Differential Indistinguishability:
Cryptography with Imperfect Randomness

Michael Backes, Aniket Kate, Sebastian Meiser, Tim Ruffing
CISPA, Saarland University

***Abstract*—Cryptographic protocols are commonly designed and their security proven under the assumption that the protocol parties have access to perfect (uniform) randomness. Physical randomness sources deployed in practical implementations of these protocols often fall short in meeting this assumption, but instead provide only a steady stream of bits with certain high entropy. Trying to ground cryptographic protocols on such imperfect, weaker sources of randomness has thus far mostly given rise to a multitude of impossibility results, including the impossibility to construct provably secure encryption, commitments, secret sharing, and zero-knowledge proofs based solely on a weak source. More generally, indistinguishability-based properties break down for such weak sources.**

In this work, we show that the loss of security induced by using a weak source can be meaningfully quantified if the source is *bounded*, e.g., for the well-studied Santha-Vazirani (SV) sources. The quantification relies on a novel relaxation of indistinguishability by a quantitative parameter. We call the resulting notion *differential indistinguishability* in order to reflect its structural similarity to differential privacy. More concretely, we prove that indistinguishability with uniform randomness implies differential indistinguishability with weak randomness. We show that if the amount of weak randomness is limited (e.g., by using it only to seed a PRG), all cryptographic primitives and protocols still achieve differential indistinguishability.

I. Introduction

Cryptographic protocols are commonly designed and their security proven under the assumption that the protocol parties have access to perfect, i.e., uniform, randomness. Actual physical randomness sources that cryptographic implementations rely on however rarely meet this assumption: instead of providing uniform randomness, they provide only a stream of bits with a certain high amount of entropy. Moreover, these so-called *weak sources*, such as the Santha-Vazirani (SV) sources [16], are often non-extractable [16], [7], i.e., it is computationally infeasible to extract more than a super-logarithmic amount of (almost) uniform randomness from them.

There have been several attempts to bridge this gap, i.e., to ground the security guarantees of cryptographic systems on such weak sources. As soon as indistinguishability-based secrecy properties are being desired, however, this line of research has mostly given rise to a multitude of impossibility results [7], [6], [13], only complemented by a few constructive results if additional assumptions are being imposed. For instance, encryption can be realized using weak sources, if one imposes strong assumptions on the entropy of encrypted messages [5], or if the weak source is restricted to the key generation algorithm and a perfect source is available for the actual encryption algorithm [8]. The plurality of impossibility results in this area, as well as the absence of comprehensive constructive results, indicates that traditional indistinguishability-based secrecy notions fall short in capturing the impact of weak randomness on cryptography. This constitutes an unsatisfactory situation, with several open questions looking for an answer:

- Is it possible to quantify the secrecy loss of cryptographic operations and primitives, if a weak source (such as an SV source) is being used?
- Imagine that today a cryptographic protocol (e.g., an e-voting system) is executed and tomorrow it turns out that the employed randomness was weak. Given that there are strong impossibility results [7], [6], [13] for indistinguishability, is all lost or can we still give quantitative guarantees about the secrecy of the system?
- Given that these quantitative guarantees will necessarily be weaker than traditional cryptographic guarantees, under which assumptions do they still provide reasonable practical security guarantees?

A. Our Contributions

1) Relaxing Indistinguishability to Quantify the Secrecy Loss: We derive quantitative guarantees for all indistinguishability-based cryptographic constructions that are used with arbitrary weak sources that are additionally *bounded* in the following sense: in addition to imposing an upper bound on the probability of each individual bitstring (i.e., requiring a sufficiently high min-entropy), one additionally imposes a lower bound on these probabilities. These *bounded weak sources* include SV sources [16] and resemble balanced sources [11].

To quantify the secrecy loss that weak randomness imposes on cryptography, we define *differential indistinguishability*, a quantitative relaxation of cryptographic indistinguishability in the spirit of differential privacy [9], [14] and pseudodensity [15]. The necessity of a new, relaxed

A full version of this work is available online [4].

notion arises from the impossibility result of Dodis et al. [7] who showed that whenever only weak sources of randomness are available, traditional indistinguishability is provably impossible for cryptographic primitives that have a secrecy requirement, e.g., encryption, commitments, and zero-knowledge proofs. More concretely, one cannot ensure that the advantage in distinguishing two challenger machines X_0 and X_1 is negligible for every probabilistic polynomial-time adversary. However, it might still be the case that no adversary has a non-negligible advantage in performing a practical attack that breaks the security *entirely*, e.g., by reaching a state in which it is *certain* whether it interacts with X_0 or X_1. The notion of differential indistinguishability consequently aims at quantifying the resulting loss of secrecy without overestimating the adversary's power to break the scheme entirely: Two games, i.e., interactions with two machines X_0 and X_1, are (ε, δ)-differentially indistinguishable if for all interactive distinguisher machines A, the output probabilities for all outputs x are related by

$$\Pr\left[\langle \mathsf{A}|\mathsf{X}_0\rangle = x\right] \leq 2^{\varepsilon} \cdot \Pr\left[\langle \mathsf{A}|\mathsf{X}_1\rangle = x\right] + \delta,$$

where x is a possible output of A.[1] Here $\varepsilon \geq 0$ is a reasonably small constant or a decreasing function such as $1/p(\cdot)$ for a polynomial p. We allow only a negligible function for δ, which corresponds to a negligible probability to break the security of the scheme entirely. Differential indistinguishability thus offers quantitative parameters to reason about the loss of secrecy incurred by the use of imperfect randomness.

2) Guarantees for Cryptographic Primitives Using Weak Sources: As our main contribution we show that traditional indistinguishability (given a uniform randomness source) suffices to guarantee differential indistinguishability if the uniform source is replaced by an arbitrary bounded weak source. This result immediately entails meaningful quantitative lower security bounds in cases where indistinguishability-based definitions are provably impossible to achieve [7].

In particular, our methodology can be applied in hindsight and produces meaningful quantitative guarantees for all cryptographic primitives and protocols, provided that the amount of used imperfect randomness is bounded; there is no need for new cryptographic constructions for any of the existing primitives whose security is defined and proven by means of indistinguishability, including simulator-based notions. Moreover, we show that if the bounded weak randomness is used only to seed a secure PRG, differential indistinguishability suffers only a negligible quantitative (additional) security loss under composition – just as traditional indistinguishability.

3) Connection to Differential Privacy: We analyze the relation between differential indistinguishability and the well-studied notion of differential privacy [9], [14], especially in terms of composition. Similar to the privacy loss in differential privacy when the privacy of several users is analyzed, differential indistinguishability suffers from a commensurate loss of entropy, which consequently leads to a secrecy loss in cases where several users use weak, potentially even dependent randomness.

II. Preliminaries and Notation

We denote sampling an element r from a distribution D by $r \leftarrow D$. The probability of the event $F(r)$, where r is sampled from the distribution D, is denoted by $\Pr\left[F(r) \mid r \leftarrow D\right]$ or more compactly by $\Pr\left[F(D)\right]$. To keep the notation simple, we write f_k for the value of a function $f(\cdot)$ applied to k, where k is typically the security parameter. We drop the explicit dependence of parameters and security bounds $(\alpha, \beta, \varepsilon, \gamma)$ on k whenever it is clear from the context. We denote by $\{D_k\}_{k\in\mathbb{N}}$ a family of distributions such that for each $k \in \mathbb{N}$ the distribution D_k samples elements from $\{0,1\}^k$. In particular, $\{U_k\}_{k\in\mathbb{N}}$ is the family of uniform distributions, where U_k is the uniform distribution over $\{0,1\}^k$.

Throughout the paper we consider (possibly interactive) Turing machines X that always have implicitly access to a *random tape with an infinite sequence of uniformly distributed random bits*, even if the machines get an additional input drawn from some random source. Unless we mention that they run in probabilistic polynomial time (ppt) in the length of their first input, those machines are *not bounded.* The distribution on the outputs of X when run on input x is denoted by $\mathsf{X}(x)$. Similarly, we write $\langle \mathsf{X}(x)|\mathsf{Y}(y)\rangle$ to denote the distribution on the output of the machine X on input x in an interaction with the machine Y on input y. We write $\log := \log_2$ for the logarithm to base 2.

4) Randomness Sources: In addition to the commonly used min-entropy, we make use of a symmetrically defined counterpart, coined *max-entropy* [11].

Definition 1. *Let D be a distribution over the set S. The* min-entropy *of D is $H_{min}(D) := \min_{x\in S}(-\log \Pr[D = x])$; the* max-entropy *of D is $H_{max}(D) := \max_{x\in S}(-\log \Pr[D = x])$.*

Definition 2. *A family of distributions $\{D_n\}_{n\in\mathbb{N}}$, each over the set $\{0,1\}^n$ of bitstrings of length n, is a (α, β)-*bounded weak source*, if every D_n satisfies the following entropy requirements:*

(i) D_n has min-entropy at least $n - \alpha$, and
(ii) D_n has max-entropy at most $n + \beta$.

III. Differential Indistinguishability

In this section we present our main results, which can be applied to a variety of cryptographic notions. Traditional cryptography defines two machines X_0 and X_1 to be *indistinguishable* for a certain class of distinguishers $\mathcal{A}$ if no distinguisher $\mathsf{A} \in \mathcal{A}$ in this class is able to notice a difference between an interaction with X_0 and an interaction with X_1. Formally, the concept of "noticing a difference" is captured by requiring that any possible view of a distinguisher is (almost) equally likely for both

[1] In contrast to differential privacy and pseudodensity, we use 2 instead of e as a base for the exponential function, because the base 2 fits standard definitions of entropy better.

X_0 and X_1, i.e., the difference between the probability that A outputs any given view in the interaction with X_0 and the probability that A outputs the same view in the interaction with X_1 is negligible. We consider a variant of indistinguishability that allows these probabilities to be also related by a multiplicative factor $2^\varepsilon > 1$, similar to mutual pseudodensity [15] and differential privacy [9], [14].

Definition 3 (Differential Indistinguishability)**.** *Two probabilistic machines X_0 and X_1 are (ε,δ)-differentially indistinguishable for a distribution $\{D_\ell\}_{\ell\in\mathbb{N}}$ over $\{0,1\}^\ell$ for a positive polynomial ℓ and a class $\mathcal{A}$ of adversaries (probabilistic machines) if for all $\mathsf{A} \in \mathcal{A}$, for all sufficiently large k, for all possible outputs x of A, and for all $b \in \{0,1\}$,*

$$\begin{aligned}&\Pr\left[\left\langle \mathsf{A}(1^k)\middle|\mathsf{X}_b(1^k, D_\ell)\right\rangle = x\right]\\ &\leq 2^\varepsilon \Pr\left[\left\langle \mathsf{A}(1^k)\middle|\mathsf{X}_{1-b}(1^k, D_\ell)\right\rangle = x\right] + \delta_k.\end{aligned}$$

This definition allows to express many of the traditional cryptographic indistinguishability notions [10], [12]. We discuss the impact of the multiplicative factor, that can (and must) be interpreted carefully, in Section V. For the traditional case of $\varepsilon = 0$ we speak of δ-indistinguishability. The definition covers interactive and non-interactive notions, as well as simulation-based notions. For perfect (information-theoretic) indistinguishability, the class of adversaries is the class $\mathcal{A}_\infty$ of all probabilistic (possibly unbounded) machines and we have $\delta = 0$.[2] Statistical indistinguishability can be expressed with the same class of adversaries for $\delta > 0$. Cryptographic (computational) indistinguishability can be achieved with the class $\mathcal{A}_{ppt}$ of ppt machines with δ being a negligible function.

A. Main Result

Traditional indistinguishability for uniform randomness directly implies differential indistinguishability for (α,β)-bounded weak sources. This is captured by the following theorem. It allows us to easily give guarantees for cryptographic primitives whenever their security notions can be expressed in terms of Definition 3.

Theorem 1. *If two probabilistic machines X_0 and X_1 are δ-indistinguishable for a class of probabilistic machines $\mathcal{A}$ and the family of uniform sources $\{U_n\}_{n\in\mathbb{N}}$ over $\{0,1\}^n$, then X_0 and X_1 are also $(\alpha+\beta, 2^\alpha \cdot \delta)$-differentially indistinguishable for $\mathcal{A}$ and any (α,β)-bounded weak source over $\{0,1\}^n$.*

B. Computational Differential Indistinguishability

In the computational setting where adversaries are ppt machines, we can achieve a stronger result: If we rely on a pseudorandom generator (PRG), we can expand a short seed from a randomness source to polynomially many bits of pseudorandomness. This well-known property is especially interesting here, as it allows us to apply Theorem 1 in a much broader form: Virtually every classically secure protocol is differentially secure when only a short random seed has been drawn from a bounded weak source and then expanded via a PRG, as this puts a limit on the entropy loss imposed by the actual bounded weak source. We formalize this observation in the following corollary, which is central to our work.

Corollary 1. *If two probabilistic machines X_0 and X_1 are computationally indistinguishable for a class of ppt machines $\mathcal{A}$ and uniform randomness, then X_0 and X_1 are also $(\alpha+\beta, 2^\alpha \cdot \delta)$-differentially indistinguishable for $\mathcal{A}$ and for a negligible function δ, if they draw their randomness from a PRG that is seeded with a (α,β)-bounded weak source.*

IV. Case Study: Public-Key Encryption

We apply differential indistinguishability to a common secrecy definition, namely indistinguishability under chosen ciphertext attacks for public-key encryption. This definition serves as example for how to instantiate the notion and how to apply our main results to quantify the secrecy loss under imperfect randomness. We demonstrate the applicability of our results by proving that public-key encryption achieves differential indistinguishability if it is used with bounded weak sources.

As an example, we relax *indistinguishability under adaptive chosen ciphertext attack* ($\mathsf{IND\text{-}CCA}$) [10] to use differential indistinguishability instead of traditional indistinguishability.

Definition 4 ((ε,δ)-$\mathsf{DIF\text{-}IND\text{-}CCA}$)**.** *A pair $\mathsf{A} = (\mathsf{A}_0, \mathsf{A}_1)$ of ppt oracle machines is an $\mathsf{IND\text{-}CCA}$ adversary if A_0 outputs two messages x_0, x_1 of the same length together with a state s, A_1 outputs a bit, and both A_0 and A_1 have access to decryption oracles as defined below. A PKE scheme $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ has (ε,δ)-*differentially indistinguishable encryptions under adaptive chosen ciphertext attack for a randomness source $\{D_n\}_{n\in\mathbb{N}}$ *if for all $\mathsf{IND\text{-}CCA}$ adversaries and for all sufficiently large k and bitstrings z of polynomial length in k, it holds that* $\Pr\left[\mathsf{P}^{(0)}_{k,z} = 1\right] < 2^\varepsilon \Pr\left[\mathsf{P}^{(1)}_{k,z} = 1\right] + \delta$*, where $P^{(i)}_{k,z}$ is the following probabilistic machine:*

$$\mathsf{P}^{(i)}_{k,z} := (e,d) \leftarrow \mathsf{Gen}(1^k);\ ((x_0,x_1),s) \leftarrow \mathsf{A}_0^{\mathsf{Dec}(d,\cdot)}(1^k, e, z)$$

$$c \leftarrow \mathsf{Enc}(e, x_i; D_n);\ \text{output } \mathsf{A}_1^{\mathsf{Dec}_c(d,\cdot)}(1^k, s, c)$$

Here, $\mathsf{Dec}_c(d,\cdot)$ denotes a decryption oracle that answers on all ciphertexts except for c, where it returns an error symbol $\perp$. The randomness used by the encryption algorithm Enc is drawn from D_n.

Note that $(0,\delta)$-$\mathsf{DIF\text{-}IND\text{-}CCA}$ security is equivalent to traditional δ-$\mathsf{IND\text{-}CCA}$ security.

1) Encryption with Imperfect Randomness: Both the encryption algorithm and the key generation algorithm require randomness. Dodis and Yu [8] show that even if weak sources are used for the key generation of $\mathsf{IND\text{-}CCA}$ secure encryption schemes, the security is preserved. However, this result does not apply when imperfect randomness is used by the *encryption algorithm.* The next theorem, an application of Theorem 1, quantifies the secrecy loss

[2] We additionally drop the formulation "for sufficiently large k" in the case of information-theoretic security.

whenever the encryption algorithm has only access to an (α, β)-bounded weak source.

Theorem 2. *Let $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ be any PKE scheme that is δ-*IND-CCA *secure under the assumption that* Enc *consumes at most n bits of uniform randomness. Then $\mathcal{E}$ is $(\alpha + \beta, 2^{\alpha}\delta)$-*DIF-IND-CCA *secure if* Enc *uses an (α, β)-bounded weak source $\{D_n\}_{n\in\mathbb{N}}$ instead of a uniform source.*

2) Discussion: Theorem 2 enables us to provide meaningful guarantees if an IND-CCA secure encryption scheme relies on an imperfect randomness, as long as the randomness used to encrypt the ciphertext in question is drawn from a bounded weak source. If an encryption scheme is $(\varepsilon$,$\delta)$-DIF-IND-CCA secure, the adversary may learn that the probability that a ciphertext contains a particular message m_0 is 2^{ε} times higher than the probability that it contains another message m_1. However, if ε is reasonably small, e.g., $\varepsilon = 0.001$ (and thus $2^{\varepsilon} \approx 1.001$), both m_0 and m_1 are a plausible content of the ciphertext. In particular, the adversary cannot reasonably believe or even convince a third party that m_0 is the value that has been encrypted. Moreover, the encryptor retains (a weak form of) deniability: She could indeed have encrypted any message.

3) Multiple Encryptions: Theorem 2 states a guarantee only for a single encryption (namely the encryption of one challenge message). However, it can be extended to the encryption of a message vector. In particular, if a PRG is used (and thus the amount of bounded weak randomness is limited to the seed of the PRG), Corollary 1 yields immediately an differential indistinguishability guarantee with ε being independent of the number of encrypted messages. If however, the encryption algorithm Enc is run several times with (fresh) imperfect randomness, the entropy loss of the randomness can increase linearly in the number of messages, and consequently, ε increases significantly.

A. Composability

Traditional indistinguishability with a negligible function δ and $\varepsilon = 0$ allows for polynomially many compositions as a polynomial factor for the advantage of an adversary that might come from from seeing multiple samples does not help the adversary substantially (the advantage remains negligible). This is not true for differential indistinguishability in general, because the (non-negligible) multiplicative factors can be accumulated as well.

For individual users we have shown that sequential composition of one or more primitives is possible without an (additional) loss of secrecy if a PRG is used (Corollary 1). If, however, several users within a protocol use imperfect randomness, the secrecy can degrade. Interestingly, we can give a bound on the loss of secrecy that is similar to the composition that occurs for differential privacy. We formulate a general composition lemma that we can instantiate to cope with several situations.

Lemma 1. *Let $\mathcal{A}$ be a class of adversaries. If X_0 and X_1 are (ε, δ)-differentially indistinguishable for $\mathcal{A}$, and X_1 and X_2 are (ε', δ')-differentially indistinguishable for $\mathcal{A}$, then X_0 and X_2 are $(\varepsilon'', \delta'')$-differentially indistinguishable for $\mathcal{A}$ where $\varepsilon'' = \varepsilon + \varepsilon'$ and $\delta'' = 2^{\varepsilon'}\delta + 2^{\varepsilon}\delta'$.*

A direct application of the lemma is the above described scenario in which multiple users (sequentially or concurrently) contribute to a protocol and use bad randomness. In this case, the machine X_1 can express an intermediate scenario that is used in a straightforward hybrid argument, where for two users X_1 is the only hybrid. Moreover, the lemma is applicable to scenarios where an individual user draws from a random source several times (for several primitives or protocols) instead of using a PRG, and also to compositions of differential indistinguishability guarantees in information-theoretical settings, where a PRG cannot be employed in the first place.

V. Interpretation and Analysis

In this section, we analyze and interpret the security guarantees provided by differential indistinguishability. In particular, we discuss the relation between differential indistinguishability and differential privacy.

A. Impact of a Multiplicative Factor

Similar to differential privacy, differential indistinguishability adds a multiplicative factor to the inequality used in the traditional indistinguishability notion. We observe that a multiplicative bound may express properties that are inexpressible by an additive bound. While every multiplicative bound of the form $\Pr[A] \leq 2^{\varepsilon}\Pr[B] + \delta$ implies a purely additive bound $\Pr[A] \leq \Pr[B] + \delta + 2^{\varepsilon} - 1 \approx \Pr[B] + \delta + \varepsilon$, the converse does not hold in general. No matter which additive bound can be shown between two probabilistic events, there does not necessarily exist a multiplicative bound. In particular, there are machines that are δ-indistinguishable for some δ but not (ε, δ')-indistinguishable for any ε such that $\delta' < \delta$.

For secrecy properties, traditional indistinguishability intuitively states that no adversary can learn any information about the secret, except with negligible probability. The multiplicative factor generalizes indistinguishability to additionally allow the adversary to learn information about the secret with more than a negligible probability, as long as the loss of secrecy is bounded; e.g., if ε is a small constant then differential indistinguishability ensures that the owner of the secret retains deniability by introducing doubt for the adversary.

Besides differential privacy, a multiplicative factor has also been used to achieve a specialized relaxation of semantic security in the presence of efficient adversaries that may tamper with an SV source [1, App B.4], and additionally for a security analysis of anonymous communication protocols [2], [3].

1) Example: Let us assume that Alice participates in an e-voting protocol based on, e.g., a commitment scheme. If the random source that she uses to seed her PRG turns out to be an (α, β)-bounded weak source, the commitments are still ε-*differently hiding* (this can be made formal), where $\varepsilon = \alpha + \beta$ is a small constant. Assume that Alice can vote for one of two popular candidates, say, Bob and Charlie, and she chooses to vote for Bob. In the traditional indistinguishability case, a non-negligible additive difference in the guarantee could result from a non-negligible probability of leaking the vote, which is highly unsatisfactory. The multiplicative factor 2^{ε}, however, allows us to guarantee that both cases will still maintain non-zero probability and no distinguisher can be sure whether Alice voted for Bob or for Charlie. Consider a distinguisher that only outputs, say '1' if it is certain that the vote was cast for Bob, and '0' in all other cases. Such a distinguisher is affected by the multiplicative bound as the output '1' is almost equally probable in all cases. Moreover, if the probability of outputting '1' is zero when the vote was cast for Charlie, then differential indistinguishability implies that the probability of outputting '1' is zero when the vote was cast for Bob.

Notice that the same analysis applies if a negligible additive value $\delta \neq 0$ is present. In this case, there might be a negligible chance for the adversary to be certain about the vote, but in all other cases, deniability is preserved.

B. *Relation to Differential Privacy and Sensitivity*

Differential privacy [9] quantifies the privacy provided by database query mechanisms: Intuitively, differential privacy requires that the output of a query mechanism should not allow to distinguish similar databases better than with a small multiplicative factor. Both in terms of the definition and in terms of the small, but usually non-negligible multiplicative factor, differential privacy and differential indistinguishability are closely related. We find this relation to be helpful for interpreting the guarantees and for understanding the drawbacks of differential indistinguishability. Differential privacy is influenced by the *sensitivity* of a statistical query, i.e., the amount of influence individual database records can have on the output of the query. Typical differential private mechanisms sanitize their output by adding random noise to guarantee a certain ε-level of privacy; the amount of added noise directly depends on the sensitivity.

Although there are neither databases nor the concept of utility (in the same sense as differential privacy) in our setting, the fact that a bounded weak source is differentially indistinguishable from a uniform source is analogous to the differential privacy of a query mechanism. From this point of view, the missing entropy of the weak source corresponds directly to the sensitivity in differential privacy.

This relation between sensitivity and entropy is interesting for sources that can be analyzed in a block-by-block manner, e.g., (n, γ)-SV sources. For such a source the entropy loss and thus the "sensitivity" is directly associated with the parameter γ and the amount of blocks that are drawn from this source. The higher the sensitivity, i.e., the more randomness is drawn by honest parties, the smaller γ must be to allow for guaranteeing ε-differential indistinguishability for a given value of ε. Clearly, the bias and thus the entropy loss in a $(1, \gamma)$-SV source can be arbitrarily increased, e.g., by drawing more random bits and taking the majority vote over them. Although this amplification does not make a difference for uniform randomness, it may increase the bias of the bits for SV sources. Therefore, for SV sources, the amount of randomness is a necessary parameter that influences the security.

References

[1] P. Austrin, K. Chung, M. Mahmoody, R. Pass, and K. Seth. On the impossibility of cryptography with tamperable randomness. In *Proc. of the 34th International Cryptology Conference (CRYPTO'14)*, pages 462–479, 2014.

[2] M. Backes, A. Kate, P. Manoharan, S. Meiser, and E. Mohammadi. AnoA: A framework for analyzing anonymous communication protocols. In *Proc. of the 26th IEEE Computer Security Foundations Symposium (CSF'13)*, pages 163–178, 2013.

[3] M. Backes, A. Kate, S. Meiser, and E. Mohammadi. (Nothing else) MATor(s): monitoring the anonymity of Tor's path selection. In *Proc. of the 21st Conference on Computer and Communications Security (CCS'14)*. ACM, 2014.

[4] M. Backes, A. Kate, S. Meiser, and T. Ruffing. Differential indistinguishability for cryptography with (bounded) weak sources. IACR Cryptology ePrint Archive, Report 2013/808, 2013. Technical Report.

[5] M. Bellare, Z. Brakerski, M. Naor, T. Ristenpart, G. Segev, H. Shacham, and S. Yilek. Hedged public-key encryption: How to protect against bad randomness. In *Proc. of the 15th International Conference on the Theory and Application of Cryptology and Information Security (ASIACRYPT'09)*, pages 232–249. Springer, 2009.

[6] C. Bosley and Y. Dodis. Does privacy require true randomness? In *Proc. of the 4th Theory of Cryptography Conference (TCC'07)*, pages 1–20. Springer, 2007.

[7] Y. Dodis, S. J. Ong, M. Prabhakaran, and A. Sahai. On the (im)possibility of cryptography with imperfect randomness. In *Proc. of the 45th Symposium on Foundations of Computer Science (FOCS'04)*, pages 196–205. IEEE, 2004.

[8] Y. Dodis and Y. Yu. Overcoming weak expectations. In *Proc. of the 10th Theory of Cryptography Conference (TCC'13)*, pages 1–22. Springer, 2013.

[9] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Proc. of the 3rd Theory of Cryptography Conference (TCC'06)*, pages 265–284. Springer, 2006.

[10] O. Goldreich. *Foundations of Cryptography: Volume 1, Basic Tools*. Foundations of Cryptography. Cambridge University Press, 2001.

[11] I. Haitner, O. Horvitz, J. Katz, C.-Y. Koo, R. Morselli, and R. Shaltiel. Reducing complexity assumptions for statistically-hiding commitment. In *Proc. of the 24th International Conference on Theory and Applications of Cryptographic Techniques (EUROCRYPT'05)*, pages 58–77. Springer, 2005.

[12] J. Katz and Y. Lindell. *Introduction to Modern Cryptography*. CRC Press, 2007.

[13] J. L. McInnes and B. Pinkas. On the impossibility of private key cryptography with weakly random keys. In *Proc. of the 10th International Cryptology Conference (CRYPTO'90)*, pages 421–435. Springer, 1990.

[14] I. Mironov, O. Pandey, O. Reingold, and S. Vadhan. Computational differential privacy. In *Proc. of the 29th International Cryptology Conference (CRYPTO'09)*. Springer, 2009.

[15] O. Reingold, L. Trevisan, M. Tulsiani, and S. Vadhan. Dense subsets of pseudorandom sets. In *Proc. of the 49th Symposium on Foundations of Computer Science (FOCS'08)*, pages 76–85. IEEE, 2008.

[16] M. Santha and U. V. Vazirani. Generating quasi-random sequences from slightly-random sources. In *Proc. of the 25th Symposium on Foundations of Computer Science (FOCS'84)*, pages 434–440. IEEE, 1984.