime fields, [2] proved that the additive secret-sharing scheme is $2.47 \cdot (2/\pi)^k$-secure against *any* local one-bit leakage attack (i.e., the

[3] This bound proves that the discrepancy of the Irwin-Hall distribution is non-zero and is an integer multiple of $1/2^k(k-1)!$. Next, it transfers this lower bound to the distinguishing advantage of the parity-of-parity attacker against the additive secret-sharing scheme over finite prime fields.

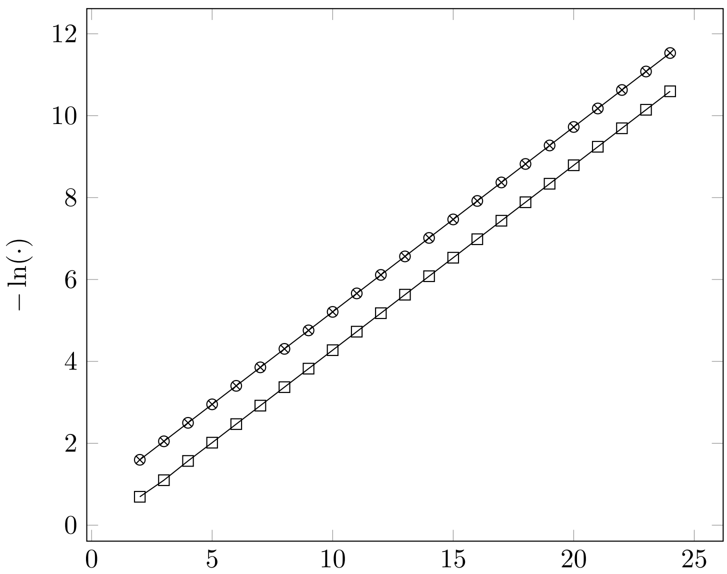

Figure 1 The horizontal axis represents the number of shares k in the additive secret-sharing scheme. The vertical axis represents the $-\ln(\cdot)$ of the distinguishing advantage of the parity-of-parity attack introduced by Maji et al. [15]. The squared points represent the empirically computed value for small k over a large enough field F as presented in [15] . The circled points represents the lower bound we prove in this work.

leakage function $\tau_i : F \to \{0, 1\}$ is arbitrary and *need not be a physical-bit probing* leakage). Consequently, the reconstruction threshold of the additive secret-sharing scheme must satisfy $k = \omega(\log \lambda)$ to be leakage-resilient to one physical-bit leakage from every secret share. Our result improves the previous best-known lower bound of $k = \omega(\log \lambda / \log \log \lambda)$ for additive secret-sharing schemes using the leakage attack presented in [2, 1].

To better bound the effectiveness of the parity-of-parity attack, Maji et al. [15] proposed the notion: *discrepancy of the Irwin-Hall distribution.* The first Irwin-Hall distribution IH_1 is the uniform distribution over $[0, 1)$. The i-th Irwin-Hall distribution IH_i is the convolution of the $(i-1)$-th Irwin-Hall distribution IH_{i-1} with the uniform distribution over $[0, 1)$. The discrepancy of the k-th Irwin-Hall distribution $\mathsf{disc}(k)$ is defined as

$$\mathsf{disc}(k) := \sup_y \left| \int_{-\infty}^{\infty} (-1)^{\lceil x-y \rceil} \cdot \mathsf{IH}_k(x)\, \mathrm{d}x \right|. \quad (1)$$

Appendix A provides a pictorial illustration of this notion. We refer the readers to [1] for more discussion on why this measure represents the effectiveness of the parity-of-parity attack. In particular, they proved that $\mathsf{disc}(k-1)$ is $\Theta\big(k^2/p\big)$-close to the effectiveness of the parity-of-parity attack on additive secret-sharing among k parties over prime field F of order p. Consequently, our result implies that the discrepancy of the Irwin-Hall distribution is also exponential in k, improving upon the previous best lower bound $1/(2^k \cdot k!)$ [1].

▶ **Corollary 3.** *For $k \in \{1, 2, \dots\}$, let $\mathsf{disc}(k)$ represents the discrepancy of the k-th Irwin-Hall distribution. Then, it holds that $\mathsf{disc}(k) = \Theta\left(\left(\frac{2}{\pi}\right)^k\right)$.*

Finally, motivated by applications in leakage-resilient secure computation, observe that our result extends to a stronger adversary who obtains some secret shares in the clear and performs local leakage attacks on the remaining secret shares.[4] We have the following theorem for such insider attackers.

▶ **Corollary 4.** *Consider the additive secret sharing scheme with k parties over the prime field F. Suppose a more general adversary obtains θ secret shares and gets the least significant bit from other shares. Then, there exist two secrets such that the adversary's advantage of distinguishing the two secrets is at least $\frac{1}{2} \cdot \left(\frac{2}{\pi}\right)^{k-\theta}$.*

Shamir's secret-sharing scheme. Let $\mathsf{ShamirSS}(n, k, \vec{X})$ represent Shamir's secret-sharing scheme among n parties, reconstruction threshold k, and evaluation places $\vec{X} = (X_1, \ldots, X_n)$. The evaluation places $X_1, \ldots, X_n$ are distinct elements of F^*. Let $s \in F$ be the secret. The secret-sharing scheme picks a random polynomial $f(Z) \in F[Z]/Z^k$ conditioned on the fact that $f(0) = s$. For $i \in \{1, \ldots, n\}$, the i-th secret share is $f(X_i)$.

Maji et al. [15] show a set of evaluation places such that one could perform the parity-of-parity attack on the first k secret shares to get the same advantage as the attack on the additive secret-sharing scheme. Hence, our result implies the following theorem.

▶ **Theorem 5.** *Let F be a prime field of order p such that $p = 1 \mod k$. Let $\alpha \in F^*$ be such that $\{\alpha, \alpha^2, \ldots, \alpha^k = 1\} \subseteq F^*$ is the set of k roots of the equation $Z^k - 1 = 0$. Suppose there exists $\beta \in F^*$ such that $\{\beta\alpha, \beta\alpha^2, \ldots, \beta\alpha^k = \beta\}$ is a subset of the evaluation places $\vec{X}$. One can perform the parity-of-parity attack on the secret shares corresponding to evaluation places $\{\beta\alpha, \beta\alpha^2, \ldots, \beta\alpha^k = \beta\}$ to get a distinguishing advantage of $\geqslant \frac{1}{2} \cdot (2/\pi)^k$. Therefore, if $\mathsf{ShamirSS}(n, k, \vec{X})$ is $\mathsf{negl}(\lambda)$-locally leakage-resilient secret-sharing scheme against one physical-bit leakage from each secret share, then it must be the case that $k = \omega(\log \lambda)$.*

Extension to arbitrary local leakage attacks. The following result extends the parity-of-parity attack to a local leakage attack to Massey secret-sharing scheme and Shamir's secret-sharing scheme. Given a linear code $C \subseteq F^{(n+1)}$, the Massey secret-sharing scheme [18] corresponding to a code C, is defined as follows. For a secret $s \in F$, one samples a random codeword $(s_0, s_1, \ldots, s_n) \in C$ such that $s_0 = s$. For $i \in \{1, 2, \ldots, n\}$, the i^{th} secret share is $s_i \in F$.

▶ **Theorem 6.** *Let F be a prime order field. For any Massey secret-sharing scheme corresponding to an $[n+1, k]_F$-linear code C or any $\mathsf{ShamirSS}(n, k, \vec{X})$ with arbitrary evaluation places $\vec{X}$ over F, there is a one-bit local leakage attack such that the distinguishing advantage is at least $\frac{1}{2} \cdot \left(\frac{2}{\pi}\right)^k$.*

To see why our results imply this theorem, assume the secret could be reconstructed from the first k shares as $s = \sum_{i=1}^{k} \alpha_i \cdot s_i$, where $\alpha_1, \ldots, \alpha_k$ are some fixed field elements (determined by the $[n+1, k]_F$ linear code). One can leak the least significant bit of $\alpha_i \cdot s_i$ from the i-th secret share s_i. It is easy to see that the advantage of this adversary is identical to the advantage of the parity-of-parity attack on the additive secret-sharing scheme.

However, we clarify that this leakage is *not* the physical-bit leakage because the local leakage involves field multiplication. As a consequence of Theorem 6, we obtain a similar lower bound for the reconstruction threshold against *arbitrary* local leakage.

▶ **Corollary 7.** *Fix $n, k \in \mathbb{N}$ and a prime order field F. If the Massey secret-sharing scheme corresponding to an $[n+1, k]_F$-linear code or $\mathsf{ShamirSS}(n, k, \vec{X})$ over F with arbitrary*

[4] For example, in secret-sharing based multi-party computation protocols [5], an adversary can corrupt some parties and get their entire secret shares in the clear. Additionally, the adversary may perform leakage attacks on the secret shares of the remaining honest parties.

evaluation places $\vec{X}$ is $\mathsf{negl}(\lambda)$*-locally leakage-resilient against one-bit local leakage, then it must hold that* $k = \omega(\log(\lambda))$*.*

We clarify that a physical-bit leakage analog for this result does not hold. [15] proved that with close-to-one-probability the $\mathsf{ShamirSS}(n, k, \vec{X})$ with *random evaluation places* $\vec{X}$ is $\mathsf{negl}(\lambda)$-locally leakage-resilient even for $k = 2$. Our result shows that the lower bound on the reconstruction threshold of Shamir's secret-sharing scheme is $k = \omega(\log \lambda)$ when the number of parties is $n = \mathcal{O}(\lambda \log \lambda)$. Before our work, the lower bound was (1) $k = \omega(\log \lambda / \log \log \lambda)$, if $n = \mathcal{O}(\lambda \log \lambda / \log \log \lambda)$ [1], and (2) $k \geqslant n/(\lambda + 1)$, if $n = \omega(\lambda \log \lambda / \log \log \lambda)$ [19].

▶ Remark 8. Our analysis also extends to the thermal noise leakage model in which the adversary obtains a noisy version of the leakage bits as considered in [1]. In this model, instead of obtaining the leakage $\tau(s) = (\tau_1(s_1), \tau_2(s_2), \ldots, \tau_n(s_n))$, the adversary receives a noisy leakage $\tau'(s) = (\tau'_1(s_1), \tau'_2(s_2), \ldots, \tau'_n(s_n))$, where every $\tau'_i(s_i)$ is ρ_i-correlated with $\tau_i(s_i)$.[5] The distinguishing advantage is reduced by a (multiplicative) factor of $\rho = \rho_1 \rho_2 \cdots \rho_n \leqslant 1$. For instance, the distinguishing advantage of the parity-of-parity attack in the presence of $(\rho_1, \ldots, \rho_n)$ noise would be

$$\rho_1 \cdot \rho_2 \cdots \rho_n \cdot \frac{1}{2} \cdot \left(\frac{2}{\pi}\right)^n.$$

This observation follows from facts of convolution.

3 Technical Overview

This section presents an overview of our technical approach. Let F be a prime field of order p. Consider the additive secret-sharing scheme over F. Let τ be the leakage attack that leaks the least significant bit from every share.

We refer the readers to Section 4.1 for an introduction to Fourier analysis. By the Fourier-analytic approach from prior works [2, 15, 16], for any two secrets $s^{(0)}$ and $s^{(1)}$, we have

$$\mathsf{SD}\left(\tau\left(s^{(0)}\right), \tau\left(s^{(1)}\right)\right) = \frac{1}{2} \cdot \sum_{\ell \in \{0,1\}^n} \left| \sum_{\alpha \in F^*} \left(\prod_{i=1}^{n} \widehat{\mathbb{1}_{\ell_i}}(\alpha)\right) \left(\omega^{\alpha \cdot s^{(0)}} - \omega^{\alpha \cdot s^{(1)}}\right) \right|,$$

where $\omega = \exp(2\pi\imath/p)$ is the p^{th} root of unity. Furthermore, $\mathbb{1}_0$ is the indicator function for the set $S_0 := \{0, 2, \ldots, p-1\}$ and, similarly, $\mathbb{1}_1$ is the indicator function for the set $S_1 := \{1, 3, \ldots, p-2\}$. That is, S_b is the set of field elements whose least significant bit is b.

Note that the above expression is an *identity*. Our first observation is that, for any $\ell \in \{0,1\}^n$, the *magnitude* of the expression

$$U(\alpha) := \prod_{i=1}^{n} \widehat{\mathbb{1}_{\ell_i}}(\alpha)$$

is exponentially decaying as α goes from the central points $\frac{p-1}{2}$ and $\frac{p+1}{2}$ to the end points 1 and $p-1$ (refer to Figure 2). Informally, it holds that

$$|U(\alpha)| \approx \left(\frac{2}{\pi}\right)^n \cdot \left(\frac{1}{|2\alpha - p|}\right)^n.$$

[5] For any $\rho \in [0,1]$, a bit b is ρ-correlated with another bit b' if $b = b'$ with probability ρ, and b is an independent and uniformly random bit with probability $1 - \rho$.

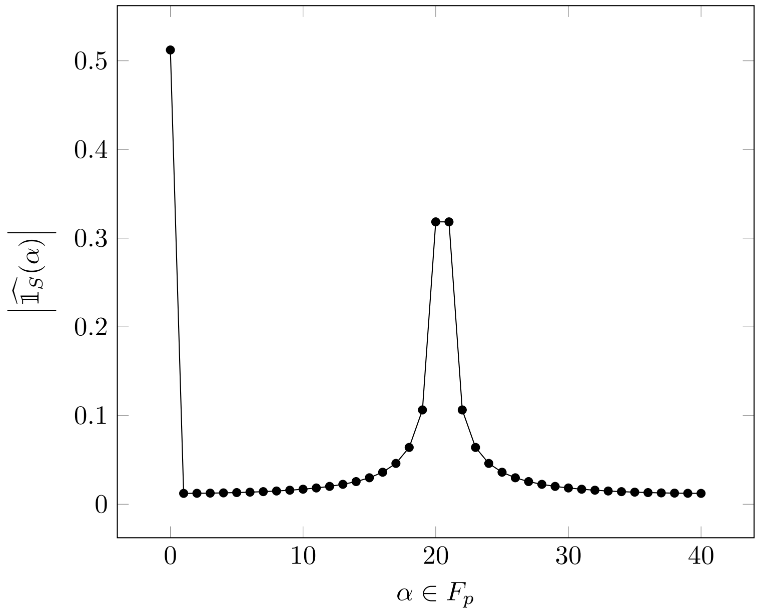

Figure 2 For the representative case of $p = 41$, the Fourier spectrum of the indicator function $\mathbb{1}_S$, where $S = \{0, 2, \ldots, 40\} \subset F_p$ is the subset of all "even elements."

For the magnitude of the other term

$$V(\alpha) := \left(\omega^{\alpha \cdot s^{(0)}} - \omega^{\alpha \cdot s^{(1)}}\right),$$

we use the naive triangle inequality to upper bound it by 2 for the *non-central terms* (i.e., $\alpha \neq \frac{p-1}{2}, \frac{p+1}{2}$). And we argue that there exists two secrets $s^{(0)}$ and $s^{(1)}$ such that $V\left(\frac{p-1}{2}\right)$ and $V\left(\frac{p+1}{2}\right)$ are large (e.g., $\geqslant 3/2$).

Together, these observations enable us to lower bound the statistical distance by (approximately) the magnitude of the dominant term $U\left(\frac{p-1}{2}\right)$ and $U\left(\frac{p+1}{2}\right)$, which are $\Theta\left(\left(\frac{2}{\pi}\right)^n\right)$.

Finally, observe that the two distributions $\tau\left(s^{(0)}\right)$ and $\tau\left(s^{(1)}\right)$ are $(n-1)$-indistinguishable. That is, these two distributions restricted to any proper subset of their coordinates are identical. Therefore, by standard techniques, parity is the optimal distinguisher for these two distributions (we provide a formal discussion on this in Appendix B). Consequently, the parity-of-parity attack [15] has an distinguishing advantage of $\Theta\left(\left(\frac{2}{\pi}\right)^n\right)$.

▶ Remark 9. Due to the form of our lower bound expression, it is tempting to naïvely argue that the advantage of the parity-of-parity attack correctly predicting the secret's parity is (some form of a) "k-fold convolution of a $(2/\pi)$-biased predictor." This intuition is (seriously) technically flawed. The least significant bit of the first $(k-1)$ secret shares are each $1/p$-biased and independent of the secret.

4 Analysis of the Parity-of-parity Attack on Additive Secret-sharing Schemes

Maji et al. [15] proposed the following parity-of-parity attack. Suppose the field elements are stored in their natural binary representation. The adversary leaks the least significant bit (LSB) as the local leakage of every secret share. Finally, the adversary outputs the parity of the LSB from every secret share as the prediction of the secret. Adams et al. [1] proved that the distinguishing advantage of this adversary is at least $\Omega(1/n!)$. In this section, we shall present a tight analysis of this attack. In particular, we shall show that the distinguishing advantage is $\exp(-\mathcal{O}(n))$.

This lower bound we prove is tight up to a small constant, as Benhamouda et al. [2] prove that the distinguishing advantage of the adversary is upper-bounded by $\left(\frac{2}{\pi}\right)^{n-2}$. Note that the upper bound of [2] holds for any local leakage attack on the additive secret-sharing scheme. Therefore, our result also demonstrates that *the "parity-of-parity" attack is the optimal attack.*

Formally, let $\mathsf{AddSS}(s)$ represent the distribution of the additive secret shares of the secret s. That is, $\mathsf{AddSS}(s) = (s_1, \ldots, s_n)$ is sampled uniformly at random conditioned on that $s_1 + s_2 + \cdots + s_n = s$. For any $x \in F$, let $\mathsf{lsb}(x)$ represent the least significant bit of x.[6]

Let τ represent the local leakage function that leaks the LSB of every secret share. That is, $\tau(\mathsf{AddSS}(s)) := \Big(\mathsf{lsb}(s_1), \mathsf{lsb}(s_2), \ldots, \mathsf{lsb}(s_n)\Big)$. We prove the following theorem.

▶ **Theorem 10.** *There exists two secrets $s^{(0)}$ and $s^{(1)}$ such that*

$$\mathsf{SD}\left(\tau(\mathsf{AddSS}(s^{(0)}))\ ,\ \tau(\mathsf{AddSS}(s^{(1)}))\right) \geqslant \frac{1}{2} \cdot \left(\frac{2}{\pi}\right)^n .$$

In particular, to ensure that the adversary has a negligible distinguishing advantage $\mathsf{negl}(\lambda)$*, it must hold that $n = \omega(\log \lambda)$.*

▶ Remark 11 (On the characteristics of the field). We emphasize that our lower bound holds for arbitrarily large characteristics. Intuitively, as the characteristic of the field increases, one expects the advantage of the adversary to decrease. However, our result shows that the advantage of the adversary is guaranteed to be higher than $\frac{1}{2} \cdot \left(\frac{2}{\pi}\right)^n$even when the characteristic of the field tends to infinity.

Finally, observe that $\tau(\mathsf{AddSS}(s^{(0)}))$ and $\tau(\mathsf{AddSS}(s^{(1)}))$ are $(n-1)$-indistinguishable distributions since the additive secret sharing is $(n-1)$-private. By standard techniques in Fourier analysis, the parity of all the bits is the best distinguisher (up to a small constant) for any two $(n-1)$-indistinguishable distributions. For completeness, we provide formal proof of this in Appendix B. This observation, together with the theorem, implies the optimality of the parity-of-parity attack.

Surprisingly, our proof of Theorem 10 is based on Fourier analysis. Typically, Fourier analytic approach is employed to *upper bound* the distinguishing advantage of the adversary. However, we shall use it to prove a *lower bound* result.

We start by introducing some notations and basics of Fourier analysis that suffice for our purposes. Next, we present the proof of Theorem 10.

4.1 Preliminaries on Fourier Analysis

Let F be a prime field of order p. For any complex number $x \in \mathbb{C}$, let $\overline{x}$ represent its conjugate. For any two functions $f, g \colon F \to \mathbb{C}$, their *inner product* is

$$\langle f, g \rangle := \frac{1}{p} \cdot \sum_{x \in F} f(x) \cdot \overline{g(x)}.$$

[6] This section restricts our discussion to a field F of prime order. If E is an degree t extension field of the field F, then every element α of E can be seen as a polynomial $a_{t-1}X^{t-1} + \cdots + a_1 X + a_0$ in $F[X]$. We shall call a_0 the least significant symbol of α. Observe that, for an additive secret sharing of the secret s over E, the least significant symbol of every secret share forms an additive secret sharing of the least significant symbol of s over F. Therefore, the result for prime order fields naturally extends to composite order fields when the attacker leaks the LSB of the least significant symbol of every share.

Let $\omega = \exp(2\pi\imath/p)$ be the p^{th} root of unity. For all $\alpha \in F$, the function χ_α is defined to be

$$\chi_\alpha(x) := \omega^{\alpha \cdot x},$$

and the respective Fourier coefficient $\widehat{f}(\alpha)$ is defined as

$$\widehat{f}(\alpha) := \langle f, \chi_\alpha \rangle .$$

Our proof relies on the following lemma. We refer the readers to [2] for a proof.

▶ **Lemma 12** (Poisson Summation Formula). *Let $C \subseteq F^n$ be a linear code with dual code $C^\perp$. For all $i \in \{1, 2, \ldots, n\}$, let $f_i : F \to \mathbb{C}$ be an arbitrary function. It holds that*

$$\mathop{\mathrm{E}}_{\vec{x} \leftarrow C}\left[\prod_{i=1}^{n} f_i(x_i)\right] = \sum_{\vec{y} \in C^\perp}\left(\prod_{i=1}^{n} \widehat{f_i}(y_i)\right).$$

The following claims will also be useful, which follows directly from the definition.

▶ **Claim 1.** *Let $S, T \subseteq F$ be a partition of F. For all $\alpha \in F$,*

$$\widehat{\mathbb{1}_S}(\alpha) = -\widehat{\mathbb{1}_T}(\alpha).$$

▶ **Claim 2.** *For all $S \subseteq F$ and $x \in F$, it holds that*

$$\widehat{\mathbb{1}_{x+S}}(\alpha) = \widehat{\mathbb{1}_S}(\alpha) \cdot \omega^{-\alpha \cdot x}.$$

The *statistical distance* (a.k.a, total variation distance) between two distributions P and Q over a finite sample space Ω is defined as $\mathsf{SD}(P, Q) = \frac{1}{2}\sum_{x \in \Omega}|P(x) - Q(x)|$. For any code $C \subseteq F^n$ and any vector $x \in F^n$, we define $x + C := \{x + c : c \in C\}$.

4.2 Proof of Theorem 10

We start by introducing some notations and facts. Define a bipartition of F as

$$S_0 := \{0, 2, \ldots, p-1\} \qquad \text{and} \qquad S_1 := \{1, 3, \ldots, p-2\}.$$

That is, S_b is the set of field elements on which the LSB function will output b.

▶ **Claim 3.** *For $\alpha \in F^*$, it holds that*

$$\widehat{\mathbb{1}_{S_0}}(\alpha) = \frac{1}{2p} \cdot \frac{1}{\cos(\pi\alpha/p)} \cdot \omega^{\alpha/2}, \text{ and } \widehat{\mathbb{1}_{S_1}}(\alpha) = -\frac{1}{2p} \cdot \frac{1}{\cos(\pi\alpha/p)} \cdot \omega^{\alpha/2}.$$

Furthermore,

$$\left|\widehat{\mathbb{1}_{S_0}}(\alpha)\right| = \left|\widehat{\mathbb{1}_{S_1}}(\alpha)\right| = \frac{1}{2p} \cdot \frac{1}{|\cos(\pi\alpha/p)|}.$$

Proof of Claim 3 . By definition, we have

$$\begin{aligned}\widehat{\mathbb{1}_{S_0}}(\alpha) = \langle \mathbb{1}_{S_0}, \chi_\alpha \rangle = \frac{1}{p}\sum_{x \in S_0} \omega^{-\alpha \cdot x} = \frac{1}{p} \cdot \sum_{j=0}^{(p-1)/2} \omega^{-\alpha \cdot (2j)} \\ = \frac{1}{p} \cdot \frac{1 - \omega^{-(2\alpha)\cdot(p+1)/2}}{1 - \omega^{-2\alpha}} = \frac{1}{p} \cdot \frac{1 - \omega^{-\alpha}}{1 - \omega^{-2\alpha}}.\end{aligned}$$

One could verify that $1 - \omega^{-\alpha} = 2\sin(\pi\alpha/p)\cdot\omega^{\frac{p}{4}-\frac{\alpha}{2}}$. Hence,

$$\widehat{\mathbb{1}_{S_0}}(\alpha) = \frac{1}{p}\cdot\frac{2\sin(\pi\alpha/p)\cdot\omega^{\frac{p}{4}-\frac{\alpha}{2}}}{2\sin(\pi(2\alpha)/p)\cdot\omega^{\frac{p}{4}-\frac{2\alpha}{2}}} = \frac{1}{2p}\cdot\frac{1}{\cos(\pi\alpha/p)}\cdot\omega^{\alpha/2}.$$

By Claim 1, we have

$$\widehat{\mathbb{1}_{S_1}}(\alpha) = -\frac{1}{2p}\cdot\frac{1}{\cos(\pi\alpha/p)}\cdot\omega^{\alpha/2}.$$

Finally, since $\left|w^{\alpha/2}\right| = 1$, it is easy to see that

$$\left|\widehat{\mathbb{1}_{S_0}}(\alpha)\right| = \left|\widehat{\mathbb{1}_{S_1}}(\alpha)\right| = \frac{1}{2p}\cdot\frac{1}{|\cos(\pi\alpha/p)|},$$

which completes the proof. ◀

Let C be the parity code. That is, $(c_1,\dots,c_n)\in C$ if $c_1+\cdots+c_n = 0$. The secret shares of a secret s is uniformly distributed over the set $(s,0,\dots,0)+C$; or equivalently, it is uniformly distributed over $(n^{-1}\cdot s,\dots,n^{-1}\cdot s)+C$. For ease of presentation, we use the latter form. Additionally, the dual code of C, denoted by $C^\perp$, is simply the repetition code, i.e., $C^\perp = \{(\alpha,\dots,\alpha)\colon \alpha\in F\}$.

We are ready to prove Theorem 10 as follows. We shall abuse notation and write $\mathbb{1}_b$ for $\mathbb{1}_{S_b}$. Observe that

$$\begin{aligned}
&\mathsf{SD}\left(\tau\left(\mathsf{AddSS}(s^{(0)})\right),\tau\left(\mathsf{AddSS}(s^{(1)})\right)\right)\\
=&\frac{1}{2}\cdot\sum_{\ell\in\{0,1\}^n}\left|\mathop{\mathrm{E}}_{\vec{x}\leftarrow C}\left[\prod_{i=1}^n \mathbb{1}_{\ell_i}(x_i+n^{-1}\cdot s^{(0)})\right] - \mathop{\mathrm{E}}_{\vec{x}\leftarrow C}\left[\prod_{i=1}^n \mathbb{1}_{\ell_i}(x_i+n^{-1}\cdot s^{(1)})\right]\right| && \text{(By definition of SD)}\\
=&\frac{1}{2}\cdot\sum_{\ell\in\{0,1\}^n}\left|\sum_{\vec{y}\in C^\perp}\left(\prod_{i=1}^n \widehat{\mathbb{1}_{\ell_i}}(y_i+n^{-1}\cdot s^{(0)})\right) - \sum_{\vec{y}\in C^\perp}\left(\prod_{i=1}^n \widehat{\mathbb{1}_{\ell_i}}(y_i+n^{-1}\cdot s^{(1)})\right)\right| && \text{(Lemma 12)}\\
=&\frac{1}{2}\cdot\sum_{\ell\in\{0,1\}^n}\left|\sum_{\alpha\in F}\left(\prod_{i=1}^n \widehat{\mathbb{1}_{\ell_i}}(\alpha+n^{-1}\cdot s^{(0)})\right) - \sum_{\alpha\in F}\left(\prod_{i=1}^n \widehat{\mathbb{1}_{\ell_i}}(\alpha+n^{-1}\cdot s^{(1)})\right)\right| && \text{(By the definition of } C^\perp)\\
=&\frac{1}{2}\cdot\sum_{\ell\in\{0,1\}^n}\left|\sum_{\alpha\in F}\left(\prod_{i=1}^n \widehat{\mathbb{1}_{\ell_i}}(\alpha)\right)\left(\omega^{\alpha\cdot s^{(0)}}-\omega^{\alpha\cdot s^{(1)}}\right)\right| && \text{(Claim 2)}\\
=&\frac{1}{2}\cdot\sum_{\ell\in\{0,1\}^n}\left|\sum_{\alpha\in F^*}\left(\prod_{i=1}^n\left((-1)^{\ell_i}\frac{1}{2p}\cdot\frac{1}{\cos(\pi\alpha/p)}\cdot\omega^{\alpha/2}\right)\right)\left(\omega^{\alpha\cdot s^{(0)}}-\omega^{\alpha\cdot s^{(1)}}\right)\right| && \text{(Claim 3)}\\
=&2^{n-1}\cdot\left|\sum_{\alpha\in F^*}\left(\frac{1}{2p}\cdot\frac{1}{\cos(\pi\alpha/p)}\cdot\omega^{\alpha/2}\right)^n\left(\omega^{\alpha\cdot s^{(0)}}-\omega^{\alpha\cdot s^{(1)}}\right)\right| && \text{(Identity transformation)}
\end{aligned}$$

Note that the proof so far has not used any inequalities. The expression above is identical to the statistical distance. For brevity, let us define

$$U(\alpha) := \left(\frac{1}{2p}\cdot\frac{1}{\cos(\pi\alpha/p)}\cdot\omega^{\alpha/2}\right)^n,\ \text{and}\ \ V(\alpha) := \omega^{\alpha\cdot s^{(0)}} - \omega^{\alpha\cdot s^{(1)}}.$$

Additionally, let $W(\alpha) := U(\alpha) \cdot V(\alpha)$. Intuitively, we shall prove that the magnitude of $\sum_{\alpha \in F^*} W(\alpha)$ is approximately the magnitude of its leading term $W((p-1)/2)$ and $W((p+1)/2)$. In particular, we prove the following claims.

▶ **Claim 4.** *There exists a universal constant $\mu \geqslant 3/2$ and two secrets $s^{(0)}, s^{(1)} \in F$ such that*

$$\left| W\left(\frac{p-1}{2}\right) + W\left(\frac{p+1}{2}\right) \right| \geqslant \mu \cdot \pi^{-n}.$$

▶ **Claim 5.** *For all secrets $s^{(0)}, s^{(1)}$, we have*

$$\left| \sum_{\alpha \in F^* \setminus \{\frac{p-1}{2}, \frac{p+1}{2}\}} W(\alpha) \right| \leqslant \exp(-\Theta(n)) \cdot \pi^{-n}.$$

Using Claim 4 and Claim 5, the proof of the Theorem 10 follows from the fact that

$$\begin{aligned}
\mathsf{SD}\left(\tau\left(\mathsf{AddSS}(s^{(0)})\right), \tau\left(\mathsf{AddSS}(s^{(1)})\right)\right) &\geqslant 2^{n-1} \cdot (\mu - \exp(-\Theta(n))) \cdot \pi^{-n}, \\
&\geqslant 2^{n-1} \cdot \left(\frac{3}{2} - \mathrm{o}(1)\right) \cdot \pi^{-n}, \\
&\geqslant \frac{1}{2} \cdot 1 \cdot \left(\frac{2}{\pi}\right)^n \qquad \text{(for large enough } n.)
\end{aligned}$$

Consequently, it suffices to prove Claim 4 and Claim 5 to complete the proof of Theorem 10.

Proof of Claim 4 . Observe that

$$\begin{aligned}
W\left(\frac{p-1}{2}\right) &= \left(\frac{1}{2p} \cdot \frac{1}{\cos\left(\pi \cdot \frac{p-1}{2p}\right)} \cdot \omega^{\frac{p-1}{4}}\right)^n \cdot V\left(\frac{p-1}{2}\right) \\
&= \left(\frac{1}{2p} \cdot \frac{1}{\sin\left(\pi \cdot \frac{1}{2p}\right)}\right)^n \cdot \omega^{n \cdot \frac{p-1}{4}} \cdot V\left(\frac{p-1}{2}\right)
\end{aligned}$$

and

$$\begin{aligned}
W\left(\frac{p+1}{2}\right) &= \left(\frac{1}{2p} \cdot \frac{1}{\cos\left(\pi \cdot \frac{p+1}{2p}\right)} \cdot \omega^{\frac{p+1}{4}}\right)^n \cdot V\left(\frac{p+1}{2}\right) \\
&= \left(\frac{1}{2p} \cdot \frac{1}{\sin\left(\pi \cdot \frac{1}{2p}\right)}\right)^n \cdot (-1)^n \cdot \omega^{n \cdot \frac{p+1}{4}} \cdot V\left(\frac{p+1}{2}\right)
\end{aligned}$$

Therefore,

$$\begin{aligned}
&\left| W\left(\frac{p-1}{2}\right) + W\left(\frac{p+1}{2}\right) \right| \\
=& \left| \left(\frac{1}{2p} \cdot \frac{1}{\sin\left(\pi \cdot \frac{1}{2p}\right)}\right)^n \right| \cdot \left| \omega^{n \cdot \frac{p-1}{4}} \cdot V\left(\frac{p-1}{2}\right) + (-1)^n \cdot \omega^{n \cdot \frac{p+1}{4}} \cdot V\left(\frac{p+1}{2}\right) \right| \\
=& \left| \left(\frac{1}{2p} \cdot \frac{1}{\sin\left(\pi \cdot \frac{1}{2p}\right)}\right)^n \right| \cdot \left| V\left(\frac{p-1}{2}\right) + (-1)^n \cdot \omega^{\frac{n}{2}} \cdot V\left(\frac{p+1}{2}\right) \right|
\end{aligned}$$

Note that $x \cdot \sin(1/x)$ is strictly increasing as x increases and tends to 1 as $x \to \infty$.[7] Therefore,

$$\left| W\left(\frac{p-1}{2}\right) + W\left(\frac{p+1}{2}\right) \right| \geqslant \pi^{-n} \cdot \left| V\left(\frac{p-1}{2}\right) + (-1)^n \cdot \omega^{\frac{n}{2}} \cdot V\left(\frac{p+1}{2}\right) \right|.$$

It remains to prove that there exist secrets $s^{(0)}$ and $s^{(1)}$ such that $V\left(\frac{p-1}{2}\right)$ and $(-1)^n \cdot \omega^{\frac{n}{2}} \cdot V\left(\frac{p+1}{2}\right)$ does not cancel each other to be too small. More formally, for any p and n, we shall show that there exist a universal constant μ and secrets $s^{(0)}$ and $s^{(1)}$ such that

$$\left| \left(\omega^{\frac{p-1}{2} \cdot s^{(0)}} - \omega^{\frac{p-1}{2} \cdot s^{(1)}}\right) + (-1)^n \cdot \omega^{\frac{n}{2}} \cdot \left(\omega^{\frac{p+1}{2} \cdot s^{(0)}} - \omega^{\frac{p+1}{2} \cdot s^{(1)}}\right) \right| \geqslant \mu.$$

Let $f(s^{(0)})$ (resp., $g(s^{(1)})$) denote the terms involving $s^{(0)}$ (resp., $s^{(1)}$) in the above expression. And we are interested in $\left|f(s^{(0)}) + g(s^{(1)})\right|$. Observe that

$$\sum_{s^{(1)} \in F} g(s^{(1)}) = 0.$$

Therefore, we have

$$\begin{aligned} \max_{s^{(1)}} \left|f(s^{(0)}) + g(s^{(1)})\right| &\geqslant \frac{1}{p} \sum_{s^{(1)} \in F} \left|f(s^{(0)}) + g(s^{(1)})\right| \\ &\geqslant \frac{1}{p} \left| \sum_{s^{(1)} \in F} \left(f(s^{(0)}) + g(s^{(1)})\right) \right| = \left|f(s^{(0)})\right|. \end{aligned}$$

Hence, it suffices to show that there exists an $s^{(0)}$ such that $\left|f(s^{(0)})\right|$ is sufficiently large. That is,

$$\max_{s^{(0)}} \left| \omega^{\frac{p-1}{2} \cdot s^{(0)}} + (-1)^n \cdot \omega^{\frac{n}{2}} \cdot \omega^{\frac{p+1}{2} \cdot s^{(0)}} \right| \geqslant \mu,$$

which is equivalent to

$$\max_{s^{(0)}} \left| 1 + (-1)^n \cdot \omega^{\frac{n}{2}} \cdot \omega^{s^{(0)}} \right| \geqslant \mu.$$

It is easy to see that the phase of $\omega^{s^{(0)}}$ could be an arbitrary multiple of $2\pi/p$. Hence, there must exist an $s^{(0)}$ such that the above expression has magnitude $\geqslant 3/2$.[8] This completes the proof. ◀

[7] Intuitively, the advantage of the adversary decreases as the characteristic of the field increases.
[8] In fact, as p tends to infinity, the maximum gets arbitrarily close to 2.

Proof of Claim 5 . By a simple triangle inequality, we have $|V(\alpha)| \leqslant 2$. Hence,

$$
\begin{aligned}
&\left|\sum_{\alpha \in F^* \setminus \{\frac{p-1}{2}, \frac{p+1}{2}\}} W(\alpha)\right| \\
&\leqslant \sum_{\alpha \in F^* \setminus \{\frac{p-1}{2}, \frac{p+1}{2}\}} |W(\alpha)| && \text{(Triangle inequality)} \\
&\leqslant 2 \cdot \sum_{\alpha \in F^* \setminus \{\frac{p-1}{2}, \frac{p+1}{2}\}} |U(\alpha)| && \text{(Triangle inequality)} \\
&= 2 \cdot \sum_{\alpha \in F^* \setminus \{\frac{p-1}{2}, \frac{p+1}{2}\}} \left|\frac{1}{2p} \cdot \frac{1}{\cos(\pi\alpha/p)}\right|^n && \text{(Identity transformation)} \\
&= 4 \cdot \sum_{j=1}^{(p-3)/2} \left(\frac{1}{2p} \cdot \frac{1}{\cos(\pi j/p)}\right)^n && \text{(Identity transformation)} \\
&= 4 \cdot \sum_{j=1}^{(p-3)/2} \left(\frac{1}{2p} \cdot \frac{1}{\sin(\pi(p-2j)/(2p))}\right)^n && \text{(Identity transformation)}
\end{aligned}
$$

Observe that $\sin(x) \geqslant x/2$ for every $x \in (0, \pi/2)$. Hence,

$$
\begin{aligned}
&\left|\sum_{\alpha \in F^* \setminus \{\frac{p-1}{2}, \frac{p+1}{2}\}} W(\alpha)\right| \\
&\leqslant 4 \cdot \sum_{j=1}^{(p-3)/2} \left(\frac{1}{2p} \cdot \frac{2}{\pi(p-2j)/(2p)}\right)^n \\
&= \pi^{-n} \cdot 4 \cdot \sum_{j=1}^{(p-3)/2} \left(\frac{2}{p-2j}\right)^n \\
&\leqslant \pi^{-n} \cdot 4 \cdot \left(\left(\frac{2}{3}\right)^n + \int_3^\infty \left(\frac{1}{x}\right)^n \mathrm{d}x\right) \\
&= \pi^{-n} \cdot 4 \cdot \left(\left(\frac{2}{3}\right)^n + \frac{1}{n+1}\left(\frac{1}{3}\right)^{n+1}\right) \\
&= \pi^{-n} \cdot \exp(-\Theta(n)).
\end{aligned}
$$

This completes the proof. ◀

References

1. Donald Q. Adams, Hemanta K. Maji, Hai H. Nguyen, Minh L. Nguyen, Anat Paskin-Cherniavsky, Tom Suad, and Mingyuan Wang. Lower bounds for leakage-resilient secret sharing schemes against probing attacks. In *IEEE International Symposium on Information Theory ISIT 2021*, 2021. URL: `https://www.cs.purdue.edu/homes/hmaji/papers/AMNNPSW21.pdf`.
2. Fabrice Benhamouda, Akshay Degwekar, Yuval Ishai, and Tal Rabin. On the local leakage resilience of linear secret sharing schemes. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018, Part I*, volume 10991 of *Lecture Notes in Computer Science*, pages 531–561, Santa Barbara, CA, USA, August 19–23, 2018. Springer, Heidelberg, Germany. `doi:10.1007/978-3-319-96884-1_18`.
3. Hoang Dau, Iwan M. Duursma, Han Mao Kiah, and Olgica Milenkovic. Repairing reed-solomon codes with multiple erasures. *IEEE Trans. Inf. Theory*, 64(10):6567–6582, 2018. `doi:10.1109/TIT.2018.2827942`.
4. Alexandre Duc, Stefan Dziembowski, and Sebastian Faust. Unifying leakage models: From probing attacks to noisy leakage. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology – EUROCRYPT 2014*, volume 8441 of *Lecture Notes in Computer Science*, pages 423–440, Copenhagen, Denmark, May 11–15, 2014. Springer, Heidelberg, Germany. `doi:10.1007/978-3-642-55220-5_24`.
5. Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In Alfred Aho, editor, *19th Annual ACM Symposium on Theory of Computing*, pages 218–229, New York City, NY, USA, May 25–27, 1987. ACM Press. `doi:10.1145/28395.28420`.
6. Vipul Goyal and Ashutosh Kumar. Non-malleable secret sharing. In Ilias Diakonikolas, David Kempe, and Monika Henzinger, editors, *50th Annual ACM Symposium on Theory of Computing*, pages 685–698, Los Angeles, CA, USA, June 25–29, 2018. ACM Press. `doi:10.1145/3188745.3188872`.
7. Venkatesan Guruswami and Ankit Singh Rawat. MDS code constructions with small sub-packetization and near-optimal repair bandwidth. In Philip N. Klein, editor, *28th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2109–2122, Barcelona, Spain, January 16–19, 2017. ACM-SIAM. `doi:10.1137/1.9781611974782.137`.
8. Venkatesan Guruswami and Mary Wootters. Repairing reed-solomon codes. In Daniel Wichs and Yishay Mansour, editors, *48th Annual ACM Symposium on Theory of Computing*, pages 216–226, Cambridge, MA, USA, June 18–21, 2016. ACM Press. `doi:10.1145/2897518.2897525`.
9. Venkatesan Guruswami and Mary Wootters. Repairing reed-solomon codes. *IEEE Trans. Inf. Theory*, 63(9):5684–5698, 2017. `doi:10.1109/TIT.2017.2702660`.
10. Yuval Ishai, Manoj Prabhakaran, Amit Sahai, and David Wagner. Private circuits II: Keeping secrets in tamperable circuits. In Serge Vaudenay, editor, *Advances in Cryptology – EUROCRYPT 2006*, volume 4004 of *Lecture Notes in Computer Science*, pages 308–327, St. Petersburg, Russia, May 28 – June 1, 2006. Springer, Heidelberg, Germany. `doi:10.1007/11761679_19`.
11. Yuval Ishai, Amit Sahai, and David Wagner. Private circuits: Securing hardware against probing attacks. In Dan Boneh, editor, *Advances in Cryptology – CRYPTO 2003*, volume 2729 of *Lecture Notes in Computer Science*, pages 463–481, Santa Barbara, CA, USA, August 17–21, 2003. Springer, Heidelberg, Germany. `doi:10.1007/978-3-540-45146-4_27`.
12. Yael Tauman Kalai and Leonid Reyzin. A survey of leakage-resilient cryptography. Cryptology ePrint Archive, Report 2019/302, 2019. `https://eprint.iacr.org/2019/302`.
13. Paul C. Kocher. Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In Neal Koblitz, editor, *Advances in Cryptology – CRYPTO'96*, volume 1109 of *Lecture Notes in Computer Science*, pages 104–113, Santa Barbara, CA, USA, August 18–22, 1996. Springer, Heidelberg, Germany. `doi:10.1007/3-540-68697-5_9`.

14 Paul C. Kocher, Joshua Jaffe, and Benjamin Jun. Differential power analysis. In Michael J. Wiener, editor, *Advances in Cryptology – CRYPTO'99*, volume 1666 of *Lecture Notes in Computer Science*, pages 388–397, Santa Barbara, CA, USA, August 15–19, 1999. Springer, Heidelberg, Germany. doi:10.1007/3-540-48405-1_25.

15 Hemanta K. Maji, Hai H. Nguyen, Anat Paskin-Cherniavsky, Tom Suad, and Mingyuan Wang. Leakage-resilience of the shamir secret-sharing scheme against physical-bit leakages. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology – EUROCRYPT 2021, Part II*, volume 12697 of *Lecture Notes in Computer Science*, pages 344–374, Zagreb, Croatia, October 17–21, 2021. Springer, Heidelberg, Germany. doi:10.1007/978-3-030-77886-6_12.

16 Hemanta K. Maji, Anat Paskin-Cherniavsky, Tom Suad, and Mingyuan Wang. Constructing locally leakage-resilient linear secret-sharing schemes. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology – CRYPTO 2021, Part III*, volume 12827 of *Lecture Notes in Computer Science*, pages 779–808, Virtual Event, August 16–20, 2021. Springer, Heidelberg, Germany. doi:10.1007/978-3-030-84252-9_26.

17 Jay Mardia, Burak Bartan, and Mary Wootters. Repairing multiple failures for scalar MDS codes. *IEEE Trans. Inf. Theory*, 65(5):2661–2672, 2019. doi:10.1109/TIT.2018.2876542.

18 James L. Massey. Some applications of coding theory in cryptography. *Mat. Contemp*, 21(16):187–209, 2001.

19 Jesper Buus Nielsen and Mark Simkin. Lower bounds for leakage-resilient secret sharing. In Anne Canteaut and Yuval Ishai, editors, *Advances in Cryptology – EUROCRYPT 2020, Part I*, volume 12105 of *Lecture Notes in Computer Science*, pages 556–577, Zagreb, Croatia, May 10–14, 2020. Springer, Heidelberg, Germany. doi:10.1007/978-3-030-45721-1_20.

20 Ryan O'Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 2014.

21 Itzhak Tamo, Min Ye, and Alexander Barg. Optimal repair of reed-solomon codes: Achieving the cut-set bound. In Chris Umans, editor, *58th Annual Symposium on Foundations of Computer Science*, pages 216–227, Berkeley, CA, USA, October 15–17, 2017. IEEE Computer Society Press. doi:10.1109/FOCS.2017.28.

A The Discrepancy of the Irwin-Hall Distribution

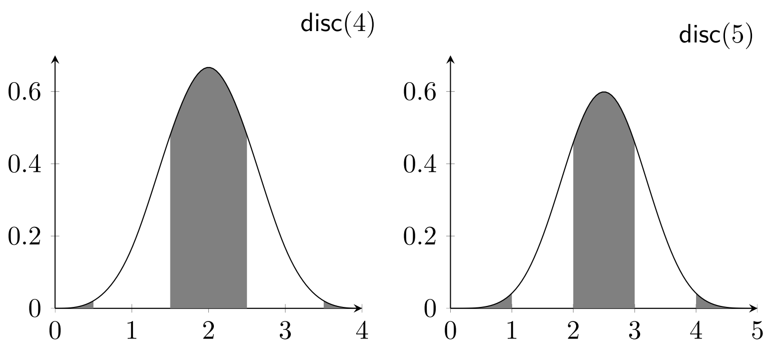

Figure 3 The plot of the fourth (left) and fifth (right) Irwin-Hall distribution. Intuitively, the discrepancy of the Irwin-Hall distribution is the difference between the total probability mass inside the black bands and the total probability mass outside the black bands. In particular, we are interested in the maximum difference as the black bands shift along the x-axis. Equation 1 provides a precise definition. This maximum difference is defined as the discrepancy of the k-th Irwin-Hall distribution, denoted by $\mathsf{disc}(k)$.

B On the Optimality of the Parity Distinguisher

Let $\mathcal{D}^{(0)}$ and $\mathcal{D}^{(1)}$ be two distributions over the universe $\{0,1\}^n$. Suppose $\mathcal{D}^{(0)}$ and $\mathcal{D}^{(1)}$ are $(n-1)$-indistinguishable.[9] That is, for any proper subset $S \subset \{1,2,\ldots,n\}$, we have

$$\mathsf{SD}\left(\left\{\begin{array}{c}\vec{x} \leftarrow \mathcal{D}^{(0)} \\ \text{Output } \vec{x}_S\end{array}\right\}, \left\{\begin{array}{c}\vec{x} \leftarrow \mathcal{D}^{(1)} \\ \text{Output } \vec{x}_S\end{array}\right\}\right) = 0.$$

For a distribution $\mathcal{D}$ and any set $S \subseteq 1,2,\ldots,n$, define the bias of $\mathcal{D}$ over S as

$$\mathsf{bias}(\mathcal{D}, S) := \mathop{\mathrm{E}}_{\vec{x} \leftarrow \mathcal{D}}\left[(-1)^{\sum_{i \in S} x_i}\right].$$

The following fact about the bias shall be useful. We refer the readers to [20] for a proof.

▶ **Lemma 13.**

$$\mathsf{SD}\left(\mathcal{D}^{(0)}, \mathcal{D}^{(1)}\right) \leqslant \frac{1}{2} \cdot \sqrt{\sum_{S \in \Omega} \left(\mathsf{bias}(\mathcal{D}^{(0)}, S) - \mathsf{bias}(\mathcal{D}^{(1)}, S)\right)^2},$$

where Ω is the power set of $\{1,2,\ldots,n\}$.

Observe that $\mathcal{D}^{(0)}$ and $\mathcal{D}^{(1)}$ are $(n-1)$-indistinguishable implies that

$$\mathsf{bias}(\mathcal{D}^{(0)}, S) = \mathsf{bias}(\mathcal{D}^{(1)}, S)$$

for all proper subsets $S \subset \{1,2,\ldots,n\}$.

Therefore, this lemma implies that

$$\mathsf{SD}\left(\mathcal{D}^{(0)}, \mathcal{D}^{(1)}\right) \leqslant \frac{1}{2} \cdot \left|\mathsf{bias}\left(\mathcal{D}^{(0)}, \{1,2,\ldots,n\}\right) - \mathsf{bias}\left(\mathcal{D}^{(1)}, \{1,2,\ldots,n\}\right)\right|.$$

This shows that the parity is the optimal distinguisher up to a constant as the right hand side is exactly the advantage of the parity distinguisher.

C Massey's Secret-sharing Schemes

For completeness, we recall Massey's Secret-sharing scheme. The following is taken verbatim from [16].

A *linear code* C (over the finite field F) of *length* $(n+1)$ and *rank* $(k+1)$ is a $(k+1)$-dimension vector subspace of F^{n+1}, referred to as an $[n+1,k+1]_F$-code. The *generator matrix* $G \in F^{(k+1)\times(n+1)}$ of an $[n+1,k+1]_F$ linear code C ensures that every element in C can be expressed as $\vec{x} \cdot G$, for an appropriate $\vec{x} \in F^{k+1}$. Given a generator matrix G, the row-span of G, i.e., the code generated by G, is represented by $\langle G \rangle$. A generator matrix G is in the *standard form* if $G = [I_{k+1}|P]$, where $I_{k+1} \in F^{(k+1)\times(k+1)}$ is the identity matrix and $P \in F^{(k+1)\times(n-k)}$ is the parity check matrix. In this work, we always assume that the generator matrices are in their standard form.

Massey Secret-sharing Schemes. Let $C \subseteq F^{n+1}$ be a linear code. Let $s \in F$ be a secret. The Massey secret-sharing scheme corresponding to C picks a random element

[9] We do not use the term $(n-1)$-independent since the LSB of a uniformly random field element is not exactly uniform over $\{0,1\}$.

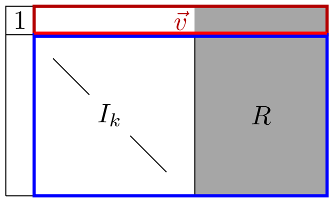

Figure 4 A pictorial summary of the generator matrix $G^+ = [I_{k+1} \mid P]$, where P is the shaded matrix. The indices of rows and columns of G^+ are $\{0, 1, \ldots, k\}$ and $\{0, 1, \ldots, n\}$, respectively. The (blue) matrix $G = [I_k \mid R]$ is a submatrix of G^+. In particular, the secret shares of secret $s = 0$ form the code $\langle G \rangle$. The (red) vector is $\vec{v}$. In particular, for any secret s, the secret shares of s form the affine subspace $s \cdot \vec{v} + \langle G \rangle$.

$(s, s_1, \ldots, s_n) \in C$ to share the secret s. The secret shares of parties $1, \ldots, n$ are $s_1, \ldots, s_n$, respectively.

Recall that the set of all codewords of the linear code generated by the generator matrix $G^+ \in F^{(k+1)\times(n+1)}$ is

$$\left\{ \vec{y} : \ \vec{x} \in F^{k+1}, \vec{x} \cdot G^+ =: \vec{y} \right\} \subseteq F^{n+1}.$$

For such a generator matrix, its rows are indexed by $\{0, 1, \ldots, k\}$ and its columns are indexed by $\{0, 1, \ldots, n\}$. Let $s \in F$ be the secret. The secret-sharing scheme picks independent and uniformly random $r_1, \ldots, r_k \in F$. Let

$$(y_0, y_1, \ldots, y_n) := (s, r_1, \ldots, r_k) \cdot G^+.$$

Observe that $y_0 = s$ because the generator matrix G^+ is in the standard form. The secret shares for the parties $1, \ldots, n$ are $s_1 = y_1, s_2 = y_2, \ldots, s_n = y_n$, respectively. Observe that every party's secret share is an element of the field F. Of particular interest will be the set of all secret shares of the secret $s = 0$. Observe that the secret shares form an $[n, k]_F$-code that is $\langle G \rangle$, where $G = G^+_{\{1,\ldots,k\}\times\{1,\ldots,n\}}$. Note that the matrix G is also in the standard form. The secret shares of $s \in F^*$ form the affine space $s \cdot \vec{v} + \langle G \rangle$, where $\vec{v} = G^+_{0,\{1,\ldots,n\}}$. Refer to Figure 4 for a pictorial summary.

Suppose parties $i_1, \ldots, i_t \in \{1, \ldots, n\}$ come together to reconstruct the secret with their, respective, secret shares $s_{i_1}, \ldots, s_{i_t}$. Let $G^+_{*,i_1}, \ldots, G^+_{*,i_t} \in F^{(k+1)\times 1}$ represent the columns indexed by $i_1, \ldots, i_t \in \{1, \ldots, n\}$, respectively. If the column $G^+_{*,0} \in F^{(k+1)\times 1}$ lies in the span of $\left\{G^+_{*,i_1}, \ldots, G^+_{*,i_t}\right\}$ then these parties can reconstruct the secret s using a linear combination of their secret shares. If the column G^+_{*0} *does not* lie in the span of $\left\{G^+_{*,i_1}, \ldots, G^+_{*,i_t}\right\}$ then the secret remains *perfectly hidden* from these parties.