On the Relationship between Statistical Zero-Knowledge and Statistical Randomized Encodings

Benny Applebaum and Pavel Raykov

School of Electrical Engineering, Tel-Aviv University
{bennyap,pavelraykov}@post.tau.ac.il

Abstract. *Statistical Zero-knowledge proofs* (Goldwasser, Micali and Rackoff, SICOMP 1989) allow a computationally unbounded server to convince a computationally limited client that an input x is in a language Π without revealing any additional information about x that the client cannot compute by herself. *Randomized encoding* (RE) of functions (Ishai and Kushilevitz, FOCS 2000) allows a computationally limited client to publish a single (randomized) message, $\mathsf{Enc}(x)$, from which the server learns whether x is in Π and nothing else.

It is known that $\mathcal{SRE}$, the class of problems that admit statistically private randomized encoding with polynomial-time client and computationally unbounded server, is contained in the class $\mathcal{SZK}$ of problems that have statistical zero-knowledge proof. However, the exact relation between these two classes, and, in particular, the possibility of equivalence was left as an open problem.

In this paper, we explore the relationship between $\mathcal{SRE}$ and $\mathcal{SZK}$, and derive the following results:

- In a non-uniform setting, statistical randomized encoding with one-side privacy ($1\mathcal{RE}$) is equivalent to non-interactive statistical zero-knowledge ($\mathcal{NISZK}$). These variants were studied in the past as natural relaxation/strengthening of the original notions. Our theorem shows that proving $\mathcal{SRE} = \mathcal{SZK}$ is equivalent to showing that $1\mathcal{RE} = \mathcal{SRE}$ and $\mathcal{SZK} = \mathcal{NISZK}$. The latter is a well-known open problem (Goldreich, Sahai, Vadhan, CRYPTO 1999).
- If $\mathcal{SRE}$ is non-trivial (not in $\mathcal{BPP}$), then infinitely-often one-way functions exist. The analog hypothesis for $\mathcal{SZK}$ yields only *auxiliary-input* one-way functions (Ostrovsky, Structure in Complexity Theory, 1991), which is believed to be a significantly weaker implication.
- If there exists an average-case hard language with *perfect randomized encoding*, then collision-resistance hash functions (CRH) exist. Again, a similar assumption for $\mathcal{SZK}$ implies only constant-round statistically-hiding commitments, a primitive which seems weaker than CRH.

We believe that our results sharpen the relationship between $\mathcal{SRE}$ and $\mathcal{SZK}$ and illuminates the core differences between these two classes.

1 Introduction

Consider a "computationally-weak" client, Alice, which holds an input $x \in \{0,1\}^n$ to a language, or promise problem, Π which is beyond her computational power. We will be interested in the following two related scenarios.

- Alice contacts a computationally-strong server Bob, and asks him to prove that x is a yes-instance of Π. The server wishes to do so without revealing any additional information about x that Alice cannot compute by herself. That is, we are interested in an interactive proof system in which, for every yes-instance, the client is able to simulate her view without any interaction with the server.
- Alice would like to send to the server Bob a single (randomized) message $\mathsf{Enc}(x)$ which allows Bob to tell whether x is a yes-instance or a no-instance but hides any other information about x. That is, the message $\mathsf{Enc}(x)$ should be *private* in the sense that all yes-instances (resp., no-instances) are mapped by $\mathsf{Enc}(x)$ to the same universal yes-distribution $\mathsf{Sim}_{\text{YES}}$ (resp., no-distribution Sim_{NO}); In addition, $\mathsf{Enc}(x)$ should be *correct* (i.e., it should be possible to decode membership in Π) and so the yes-distribution is required to be statistically-far from the no-distribution.

The first setting is captured by the notion of *zero-knowledge* (ZK) proofs introduced in [GMR89], while the second is captured by the notion of *randomized encoding* (RE) of functions [IK00,AIK04]. In this paper, we model the client as a polynomial-time machine, the server as a computationally-unbounded party, and ask for information-theoretic security.[1] Problems that admit such a statistical zero-knowledge proofs (resp., such statistical randomized encodings) give rise to the complexity class $\mathcal{SZK}$ (resp., $\mathcal{SRE}$).

The class $\mathcal{SZK}$ and its variants were extensively studied and we have relatively rich insights about its power and structure including non-trivial upper-bounds (e.g., $\mathcal{SZK} \subseteq \mathcal{AM} \cap \text{co-}\mathcal{AM}$ [AH87]), complete problems [SV03,GV99], and closure properties [Oka00,Vad99]. Unfortunately, the status of $\mathcal{SRE}$ is very different. Although randomized encoding are extensively used in cryptography (see the surveys [App11,Ish13]), the class $\mathcal{SRE}$ was left relatively unexplored. The main known result (observed in [App14]) is that

$$\mathcal{SRE} \subseteq \mathcal{SZK}.$$

That is, a statistical randomized encoding for a problem Π can be transformed into a statistical zero knowledge proof system for the same problem. The exact relation between $\mathcal{SRE}$ and $\mathcal{SZK}$, and, in particular, the intriguing possibility that these two classes are actually equivalent was left as an open problem. This question was recently addressed by Agrawal, Ishai, Khurana, and

[1] The literature contains many other natural choices for security (e.g., computational [AIK05]) and efficiency (e.g., client with low parallel complexity and polynomial-time server [AIK04]). Following Agrawal, Ishai, Khurana, and Paskin-Cherniavsky [AIKP15], we view the current choice as a natural starting point for a complexity-theoretic treatment.

Paskin-Cherniavsky [AIKP15] who provided an oracle separation between the two classes, in addition to candidates for problems in $\mathcal{SRE}$ that are not solvable in (non-uniform) polynomial-time. As usual, an oracle separation tells us that equivalence cannot be established via relativized techniques, and so it essentially addresses the *proof of equivalence* (or technical barriers against it). However, such separations tell us very little on the statement itself ($\mathcal{SRE} = \mathcal{SZK}$) and its potential implications on the landscape of computational complexity.[2]

1.1 Our Results

In this paper, we continue the complexity theoretic study of $\mathcal{SRE}$, as advocated by [AIKP15], and further explore the exact relationship between $\mathcal{SRE}$ and $\mathcal{SZK}$. We study variants of these classes, prove their equivalence, and sharpen the difference between $\mathcal{SRE}$ and $\mathcal{SZK}$. We also point out several interesting complexity-theoretic implications of an equivalence between $\mathcal{SRE}$ and $\mathcal{SZK}$. Overall, we believe that our results shed light on the causes for which $\mathcal{SZK}$ is (seemingly) more powerful than $\mathcal{SRE}$.

Non-interactive ZK is equivalent to Semi-private RE Zero-knowledge proofs differ from randomized-encoding in many aspects. Most notably, the flow of information is reversed (Server-to-Client for ZK-proofs vs. Client-to-Sever for encodings). Let us ignore this major difference and focus on two seemingly less important syntactic differences. First, recall that REs are non-interactive while zero-knowledge proofs are allowed to use interaction. Secondly, the privacy condition of REs should hold for both yes and no-instances, whereas the ZK condition is defined only with respect to yes-instances. In an attempt to make a "fair" comparison between these two notions, we consider *non-interactive* zero-knowledge proofs (NISZK) [BFM88] and statistical randomized encoding with *one-sided privacy* (1RE) [AIK04,AIK15].

The NISZK model, introduced by Blum, Feldman and Micali [BFM88], restricts the prover to send a single message to the verifier at the expense of allowing the parties to share a common reference string that was pre-sampled by a trusted (efficient) dealer.[3] The notion of statistical randomized encoding with one-sided privacy was introduced by Applebaum, Ishai, and Kushilevitz [AIK04] (under the term semi-private encoding) as a relaxation of REs in which the privacy condition should hold only for yes-instances.

[2] Moreover, there are examples for classes which are separated relative to some oracle, but, without an oracle, are actually equal. (E.g., $\mathcal{IP}$ vs. $\mathcal{PSPACE}$; see the discussion in [CCG+94]).

[3] Our description corresponds to the *public-parameter model*, which is widely used in the literature (see [PS05] and references therein). This setting generalizes the original *common random string* (crs) model proposed by Blum et al. [BFM88], in which the reference string is simply a uniformly random string of polynomial length. Following [CCKV08], we use the superscripts PUB and CRS to distinguish between these two variants. Observe that $\mathcal{NISZK}^{\text{CRS}} \subseteq \mathcal{NISZK}^{\text{PUB}}$.

We show that the corresponding complexity classes $\mathcal{NISZK}^{\text{PUB}}$ and $\mathit{1}\mathcal{RE}$ are essentially equivalent.

Theorem 1. *It holds that* $\mathcal{NISZK}^{\text{PUB}} \subseteq \mathit{1}\mathcal{RE}$ *and, in the non-uniform setting,* $\mathit{1}\mathcal{RE} \subseteq \mathcal{NISZK}^{\text{PUB}}$.

The "non-uniform" setting refers to the case where all efficient entities (the client, the dealer, and the RE/SZK simulators) are modeled by polynomial-size circuits. The theorem shows that, non-uniformly, the class $\mathcal{NISZK}^{\text{PUB}}$ is *equivalent* to the class $\mathit{1}\mathcal{RE}$. It is known that $\mathcal{NISZK}^{\text{PUB}} \subseteq \mathcal{SZK}$ [PS05] and, by definition, we have that $\mathcal{SRE} \subseteq \mathit{1}\mathcal{RE}$. Hence, together with Theorem 1, we derive the following interesting picture (in the non-uniform setting):

$$\mathcal{SRE} \subseteq \mathit{1}\mathcal{RE} = \mathcal{NISZK}^{\text{PUB}} \subseteq \mathcal{SZK}.$$

Note that if $\mathcal{SZK}$ collapses to $\mathcal{SRE}$ then all intermediate classes also collapse. This means that the question of putting $\mathcal{SZK}$ inside $\mathcal{SRE}$ boils down to two separate questions: "Can statistical zero-knowledge be made non-interactive?" ($\mathcal{NISZK}^{\text{PUB}} = \mathcal{SZK}$?) and "Can one-side privacy be upgraded to full privacy?" ($\mathcal{SRE} = \mathit{1}\mathcal{RE}$?). Nicely, each of these well motivated questions is "pure" in the sense that it only addresses one object (either randomized encoding or zero-knowledge proofs). We further mention that the first question ($\mathcal{NISZK} = \mathcal{SZK}$?) is a well-known open problem that was studied before by [GSV99].[4]

Consequences of Randomized Encoding for Intractable Problems Another way to compare $\mathcal{SZK}$ to $\mathcal{SRE}$ is by asking what are the consequences of the existence of computationally-intractable problems in the class. For example, the following theorem was proven by Ostrovsky.

Theorem 2 ([Ost91]). *If* $\mathcal{SZK}$ *is not in* $\mathcal{BPP}$*, then Auxiliary-Input One-way functions exist.*[5]

Auxiliary-input one-way functions (ai-OWF) are keyed functions that achieve a very weak form of one-wayness. Roughly speaking, for each adversary there exists a set of hard keys on which the adversary fails to invert the function. (See [Gol01] for definition.) However, it may be the case that there is no universal set of keys which is simultaneously hard for all efficient adversaries.

For $\mathcal{SRE}$ we prove (Section 6) the following stronger implication:

Theorem 3. *If* $\mathcal{SRE}$ *is not in* $\mathcal{BPP}$*, then infinitely-often one-way functions exist.*

[4] More precisely, [GSV99] focused on the CRS model, and provided several necessary and sufficient conditions for the equality $\mathcal{NISZK}^{\text{CRS}} = \mathcal{SZK}$.

[5] This theorem, and all the other results in this section, is formulated in the uniform setting. If one considers a non-uniform variant of $\mathcal{SZK}$, then the theorem holds by changing $\mathcal{BPP}$ to $\mathcal{P}/\text{poly}$ and by relaxing the notion of AIOWFs to be computable by polynomial-size circuits. Similar modifications can be applied to the other theorems of this section.

Infinitely-often one-way functions (io-OWFs) are essentially standard one-way functions except that their hardness holds over a (universal) set of infinitely many input lengths. This notion is considered to be significantly stronger than ai-OWFs. For example, while it is possible to construct ai-OWFs based on the worst-case hardness of graph-isomorphism (GI), it is unknown how to obtain io-OWF from such an assumption. By Theorem 3, such a GI-based io-OWF would follow from the equivalence of $\mathcal{SZK}$ and $\mathcal{SRE}$. More generally, a proof of such an equivalence would allow us to base io-OWFs on worst-case hardness in $\mathcal{SZK}$ improving the 25-year old classical result of [Ost91].

Theorem 3 also explains why all the candidates of Agrawal et al. [AIKP15] for computationally-hard problems in $\mathcal{SRE}$ imply the existence of one-way functions – Such an assumption is inherently necessary to separate $\mathcal{SRE}$ from $\mathcal{BPP}$.

We can further ask what are the implications of an average-case hard problem in these complexity classes. Roughly speaking, a promise problem Π is average-case hard if it is equipped with a probability distribution D such that no efficient algorithm can classify correctly an instance x sampled from D with probability significantly better than $1/2$. Ostrovsky's result can be used to prove that the existence of an average-case hard language in $\mathcal{SZK}$ implies the existence of a one-way function. The following (stronger) theorem is implicit in the work of Ong and Vadhan [OV08].

Theorem 4 (implicit in [OV08]). *If there exists an average-case hard language in $\mathcal{SZK}$ then a constant-round statistically-hiding commitments (CRSC) exists.*

As a general primitive, CRCS implies the existence of one-way functions, and is believed to be strictly stronger due to the black-box separation of [HHRS15]. We derive a stronger implication if we have randomized encoding for an average-case hard problem. Specifically, we consider the class $\mathcal{PRE}$ of problems that admit *perfect randomized encoding* [AIK04] – a stronger variant of $\mathcal{SRE}$ which achieves perfect correctness (zero-decoding error), perfect privacy (the simulators perfectly simulate the encoding) and enjoys some additional syntactic properties. (See Section 4 for a formal definition.)

Theorem 5. *If there exists an average-case hard language in $\mathcal{PRE}$ then collision-resistance hash functions (CRH) exist.*

The proof of the theorem is sketched in Section 7. CRH imply CRSC but the converse is not known to be true. Hence, this implication is seemingly stronger than the one proven in [OV08]. Extending this theorem to the case of $\mathcal{SRE}$ is left as an interesting open problem.

2 Our Techniques

Let us outline the main ideas behind the proofs of Theorems 1, 3 and 5.

Proof of Theorem 1 We begin with the equivalence of $\mathit{1}\mathcal{RE}$ and $\mathcal{NISZK}^{\text{PUB}}$. It is instructive to note that all the complexity classes $\mathcal{SZK}, \mathcal{NISZK}$, $\mathit{1}\mathcal{RE}$ and $\mathcal{SRE}$ essentially capture different variants of "statistical-distance" problems. Indeed, as we already saw, for a $\mathcal{SRE}$-problem Π, the membership of x boils down to determining whether the distribution $\mathsf{Enc}(x)$ is close to one of two distributions $\mathsf{Sim}_{\text{YES}}$ and Sim_{NO} which are statistically-far apart from each other. Notably, these distributions are *universal* and they depend only on the problem Π (and not on the input x). The work of [SV03] also shows that, for any $\mathcal{SZK}$-problem Π, there exists an efficient mapping from an instance x to a pair of distributions (A_x, B_x) which are statistically-close if x is a yes-instance and statistically-far otherwise. However, in contrast to the case of SREs, the distributions (A_x, B_x) are *instance dependent*. In particular, two different yes-instances x and x' may be mapped to completely different pairs of distributions (A_x, B_x) and $(A_{x'}, B_{x'})$.

In the intermediate notion of $\mathcal{NISZK}$, one of the distributions, say B, corresponds to the dealer's distribution and so it becomes universal [SCPY98,GSV99].[6] Correspondingly, all yes-instances x are mapped to this single universal distribution, i.e., $A_x \approx B$. (A_x essentially corresponds to the simulated version of the public-parameter). For no-instances, the distribution A_x may be instance-dependent. Similarly, for $\mathit{1}\mathcal{RE}$, only yes-instances are mapped by $\mathsf{Enc}(x)$ to some universal yes-distribution $\mathsf{Sim}_{\text{YES}}$, whereas the encoding of a no-instance $\mathsf{Enc}(x)$ may be instance-dependent. Overall, the privacy properties of $\mathit{1}\mathcal{RE}$ and the zero-knowledge properties of $\mathcal{NISZK}$ match nicely. Still, there is one technical difference with respect to the requirements on the distributions of no-instances.

In $\mathit{1}\mathcal{RE}$, correctness requires the existence of a single decoder that distinguishes between the yes-distribution $\mathsf{Sim}_{\text{YES}}$ and *all* possible no-distributions $\{\mathsf{Enc}(x)\}_{x \in \Pi_{\text{NO}}}$. This means that $\mathsf{Sim}_{\text{YES}}$ is "universally-far" from all the no-distributions. In contrast, the soundness property of $\mathcal{NISZK}$ requires from every no-distribution A_x to be "disjoint" from B in the following sense: A random sample from the universal distribution $b \stackrel{R}{\leftarrow} B$ should fall, with high probability, outside the support of A_x. To prove Theorem 1 we should be able to move from "universal-farness" to "disjointness" and vice versa. While it is relatively straightforward to convert disjointness to universal-farness (e.g., via parallel-repetition), the converse direction requires some work.

As a concrete (and somewhat simplified) example, imagine the case where we have a single pair of distributions X and Y, where X outputs, with probability $1-\varepsilon$, a random n-bit string whose first bit is 1, and, with probability ε, a random n-bit string whose first bit is 0. Assume that Y does exactly the opposite. These distributions are $(1-2\varepsilon)$-far in statistical distance, but they do not satisfy the disjointness property as their supports are equal. The key observation is to note that a typical $y \stackrel{R}{\leftarrow} Y$ value, has much larger weight under Y compared to its weight under X. When these distributions are implemented by circuits that use m random bits as inputs, this means that the set of preimages $Y^{-1}(y)$ is likely to be significantly larger than the set $X^{-1}(y)$. In other words, the entropy e_1

[6] Interestingly, in the CRS model, this distribution is simply the uniform distribution and it is therefore also *problem-independent*.

of the conditional distribution $[r|Y(r) = y]$ is larger than the entropy e_2 of the conditional distribution $[r|X(r) = y]$. Following the approach of [GSV99], we can turn these distributions to be disjoint by hashing out about $e_1 \ll e \ll e_2$ random bits from r, and appending the result $h(r)$ to the output. That is, we define a pair of new distributions by $X' = (X(r), h, h(r))$ and $Y' = (Y(r), h, h(r))$ where h is sampled from a 2-universal family of hash functions.[7] One can now show that for a typical $y \stackrel{R}{\leftarrow} Y$ (and most h's), the conditional distribution $[h(r)|Y(r) = y]$ is almost uniform, whereas the conditional distribution $[h(r)|X(y) = y]$ has small support. This means that a random sample from Y' is likely to land out of the support of X', as required.

The actual construction introduces some additional technicalities. Most notably, it requires an estimation on the amount of entropy of the distribution which is sampled by $\mathsf{Sim}_{\text{YES}}$, the simulator of the original encoding. We overcome this problem by treating this value as a non-uniform advice. We note that this advice is short (of logarithmic length) and so one may hope to simply try all possible values. The problem is that some of these values will violate the zero-knowledge property, while others would violate soundness. Unfortunately, we do not know how to "combine" together several faulty NISZK protocol into a single good protocol. The question of finding a way around this problem and achieving a fully uniform reduction is left for future research.

Proof of Theorem 3 Recall that Theorem 3 asserts that if infinitely-often one-way functions do not exist, then any language Π in $\mathcal{SRE}$ can be decided by some $\mathcal{BPP}$ algorithm A. The proof is based on the following observation: Given an instance x, one can probabilistically decide if $x \in \Pi$ by first sampling an encoding $y = \mathsf{Enc}(x)$, and then outputting "yes" if the weight of y under the distribution $\mathsf{Sim}_{\text{YES}}$ is larger than its weight under Sim_{NO}. Note that the latter problem can be reduced to the following "distributional inversion" problem. Define the function

$$g(r, b) = \begin{cases} \mathsf{Sim}_{\text{NO}}(r), & \text{if } b = 0, \\ \mathsf{Sim}_{\text{YES}}(r), & \text{if } b = 1; \end{cases}$$

sample a random preimage (r, b) of y under g, and output the bit b. (I.e., when $b = 0$ the instance x is classified as a no-instance, and if $b = 1$ then x is classified as a yes-instance.) It can be shown, based on the privacy and the correctness guaranties of the encoding, that b is likely to classify x correctly. By the results of Impagliazzo and Luby [IL89], the distributional inversion problem can be efficiecntly solved (up to small, inverse-polynomial, deviation error), assuming that infinitely-often one-way functions do not exist.

It is instructive to compare the above to the SZK setting. The RE simulators give rise to a universal function g (independent of the instance x) whose inversion is as hard as deciding Π. In contrast, in the SZK setting, the corresponding distributions depend on x, and so deciding $x \in \Pi$ reduces to inverting an

[7] More generally, we could use any seeded randomness extractor that extracts e almost uniform bits from any e_2-bit source.

instance-dependent function g_x. Correspondingly, the intractability of Π yields only auxiliary-input one-way functions.

Proof of Theorem 5 In Theorem 5 we show that if an average-case hard language Π admits a prefect RE then CRH exist. The notion of perfect encoding guarantees that the image of the encoder Enc can be partitioned into two equal sets Y and N and that for any yes-instance (resp., no-instance) x, the mapping $\mathsf{Enc}(x;r)$ is a bijection from the randomness space to Y (resp., N). Similarly both simulators, $\mathsf{Sim}_{\text{YES}}(r)$ and $\mathsf{Sim}_{\text{NO}}(r)$, form a bijective mapping from the randomness space to Y and N, respectively. Let us define a pair of functions, keyed by instances x, y,

$$h^0_x(r,b) = \begin{cases} g(x;r), \text{ if } b = 0, \\ \mathsf{Sim}_{\text{NO}}(r), \text{ otherwise;} \end{cases} \qquad h^1_y(r,b) = \begin{cases} g(y;r), \text{ if } b = 0, \\ \mathsf{Sim}_{\text{YES}}(r), \text{ otherwise;} \end{cases}$$

Since the encoding is perfect, h^0_x and h^1_y are permutations if x is a yes-instance and y is a no-instance; on the other hand, if x is a no-instance and y is a yes-instance the images of the functions are disjoint. Suppose that there exists an efficiently samplable distribution $\mathcal{Y}$ over yes-instances which is indistinguishable from some efficiently samplable distribution $\mathcal{N}$ over no-instances. Then, we can sample a pair of yes/no instances $(x,y) \xleftarrow{R} \mathcal{Y} \times \mathcal{N}$ which is indistinguishable from a pair of no/yes instances $(x',y') \xleftarrow{R} \mathcal{N} \times \mathcal{Y}$. This means that, although the functions h^0_x, h^1_y are permutations with identical images, it is computationally hard to find a pair (u,v) which forms a "claw", i.e., $h^0_x(u) = h^1_y(u)$. (Indeed, a claw-finder can be used to distinguish (x,y) from (x',y').) Such *claw-free permutations* [Dam87,GMR88] imply the existence of CRH. The argument extends to the case where there exists only a single "hard" distribution over yes/no instances of Π (as opposed to a pair of "pure" distributions). In this case, we get claw-free *pseudo-permutations* [Rus95], whose existence still implies CRH.

2.1 A Broader Perspective

So far we emphasized the differences between $\mathcal{SRE}$ and $\mathcal{SZK}$, however, from a broader point of view, our results may be interpreted as saying that the two classes are actually close variants of each other. This is similar in spirit to a recent result [AR16] that reveals a close connection between *private simultaneous message protocols* (PSM) [FKN94] and *Zero-Information Arthur-Merlin* (ZAM) protocols [GPW15]. PSMs and ZAMs can be viewed as the communication-complexity analog of Randomized Encodings and Zero-Knowledge proofs, where instead of limiting the computational power of the client, we split it into two non-communicating (computationally-unbounded) parties Alice and Bob each holding different parts of the input $x = (x_A, x_B)$. It is shown in [AR16] that the communication complexity of ZAM protocols is closely related to the randomness complexity of (variants of) PSMs, and vice versa. This is conceptually similar to some of the current results (e.g., $\mathit{1}\mathcal{RE} = \mathcal{NISZK}^{\text{PUB}}$) though the computational

setting introduces different technical challenges, and correspondingly it requires a significantly different approach.

Organization. We begin with some standard preliminaries in Section 3. In Section 4 we provide formal definitions of statistical zero knowledge proofs, statistical randomized encoding and their variants. Theorem 1 is proved in Section 5, Theorem 3 in Section 6 and Theorem 5 in Section 7.

3 Preliminaries

Basic Definitions. For a finite set S, let $s \xleftarrow{R} S$ denote an element that is sampled uniformly at random from S, and let $U(S)$ denote the corresponding random variable. The uniform distribution over n-bit strings is denoted by U_n. The *support* of a random variable X is the set $\mathrm{supp}(X) := \{x \mid \Pr[X = x] > 0\}$. The Shannon *entropy* of X is $H(X) := -\sum_z \Pr[X = z] \log \Pr[X = z]$. For a distribution D, we let $\otimes^k D$ be the probability distribution over k-tuples where each element is sampled independently according to D. Similarly, for a randomized algorithm $F(x)$, we let $\otimes^k F(x)$ be a k-tuple of k independent samples of $F(x)$. We sometimes make the coins of a randomized algorithm F explicit by writing $F(x; r)$ where $r \xleftarrow{R} U_{s(x)}$ denotes the random coins used on an input x and $s(x)$ denotes the randomness complexity of F on an input x, which, by default, is assumed to solely depend on the length of x.

Statistical Distance. The *statistical distance* between a pair of random variables X and Y distributed over the set Z is defined as

$$\Delta(X; Y) := \frac{1}{2} \sum_{z \in Z} \left| \Pr[X = z] - \Pr[Y = z] \right|.$$

Equivalently, $\Delta(X; Y) = \max_A |\Pr[A(X) = 1] - \Pr[A(Y) = 1]|$ where the maximum ranges over all Boolean functions $A : Z \to \{0, 1\}$. We write

$$\Delta_{x_1 \xleftarrow{R} D_1, \ldots, x_k \xleftarrow{R} D_k} (F(x_1, \ldots, x_k); G(x_1, \ldots, x_k))$$

to denote the statistical distance between two random variables obtained as a result of sampling x_i's from D_i's and applying the functions F and G to $(x_1, \ldots, x_k)$, respectively. We will use the following properties of statistical distance and entropy.

Fact 1. *Let X and Y be a pair of random variables. Then the following holds:*

1. [Vad99, Fact 3.2.2] *For every (possibly randomized) function F, we have that $\Delta(F(X); F(Y)) \leq \Delta(X; Y)$.*
2. [Vad99, Fact 3.3.9] *Let D be the range of X and Y, then $|H(X) - H(Y)| \leq (\log |D|) \cdot \Delta(X; Y) + 1$.*

3. [Vad99, Lemma 3.1.15] *For any integer $q > 0$, we have that* $1 - 2\exp(-q(\Delta(X;Y))^2/2) \leq \Delta(\otimes^q X; \otimes^q Y) \leq q\,\Delta(X;Y)$.
4. [SV03, Fact 2.5] *Suppose that $X = (X_1, X_2)$ and $Y = (Y_1, Y_2)$ are distributed over a set $D \times E$ such that: (a) X_1 and Y_1 are identically distributed; and (b) with probability greater than $1 - \varepsilon$ over $x \xleftarrow{R} X_1$, we have $\Delta(X_2|_{X_1=x}, Y_2|_{Y_1=x}) \leq \delta$. Then $\Delta(X,Y) \leq \varepsilon + \delta$.*
5. (cf. Appendix A.1) *If $\Delta(X;Y) \geq 1 - \varepsilon$, then, for any $t > 1$, it holds that* $\Pr_{x \xleftarrow{R} X}[\Pr[X = x] < t \cdot \Pr[Y = x]] \leq \varepsilon t$.

Flattening. We will use the following notion of Δ-flat distributions from [GSV99].

Definition 1 (Flat Distributions). *Let X be a distribution. An element x of* $\mathrm{supp}(X)$ *is called ε-typical if $|\log(1/\Pr[X = x]) - H(X)| \leq \varepsilon$. We say that X is Δ-flat if for every $t > 0$ the probability that an element chosen from X is $(t \cdot \Delta)$-typical is at least $1 - 2^{-t^2+1}$.*

A 0-flat distribution is uniform on its support, and is simply referred to as a *flat* distribution. A natural way to flatten a distribution is via parallel repetition.

Lemma 1 (Flattening Lemma [Vad99,GSV99]). *Let D be a distribution such that for all x from* $\mathrm{supp}(D)$ *we have that $D(x) \geq 2^{-m}$. Then, for any $k \in \mathbb{N}$, the distribution $\otimes^k D$ is $(\sqrt{k} \cdot m)$-flat.*

Hashing. A family $\mathcal{H}$ of functions mapping a domain $\mathcal{D}$ to a range $\mathcal{R}$ is 2-*universal* [CW79] if for every two elements $x \neq y$ from $\mathcal{D}$ and a, b from $\mathcal{R}$ it holds that $\Pr_{h \xleftarrow{R} \mathcal{H}}[h(x) = a \wedge h(y) = b] = \frac{1}{|\mathcal{R}|^2}$. We write $\mathcal{H}_{n,m}$ to denote a 2-universal family from $\{0,1\}^n$ to $\{0,1\}^m$. There are efficient constructions of 2-universal families of hash functions $\mathcal{H}_{n,m}$ that can be evaluated and sampled in $\mathrm{poly}(n,m)$ time [CW79].

Lemma 2 (Leftover Hash Lemma [ILL89,GSV99]). *Let $\mathcal{H}$ be a 2-universal family of hash functions mapping a domain $\mathcal{D}$ to a range $\mathcal{R}$. Let X be a flat distribution on $\mathcal{D}$ such that for all $x \in \mathrm{supp}(X)$ we have that* $\Pr[X = x] \leq \alpha/|\mathcal{R}|$. *Then*

$$\underset{h \xleftarrow{R} \mathcal{H}}{\Delta}((h, h(X)); (h, U(\mathcal{R}))) \leq O(\alpha^{1/3}).$$

Sampling distributions via circuits. Let X be a circuit with m input and n output gates. We will sometimes abuse notation and use X to denote the random variable $X(U_m)$ which corresponds to the output distribution of the circuit induced by "feeding" a uniformly chosen n-bit input. We let $X^{-1}(x)$ denote the set of preimages of x under X, i.e., $X^{-1}(x) := \{r \in \{0,1\}^m \mid X(r) = x\}$. Observe that $\Pr[X = x] = 2^{-m} \cdot |X^{-1}(x)|$.

4 $\mathcal{NISZK}$ and $\mathcal{SRE}$

A *promise problem* [ESY84] Π is a pair of two non-intersecting sets of strings $(\Pi_{\text{YES}}, \Pi_{\text{NO}})$. The strings in Π_{YES} are called *yes-instances* and the strings in Π_{NO}

are called *no-instances*. Let $\chi_{\Pi}(x)$ be the characteristic function of Π which outputs 1 on yes-instances and 0 on no-instances. Note that a promise problem is a generalization of a language $L \subseteq \{0,1\}^*$, i.e., L is translated into a promise problem Π_L where L corresponds to the set of yes-instances and $\{0,1\}^* \setminus L$ corresponds to the set of no-instances. (See [Gol06] for a survey.)

Definition 2 (statistical randomized encoding [IK00,AIK04]). *We say that an efficient randomized algorithm* Enc *is a* ε*-private and* δ*-correct* statistical randomized encoding *of a promise problem* $\Pi = (\Pi_{\text{YES}}, \Pi_{\text{NO}})$ *(abbreviated* (ε,δ)-SRE*), if the following holds:*

ε-PRIVACY FOR YES-INSTANCES: *There exists an efficient simulator* $\mathsf{Sim}_{\text{YES}}$ *such that for every yes-instance* x_{YES} *of length* n *from* Π*,*

$$\Delta(\mathsf{Sim}_{\text{YES}}(1^n); \mathsf{Enc}(x_{\text{YES}})) \leq \varepsilon(n).$$

ε-PRIVACY FOR NO-INSTANCES: *There exists an efficient simulator* Sim_{NO}*, such that for every no-instance* x_{NO} *of length* n *from* Π*,*

$$\Delta(\mathsf{Sim}_{\text{NO}}(1^n); \mathsf{Enc}(x_{\text{NO}})) \leq \varepsilon(n).$$

δ-CORRECTNESS: *There exists a computationally-unbounded decoder* Dec*, such that for every instance* $x \in (\Pi_{\text{YES}} \cup \Pi_{\text{NO}})$ *of length* n*,*

$$\Pr[\mathsf{Dec}(\mathsf{Enc}(x)) \neq \chi_{\Pi}(x)] \leq \delta(n).$$

By default, $\varepsilon(n)$ *and* $\delta(n)$ *are required to be negligible functions.*

Perfect Encoding [AIK04]. A randomized encoding which is 0-private (resp., 0-correct) is called *perfectly* private (resp., *perfectly correct*). For an input of length n, let $s(n)$ denote the length of the random strings used by Enc and let $t(n)$ be the output length of the encoding. A perfectly private and perfectly correct randomized encoding whose simulators $\mathsf{Sim}_{\text{YES}}$ and Sim_{NO} use $s(n)$ coins, $\text{supp}(\mathsf{Sim}_{\text{YES}}(1^n)) \cup \text{supp}(\mathsf{Sim}_{\text{NO}}(1^n)) = \{0,1\}^{t(n)}$, and $1 + s(n) = t(n)$ is called *perfect*. (See [AIK04] for an intuitive explanation of these requirements.)

One-sided Encoding [AIK04,AIK15]. A randomized encoding which is ε-private on yes-instances and δ-correct is called *one-sided* (or *semi-private*) randomized encoding (denoted with (ε,δ)-$\mathsf{1RE}$)[AIK04,AIK15]. Clearly, any (ε,δ)-SRE is also (ε,δ)-$\mathsf{1RE}$, though the converse does not necessarily hold. A *disjoint one-sided* randomized encoding is an encoding which is ε-private on yes-instances and, instead of standard correctness, it satisfies the following ρ*-disjointness* property: For every no-instance x_{NO} of length n from Π, it holds that $\Pr[\mathsf{Sim}_{\text{YES}}(1^n) \in \text{supp}(\mathsf{Enc}(x_{\text{NO}}))] \leq \rho(n)$. We refer to such an encoding as (ε,ρ)-$\mathsf{D1RE}$.

Definition 3 (non-interactive statistical zero-knowledge [BSMP91]). *A* non-interactive statistical zero-knowledge proof system *(*NISZK*) for a promise problem* $\Pi = (\Pi_{\text{YES}}, \Pi_{\text{NO}})$ *is defined by probabilistic algorithms* Prov *(prover),* Deal *(dealer),* Sim *(simulator), and a deterministic algorithm* Ver *(verifier), such that for every* n*-bit instance* x *the following holds*

α-COMPLETENESS: *If* $x \in \Pi_{\text{YES}}$ *then* $\Pr[\mathsf{Ver}(x, \sigma, \mathsf{Prov}(x, \sigma)) \neq 1] \leq \alpha(n)$*, where* $\sigma \stackrel{R}{\leftarrow} \mathsf{Deal}(1^n)$.

β-SOUNDNESS: *If* $x \in \Pi_{\text{NO}}$ *then* $\Pr[\exists p = p(x, \sigma) : \mathsf{Ver}(x, \sigma, p) = 1] \leq \beta(n)$*, where* $\sigma \stackrel{R}{\leftarrow} \mathsf{Deal}(1^n)$.

γ-ZERO-KNOWLEDGE: *If* $x \in \Pi_{\text{YES}}$ *then the pair* (σ, p) *is* $\gamma(n)$*-close in statistical distance to the pair* (σ', p') *where* $\sigma \stackrel{R}{\leftarrow} \mathsf{Deal}(1^n), p \stackrel{R}{\leftarrow} \mathsf{Prov}(x, \sigma)$ *and* $(\sigma', p') \stackrel{R}{\leftarrow} \mathsf{Sim}(x)$.

The algorithms $\mathsf{Ver}, \mathsf{Deal}$*, and* Sim *are required to be efficient, while the prover's algorithm* Prov *is allowed to be computationally unbounded. By default,* α, β *and* γ *are assumed to be negligible in* n.

Variants. In the special case where the dealer $\mathsf{Deal}(1^n)$ samples σ uniformly from the set of all strings of length $r(n)$ (for some polynomial $r(\cdot)$), the proof system is called an interactive zero-knowledge proof system in the *common random string model* and is denoted by (α, β, γ)-$\mathsf{NISZK}^{\text{CRS}}$ [BFM88]. We will focus on the more general setting (defined above) where the dealer is allowed to use any arbitrary (polynomial-time samplable) distribution. This setting is referred to as the *public parameter model* and protocols in the model are denoted by (α, β, γ)-$\mathsf{NISZK}^{\text{PUB}}$.[8]

Remark 1 (Efficiency: Uniformity vs. Non-Uniformity). Randomized encodings and non-interactive statistical-zero knowledge proof systems can be defined either in the uniform setting where all efficient entities (encoder, RE-simulator, verifier, dealer, and NISZK-simulator) are assumed to be probabilistic polynomial-time algorithms, or in the non-uniform setting where these entities are represented by probabilistic polynomial-time algorithms which take a non-uniform advice. We will emphasize this distinction only when it matters (Theorem 6), and otherwise, (when the results are insensitive to the difference) ignore it.

Definition 4 (Complexity classes). *The complexity class* $\mathcal{SRE}$ *(resp.,* $1\mathcal{RE}$*,* $\mathcal{NISZK}^{\text{PUB}}$*) is the set of all the promise problems that have an* SRE *(resp.,* $\mathsf{1RE}$*,* $\mathsf{NISZK}^{\text{PUB}}$*).*

5 $\mathcal{NISZK}^{\text{pub}} = 1\mathcal{RE}$

In this section we will prove Theorem 1. We start by showing that the notions of $\mathsf{1RE}$ and $\mathsf{D1RE}$ are equivalent in Section 5.1. Then, based on this equivalence we prove that $\mathcal{NISZK}^{\text{PUB}} = 1\mathcal{RE}$. In the first part of the proof we show that $\mathcal{NISZK}^{\text{PUB}} \subseteq 1\mathcal{RE}$ (cf. Section 5.2). In the second part of the proof we show that $1\mathcal{RE} \subseteq \mathcal{NISZK}^{\text{PUB}}$ (cf. Section 5.3).

[8] The class $\mathsf{NISZK}^{\text{PUB}}$ was implicitly considered in [BDLP88], and was later referred to as NISZK in the *auxiliary string model* [Dam00] and as *protocol-dependent* NISZK by [GB00]. Our terminology (NISZK in *public parameter model*) is taken from [PS05].

5.1 Equivalence of 1RE and D1RE

We start by showing how to convert a 1RE F for a promise problem Π into a D1RE G for the same problem. The construction is inspired by the techniques of [GSV99]. The encoding G consists of sufficiently many independent copies of F together with a hash of the randomness used to generate the copies. In order to achieve disjointness, while keeping privacy, the length of the hash is chosen such that for yes-instance the hash is close to uniform and in the case of no-instances the support of the hash output is relatively small.

We note that this construction is *non-uniform*. That is, the length of the hash is chosen using a non-uniform advice that depends on the entropy of the encoding distribution on yes-instances. It is an interesting open question whether one can give a uniform construction achieving disjointness.

Theorem 6. *If the promise problem Π has a (possibly non-uniform)* 1RE *F, then it also has a non-uniform* D1RE *G. Moreover, if F is uniform then G can be implemented based on F and an advice of $O(\log n)$ bits.*

Proof. Let Π be a promise problem that has an ε-private and δ-correct 1RE F, where ε and δ are negligible. Let Sim_F be the simulator showing the privacy of F on yes-instances. For an input length of n, let $m = m(n) = \mathrm{poly}(n)$ denote the maximum bit-length of the randomness used by Sim_F and F. We define a D1RE $G(x)$ for Π as follows:

1. Parameters: $q = 10^6 nm^2$, $m' = qm$.
2. Non-uniform advice $\ell := \lceil m' - H(S_n) - \sqrt{qn} \cdot m - 2n \rceil$.
3. Input: $x \in \{0,1\}^n$.
4. Sample randomness $r = (r_1, \ldots, r_q) \xleftarrow{R} \{0,1\}^{m'}$ (where $|r_i| = m$), and a pair-wise independent hash function $h \xleftarrow{R} \mathcal{H}_{m',\ell}$.
5. Output $((F(x;r_1), \ldots, F(x;r_q)), h, h(r))$.

To simplify notation, we let $J_x(r) = (F(x,r_1), \ldots, F(x,r_q))$ and write J_x to denote the distribution induced by a uniform choice of $r \xleftarrow{R} U_{m'}$. We let $S_n = \otimes^q \mathsf{Sim}_F(1^n)$, and let $\mathcal{H}$ denote the family $\mathcal{H}_{m',\ell}$.

We proceed with an analysis of the encoding G, starting with privacy. We define the simulator $\mathsf{Sim}_G(1^n)$ to generate the random variable $(S_n, U(\mathcal{H}), U_\ell)$. Fix some yes-instance x of length n from Π. Our goal is to show that the statistical distance $\varepsilon'(n)$ between $\mathsf{Sim}_G(1^n)$ and $G(x)$ is upper-bounded by some negligible function. First observe that, by the triangle inequality, ε' is upper-bounded by

$$\Delta(\mathsf{Sim}_G(1^n)\ ;\ (J_x, U(\mathcal{H}), U_\ell)) + \Delta((J_x, U(\mathcal{H}), U_\ell)\ ;\ G(x)). \tag{1}$$

By the ε-privacy of the original encoding and by Fact 1 item 3, the first summand satisfies

$$\begin{aligned}\Delta(\mathsf{Sim}_G(1^n)\ ;\ (J_x, U(\mathcal{H}), U_\ell)) &= \Delta((S_n, U(\mathcal{H}), U_\ell)\ ;\ (J_x, U(\mathcal{H}), U_\ell)) \\ &\leq \Delta(S_n, J_x) \\ &\leq q\varepsilon(n) = \mathrm{neg}(n).\end{aligned}$$

It is left to analyze the second summand in (1), i.e., to upper-bound the quantity

$$\Delta_{r \xleftarrow{R} \{0,1\}^{m'}, h \xleftarrow{R} \mathcal{H}} ((J_x(r), h, U_\ell)\ ;\ (J_x(r), h, h(r))). \tag{2}$$

Since the first entry is identically distributed in both distributions, it suffices to analyze the statistical distance between the two tuples conditioned on the outcome of the first entry J_x. Indeed, we prove the following claim.

Claim 1. With probability $1 - 2^{-\Omega(n)}$ over $z \xleftarrow{R} J_x$, it holds that

$$\Delta_{r \xleftarrow{R} \{0,1\}^{m'}, h \xleftarrow{R} \mathcal{H}} ([J_x(r), h, U_\ell | J_x(r) = z]\ ;\ [J_x(r), h, h(r) | J_x(r) = z]) < 2^{-\Omega(n)}. \tag{3}$$

It follows (by Fact 1 item 4) that (2) is upper-bounded by $2^{-\Omega(n)}$.

Proof (Proof of Claim 1). Recall that on any input x the encoding F uses at most m random bits, and so any element in its support has weight at least 2^{-m}. Hence, due to the Flattening Lemma 1, the distribution J_x is Δ-flat for $\Delta = \sqrt{q}m$. Since $z \xleftarrow{R} J_x$ is $(\sqrt{n}\Delta)$-typical with probability at least $1 - O(2^{-n})$, it suffices to show that (3) holds for every $(\sqrt{n}\Delta)$-typical z.

Fix some $(\sqrt{n}\Delta)$-typical z from J_x and consider the distribution $(J_x(r), h, h(r))$ conditioned on $J_x(r) = z$. The conditional distribution of r is uniform over the set $J_x^{-1}(z)$. We will show below that

$$\log(|J_x^{-1}(z)|) \geq \ell + n \tag{4}$$

Therefore we can apply the Leftover Hash Lemma 2 to the distribution of $r \xleftarrow{R} J_x^{-1}(z)$ with $\mathcal{R} = \{0,1\}^\ell$ and $\alpha = 2^{-n}$, and conclude that the distribution of $(J_x(r), h, h(r))$ conditioned on $J_x(r) = z$ is $O(2^{-n/3})$-close to the distribution $(z, U(\mathcal{H}), U_\ell)$.

It remains to prove (4). First, we show that the entropies $H(J_x)$ and $H(S_n)$ are close. Indeed, by the privacy of F, we have that $\Delta(\mathsf{Sim}_F(1^n); F(x)) \leq \varepsilon(n)$ and therefore (by Fact 1 item 3) $\Delta(J_x; S_n) \leq q\varepsilon(n)$. Hence, by Fact 1 item 2, we get that, for all sufficiently large n's,

$$|H(J_x) - H(S_n)| \leq m'q\varepsilon(n) + 1 \leq 2, \tag{5}$$

where the second inequality follows by noting that $\varepsilon(n)$ is negligible in n, and m', q are polynomials in n. Now, recall that z is $(\sqrt{n}\Delta)$-typical, and therefore

$\log(|J_x^{-1}(z)|) \geq m' - H(J_x) - \sqrt{n}\Delta$. Plugging in (5) we conclude that

$$\begin{aligned}\log(|J_x^{-1}(z)|) &\geq m' - H(S_n) - 2 - \sqrt{n}\Delta \\ &\geq \underbrace{\lceil m' - H(S_n) - \sqrt{n}\Delta - 2n \rceil}_{=\ell} + (n-3) + n \\ &\geq \ell + n,\end{aligned}$$

where the last inequality holds for $n \geq 3$. □

We move on to prove the disjointness property. Fix some no-instance x. Our goal is to upper-bound

$$\Pr\left[\mathsf{Sim}_G(1^n) \in \mathrm{supp}(G(x))\right] = \Pr\left[(S_n, U(\mathcal{H}), U_\ell) \in \mathrm{supp}(G(x))\right] \tag{6}$$

by some negligible function. For $z \xleftarrow{R} S_n$, let $\mathcal{E} = \mathcal{E}(z)$ be the event that $|J_x^{-1}(z)| \leq 2^{\ell-n}$. By marginalizing the probability, we can upper-bound (6) by

$$\Pr_{z \xleftarrow{R} S_n, h \xleftarrow{R} \mathcal{H}, w \xleftarrow{R} \{0,1\}^\ell}\left[(z,h,w) \in \mathrm{supp}(G(x)) \mid \mathcal{E}(z)\right] + \Pr_{z \xleftarrow{R} S_n}\left[\neg\mathcal{E}(z)\right].$$

We will show that both the first and second summand are negligible in n.

Claim 2. $\Pr_{z \xleftarrow{R} S_n, h \xleftarrow{R} \mathcal{H}, w \xleftarrow{R} \{0,1\}^\ell}\left[(z,h,w) \in \mathrm{supp}(G(x)) \mid \mathcal{E}(z)\right] \leq 2^{-n}$.

Proof. By definition $\mathrm{supp}(G(x)) = \{(J_x(r), h, h(r)) \mid r \in \{0,1\}^{m'}, h \in \mathcal{H}\}$. Therefore, for any fixed z and h the probability, over $w \xleftarrow{R} \{0,1\}^\ell$, that the triple (z,h,w) lands in $\mathrm{supp}(G(x))$ is exactly

$$\frac{|h(J_x^{-1}(z))|}{2^\ell} \leq \frac{|J_x^{-1}(z)|}{2^\ell},$$

which is upper-bounded by $2^{\ell-n}/2^\ell = 2^{-n}$ when we condition on $\mathcal{E}(z)$. □

We conclude the proof by showing that for $z \xleftarrow{R} S_n$ the event $\mathcal{E}(z)$ happens almost surely.

Claim 3. $\Pr_{z \xleftarrow{R} S_n}\left[\log|J_x^{-1}(z)| \leq \ell - n\right] \geq 1 - 2^{-\Omega(n)}$.

Proof. Call z *good* if

$$z \text{ is } (\sqrt{n}\Delta)\text{-typical}, \qquad \text{where } \Delta = \sqrt{q}m, \tag{7}$$

and

$$\Pr[S_n = z] \geq 2^{q/10}\Pr[J_x = z]. \tag{8}$$

We begin by showing that, except with probability $2^{-\Omega(n)}$, a random $z \xleftarrow{R} S_n$ is good. First, recall that $\mathsf{Sim}_F(1^n)$ uses at most m random bits, and so any element in its support has weight at least 2^{-m}. Hence, due to the Flattening Lemma 1,

the distribution S_n is Δ-flat for $\Delta = \sqrt{q}m$ which implies that a random $z \xleftarrow{R} S_n$ satisfies (7) with probability at least $1-2^{-\Omega(n)}$. Next, we show that, except with probability $2^{-\Omega(n)}$, a random $z \xleftarrow{R} S_n$ satisfies (8). Indeed, due to the correctness property of F, we have that $\Delta(\mathsf{Sim}_F(1^n); F(x)) \geq 1/2$ which implies (by Fact 1 item 3) that $\Delta(S_n, J_x) \geq 1 - 2\exp(-q/8)$. Applying Fact 1 item 5, we conclude that

$$\Pr_{z \xleftarrow{R} S_n}[\Pr[S_n = z] < t \Pr[J_x = z]] \leq t \cdot 2\exp(-q/8),$$

for any $t \geq 1$. Taking $t := 2^{q/10}$, and noting that

$$t \cdot 2\exp(-q/8) \leq 2t \cdot 2^{-q/8} = 2 \cdot 2^{q/10} \cdot 2^{-q/8} = 2^{-q/40+1} = 2^{-\Omega(n)},$$

we conclude that (8) holds for all but $2^{-\Omega(n)}$-fraction of the $z \xleftarrow{R} S_n$. It follows, by a union-bound, that, except with probability $2^{-\Omega(n)}$, a random $z \xleftarrow{R} S_n$ is good.

Finally, we prove that for any good z it holds that $\log|J_x^{-1}(z)| \leq \ell - n$. By definition

$$|J_x^{-1}(z)| = 2^{m'} \cdot \Pr[J_x = z]$$

and by (8) the latter is upper-bounded by

$$2^{m'-q/10} \cdot \Pr[S_n = z].$$

Recalling that $\Pr[S_n = z] \leq 2^{-H(S_n)+\sqrt{n}\Delta}$ (since z is $\sqrt{n}\Delta$-typical) we conclude that

$$|J_x^{-1}(z)| \leq 2^{m'-q/10-H(S_n)+\sqrt{n}\Delta}.$$

Hence, we get that

$$\begin{aligned}\log|J_x^{-1}(z)| &\leq m' - H(S_n) + \sqrt{n}\Delta - q/10 \\ &\leq \underbrace{\lceil (m' - H(S_n) - \sqrt{n}\Delta - 2n) \rceil}_{=\ell} - n + \underbrace{(3n + 3\sqrt{n}\Delta - q/10)}_{T}.\end{aligned}$$

Since $q = 10^6 nm^2$ the expression T is always negative, and the claim follows. □

This completes the proof of Theorem 6. □

Now we show that if we repeat a D1RE polynomially many times we preserve the privacy of the encoding on yes-instances and gain the correctness security property of 1RE.

Theorem 7. *Let Π be a promise problem that has an ε-private and ρ-disjoint* D1RE *F, where ε and ρ are negligible. Then, there exists G a* 1RE *for Π that is ε'-private and δ-correct, where ε' and δ are negligible.*

Proof. For an instance x of length n, we define a randomized encoding $G(x)$ to be $\otimes^n F(x)$. Since F is efficient, the encoding G is also efficient. We prove that G is a 1RE for Π.

PRIVACY FOR YES-INSTANCES: Let Sim_F be the simulator showing the privacy of F on yes-instances. Define $\mathsf{Sim}_G(1^n) := \otimes^n \mathsf{Sim}_F(1^n)$. Take any yes-instance x from Π. We have that

$$\Delta(\mathsf{Sim}_G(1^n); G(x)) = \Delta(\otimes^n \mathsf{Sim}_F(1^n); \otimes^n F(x)) \leq n \cdot \varepsilon(n),$$

where the last inequality holds due to Fact 1 item 3. Since $\varepsilon(n)$ is negligible, we have that $\varepsilon'(n) := n \cdot \varepsilon(n)$ is also negligible.

CORRECTNESS: Let $Z = \bigcup_{x \in \Pi_{\text{NO}}} \text{supp}(G(x))$. The decoder Dec on input s outputs 0 if $s \in Z$; and outputs 1, otherwise. Clearly, a no-instance is always decoded correctly. For a yes-instance x, we upper-bound the decoding error by showing that $\Pr[G(x) \in Z]$ is negligible. Since G is ε'-private on yes-instances, we have that

$$\Pr[G(x) \in Z] \leq \Pr[\mathsf{Sim}_G(1^n) \in Z] + \varepsilon'(n).$$

By ρ-disjointness, it holds that $\Pr[\mathsf{Sim}_F(1^n) \in \text{supp}(F(x_{\text{NO}}))] \leq \rho(n)$, for any no-instance x_{NO}. This implies that if we repeat this experiment n times we get that $\Pr[\mathsf{Sim}_G(1^n) \in \text{supp}(G(x_{\text{NO}}))] \leq \rho(n)^n$. By a union bound, we conclude that $\Pr[\mathsf{Sim}_G(1^n) \in Z] \leq 2^n \rho(n)^n$, which implies that

$$\Pr[G(x) \in Z] \leq 2^n \rho(n)^n + \varepsilon'(n) \leq \text{neg}(n).$$

The theorem follows. □

5.2 From NISZK$^{\text{pub}}$ to 1RE

In this section we prove that $\mathcal{NISZK}^{\text{PUB}} \subseteq \mathcal{1RE}$.

Theorem 8. $\mathcal{NISZK}^{\text{PUB}} \subseteq \mathcal{1RE}$.

Proof. Let Π be a promise problem with (α, β, γ)-$\mathsf{NISZK}^{\text{PUB}}$ proof system consisting of $(\mathsf{Prov}, \mathsf{Ver}, \mathsf{Deal}, \mathsf{Sim}_{\text{ZK}})$, where α, β, γ are negligible. By Theorem 7, it suffices to show that Π has a (ε, ρ)-D1RE Enc for some negligible ε and ρ. For an n-bit string x, we define a randomized encoding $\mathsf{Enc}(x)$ as follows:

1. Compute $(\sigma, p) = \mathsf{Sim}_{\text{ZK}}(x)$.
2. Compute the bit $b = \mathsf{Ver}(x, \sigma, p)$.
3. If $b = 1$ output σ, otherwise output a fixed string $z_n \notin \text{supp}(\mathsf{Deal}(1^n))$.[9]

Observe that Enc is efficient because Sim_{ZK} and Ver are efficient. We prove that Enc is a D1RE.

PRIVACY: We define $\mathsf{Sim}_{\text{YES}}(1^n) = \mathsf{Deal}(1^n)$ and prove that for any yes-instance x the distribution $\mathsf{Sim}_{\text{YES}}(1^n)$ is $\varepsilon(n)$-close to $\mathsf{Enc}(x)$ where $\varepsilon(n) = \alpha(n) + 2 \cdot \gamma(n) =$

[9] For example, such a $z(n)$ can be efficiently constructed by appending a trailing 1 to the output of $\mathsf{Deal}(1^n)$ and setting $z(n)$ to the all-zero string.

$\mathrm{neg}(n)$. Fix some yes-instance x of length n. Due to the zero-knowledge property of NISZK, we have that

$$\Delta_{\sigma \xleftarrow{R} \mathsf{Deal}(1^n)}(\mathsf{Sim}_{\mathrm{ZK}}(x), (\sigma, \mathsf{Prov}(x, \sigma))) \leq \gamma(n).$$

By the definition of the statistical distance, this implies that

$$\left| \Pr_{\sigma \xleftarrow{R} \mathsf{Deal}(n)}[\mathsf{Ver}(\sigma, x, \mathsf{Prov}(x, \sigma)) \neq 1] - \Pr_{(\sigma,p) \xleftarrow{R} \mathsf{Sim}_{\mathrm{ZK}}(x)}[\mathsf{Ver}(\sigma, x, p) \neq 1] \right| \leq \gamma(n).$$

Because of the correctness property of NISZK, we have that

$$\Pr_{\sigma \xleftarrow{R} \mathsf{Deal}(n)}[\mathsf{Ver}(\sigma, x, \mathsf{Prov}(x, \sigma)) \neq 1] \leq \alpha(n).$$

This implies that

$$\Pr_{(\sigma,p) \xleftarrow{R} \mathsf{Sim}_{\mathrm{ZK}}(x)}[\mathsf{Ver}(\sigma, x, p) \neq 1] \leq \alpha(n) + \gamma(n).$$

The latter inequality means that in the execution of $\mathsf{Enc}(x)$ the bit b equals to 1 except with the probability $\alpha(n) + \gamma(n)$. Hence, $\Delta(\mathsf{Enc}(x); \mathsf{Sim}_{\mathrm{ZK}}(x)[1]) \leq \alpha(n) + \gamma(n)$, where $\mathsf{Sim}_{\mathrm{ZK}}(x)[1]$ denotes the first component of the tuple output by the simulator. Because of the zero-knowledge property of NISZK and due to Fact 1 item 1, we have that $\Delta(\mathsf{Sim}_{\mathrm{ZK}}(x)[1]; \mathsf{Deal}(1^n)) \leq \gamma(n)$. Finally, combining the last two inequalities, we get that

$$\Delta(\mathsf{Enc}(x); \mathsf{Deal}(1^n)) \leq \alpha(n) + 2 \cdot \gamma(n) = \mathrm{neg}(n).$$

Disjointness: Let x be a no-instance of Π. Let $E \subseteq \mathrm{supp}(\mathsf{Deal}(1^n))$ denote the set of the strings admitting a proof for the no-instance x, i.e., $E := \{\sigma \in \mathrm{supp}(\mathsf{Deal}(1^n)) \mid \exists p : \mathsf{Ver}(\sigma, x, p) = 1\}$. By Enc's construction we have that $\mathrm{supp}(\mathsf{Enc}(x)) \subseteq E \cup \{z_n\}$. This implies that

$$\begin{aligned} \Pr[\mathsf{Deal}(1^n) \in \mathrm{supp}(\mathsf{Enc}(x))] &\leq \Pr[\mathsf{Deal}(1^n) \in E \cup \{z_n\}] \\ &\overset{(\star)}{=} \Pr[\mathsf{Deal}(1^n) \in E] \\ &\leq \beta(n), \end{aligned}$$

where the last inequality follows from the soundness property of NISZK, and the equality $(\star)$ holds because $z_n \notin \mathrm{supp}(\mathsf{Deal}(1^n))$.

□

5.3 From 1RE to $\mathsf{NISZK}^{\mathrm{pub}}$

Theorem 9. *If the promise problem Π has a (possibly non-uniform) $\mathsf{1RE}$ F, then it also has a non-uniform $\mathsf{NISZK}^{\mathrm{PUB}}$ proof system. Moreover, if F is uniform then the $\mathsf{NISZK}^{\mathrm{PUB}}$ proof system can be implemented based on F and an advice of $O(\log n)$ bits.*

Proof. Let $\Pi \in 1\mathcal{RE}$. Due to Theorem 6, there exists a non-uniform (ε, ρ)-D1RE Enc for Π such that ε and ρ are negligible. Let $s(n)$ denote the randomness complexity of the encoding Enc when it is applied to an n-bit input x, and let Sim_{RE} be the simulator showing the privacy of Enc on yes-instances. We construct a proof system $(\mathsf{Prov}, \mathsf{Ver}, \mathsf{Deal}, \mathsf{Sim}_{\text{ZK}})$ for Π as follows:

- Deal: Given 1^n, the dealer outputs $\mathsf{Sim}_{\text{RE}}(1^n)$.
- Prov: Given an n-bit input x and a string σ from Deal, the prover samples a random $r \in \{0,1\}^{s(n)}$ subject to $\mathsf{Enc}(x, r) = \sigma$, and sends it to the verifier. If no such r exists the prover sends some arbitrary message.
- Ver: Given (x, σ, r), the verifier accepts if $\mathsf{Enc}(x, r) = \sigma$, and other rejects.
- Sim_{ZK}: Given x, the simulator Sim_{ZK} samples a random r and outputs the pair $(\mathsf{Enc}(x, r), r)$.

We show that $(\mathsf{Prov}, \mathsf{Ver}, \mathsf{Deal}, \mathsf{Sim}_{\text{ZK}})$ forms a NISZK for Π.

Completeness: Consider some yes-instance x of length n. Recall that, by the privacy of $\mathsf{D1RE}$, the simulator's distribution $\mathsf{Sim}_{\text{RE}}(1^n)$ is $\varepsilon(n)$-close to $\mathsf{Enc}(x)$, which implies that

$$\Pr[\mathsf{Sim}_{\text{RE}}(1^n) \in \text{supp}(\mathsf{Enc}(x))] \geq 1 - \varepsilon(n).$$

Hence, except with probability $\varepsilon(n)$, for a string σ generated by $\mathsf{Sim}_{\text{RE}}(1^n)$, the prover Prov can find r, such that $\mathsf{Enc}(x, r) = \sigma$.

Soundness: For all no-instances x of Π, we have that

$$\Pr_{\sigma \xleftarrow{R} \mathsf{Deal}(1^n)}[\exists p : V(x, \sigma, p) = 1] = \Pr_{\sigma \xleftarrow{R} \mathsf{Sim}_{\text{RE}}(1^n)}[\sigma \in \text{supp}(\mathsf{Enc}(x))] \leq \delta(n),$$

where the last inequality follows from the disjointness property of Enc.

Zero Knowledge: For all yes-instances x of Π, we have that

$$\begin{aligned}
&\Delta_{\sigma \xleftarrow{R} \mathsf{Deal}(1^n)}(\mathsf{Sim}_{\text{ZK}}(x) \ ; \ (\sigma, \mathsf{Prov}(x, \sigma))) = \\
&\Delta_{\sigma \xleftarrow{R} \mathsf{Sim}_{\text{RE}}(1^n), r \xleftarrow{R} \{0,1\}^{s(n)}}((\mathsf{Enc}(x, r), r) \ ; \ (\sigma, \mathsf{Prov}(x, \sigma))) = \\
&\Delta_{\sigma \xleftarrow{R} \mathsf{Sim}_{\text{RE}}(1^n), r \xleftarrow{R} \{0,1\}^{s(n)}}((\mathsf{Enc}(x, r), \mathsf{Prov}(x, \mathsf{Enc}(x, r))) \ ; \ (\sigma, \mathsf{Prov}(x, \sigma))) \leq \\
&\Delta_{\sigma \xleftarrow{R} \mathsf{Sim}_{\text{RE}}(1^n), r \xleftarrow{R} \{0,1\}^{s(n)}}(\mathsf{Enc}(x, r); \sigma) \leq \\
&\varepsilon(n),
\end{aligned}$$

where the second equality follows by recalling that $\mathsf{Prov}(\sigma)$ samples a random r subject to $\mathsf{Enc}(x, r) = \sigma$ and so $(\mathsf{Enc}(x, r), r)$ is identically distributed to $(\mathsf{Enc}(x, r), \mathsf{Prov}(x, \mathsf{Enc}(x, r)))$, and the first inequality follows from Fact 1 item 1.

□

6 If $\mathcal{SRE}$ is non-trivial then one-way functions exist

In this section we prove Theorem 3:

Theorem 3 (Restated). *If $\mathcal{SRE}$ is not in $\mathcal{BPP}$, then infinitely-often one-way functions exist.*

Proof. Assume that infinitely-often one-way functions do not exist. Impagliazzo and Luby [IL89] showed that in this case every efficiently computable function $g(x)$ can be "distributionally-inverted" in the following sense: For every inverse polynomial $\alpha(\cdot)$, there exists an efficient adversary A such that, for random $x \in \{0,1\}^n$, the pair $(x, g(x))$ is $\alpha(n)$-close to the pair $(A(g(x)), g(x))$. In other words, for most x's, A finds an almost uniform preimage of $g(x)$. We refer to α as the *deviation* of the inverter and set it to $1/10$.

We will show that such an inverter allows to put $\mathcal{SRE}$ in $\mathcal{BPP}$. Let Π be a promise problem in $\mathcal{SRE}$ with ε-private δ-correct statistical encoding Enc for some negligible ε and δ. Let $\mathsf{Sim}_{\text{YES}}$ and Sim_{NO} be the simulators of the encoding and define $\mathsf{Sim}(b, r)$ to be a "joint" simulator which takes as an input a single bit $b \in \{0,1\}$ and random string r and outputs a sample from $\mathsf{Sim}_{\text{YES}}(r)$ if $b = 1$ and from $\mathsf{Sim}_{\text{NO}}(r)$ if $b = 0$.[10] We decide Π via the following $\mathcal{BPP}$ procedure B: Given a string $x \in \{0,1\}^n$, sample an encoding $y \xleftarrow{R} \mathsf{Enc}(x)$ and α-distributionally invert the simulator Sim on the string y. Take the resulting preimage (b, r) (where r is the coins of the simulator) and output the bit b. We analyze the success probability of deciding Π with this procedure.

Claim 4. The procedure B decides Π with error probability of at most $1/6 + 5\delta + \varepsilon + \alpha$.

Proof. Let us focus on the case where $x \in \{0,1\}^n$ is a yes-instance (the other case is symmetric). First consider an "ideal" version B' of the algorithm B in which (1) the string y is sampled from $\mathsf{Sim}_{\text{YES}}(r)$ and (2) the distributional inversion algorithm is perfect and has zero deviation. Observe that the gap between the error probability of the real algorithm B to the error probability of the ideal algorithm B' is at most $\varepsilon + \alpha$ (this is due to ε-privacy and to α-deviation of the actual inverter). Hence, it suffices to show that the ideal version errs with probability of at most $1/6 + 5\delta$.

For a given encoding y, the ideal algorithm outputs the right answer $b = 1$ with probability $\frac{p_1(y)}{p_0(y)+p_1(y)}$ where $p_0(y)$ denotes the weight of y under the distribution sampled by Sim_{NO} and $p_1(y)$ denotes the weight of y under $\mathsf{Sim}_{\text{YES}}$. By the δ-correctness of the encoding and by Fact 1 item 5 (instantiated with

[10] We omit the unary input 1^n of the simulators, and assume that the randomness complexity $m(n)$ of the simulators uniquely determines the instance length n. Similarly, we assume that the output of $\mathsf{Sim}(b, r)$ uniquely determines n. Both requirements can be achieved without loss of generality via standard padding conventions. (E.g., pad the randomness r and concatenate the input length 1^n to the encoding and to the output of Sim.)

$t = 5$), it holds that, except with probability at most 5δ over $y \xleftarrow{R} \mathsf{Sim}_{\text{YES}}$, we have that $p_1(y) \geq 5p_0(y)$. It follows, by a union bound, that the ideal algorithm errs with probability of at most $5\delta + 1/6$, as required. □

It remains to notice, that since δ and ε are negligible and α is an inverse polynomial, we have that Π can be decided with success probability at least $2/3$. □

7 If $\mathcal{PRE}$ is hard on the average then CRH exist

In this section we will study the consequences of the existence of an average-case hard problem $\Pi \in \mathcal{PRE}$.

Definition 5. *We say that a promise problem $\Pi = (\Pi_{\text{YES}}, \Pi_{\text{NO}})$ is* hard on average *if there exists an efficient sampler S that given 1^n outputs an n-bit instance of Π such that for every non-uniform efficient algorithm A,*

$$\left| \Pr_{x \xleftarrow{R} S(1^n)} [A(x) = \chi_{\Pi}(x)] - 1/2 \right| < \text{neg}(n).$$

We say that the problem has efficient Yes/No samplers *if it is possible to efficiently sample from the conditional Yes distribution $Y_n = [S(1^n)|S(1^n) \in \Pi_{\text{YES}}]$ and from the conditional No distribution $N_n = [S(1^n)|S(1^n) \in \Pi_{\text{NO}}]$.*

A collection of *claw-free pseudo-permutations* (CFPP) [Dam87,GMR88,Rus95] is a set of pairs of efficiently computable functions $f^0, f^1 : \{0,1\}^n \to \{0,1\}^n$ for which it is hard to find a pair (u, v) which forms a *claw*, i.e., $f^0(u) = f^1(v)$, or a *collapse*, i.e., $f^b(u) = f^b(v)$ and $u \neq v$ for some bit b. Collections of claw-free *permutations* (CFPs) correspond to the special case where f_0 and f_1 are permutations and so collapses simply do not exist.

Definition 6 (claw-free functions). *A* collection of pairs of functions *consists of an infinite set of indices, denoted $\overline{I}$, finite sets D_i for each $i \in \overline{I}$, and two functions f_i^0 and f_i^1 mapping D_i to D_i, respectively. Such a collection is called a* claw-free pseudo-permutations *if there exist three probabilistic polynomial-time algorithms I, D, and F such that the following conditions hold:*

Easy to sample and compute: *The random variable $I(1^n)$ is assigned values in the set $\overline{I} \cap \{0,1\}^{p(n)}$ for some polynomial $p(\cdot)$. For each $i \in \overline{I}$, the random variable $D(i)$ is distributed uniformly over D_i. For each $i \in \overline{I}$, $b \in \{0,1\}$ and $x \in D_i$, $F(b, i, x) = f_i^b(x)$.*

Hard to form claws: *A pair (x, y) satisfying $f_i^0(x) = f_i^1(y)$ is called a* claw *for index i. Let C_i denote the set of claws for index i. It is required that for every probabilistic polynomial-time algorithm A,*

$$\Pr_{i \xleftarrow{R} I(1^n)} [A(i) \in C_i] < \text{neg}(n).$$

HARD TO FORM COLLAPSES: *A pair (x, y) satisfying $f_i^b(x) = f_i^b(y)$ is called a* collapse *for an index i and a bit b. Let $T_{i,b}$ denote the set of collapses for (i, b). It is required that for every probabilistic polynomial-time algorithm A and every $b \in \{0,1\}$,*

$$\Pr_{i \xleftarrow{R} I(1^n)}[A(i) \in T_{i,b}] < \mathrm{neg}(n).$$

If the last item holds for unbounded adversaries, i.e., f_i^0 and f_i^1 are permutations over D_i, then the collection is called a collection of claw-free permutations.

It is known that CFPP's imply Collision-Resistant Hash functions (CRH) [Rus95]. We will show that the existence of an average-case hard problem $\Pi \in \mathcal{PRE}$ implies the existence of CFPPs. We begin with the simpler case in which Π has an *efficient Yes/No samplers* and show that, in this case, we obtain a collection of claw-free permutations.

Theorem 10. *If there exists an average-case hard language in $\mathcal{PRE}$ with efficient Yes/No samplers then CFPs exist.*

We will need the following simple claim.

Claim 5. Let Π be a promise problem with perfect randomized encoding g whose simulators are $\mathsf{Sim}_{\text{YES}}$ and Sim_{NO}. Define the functions h_x^0, h_y^1 which are indexed by a pair of instances (x, y) of Π as follows:

$$h_x^0(r,b) = \begin{cases} g(x;r), \text{ if } b = 0, \\ \mathsf{Sim}_{\text{NO}}(r), \text{ otherwise;} \end{cases} \qquad h_y^1(r,b) = \begin{cases} g(y;r), \text{ if } b = 0, \\ \mathsf{Sim}_{\text{YES}}(r), \text{ otherwise;} \end{cases} \tag{9}$$

Then the following holds for any n-bit strings x and y:

1. If $x \in \Pi_{\text{YES}}$, then h_x^0 is a permutation.
2. If $y \in \Pi_{\text{NO}}$, then h_y^1 is a permutation.
3. If $(x, y) \in \Pi_{\text{NO}} \times \Pi_{\text{YES}}$ then $\mathrm{Im}\left(h_x^0\right) \cap \mathrm{Im}\left(h_y^1\right) = \emptyset$.

Proof. Let R_0 and R_1 denote $\mathrm{Im}(\mathsf{Sim}_{\text{NO}})$ and $\mathrm{Im}(\mathsf{Sim}_{\text{YES}})$, respectively. Let $s(n)$ denote the randomness complexity of g and let $t(n)$ denote the output length of g. Since g is a perfect randomized encoding, we have that $R_0 \cap R_1 = \emptyset$, $R_0 \cup R_1 = \{0,1\}^{t(n)}$, and $t(n) = s(n) + 1$. Consider the case where $x \in \Pi_{\text{YES}}$. Then $h_x^0(\cdot, 0) : \{0,1\}^{s(n)} \to R_1$ is a bijection and $h_x^0(\cdot, 1) : \{0,1\}^{s(n)} \to R_0$. Since $R_0 \cap R_1 = \emptyset$, the function $h_x^0(\cdot,\cdot)$ is a permutation on $R_0 \cup R_1 = \{0,1\}^{t(n)}$. Similarly, if $y \in \Pi_{\text{NO}}$, the function $h_y^1(\cdot,\cdot)$ is a permutation on $\{0,1\}^{t(n)}$.

In order to prove the third item, we observe that if $x \in \Pi_{\text{NO}}$, then $\mathrm{Im}\left(h_x^0\right) = R_0$; and if $y \in \Pi_{\text{YES}}$, then $\mathrm{Im}\left(h_y^1\right) = R_1$. This implies that for all $(x, y) \in \Pi_{\text{NO}} \times \Pi_{\text{YES}}$ it holds that $\mathrm{Im}\left(h_x^0\right) \cap \mathrm{Im}\left(h_y^1\right) = R_0 \cap R_1 = \emptyset$. □

We can now prove Theorem 10.

Proof (Proof of Theorem 10). Let Π be an average-case hard language with efficient Yes/No samplers (Y_n, N_n), and let g be a perfect randomized encoding for Π. For a pair of inputs (x, y) from Π, we say that (x, y) is a (YES, NO)-instance (resp., (NO, YES)), if x is a yes-instance and y is a no-instance (resp., if x is a no-instance and y is a yes-instance).

We construct a CFP family which is indexed by pairs $(x, y) \in \Pi_{\text{YES}} \times \Pi_{\text{NO}}$. Given a security parameter 1^n, an index (x, y) is chosen by sampling $x \xleftarrow{R} Y_n$ and $y \xleftarrow{R} N_n$. For each index (x, y) we let $f^0_{(x,y)} \equiv h^0_x$ and $f^1_{(x,y)} \equiv h^1_y$, where h^0_x and h^1_x are defined as in (9). Recall that the domain and range of $f^b_{x,y}$ are $\{0,1\}^{t(n)}$ where $t(n)$ is the output length of g's output. Clearly this collection is efficiently samplable and efficiently computable. Moreover, since our sampler always samples a (YES, NO)-instance (x, y), it holds, due to Claim 5, that $f^0_{(x,y)} \equiv h^0_x$ and $f^1_{(x,y)} \equiv h^1_y$ are permutations on $\{0,1\}^{t(n)}$. We complete the proof by showing that claws are hard to find.

Recall that we assume that the distribution ensemble $\{Y_n\}$ is computationally indistinguishable from $\{N_n\}$. By a standard hybrid argument, it follows that the pair (Y_n, N_n) is computationally indistinguishable from the pair (Y_n, Y_n) which, in turn, is computationally indistinguishable from the pair (N_n, Y_n). Now assume, for the sake of contradiction, that there exists an efficient algorithm A that given $(x, y) \xleftarrow{R} (Y_n, N_n)$ can find claws with non-negligible probability ε. We can use A to distinguish (Y_n, N_n) from (N_n, Y_n) as follows: Given (x, y) call $A(x, y)$ and output 1 if A's output (u, v) forms a collision under h^0_x and h^1_y. By assumption, the resulting distinguisher outputs 1 when $(x, y) \xleftarrow{R} (Y_n, N_n)$ with probability ε. In contrast, when $(x, y) \xleftarrow{R} (N_n, Y_n)$, the distinguisher never finds a claw since claws do not exist (due to Claim 5). Hence the distinguisher has a noticeable advantage of ε, in contradiction to our assumption. □

We continue by considering the more general case where Π is hard on average but does not admit efficient Yes/No samplers, and obtain, in this case, claw-free *pseudo-permutations* (whose existence still implies collision-resistance hash functions).

Theorem 11. *If there exists an average-case hard language in $\mathcal{PRE}$ then claw-free pseudo-permutations (CFPP) exist.*

Proof. The construction is identical to the one presented in Theorem 10, except that the index $(x, y) \in \Pi \times \Pi$ is chosen by sampling both x and y independently from the distribution $S(1^n)$ over which Π is average-case hard. By definition, the collection $f^b_{(x,y)} = h^b_x$, where h is defined as in (9), is efficiently samplable and efficiently computable. We verify that it is CFPP.

We begin by showing that $f^0_{(x,y)} = h^0_x$ is a pseudo-permutation (the case of $f^1_{(x,y)}$ is analogous). Assume for the sake of contradiction that there is an algorithm A that can find collapses for $f^0_{(x,y)}$ with a non-negligible probability ε. Using A we construct a new algorithm A' that has a non-negligible advantage

in guessing $\chi_{\Pi}(x)$ for $x \xleftarrow{R} S(1^n)$. Given an input $x \xleftarrow{R} S(1^n)$, the algorithm A' samples $y \leftarrow S(1^n)$, and then invokes $A(x,y)$ to find a collapse (u,v) for $f^0_{(x,y)} = h^0_x$. If A finds a valid collapse (i.e., $u \neq v$ and $h^0_x(u) = h^0_x(v)$), the algorithm A' classifies the input x as a no-instance and outputs 0; otherwise A' outputs a random bit. Recall that when x is a yes-instance the function h^0_x is a permutation, and so it does not have collapses. Hence, A' outputs a correct answer whenever A finds a collapse. Also, when a collapse is not found, the success probability of A' is $1/2$. Hence, the overall success probability of A' is

$$\Pr_{x \xleftarrow{R} S(1^n)}[A'(x) = \chi_{\Pi}(x)] = 1/2 \cdot (1-\varepsilon) + 1 \cdot \varepsilon = 1/2 + \varepsilon/2,$$

in contradiction to the average-case hardness of Π.

We move on to show that it is hard to find claws. Assume for the sake of contradiction that there exists an efficient algorithm A that finds claws with a non-negligible probability ε. We construct a new algorithm A' that has a non-negligible advantage in guessing $\chi_{\Pi}(x)$ for $x \xleftarrow{R} S(1^n)$. Let

$$p = \Pr_{x \xleftarrow{R} S(1^n), y \xleftarrow{R} S(1^n)}[A(x,y) \text{ finds a claw } | x \in \Pi_{\text{NO}}].$$

We distinguish between two cases based on the value of p.

First, consider the case where $p \geq \varepsilon/2$. Then, by an averaging argument, there exists some fixed no-instance x_0 for which

$$\Pr_y[A(x_0,y) \text{ finds a claw }] \geq \varepsilon/2.$$

Recall that when the index is a (NO, YES) pair there are no claws and so when A finds a claw, y must be a no-instance We can therefore construct a non-uniform algorithm that decides $y \xleftarrow{R} S(1^n)$ as follows: Call $A(x_0,y)$ and output zero ("no") if a collision is found and otherwise toss a random coin. The success probability is at least $\varepsilon/2 + (1-\varepsilon/2)/2 = 1/2 + \varepsilon/4$.

Second, consider the case where $p < \varepsilon/2$. In this case, we determine whether $x \xleftarrow{R} S(1^n)$ is a yes-instance or a no-instance via the following procedure A'. Sample $y \xleftarrow{R} S(1^n)$, and call $A(x,y)$ if A returns a valid claw, outputs 1 (classify x as a yes-instance); otherwise, output a random bit. The success probability of A' can be marginalized as follows:

$$\begin{aligned}\Pr_x[A'(x) \text{ succeeds }] &= \Pr_{x,y}[A'(x) \text{ succeeds } | A(x,y) \text{ finds a claw }] \cdot \varepsilon \\ &\quad + \Pr_{x,y}[A'(x) \text{ succeeds } | A(x,y) \text{ doesn't find a claw }] \cdot (1-\varepsilon) \\ &= \Pr_x[x \in \Pi_{\text{YES}} | A(x,y) \text{ finds a claw }] \cdot \varepsilon + (1-\varepsilon)/2,\end{aligned}$$

Therefore, it suffices to show that

$$\Pr_x[x \in \Pi_{\text{YES}} | A(x,y) \text{ finds a claw }] \geq 2/3 \tag{10}$$

since this implies that A' succeeds with probability of at least $2/3 \cdot \varepsilon + (1 - \varepsilon)/2 = 1/2 + \varepsilon/6$. To prove (10), we upper-bound by $1/3$ the probability of the complementary event:

$$\begin{aligned}
&\Pr_x[x \in \Pi_{\text{NO}} | A(x,y) \text{ finds a claw}] = \\
&\frac{\Pr_{x,y}[A(x,y) \text{ finds a claw} | x \in \Pi_{\text{NO}}] \cdot \Pr_x[x \in \Pi_{\text{NO}}]}{\Pr[A(x,y) \text{ finds a claw }]} \leq \\
&\frac{(\varepsilon/2) \cdot (2/3)}{\varepsilon} = \\
&\frac{1}{3},
\end{aligned}$$

where the inequality follows by our assumption ($p < \varepsilon/2$) and by the fact that $\Pr_x[x \in \Pi_{\text{NO}}] < 2/3$ (since otherwise the trivial adversary that always outputs 0 breaks the average-case hardness of Π over $S(1^n)$). The proof follows. □

Acknowledgements

Research supported by the European Union's Horizon 2020 Programme (ERC-StG-2014-2020) under grant agreement no. 639813 ERC-CLC, ISF grant 1155/11, GIF grant 1152/2011, and the Check Point Institute for Information Security. This work was done in part while the first author was visiting the Simons Institute for the Theory of Computing, supported by the Simons Foundation and by the DIMACS/Simons Collaboration in Cryptography through NSF grant CNS-1523467.

References

[AH87] W. Aiello and J. Håstad. Perfect zero-knowledge languages can be recognized in two rounds. In *28th Annual Symposium on Foundations of Computer Science, Los Angeles, California, USA, 27-29 October 1987*, pages 439–448. IEEE Computer Society, 1987.

[AIK04] B. Applebaum, Y. Ishai, and E. Kushilevitz. Cryptography in NC^0. In *45th Symposium on Foundations of Computer Science (FOCS 2004), 17-19 October 2004, Rome, Italy, Proceedings*, pages 166–175. IEEE Computer Society, 2004.

[AIK05] B. Applebaum, Y. Ishai, and E. Kushilevitz. Computationally private randomizing polynomials and their applications. In *20th Annual IEEE Conference on Computational Complexity (CCC 2005), 11-15 June 2005, San Jose, CA, USA*, pages 260–274. IEEE Computer Society, 2005.

[AIK15] B. Applebaum, Y. Ishai, and E. Kushilevitz. Minimizing locality of one-way functions via semi-private randomized encodings. *Electronic Colloquium on Computational Complexity (ECCC)*, 22:45, 2015.

[AIKP15] S. Agrawal, Y. Ishai, D. Khurana, and A. Paskin-Cherniavsky. Statistical randomized encodings: A complexity theoretic view. In M. M. Halldórsson, K. Iwama, N. Kobayashi, and B. Speckmann, editors, *Automata, Languages,*

and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part I, volume 9134 of *Lecture Notes in Computer Science*, pages 1–13. Springer, 2015.

[App11] B. Applebaum. Randomly encoding functions: A new cryptographic paradigm - (invited talk). In S. Fehr, editor, *Information Theoretic Security - 5th International Conference, ICITS 2011, Amsterdam, The Netherlands, May 21-24, 2011. Proceedings*, volume 6673 of *Lecture Notes in Computer Science*, pages 25–31. Springer, 2011.

[App14] B. Applebaum. *Cryptography in Constant Parallel Time*. Information Security and Cryptography. Springer, 2014.

[AR16] B. Applebaum and P. Raykov. From private simultaneous messages to zero-information arthur-merlin protocols and back. To appear in TCC 2016A, 2016. Available as eprint report 2015/1046 at `http://eprint.iacr.org/2015/1046`.

[BDLP88] J. Brandt, I. Damgård, P. Landrock, and T. P. Pedersen. Zero-knowledge authentication scheme with secret key exchange (extended abstract). In S. Goldwasser, editor, *Advances in Cryptology - CRYPTO '88, 8th Annual International Cryptology Conference, Santa Barbara, California, USA, August 21-25, 1988, Proceedings*, volume 403 of *Lecture Notes in Computer Science*, pages 583–588. Springer, 1988.

[BFM88] M. Blum, P. Feldman, and S. Micali. Non-interactive zero-knowledge and its applications (extended abstract). In J. Simon, editor, *Proceedings of the 20th Annual ACM Symposium on Theory of Computing, May 2-4, 1988, Chicago, Illinois, USA*, pages 103–112. ACM, 1988.

[BSMP91] M. Blum, A. D. Santis, S. Micali, and G. Persiano. Noninteractive zero-knowledge. *SIAM J. Comput.*, 20(6):1084–1118, 1991.

[CCG+94] R. Chang, B. Chor, O. Goldreich, J. Hartmanis, J. Håstad, D. Ranjan, and P. Rohatgi. The random oracle hypothesis is false. *J. Comput. Syst. Sci.*, 49(1):24–39, 1994.

[CCKV08] A. Chailloux, D. F. Ciocan, I. Kerenidis, and S. P. Vadhan. Interactive and noninteractive zero knowledge are equivalent in the help model. In R. Canetti, editor, *Theory of Cryptography, Fifth Theory of Cryptography Conference, TCC 2008, New York, USA, March 19-21, 2008.*, volume 4948 of *Lecture Notes in Computer Science*, pages 501–534. Springer, 2008.

[CW79] L. Carter and M. N. Wegman. Universal classes of hash functions. *Journal of Computer and System Sciences*, 18(2):143–154, 1979. Preliminary version appeared in STOC '77.

[Dam87] I. Damgård. Collision free hash functions and public key signature schemes. In D. Chaum and W. L. Price, editors, *Advances in Cryptology - EUROCRYPT '87, Workshop on the Theory and Application of of Cryptographic Techniques, Amsterdam, The Netherlands, April 13-15, 1987, Proceedings*, volume 304 of *Lecture Notes in Computer Science*, pages 203–216. Springer, 1987.

[Dam00] I. Damgård. Efficient concurrent zero-knowledge in the auxiliary string model. In B. Preneel, editor, *Advances in Cryptology - EUROCRYPT 2000, International Conference on the Theory and Application of Cryptographic Techniques, Bruges, Belgium, May 14-18, 2000, Proceeding*, volume 1807 of *Lecture Notes in Computer Science*, pages 418–430. Springer, 2000.

[ESY84] S. Even, A. L. Selman, and Y. Yacobi. The complexity of promise problems with applications to public-key cryptography. *Inf. Control*, 61(2):159–173, May 1984.

[FKN94] U. Feige, J. Kilian, and M. Naor. A minimal model for secure computation (extended abstract). In F. T. Leighton and M. T. Goodrich, editors, *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing, 23-25 May 1994, Montréal, Québec, Canada*, pages 554–563. ACM, 1994.

[GB00] D. Gutfreund and M. Ben-Or. Increasing the power of the dealer in non-interactive zero-knowledge proof systems. In T. Okamoto, editor, *Advances in Cryptology - ASIACRYPT 2000, 6th International Conference on the Theory and Application of Cryptology and Information Security, Kyoto, Japan, December 3-7, 2000, Proceedings*, volume 1976 of *Lecture Notes in Computer Science*, pages 429–442. Springer, 2000.

[GMR88] S. Goldwasser, S. Micali, and R. L. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM J. Comput.*, 17(2):281–308, 1988.

[GMR89] Goldwasser, Micali, and Rackoff. The knowledge complexity of interactive proof systems. *SICOMP: SIAM Journal on Computing*, 18, 1989.

[Gol01] O. Goldreich. *The Foundations of Cryptography - Volume 1, Basic Techniques*. Cambridge University Press, 2001.

[Gol06] O. Goldreich. On promise problems: A survey. In O. Goldreich, A. L. Rosenberg, and A. L. Selman, editors, *Theoretical Computer Science, Essays in Memory of Shimon Even*, volume 3895 of *Lecture Notes in Computer Science*, pages 254–290. Springer, 2006.

[GPW15] M. Göös, T. Pitassi, and T. Watson. Zero-information protocols and unambiguity in arthur-merlin communication. In T. Roughgarden, editor, *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science, ITCS 2015, Rehovot, Israel, January 11-13, 2015*, pages 113–122. ACM, 2015.

[GSV99] O. Goldreich, A. Sahai, and S. P. Vadhan. Can statistical zero knowledge be made non-interactive? or on the relationship of SZK and NISZK. In M. J. Wiener, editor, *Advances in Cryptology - CRYPTO '99, 19th Annual International Cryptology Conference, Santa Barbara, California, USA, August 15-19, 1999, Proceedings*, volume 1666 of *Lecture Notes in Computer Science*, pages 467–484. Springer, 1999.

[GV99] O. Goldreich and S. P. Vadhan. Comparing entropies in statistical zero knowledge with applications to the structure of SZK. In *Proceedings of the 14th Annual IEEE Conference on Computational Complexity, Atlanta, Georgia, USA, May 4-6, 1999*, page 54. IEEE Computer Society, 1999.

[HHRS15] I. Haitner, J. J. Hoch, O. Reingold, and G. Segev. Finding collisions in interactive protocols - tight lower bounds on the round and communication complexities of statistically hiding commitments. *SIAM J. Comput.*, 44(1):193–242, 2015.

[IK00] Y. Ishai and E. Kushilevitz. Randomizing polynomials: A new representation with applications to round-efficient secure computation. In *41st Annual Symposium on Foundations of Computer Science, FOCS 2000, 12-14 November 2000, Redondo Beach, California, USA*, pages 294–304. IEEE Computer Society, 2000.

[IL89] R. Impagliazzo and M. Luby. One-way functions are essential for complexity based cryptography (extended abstract). In *30th Annual Symposium on Foundations of Computer Science, Research Triangle Park, North Carolina, USA, 30 October - 1 November 1989*, pages 230–235. IEEE Computer Society, 1989.

[ILL89] R. Impagliazzo, L. A. Levin, and M. Luby. Pseudo-random generation from one-way functions (extended abstracts). In D. S. Johnson, editor, *Proceedings of the 21st Annual ACM Symposium on Theory of Computing, May 14-17, 1989, Seattle, Washigton, USA*, pages 12–24. ACM, 1989.

[Ish13] Y. Ishai. Randomization techniques for secure computation. In M. Prabhakaran and A. Sahai, editors, *Secure Multi-Party Computation*, volume 10 of *Cryptology and Information Security Series*, pages 222–248. IOS Press, 2013.

[Oka00] T. Okamoto. On relationships between statistical zero-knowledge proofs. *J. Comput. Syst. Sci.*, 60(1):47–108, 2000.

[Ost91] R. Ostrovsky. One-way functions, hard on average problems, and statistical zero-knowledge proofs. In *Proceedings of the Sixth Annual Structure in Complexity Theory Conference, Chicago, Illinois, USA, June 30 - July 3, 1991*, pages 133–138. IEEE Computer Society, 1991.

[OV08] S. J. Ong and S. P. Vadhan. An equivalence between zero knowledge and commitments. In R. Canetti, editor, *Theory of Cryptography, Fifth Theory of Cryptography Conference, TCC 2008, New York, USA, March 19-21, 2008.*, volume 4948 of *Lecture Notes in Computer Science*, pages 482–500. Springer, 2008.

[PS05] R. Pass and A. Shelat. Unconditional characterizations of non-interactive zero-knowledge. In V. Shoup, editor, *Advances in Cryptology - CRYPTO 2005: 25th Annual International Cryptology Conference, Santa Barbara, California, USA, August 14-18, 2005, Proceedings*, volume 3621 of *Lecture Notes in Computer Science*, pages 118–134. Springer, 2005.

[Rus95] A. Russell. Necessary and sufficient condtions for collision-free hashing. *J. Cryptology*, 8(2):87–100, 1995.

[SCPY98] A. D. Santis, G. D. Crescenzo, G. Persiano, and M. Yung. Image density is complete for non-interactive-szk (extended abstract). In K. G. Larsen, S. Skyum, and G. Winskel, editors, *Automata, Languages and Programming, 25th International Colloquium, ICALP'98, Aalborg, Denmark, July 13-17, 1998, Proceedings*, volume 1443 of *Lecture Notes in Computer Science*, pages 784–795. Springer, 1998.

[SV03] A. Sahai and S. P. Vadhan. A complete problem for statistical zero knowledge. *J. ACM*, 50(2):196–249, 2003.

[Vad99] S. P. Vadhan. *A Study of Statistical Zero-Knowledge Proofs*. PhD thesis, 1999.

A Omitted Proofs

A.1 Proof of Item 5 of Fact 1

We prove that if $\Delta(X;Y) \geq 1-\varepsilon$, then, for any $t > 1$, it holds that $\Pr_{x \xleftarrow{R} X}[\Pr[X = x] < t \cdot \Pr[Y = x]] \leq \varepsilon t$.

Proof. We start by proving an additional claim:

Claim 6. For any two distributions X, Y and a subset S of their domain, it holds that:

$$\Delta(X;Y) \leq 1 - \sum_{x \in S} \min(\Pr[X = x], \Pr[Y = x]).$$

Proof.

$$\begin{aligned}
2\,\Delta(X;Y) &= \sum_x |\Pr[X=x] - \Pr[Y=x]| \\
&= \sum_{x\notin S} |\Pr[X=x] - \Pr[Y=x]| + \sum_{x\in S} |\Pr[X=x] - \Pr[Y=x]| \\
&\le \sum_{x\notin S} \Pr[X=x] + \sum_{x\notin S} \Pr[Y=x] + \sum_{x\in S} |\Pr[X=x] - \Pr[Y=x]| \\
&= \sum_x \Pr[X=x] + \sum_x \Pr[Y=x] - \\
&\quad \sum_{x\in S} (\Pr[X=x] + \Pr[Y=y] - |\Pr[X=x] - \Pr[Y=x]|) \\
&= 2 - 2\sum_{x\in S} \min(\Pr[X=x], \Pr[Y=x]).
\end{aligned}$$

The last equality holds because $\sum_x \Pr[X=x] = 1 = \sum_x \Pr[Y=x]$, and for all a, b we have that $a + b - |a-b| = 2\min(a,b)$. □

Now we proceed to the proof of the lemma. Let $S := \{x \mid \Pr[X=x] < t\cdot\Pr[Y=x]\}$. Due to the claim, we have that

$$\Delta(X;Y) \le 1 - \sum_{x\in S} \min(\Pr[X=x], \Pr[Y=x]) \tag{11}$$

We now give a lower bound for each summand $\min(\Pr[X=x], \Pr[Y=x])$. Namely, we show that

$$\forall x \in S \;\; \min(\Pr[X=x], \Pr[Y=x]) \ge \Pr[X=x]/t. \tag{12}$$

By the construction of S, we have that for any $x \in S$ $\Pr[Y=x] > \Pr[X=x]/t$. Hence, $\min(\Pr[X=x], \Pr[Y=x]) \ge \min(\Pr[X=x], \Pr[X=x]/t)$. Since $t > 1$, we have that $\min(\Pr[X=x], \Pr[X=x]/t) = \Pr[X=x]/t$. Combining inequalities 11 and REeq:part, we get that

$$\begin{aligned}
\Delta(X;Y) &\le 1 - \sum_{x\in S} \min(\Pr[X=x], \Pr[Y=x]) \\
&\le 1 - \sum_{x\in S} \Pr[X=x]/t \\
&= 1 - \Pr[X\in S]/t.
\end{aligned}$$

Recall that by assumption $1-\varepsilon \le \Delta(X;Y)$, and therefore, we conclude that $\varepsilon \ge \Pr[X\in S]/t$ implying that $\Pr[X\in S] \le \varepsilon t$. □