



Can Statistical Zero Knowledge be made Non-Interactive?

OR

On the Relationship of \mathcal{SZK} and \mathcal{NISZK}^*

Oded Goldreich[†]

Amit Sahai[‡]

Salil Vadhan[§]

May 27, 1999

Abstract

We extend the study of non-interactive statistical zero-knowledge proofs. Our main focus is to compare the class \mathcal{NISZK} of problems possessing such *non-interactive* proofs to the class \mathcal{SZK} of problems possessing *interactive* statistical zero-knowledge proofs. Along these lines, we first show that if statistical zero knowledge is non-trivial then so is non-interactive statistical zero knowledge, where by non-trivial we mean that the class includes problems which are *not* solvable in probabilistic polynomial-time. (The hypothesis holds under various assumptions, such as the intractability of the Discrete Logarithm Problem.) Furthermore, we show that if \mathcal{NISZK} is closed under complement, then in fact $\mathcal{SZK} = \mathcal{NISZK}$, i.e. all statistical zero-knowledge proofs can be made non-interactive.

The main tools in our analysis are two promise problems that are natural restrictions of promise problems known to be complete for \mathcal{SZK} . We show that these restricted problems are in fact complete for \mathcal{NISZK} and use this relationship to derive our results comparing the two classes. The two problems refer to the statistical difference, and difference in entropy, respectively, of a given distribution from the uniform one. We also consider a weak form of \mathcal{NISZK} , in which only requires that for every inverse polynomial $1/p(n)$, there exists a simulator which achieves simulator deviation $1/p(n)$, and show that this weak form of \mathcal{NISZK} actually equals \mathcal{NISZK} .

Keywords: Statistical Zero-Knowledge Proofs, Non-Interactive Zero-Knowledge Proofs.

*An extended abstract of this work appears in CRYPTO '99 [GSV99].

[†]Department of Computer Science, Weizmann Institute of Science, Rehovot, ISRAEL. E-mail: oded@wisdom.weizmann.ac.il. Work done while visiting LCS, MIT. Supported by DARPA grant DABT63-96-C-0018.

[‡]Laboratory for Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139. E-mail: amits@theory.lcs.mit.edu. Supported by DOD/NDSEG fellowship and DARPA grant DABT63-96-C-0018.

[§]Laboratory for Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139. E-mail: salil@theory.lcs.mit.edu. Supported by DOD/NDSEG fellowship and in part by DARPA grant DABT63-96-C-0018.

1 Introduction

Zero-Knowledge proofs, introduced by Goldwasser, Micali and Rackoff [GMR89], are fascinating and extremely useful constructs. Their fascinating nature is due to their seemingly contradictory nature; they are both convincing and yet yield nothing beyond the validity of the assertion being proven. Their applicability in the domain of cryptography is vast; they are typically used to force malicious parties to behave according to a predetermined protocol (which requires parties to provide proofs of the correctness of their secret-based actions without revealing these secrets). Zero-knowledge proofs come in many flavors, and in this paper we focus on two parameters: The first parameter is the underlying *communication model*, and the second is the *type of the zero-knowledge guarantee*.

The communication model. When Goldwasser, Micali, and Rackoff proposed the definition of zero-knowledge proofs, it seemed that interaction was crucial to achieving zero knowledge – that the possibility of zero knowledge arose through the power of interaction. Indeed, it was not unexpected when [GO94] showed zero knowledge to be trivial (i.e., only exists for proofs of \mathcal{BPP} statements) in the most straightforward non-interactive models. Surprisingly, however, Blum, Feldman, and Micali [BFM88], showed that by changing the model slightly, it is possible to achieve zero knowledge in a non-interactive setting (i.e. where only unidirectional communication can occur). Specifically, they assume that both Prover and Verifier have access to a shared truly random string, called the *reference string*. Aside from this assumption, all communication consists of one message, the “proof,” which is generated by the Prover (based on the assertion being proved and the reference string) and sent from the Prover to the Verifier.

Non-interactive zero-knowledge proofs, on top of being more communication-efficient by definition, have several applications not offered by ordinary interactive zero-knowledge proofs. They have been used, among other things, to build digital signature schemes secure against adaptive chosen message attack [BG89], public-key cryptosystems secure against chosen-ciphertext attack [NY90, DDN91], and non-malleable cryptosystems [DDN91].

The zero-knowledge guarantee. For ordinary *interactive* zero-knowledge proofs, the zero-knowledge requirement is formulated by saying that the transcript of the Verifier’s interaction with the Prover can be *simulated* by the Verifier itself. Similarly, for the *non-interactive* setting described above, the zero-knowledge condition is formulated by requiring that one can produce, knowing only the statement of the assertion, a random reference string *along with* a “proof” that works for the reference string. More precisely, we require that there exists an efficient procedure that on input a valid assertion produces a distribution which is “similar” to the joint distribution of random reference strings and proofs generated by the Prover. The key parameter is the interpretation of “similarity.” Two notions have been commonly considered in the literature (cf., [GMR89, GMW91, For89, BDMP91, BR90]). *Statistical zero knowledge* requires that these distributions be statistically close (i.e., the statistical difference between them is negligible). *Computational zero knowledge* instead requires that these distributions are computationally indistinguishable (cf., [GM84, Yao82]). In this work, we focus on the stronger security requirement of statistical zero knowledge.

Since its introduction in [BFM88], most work on non-interactive zero knowledge has focused on the computational type (cf., [BFM88, DMP87, DMP88, BDMP91, FLS90, KP98]). With non-interactive statistical zero knowledge, the main objects of investigation have been the specific proof

system for **Quadratic Nonresiduosity** and variants [BDMP91, DDP94, DDP97].¹ Recently, De Santis *et. al.* [DDPY98] opened the door to a general study of non-interactive statistical zero-knowledge by showing that it contains a complete (promise²) problem.

Notation. Throughout the paper, \mathcal{SZK} denotes the class of promise problems having statistical zero-knowledge interactive proof systems (defined in Appendix A), and \mathcal{NISZK} denotes the class of promise problems having non-interactive statistical zero-knowledge proof systems (defined in Section 1.1).

Our Contribution.

In this work, we seek to understand what, if any, additional power interaction gives in the context of statistical zero knowledge. Thus, we continue the investigation of \mathcal{NISZK} , focusing on its relationship with \mathcal{SZK} . Our first result is that the non-triviality of \mathcal{SZK} implies non-triviality of \mathcal{NISZK} , where by non-trivial we mean that a class includes problems which are *not* solvable in probabilistic polynomial-time. The hypothesis holds under various assumptions, such as the intractability of Discrete Logarithm Problem [GK93] (or Quadratic Residuosity [GMR89] or Graph Isomorphism [GMW91]), but variants of these last two problems are already known to be in \mathcal{NISZK} [BDMP91, BR90]).

Furthermore, we show that if \mathcal{NISZK} is closed under complement, then in fact $\mathcal{SZK} = \mathcal{NISZK}$ — i.e., all statistical zero-knowledge proofs can be made non-interactive. We note that [DDPY98] have claimed that \mathcal{NISZK} is closed under complement (and OR), but these claims have been retracted [DDPY99].

We also show the equivalence of \mathcal{NISZK} with a variant in which the statistical zero knowledge requirement is weakened somewhat.

Complete Problems. Central to our methodology is the use of simple and natural complete problems to understand classes, such as \mathcal{SZK} and \mathcal{NISZK} , whose definitions are rather complicated. In particular, we exhibit two natural promise problems and prove that they are complete for \mathcal{NISZK} . The two problems refer to the “distance” (in two different senses) of a given distribution from the uniform one. These two problems are natural restrictions of two promise problems shown complete for \mathcal{SZK} , in [SV97] and [GV99], respectively. Indeed, our results about the relationship between \mathcal{SZK} and \mathcal{NISZK} come from relating the corresponding complete problems. This general theme of using completeness to simplify the study of a class, rather than as evidence for computational intractability (as is the traditional use of \mathcal{NP} -completeness), has been evidenced in a number of recent works (cf., [GMW91, LFKN90, Sha90, ALM⁺92, AS92]) and has been particularly useful in understanding statistical zero knowledge (cf., [SV97, SV99, DDPY98, GV99]).

1.1 The non-interactive model

Let us recall the definition of a non-interactive statistical zero-knowledge proof system from [BDMP91].³ We will adapt the definition to promise problems. Note that our definition will capture what

¹The only exception is an unpublished manuscript of Bellare and Rogaway [BR90] who proved some basic results about non-interactive perfect zero-knowledge and showed a non-interactive perfect zero-knowledge proof for the language of graphs with trivial automorphism group.

²A *promise problem* Π is a pair $\Pi = (\Pi_{\text{YES}}, \Pi_{\text{NO}})$ of disjoint sets of strings, corresponding to YES and NO instances of a decision problem.

³Actually, only non-interactive *perfect* and computational zero-knowledge proofs were defined in [BDMP91]. The definition we are using, previously given in [BR90, DDPY98], is the natural non-interactive analogue of (interactive)

[BDMP91] call a *bounded proof system*, in that each shared reference string can only be used once. In contrast to non-interactive *computational* zero knowledge (cf., [BDMP91, FLS90]), it is unknown whether every problem that has such a (bounded) non-interactive *statistical* zero-knowledge proof system also has one in which the shared reference string can be used an unbounded (polynomial) number of times.

A *non-interactive statistical zero-knowledge* proof system for a promise problem Π is defined by a triple of probabilistic machines P , V , and S , where V and S are polynomial-time and P is computationally unbounded, and a polynomial $r(n)$ (which will give the size of the random *reference string* σ), such that:

1. (Completeness:) For all $x \in \Pi_{\text{YES}}$, the probability that $V(x, \sigma, P(x, \sigma))$ accepts is at least $2/3$.
2. (Soundness:) For all $x \in \Pi_{\text{NO}}$, the probability that $V(x, \sigma, P(x, \sigma))$ accepts is at most $1/3$.
3. (Zero Knowledge:) For all $x \in \Pi_{\text{YES}}$, the statistical deviation between the following two distributions is at most $\beta(|x|)$:
 - (A) Choose σ uniformly from $\{0, 1\}^{r(|x|)}$, sample p from $P(x, \sigma)$, and output (p, σ) .
 - (B) $S(x)$ (where the coins for S are chosen uniformly at random.)

where $\beta(n)$ is a negligible function,⁴ termed the *simulator deviation*, and the probabilities in Conditions 1 and 2 are taken over the random coins of V and P , and the choice of σ uniformly from $\{0, 1\}^{r(n)}$. Note that non-interactive statistical zero knowledge is closed under parallel repetition, so the completeness and soundness errors (i.e. the probability of rejection (resp., acceptance) for YES (resp., NO) instances) can be made exponentially small in $|x|$.

We also define a weaker notion of zero knowledge, known as a *weak non-interactive statistical zero-knowledge proof system*, where we ask only that for every polynomial $g(n)$, there exists a probabilistic polynomial-time simulator S_g (whose running time may depend on g), such that the simulator deviation as defined above is at most $1/g(|x|)$. This is the natural analogue of a notion defined in the interactive setting for statistical zero knowledge [DOY97] as well as concurrent zero knowledge [DNS98].

The class of promise problems that possess non-interactive statistical zero-knowledge proof systems is denoted \mathcal{NISZK} , and we denote by *weak-NISZK* the class of promise problems that possess weak non-interactive statistical zero-knowledge proof systems. Note that by definition, $\mathcal{NISZK} \subset \text{weak-NISZK}$. De Santis *et. al.* [DDPY98] recently began a general investigation of the class \mathcal{NISZK} . They introduced a promise problem, called **Image Density**, and claimed that is complete for \mathcal{NISZK} and that the latter class is closed under OR and complement. We were able to verify that some variants of **Image Density** are \mathcal{NISZK} -complete, and indeed the ideas used towards this goal are important to our work. However, they have retracted their claims that \mathcal{NISZK} is closed under OR and complement [DDPY99].

In this paper, in addition to examining \mathcal{NISZK} on its own, we also consider the relationship non-interactive statistical zero-knowledge proofs have with *interactive* statistical zero-knowledge proofs. In the context of interactive zero-knowledge proofs, another issue that arises in the zero-knowledge condition is the behavior of the verifier. The general definition of zero knowledge requires that the zero-knowledge requirement hold for any probabilistic polynomial-time verifier. A

statistical zero knowledge [GMR89].

⁴Recall that a function is *negligible* if it is eventually less than $1/g(n)$ for any polynomial g .

weaker requirement, called *honest-verifier zero knowledge*, requires the zero-knowledge condition to hold only if the verifier behaves honestly. However, it is known that these two conditions are equivalent for statistical zero knowledge, in the sense that every statistical zero-knowledge proof against the honest verifier can be transformed into one that is statistical zero knowledge against any verifier [GSV98]. Thus, we write \mathcal{SZK} for the class of promise problems possessing statistical zero-knowledge proofs (against any polynomial-time verifier or, equivalently, against just the honest verifier).

Note that in the case of non-interactive zero knowledge, the issue of honest verifiers does not arise since the verifier does not interact with the prover. Also, note that we can always transform a non-interactive zero-knowledge proof into an honest verifier zero-knowledge proof, since we could have the honest verifier supply a random string which can replace the common reference string required for non-interactive zero knowledge. That is, $\mathcal{NISZK} \subset \mathcal{SZK}$ (recalling the equivalence of \mathcal{SZK} with honest-verifier \mathcal{SZK}).

1.2 Our Results

The primary tools we use in our investigation are promise problems that are complete for \mathcal{SZK} or \mathcal{NISZK} . In [SV97], a promise problem called Statistical Difference (SD) was introduced and proved complete for \mathcal{SZK} , providing the first completeness result for \mathcal{SZK} . Recently, it was shown in [GV99] that another natural problem, called Entropy Difference (ED), is complete for \mathcal{SZK} as well. In this work, we show that “one-sided” versions of these problems, which we call Statistical Difference from Uniform (SDU) and Entropy Approximation (EA), are complete for \mathcal{NISZK} . To define these problems more precisely, we first recall that the *statistical difference* between two random variables X and Y on a finite set D , denoted $\Delta(X, Y)$, is defined to be

$$\Delta(X, Y) \stackrel{\text{def}}{=} \max_{S \subset D} |\Pr[X \in S] - \Pr[Y \in S]| = \frac{1}{2} \cdot \sum_{\alpha} |\Pr[X = \alpha] - \Pr[Y = \alpha]|.$$

All the promise problems we consider involve distributions which are encoded by circuits which sample from them. That is, if X is a circuit mapping $\{0, 1\}^m$ to $\{0, 1\}^n$, we identify X with the probability distribution induced on $\{0, 1\}^n$ by feeding X the uniform distribution on $\{0, 1\}^m$. Since circuits can be evaluated in time polynomial in their size, yet are rich enough to encode general (e.g., Turing machine) computations, they effectively capture the notion of an “efficiently sampleable distribution.”

Definition 1.1 (Problems involving statistical difference): *The promise problem Statistical Difference, denoted $\text{SD} = (\text{SD}_{\text{YES}}, \text{SD}_{\text{NO}})$, consists of*

$$\begin{aligned} \text{SD}_{\text{YES}} &\stackrel{\text{def}}{=} \{(X, Y) : \Delta(X, Y) < 1/3\} \\ \text{SD}_{\text{NO}} &\stackrel{\text{def}}{=} \{(X, Y) : \Delta(X, Y) > 2/3\} \end{aligned}$$

where X and Y are distributions encoded as circuits which sample from them. The promise problem Statistical Difference from Uniform, denoted $\text{SDU} = (\text{SDU}_{\text{YES}}, \text{SDU}_{\text{NO}})$, consists of

$$\begin{aligned} \text{SDU}_{\text{YES}} &\stackrel{\text{def}}{=} \{X : \Delta(X, U) < 1/n\} \\ \text{SDU}_{\text{NO}} &\stackrel{\text{def}}{=} \{X : \Delta(X, U) > 1 - 1/n\} \end{aligned}$$

where X is a distribution encoded as a circuit outputting n bits, and U is the uniform distribution on n bits.

For the two problems related to entropy, we recall that the (Shannon) entropy of a random variable X , denoted $H(X)$, is defined as

$$H(X) \stackrel{\text{def}}{=} \sum_{\alpha} \Pr[X = \alpha] \cdot \log_2(1/\Pr[X = \alpha])$$

Definition 1.2 (Problems involving entropy): *The promise problem Entropy Difference, denoted $ED = (ED_{\text{YES}}, ED_{\text{NO}})$, consists of*

$$\begin{aligned} ED_{\text{YES}} &\stackrel{\text{def}}{=} \{(X, Y) : H(X) > H(Y) + 1\} \\ ED_{\text{NO}} &\stackrel{\text{def}}{=} \{(X, Y) : H(Y) > H(X) + 1\} \end{aligned}$$

The promise problem Entropy Approximation, denoted $EA = (EA_{\text{YES}}, EA_{\text{NO}})$, consists of

$$\begin{aligned} EA_{\text{YES}} &\stackrel{\text{def}}{=} \{(X, k) : H(X) > k + 1\} \\ EA_{\text{NO}} &\stackrel{\text{def}}{=} \{(X, k) : H(X) < k - 1\} \end{aligned}$$

In these problems, k is a positive integer and X and Y are distributions encoded as circuits which sample from them.

Our first theorem, which is the starting point for our other results, is:

Theorem 1.3 (EA and SDU are \mathcal{NISZK} -complete) *The promise problems EA and SDU are complete for \mathcal{NISZK} . That is, $EA, SDU \in \mathcal{NISZK}$ and for every promise problem $\Pi \in \mathcal{NISZK}$, there is a polynomial-time Karp (many-one) reduction from Π to EA and another from Π to SDU.*

From the proof of this theorem, we also obtain a method for transforming weak non-interactive statistical zero knowledge proofs into standard ones.

Theorem 1.4 *weak- $\mathcal{NISZK} = \mathcal{NISZK}$.*

Armed with our complete problems, we then begin the work of comparing \mathcal{SZK} and \mathcal{NISZK} . First we show that the non-triviality of \mathcal{NISZK} is equivalent to the non-triviality of \mathcal{SZK} . This is shown by giving a Cook reduction from ED to EA.

Theorem 1.5 (non-triviality of \mathcal{NISZK}) *$\mathcal{SZK} \neq \mathcal{BPP} \iff \mathcal{NISZK} \neq \mathcal{BPP}$.*

In this theorem (and throughout the paper), \mathcal{BPP} denotes the class of *promise problems* solvable in probabilistic polynomial time.

In fact, it turns out that the type of Cook reduction we use is a special one, and by examining it further, we are able to shed more light on the \mathcal{SZK} vs. \mathcal{NISZK} question. Specifically, we observe that the reduction we give from ED to EA is an \mathcal{AC}^0 *truth-table reduction*. That is, it is a nonadaptive Cook reduction in which the postprocessing is done in \mathcal{AC}^0 . (Formal definitions are given in Section 5.2.) Further, we can prove that if \mathcal{NISZK} is closed under complement, then \mathcal{NISZK} is closed under \mathcal{AC}^0 truth-table reductions. Thus we deduce that \mathcal{NISZK} being closed under complement implies that $\mathcal{NISZK} = \mathcal{SZK}$. In fact, we can show that closure under complement and a number of other natural conditions are equivalent to $\mathcal{SZK} = \mathcal{NISZK}$:

Theorem 1.6 (conditions for $\mathcal{SZK} = \mathcal{NISZK}$) *The following are equivalent:*

1. $SZK = NISZK$.
2. $NISZK$ is closed under complement.
3. $NISZK$ is closed under \mathcal{NC}^1 truth-table reductions.
4. ED (resp., SD) Karp-reduces to EA (resp., SDU). (“general versions reduce to one-sided ones”)
5. EA (resp., SDU) Karp-reduces to its complement. (“one-sided versions reduce to their complements”)

Theorem 1.6 can be interpreted as saying that if $NISZK$ has a relatively weak closure property (closure under complement), then the class is surprisingly rich (equals SZK) and has a much stronger closure property (closure under \mathcal{NC}^1 truth-table reductions.) At first, it might seem implausible that a class like $NISZK$ with such an asymmetric definition would be closed under complement. But SZK , which has a similarly asymmetric definition, is known to be closed under complement [Oka96]. In light of this, the closure of $NISZK$ under complement would not be quite as unexpected, and Theorem 1.6 illustrates that proving it would have wider consequences.

The last two conditions in Theorem 1.6 show that these questions about non-interactive versus interactive statistical zero-knowledge proofs are actually equivalent to basic questions about relationships between natural computational problems whose definitions have no *a priori* relationship to zero-knowledge proofs.

The equality of SZK and $NISZK$ has interesting consequences not just for $NISZK$, but also for SZK . Currently, the best known generic protocol for SZK (against cheating verifiers, making no computational assumptions) requires a polynomial number of rounds [Oka96, GSV98].⁵ For $NISZK$, however, by [DGW94], it is known that every problem in $NISZK$ has a *constant round* statistical zero-knowledge proof system (against general, cheating verifiers) with inverse polynomial soundness error. Whether every problem in SZK has such a proof system is still an open question, which would be resolved in the positive if $SZK = NISZK$.

1.3 A wider perspective

The study of non-interactive *statistical* (rather than *computational*) zero-knowledge proofs may be of interest for two reasons. Firstly, *statistical* zero-knowledge proofs provide an almost absolute level of security, whereas *computational* zero-knowledge proofs only provide security relative to computational abilities (and typically under complexity theoretic assumptions). Secondly, by analogy from the study of zero-knowledge *interactive* proofs, we believe that techniques developed for the “cleaner” statistical model can be applied or augmented to yield results for computational zero-knowledge: The proof that one-way functions are necessary for SZK to be non-trivial [Ost91] was later generalized to \mathcal{CZK} [OW93]. More recently, the transformations of honest-verifier zero knowledge to general zero knowledge, presented in [Dam94, DGW94, DGOW95, GSV98], apply both to statistical and computational zero knowledge (whereas the original motivation was the study of statistical zero knowledge). It is our hope that the current study of $NISZK$ will eventually lead to a better understanding of $NICZK$, where there are still important open questions such as the minimal conditions under which \mathcal{NP} has $NICZK$ proofs.

⁵Under the assumption that the Discrete Logarithm is hard, however, there is a constant round, cheating verifier SZK proof system with inverse polynomial soundness error for all of SZK [Oka96, BMO90].

2 Preliminaries

Recall that a *promise problem* Π is a pair $\Pi = (\Pi_{\text{YES}}, \Pi_{\text{NO}})$ of disjoint sets of strings, corresponding to the following decision problem: Given a string $x \in \Pi_{\text{YES}} \cup \Pi_{\text{NO}}$, decide whether it is in Π_{YES} (i.e. is a YES instance) or in Π_{NO} (i.e. is a NO instance). A string in $\Pi_{\text{YES}} \cup \Pi_{\text{NO}}$ is said to *satisfy the promise*, and all other strings are said to *violate the promise*. A function f is said to be a *Karp (or polynomial-time many-one) reduction* from a promise problem Π to a promise problem Γ if f is polynomial-time computable, $x \in \Pi_{\text{YES}} \Rightarrow f(x) \in \Gamma_{\text{YES}}$, and $x \in \Pi_{\text{NO}} \Rightarrow f(x) \in \Gamma_{\text{NO}}$. If such a reduction exists, we write $\Pi \leq_{\text{Karp}} \Gamma$.

Recall that all the promise problems we are considering involve distributions which are encoded by circuits which sample from them. That is, if X is a circuit mapping $\{0, 1\}^m$ to $\{0, 1\}^n$, we identify X with the probability distribution induced on $\{0, 1\}^n$ by feeding X the uniform distribution on $\{0, 1\}^m$. The *support* of X is the set of strings in $\{0, 1\}^n$ which have nonzero probability under X , i.e. $\{y \in \{0, 1\}^n : \exists r \in \{0, 1\}^m \text{ s.t. } X(r) = y\}$. For any distribution X on a set D , we write $\otimes^k X$ to denote the distribution on D^k consisting of k independent copies of X .

3 EA is in \mathcal{NISZK}

In this section, we show that EA has a non-interactive statistical zero-knowledge proof system. The proof is given in Subsection 3.1, assuming a certain lemma (Lemma 3.2). Subsections 3.2 to 3.4 are devoted to the proof Lemma 3.2, which is technically somewhat involved. Therefore, the reader is encouraged to read only Subsection 3.1 from this section on first reading.

3.1 The proof system

Our aim in this section is to prove the following:

Lemma 3.1 *EA $\in \mathcal{NISZK}$. Moreover, there is a non-interactive statistical zero-knowledge proof system for EA in which the completeness error, soundness error, and simulator deviation are all exponentially vanishing (specifically 2^{-s} , where s is the length of the input).*

The transformation given by the following lemma (proved in Subsection 3.4) will be applied at the start of the proof system:

Lemma 3.2 *There is a polynomial-time computable function that takes an instance (X, k) of EA and a parameter s (in unary) and produces a distribution Z on $\{0, 1\}^\ell$ (encoded by a circuit which samples from it) such that*

1. *If $H(X) > k + 1$, then Z has statistical difference at most 2^{-s} from the uniform distribution on $\{0, 1\}^\ell$, and*
2. *If $H(X) < k - 1$, then the support of Z is at most a 2^{-s} fraction of $\{0, 1\}^\ell$.*

Lemma 3.2 essentially transforms an instance of **Entropy Approximation** into an instance of **Image Density**, the complete problem of [DDPY98]. Given this transformation, it is straightforward to give a noninteractive statistical zero-knowledge proof system for EA:

Non-interactive proof system for EA, on input (X, k)

1. Let Z be the distribution on $\{0, 1\}^\ell$ obtained from (X, k) as in Lemma 3.2 taking s to be the total description length of (X, k) in bits. Let $\sigma \in \{0, 1\}^\ell$ be the reference string.
2. P selects r uniformly among $\{r' : Z(r') = \sigma\}$ and sends r to V .
3. V accept if $Z(r) = \sigma$ and rejects otherwise.

It is immediate from Lemma 3.2 that the completeness error and soundness error of this proof system are 2^{-s} . For zero-knowledgeness, we consider the following probabilistic polynomial-time simulator:

Simulator for EA proof system, on input (X, k)

1. Let Z be obtained from (X, k) as in the proof system.
2. Select an input r to Z uniformly at random and let $\sigma = Z(r)$.
3. Output (σ, r) .

It follows from Part 1 of Lemma 3.2 that this simulator has statistical difference at most 2^{-s} from the distribution of transcripts of (P, V) . Thus, assuming Lemma 3.2, we have established Lemma 3.1. In fact, we need not require that s be the length of (X, k) . Instead, s can be taken to be an arbitrary security parameter, and the completeness, soundness, and simulation error will be exponentially small in s , while the running time of the protocol only depends polynomially on s . We can use this to prove the following, which will be useful to us later.

Proposition 3.3 *If any promise problem Π reduces to EA by a Karp (i.e. many-one) reduction (even if it is length-reducing), then $\Pi \in \mathcal{NISZK}$.*

Proof: A noninteractive statistical zero-knowledge proof system for Π can be given as follows: On an instance x of Π , both parties compute the image (X, k) of x under the reduction $\Pi \leq_{\text{Karp}} \text{EA}$ and execute the proof system for EA on (X, k) , except that we take s to be the length of x . Hence, the completeness and soundness errors and simulator deviation of this proof system are exponentially small in $|x|$ (rather than $|(X, k)|$ which could be shorter than x). ■

3.2 Flat distributions and the Leftover Hash Lemma

Here we discuss some standard notions and techniques that will be useful in the proof of Lemma 3.2. We use the clean formulations of these tools given in [GV99]. A distribution X is called *flat* if all strings in the support of X have the same probability. Notice that if X is flat, then by the definition of entropy, $\Pr[X = x] = 2^{-H(X)}$ for every x in the support of X . We quantify deviation from flatness as follows:

Definition 3.4 (heavy, light and typical elements): *Let X be a distribution, x an element possibly in its support, and Δ a positive real number. We say that x is Δ -heavy (resp., Δ -light) if $\Pr[X = x] \geq 2^\Delta \cdot 2^{-H(X)}$ (resp., $\Pr[X = x] \leq 2^{-\Delta} \cdot 2^{-H(X)}$). Otherwise, we say that x is Δ -typical.*

A natural relaxed definition of flatness follows. The definition links the amount of slackness allowed in “typical” elements with the probability mass assigned to non-typical elements.

Definition 3.5 (flat distributions): *A distribution X is called Δ -flat if for every $t > 0$, the probability that an element chosen from X is $t \cdot \Delta$ -typical is at least $1 - 2^{-t^2+1}$.*

By straightforward application of Hoeffding Inequality (cf., Appendix C), we have

Lemma 3.6 (flattening lemma): *Let X be a distribution, k a positive integer, and $\otimes^k X$ denote the distribution composed of k independent copies of X . Suppose that for all x in the support of X it holds that $\Pr[X = x] \geq 2^{-m}$. Then $\otimes^k X$ is $\sqrt{k} \cdot m$ -flat.*

The key point is that the entropy of $\otimes^k X$ grows linearly with k , whereas its deviation from flatness grows significantly slower (i.e., linear in \sqrt{k}) as a function of k . Note that if X is a distribution defined by a circuit with ℓ input gates, then $\Pr[X = x] \geq 2^{-\ell}$ for all x in the support of X , so the conclusion of Lemma 3.6 holds with $m = \ell$. The other main tool we will use is:

Lemma 3.7 (Leftover Hash Lemma [ILL89]) *Let \mathcal{H} be a 2-universal family of hash functions mapping a domain D to a range R . Suppose X is a distribution on D such that with probability at least $1 - \delta$ over x selected from X , $\Pr[X = x] \leq \varepsilon/|R|$. Then the statistical difference between the following two distributions is at most $O(\delta + \varepsilon^{1/3})$:*

- (A) *Choose h uniformly from \mathcal{H} and x according to X . Output $(h, h(x))$.*
- (B) *Choose h uniformly from \mathcal{H} and y uniformly from R . Output (h, y) .*

In particular, notice that if X is a Δ -flat distribution, then for any parameters $s, t > 0$, X satisfies the hypothesis of the Leftover Hash Lemma with $|R| = 2^{\mathbb{H}(X) - t\Delta - s}$, $\delta = 2^{-t^2+1}$, and $\varepsilon = 2^{-s}$. As we will be applying Lemma 3.7 to sets of strings, we define, for any pair of positive integers a and b , $\mathcal{H}_{a,b}$ to be one of the standard 2-universal families of hash functions mapping $\{0, 1\}^a$ to $\{0, 1\}^b$ (e.g., affine GF(2)-linear transformations).

3.3 Overview of Lemma 3.2

The transformation proceeds in four stages, which are roughly described below:

1. Let X' consist of many copies of X so that the entropy gap between YES and NO instances increases, and the distribution becomes quite flat relative to its entropy.
2. Hash X' so that YES instances become close to the uniform distribution while NO instances have much smaller entropy than the uniform distribution. That is, let Y be of the form $(h, h(X'))$, where h is uniformly distributed in a 2-universal family with appropriate parameters.
3. Let Y' consist of many copies of Y so that for NO instances, the entropy deficiency (as compared to the uniform distribution) becomes large and yet Y' becomes quite flat relative to its entropy; while YES instances remain close to uniform.
4. Hash the *inputs* to Y' so that NO instances have small support (rather than just small entropy), while keeping YES instances close to uniform. That is, let Z be of the form $(Y'(r), h, h(r))$ where h is uniformly distributed in a 2-universal family with appropriate parameters.

3.4 Proof of Lemma 3.2

Let (X, k) be an instance of EA, let m (resp., n) denote the number of input and output gates to X , and let s be the extra parameter in the transformation. By increasing s if necessary, we may assume that s is greater than the total description length of (X, k) . Thus, all the intermediate circuits we build will be of size $\text{poly}(s)$. Also note that it suffices for the transformation to achieve error parameters just $2^{-\Omega(s)}$ rather than 2^{-s} , as this can be compensated for by first increasing s by a linear factor.

Many copies I. The first step is to take many copies of each distribution; this has the effect of increasing the entropy gap between YES and NO instances relative to X 's deviation from flatness. Namely, let $q = 4sm^2$ and let $X' = \otimes^q X$ (i.e., X' consists of q independent copies of X). Then $H(X') = q \cdot H(X)$ and, by Lemma 3.6, X' is Δ -flat for $\Delta = \sqrt{4sm^2} \cdot m = 2\sqrt{s} \cdot m^2$. In particular, we have established

Claim 3.8

1. If $H(X) > k + 1$, then $H(X') > qk + q \geq qk + \sqrt{s}\Delta + s$.
2. If $H(X) < k - 1$, then $H(X') < qk - q < qk$.

Hashing I. Now consider the distribution Y on pairs $(h, h(x))$ induced by choosing h uniformly from $\mathcal{H}_{qn, qk+1}$ and x according to X' . Say that elements of $\mathcal{H}_{qn, qk+1}$ take $u \leq \text{poly}(qn, qk) \leq \text{poly}(s)$ bits to represent. Then Y is represented by a circuit with inputs (resp., outputs) of length $m' = u + qm$ (resp., $n' = u + qk + 1$). Y has the following properties:

Claim 3.9

1. If $H(X) > k + 1$, then Y has statistical difference at most $2^{-\Omega(s)}$ from the uniform distribution on $\{0, 1\}^{n'}$.
2. If $H(X) < k - 1$, then the entropy of Y is less than $n' - 1$.

Proof: Part 1 follows from the Δ -flatness of X and the Leftover Hash Lemma. Part 2 follows from the fact that the entropy of Y is at most the entropy of X' (which is less than qk) plus the entropy of the uniform distribution on $\mathcal{H}_{qn, qk+1}$ (which is u). ■

Many copies II. We now take many copies of Y , so that the entropy deficiency of NO instances becomes large relative to the flatness while YES instances remain close to uniform. Specifically, let $q' = 4s \cdot (m')^2$ and let $Y' = \otimes^{q'} Y$, so that Y' has $M = m'q'$ input gates, $N = n'q'$ output gates, and Y' is Δ' -flat for $\Delta' = \sqrt{4s(m')^2} \cdot m' = 2\sqrt{s} \cdot (m')^2$. Then we immediately have:

Claim 3.10

1. If $H(X) > k + 1$, then Y' has statistical difference at most $q' \cdot 2^{-\Omega(s)} = 2^{-\Omega(s)}$ from the uniform distribution on $\{0, 1\}^N$.
2. If $H(X) < k - 1$, then $H(Y') < N - q' \leq N - \sqrt{3s} \cdot \Delta' - s$.

Hashing II. The final step is to make a distribution which, for NO instances, has small support (rather than just low entropy) in the case of NO instances, while YES instances remain close to uniform. Consider a circuit Z which takes as input $r \in \{0, 1\}^M$ and a hash function $h \in \mathcal{H}_{M, M-N-s}$ and outputs $(Y'(r), h, h(r))$. Then,

Claim 3.11 Z satisfies the requirements of Lemma 3.2 (with error parameters $2^{-\Omega(s)}$ rather than 2^{-s}). That is,

1. If $H(X) > k + 1$, then Z has statistical difference at most $2^{-\Omega(s)}$ from the uniform distribution on $\{0, 1\}^N \times \mathcal{H}_{M, M-N-s} \times \{0, 1\}^{M-N-s}$, and
2. If $H(X) < k - 1$, then the support of Z is at most a $2^{-\Omega(s)}$ fraction of $\{0, 1\}^N \times \mathcal{H}_{M, M-N-s} \times \{0, 1\}^{M-N-s}$.

The intuition for this is the following: In the case of YES instances, Y' is close to the uniform distribution on $\{0, 1\}^N$, so for almost all $y \in \{0, 1\}^N$, there will be about 2^{M-N} values of r such that $Y'(r) = y$. Thus, hashing r down to $M - N - s$ bits will still result in a nearly uniform distribution.

In the case of a NO instance, Y' has large entropy deficiency and is nearly flat. From this, we can deduce that Y' lands in some small subset T of $\{0, 1\}^N$ with very high probability. Thus, points $y \notin T$ must have very low probability under Y' , i.e. there are very few inputs r such that $Y'(r) = y$. So, for each $y \notin T$, the pairs $(h, h(r))$ will only hit a small subset of the possible values. Therefore, $(Y'(r), h, h(r))$ has small support, because either the first component lands in a small set (namely T) or the last two components land in a small set.

Proof: Suppose $H(X) > k + 1$. From the fact that Y' has statistical difference at most $2^{-\Omega(s)}$ from uniform it follows that with probability at least $1 - 2^{-\Omega(s)}$ over y selected according to Y' ,

$$\Pr [Y' = y] \geq \frac{1}{2} \cdot \frac{1}{2^N}. \quad (1)$$

Fix any y satisfying Inequality 1. Conditioned on $Y'(r) = y$, r is selected uniformly from $\{r: Y'(r) = y\}$, which by Equation 1 is a set of size at least 2^{M-N-1} . Thus, by the Leftover Hash Lemma, conditioned on $Y'(r) = y$, the distribution of $(h, h(r))$ has statistical difference at most $2^{-\Omega(s)}$ from uniform. Therefore the total statistical difference of Z from uniform is $2^{-\Omega(s)}$.

Now suppose $H(X) < k - 1$. We want to show that the support S of Z is a small fraction of $D = \{0, 1\}^N \times \mathcal{H}_{M, M-N-s} \times \{0, 1\}^{M-N-s}$. To do this, we divide S into three parts, depending on the probability mass given to the y component by Y' . Recall that a “typical” y for Y' has probability mass $\approx 2^{-H(Y')} \geq 2^{-N+\sqrt{3s}\Delta'+s}$.

$$\begin{aligned} S_1 &= \{(y, h, z) \in S: \Pr[Y' = y] \leq 2^{-N-2s}\} && \text{ (“much too light”)} \\ S_2 &= \{(y, h, z) \in S: 2^{-N-2s} < \Pr[Y' = y] \leq 2^{-N+s}\} && \text{ (“too light, but not much too light”)} \\ S_3 &= \{(y, h, z) \in S: 2^{-N+s} < \Pr[Y' = y]\} && \text{ (“not too light”)} \end{aligned}$$

Clearly, $S = S_1 \cup S_2 \cup S_3$. We will show that $|S_i|/|D| \leq 2^{-s}$ for $i = 1, 2, 3$, and so $|S|/|D| \leq 3 \cdot 2^{-s} = 2^{-\Omega(s)}$.

First we bound $|S_1|$. For any y such that $\Pr[Y' = y] \leq 2^{-N-2s}$, there are at most 2^{M-N-2s} values of r such that $Y'(r) = y$. Thus, for any such y and any h , the set of z such that $(y, h, z) \in S_1$ is of size at most 2^{M-N-2s} (because each such z must be of the form $h(r)$ for some r such that $Y'(r) = y$). This implies that S_1 is at most a $2^{M-N-2s}/2^{M-N-s} = 2^{-s}$ fraction of D .

Now we bound $|S_2|$. We show that the set A of y such that $2^{-N-2s} < \Pr[Y' = y] \leq 2^{-N+s}$ is at most a 2^{-s} fraction of $\{0, 1\}^N$. From this, it follows that S_2 is at most a 2^{-s} fraction of D . Every $y \in A$ is $\sqrt{3s} \cdot \Delta'$ -light (since Y' has entropy at most $N - s - \sqrt{3s} \cdot \Delta'$). By the Δ' -flatness of Y' , $\Pr[Y' \in A]$ is at most 2^{-3s+1} . Since every y in A has probability mass at least 2^{-N-2s} under Y' , $|A|$ is at most $2^{-3s+1}/2^{-N-2s} < 2^{N-s}$, as desired.

Finally, we bound $|S_3|$. Clearly, there can be at most 2^{N-s} values of y such that $\Pr[Y' = y] \geq 2^{-N+s}$. From this it follows that $|S_3|/|D| \leq 2^{N-s}/2^N = 2^{-s}$. ■

4 EA and SDU are \mathcal{NISZK} -complete

In this section, we complete the proof of Theorem 1.3. First, we establish that $\text{SDU} \in \mathcal{NISZK}$ by showing:

Lemma 4.1 $\text{SDU} \leq_{\text{Karp}} \text{EA}$. *In particular, $\text{SDU} \in \mathcal{NISZK}$.*

Proof: Let X be an instance of SDU. We assume that $\log(n) > 5$, where n is the output length of the circuit X (otherwise, one can decide in probabilistic polynomial time whether X is a YES or NO instance of SDU by random sampling). Let U denote the uniform distribution on n bits. We claim the map $X \mapsto (X, n-3)$ is the reduction required by the lemma.

If $X \in \text{SDU}_{\text{YES}}$, then $\delta = \Delta(X, U) < 1/n$. Now we use the fact (cf., Appendix B) that for any two random variables, Y and Z , ranging over domain D it holds that

$$|\mathbb{H}(Y) - \mathbb{H}(Z)| \leq (\log |D|) \cdot \Delta(Y, Z) + \mathbb{H}_2(\Delta(Y, Z)),$$

where $\mathbb{H}_2(\theta)$ denotes the entropy of a 0-1 random variable with mean θ . Applying this with $Y = U$ and $Z = X$, we have

$$n - \mathbb{H}(X) < n \cdot 1/n + \mathbb{H}_2(1/n) < 2.$$

Hence $(X, n-3) \in \text{EA}_{\text{YES}}$.

If $X \in \text{SDU}_{\text{NO}}$, then $\Delta(X, U) \geq 1 - 1/n$. By the definition of statistical difference, this implies the existence of a set $S \subset \{0, 1\}^n$ such that $\Pr[X \in S] - \Pr[U \in S] > 1 - 1/n$. This implies that

$$\Pr[X \in S] > 1 - 1/n \quad \text{and} \quad \Pr[U \in S] < 1/n.$$

Thus, $\mathbb{H}(X) \leq \Pr[X \in S] \cdot \log(|S|) + \Pr[X \notin S] \cdot n < 1 \cdot (n - \log n) + (1/n) \cdot n < n - 4$, and we have that $(X, n-3) \in \text{EA}_{\text{NO}}$.

The “in particular” part of Lemma 4.1 follows immediately from Proposition 3.3. ■

Now, we establish both Theorem 1.3 and Theorem 1.4 by showing that all promise problems in weak- \mathcal{NISZK} (and hence all promise problems in \mathcal{NISZK}) are reducible to SDU (and hence by the previous lemma to EA).

Lemma 4.2 *Every promise problem in weak- \mathcal{NISZK} Karp-reduces to SDU.*

Proof: Let Π be any promise problem in weak- \mathcal{NISZK} . As weak- \mathcal{NISZK} is preserved under parallel repetition, we may assume that Π has a weak- \mathcal{NISZK} proof system (P, V) with completeness and soundness errors at most 2^{-n} on inputs of length n . Let $r(n) = \text{poly}(n)$ be the length of the random reference string in (P, V) , and let S be a randomized polynomial-time simulator S such that the statistical difference between the output distribution of S and the distribution of true

transcripts of P is at most $1/(3r(n))$. (Such an S is guaranteed by the *weak-NISZK* property.) Let U denote the uniform distribution on $r(n)$ bits.

Let x be an instance of Π . Define M_x to be a circuit which does the following on input s : $M_x(s)$: Simulate $S(x)$ with randomness s to obtain a transcript (σ, p) . If $V(x, \sigma, p)$ accepts, then output σ , else output $0^{r(n)}$.

We claim that the map $x \mapsto M_x$ is the reduction required by the lemma. Suppose $x \in \Pi_{\text{YES}}$. In this case, we know that the random reference string σ in the output of S has statistical difference less than $1/3r(n)$ from U . In addition, since the completeness error of protocol P is at most 2^{-n} , $S(x)$ can output rejecting transcripts with probability at most $1/(3r(n)) + 2^{-n} \leq 2/(3r(n))$. Hence, $\Delta(M_x, U) < 2/(3r(n)) + 1/(3r(n)) \leq 1/r(n)$, and $M_x \in \text{SDU}_{\text{YES}}$.

Suppose $x \in \Pi_{\text{NO}}$. Since the soundness error of protocol P is bounded by 2^{-n} , for at most a 2^{-n} fraction of reference strings σ does there exist an accepting transcript (σ, p) . Since M_x only outputs reference strings corresponding to accepting transcripts or $0^{r(n)}$, $\Delta(M_x, U) \geq 1 - (2^{-n} + 2^{-r(n)}) > 1 - 1/r(n)$. Thus, $M_x \in \text{SDU}_{\text{NO}}$. ■

Clearly, Lemmas 3.1, 4.1, and 4.2 combine to prove Theorem 1.3. Lemmas 4.2 and 4.1 show that any promise problem Π in *weak-NISZK* reduces to **EA**; by Proposition 3.3, this implies that $\Pi \in \text{NISZK}$ and establishes Theorem 1.4.

5 Comparing *NISZK* and *SZK*

Armed with *NISZK*-complete promise problems so closely related to problems known to be complete for *SZK*, we can quickly begin relating the two classes.

5.1 Nontriviality of *NISZK*

First, we establish Theorem 1.5 by giving a Cook reduction from **Entropy Difference (ED)**, complete for *SZK*, to **Entropy Approximation (EA)**, complete for *NISZK*.

Lemma 5.1 *Suppose (X, Y) is an instance of ED. Let $X' = \otimes^4 X$ (resp., $Y' = \otimes^4 Y$) consist of 4 independent copies of X (resp., Y), and let n denote the maximum of the output sizes of X' and Y' . Then,*

$$\begin{aligned} (X, Y) \in \text{ED}_{\text{YES}} &\implies \bigvee_{k=1}^n [((X', k) \in \text{EA}_{\text{YES}}) \wedge ((Y', k) \in \text{EA}_{\text{NO}})] \\ (X, Y) \in \text{ED}_{\text{NO}} &\implies \bigwedge_{k=1}^n [((X', k) \in \text{EA}_{\text{NO}}) \vee ((Y', k) \in \text{EA}_{\text{YES}})] \end{aligned}$$

Proof: Suppose $(X, Y) \in \text{ED}_{\text{YES}}$, so that $H(X') > H(Y') + 4$. Let $k = \lfloor H(X') \rfloor - 2$. Then $H(X') > k + 1$. On the other hand, $k + 3 > H(X') > H(Y') + 4$, and hence $H(Y') < k - 1$. Suppose instead $(X, Y) \in \text{ED}_{\text{NO}}$, so that $H(Y') > H(X') + 4$. Then for all $k > \lceil H(X') \rceil + 1$, we have $H(X') < k - 1$. So, for all $k \leq \lceil H(X') \rceil + 1$, we have $k + 1 < H(X') + 3 < H(Y')$. ■

From this reduction, we conclude that $\text{SZK} \neq \text{BPP} \iff \text{NISZK} \neq \text{BPP}$, which is Theorem 1.5. Again, by *BPP* we mean the class of *promise problems* solvable in probabilistic polynomial time.

Proof of Theorem 1.5. By definition, $\mathcal{NISZK} \subset \mathcal{SZK}$ (recall that \mathcal{SZK} equals honest-verifier \mathcal{SZK} [GSV98]). Hence if $\mathcal{SZK} = \mathcal{BPP}$, then $\mathcal{NISZK} = \mathcal{BPP}$.

Now suppose $\mathcal{NISZK} = \mathcal{BPP}$, so in particular there is a probabilistic polynomial-time machine M which decides EA (with exponentially small error probability). To show $\mathcal{SZK} = \mathcal{BPP}$, it suffices to show that ED $\in \mathcal{BPP}$ since ED is \mathcal{SZK} -complete. We now describe how to decide instances of ED: Let (X, Y) be an instance of ED. Letting X' and Y' be as stated in Lemma 5.1, we run $M(X', k)$ and $M(Y', k)$ for all $k \in [1, n]$. If for some k , we see that $M(X', k) = 1$ and $M(Y', k) = 0$, we output 1. Otherwise, we output 0. By Lemma 5.1, this is a correct \mathcal{BPP} algorithm for deciding ED. ■

5.2 Conditions under which $\mathcal{NISZK} = \mathcal{SZK}$

The reduction given by Lemma 5.1 is a very special type of Cook reduction, which we call an \mathcal{AC}^0 truth-table reduction. In this section, we use the special properties of this reduction to show that if \mathcal{NISZK} is closed under complement, then in fact $\mathcal{NISZK} = \mathcal{SZK}$. We now precisely define the types of reductions we are using, taking care how they are defined for promise problems.

Definition 5.2 (truth-table reduction [LLS75]): *We say a promise problem Π truth-table reduces to a promise problem Γ , written $\Pi \leq_{\text{tt}} \Gamma$, if there exists a (deterministic) polynomial-time computable function f , which on input x produces a tuple (x_1, x_2, \dots, x_k) and a circuit C , such that*

1. *If $x \in \Pi_{\text{YES}}$ then for all valid settings of b_1, b_2, \dots, b_k , $C(b_1, b_2, \dots, b_k) = 1$, and*
2. *If $x \in \Pi_{\text{NO}}$ then for all valid settings of b_1, b_2, \dots, b_k , $C(b_1, b_2, \dots, b_k) = 0$.*

where a setting for b_i is considered valid when $b_i = 1$ if $x_i \in \Gamma_{\text{YES}}$ and $b_i = 0$ if $x_i \in \Gamma_{\text{NO}}$ (and b_i is unrestricted when x_i violates the promise).

In other words, a truth-table reduction for promise problems is a non-adaptive Cook reduction which is allowed to make queries which violate the promise, but must be able to tolerate both yes and no answers in response to queries that violate the promise. We further consider the case where we restrict the complexity of computing the output of the reduction from the queries:

Definition 5.3 (\mathcal{AC}^0 and \mathcal{NC}^1 truth-table reductions): *A truth-table reduction f between promise problems is an \mathcal{AC}^0 (resp., \mathcal{NC}^1) truth-table reduction if the circuit C produced by the reduction on input x has depth bounded by a constant c_f independent of x (resp., has depth bounded by $c_f \log |x|$). If there is an \mathcal{AC}^0 (resp., \mathcal{NC}^1) truth-table reduction from Π to Γ , we write $\Pi \leq_{\mathcal{AC}^0 - \text{tt}} \Gamma$ (resp., $\Pi \leq_{\mathcal{NC}^1 - \text{tt}} \Gamma$).*

With this definition, we observe that Lemma 5.1 in fact gives an \mathcal{AC}^0 truth-table reduction, since the formula given in the lemma can be expressed as an \mathcal{AC}^0 circuit, and the statement of the lemma shows that the reduction has the robustness properties against promise violations that are required in Definition 5.3. Thus, we have:

Proposition 5.4 $\text{ED} \leq_{\mathcal{AC}^0 - \text{tt}} \text{EA}$.

We say that a class \mathcal{C} of promise problems is closed under a class of reductions \leq_* if $\Pi \leq_* \Gamma$ and $\Gamma \in \mathcal{C}$ implies that $\Pi \in \mathcal{C}$. By the above, if \mathcal{NISZK} is closed under \mathcal{AC}^0 truth-table reductions, then ED $\in \mathcal{NISZK}$ and hence $\mathcal{NISZK} = \mathcal{SZK}$. Thus, we would like to capture the minimal conditions necessary for a promise class to be closed under \mathcal{AC}^0 truth-table reductions. Here, care

must be taken to because of the possibility of promise violations. Keeping this in mind, we define the following operator on promise problems to capture the notion of an unbounded fan-in AND gate for promise problems:

Definition 5.5 (*unbounded AND*): For any promise problem Π , we define $\text{AND}(\Pi)$ to be the promise problem:

$$\begin{aligned}\text{AND}_{\text{YES}}(\Pi) &\stackrel{\text{def}}{=} \{(x_1, x_2, \dots, x_k) : k \geq 0, \forall i \in [1, k] x_i \in \Pi_{\text{YES}}\} \\ \text{AND}_{\text{NO}}(\Pi) &\stackrel{\text{def}}{=} \{(x_1, x_2, \dots, x_k) : k \geq 0, \exists i \in [1, k] x_i \in \Pi_{\text{NO}}\}\end{aligned}$$

We say a class of promise problems \mathcal{C} is closed under unbounded AND if $\Pi \in \mathcal{C}$ implies that $\text{AND}(\Pi) \in \mathcal{C}$.

We have defined AND so that it has the weakest promise condition possible to remain well-defined. In particular, we see that $\text{AND}_{\text{NO}}(\Pi)$ is defined to include x_i 's that violate Π 's promise, as long as just *one* of them is in Π_{NO} . $\Pi \in \mathcal{C}$, $\text{AND}(\Pi) \in \mathcal{C}$. We also need a way of combining two promise problems:

Definition 5.6 (*disjoint union*): For any pair of promise problems Π and Γ , we define the disjoint union of Π and Γ to be the promise problem $\text{DisjUn}(\Pi, \Gamma)$ defined as follows:

$$\begin{aligned}\text{DisjUn}_{\text{YES}}(\Pi, \Gamma) &\stackrel{\text{def}}{=} \{0\} \times \Pi_{\text{YES}} \cup \{1\} \times \Gamma_{\text{YES}} \\ \text{DisjUn}_{\text{NO}}(\Pi, \Gamma) &\stackrel{\text{def}}{=} \{0\} \times \Pi_{\text{NO}} \cup \{1\} \times \Gamma_{\text{NO}}\end{aligned}$$

We say a class of promise problems \mathcal{C} is closed under disjoint union if $\Pi, \Gamma \in \mathcal{C}$ implies that $\text{DisjUn}(\Pi, \Gamma) \in \mathcal{C}$.

With these definitions, we can give the following lemma which gives some conditions sufficient to give closure under AC^0 truth-table reductions.

Lemma 5.7 A promise class \mathcal{C} is closed under AC^0 truth-table reductions if the following conditions hold:

1. \mathcal{C} is closed under Karp (i.e., many-one) reductions.
2. \mathcal{C} is closed under unbounded AND.
3. \mathcal{C} is closed under disjoint union.
4. \mathcal{C} is closed under complementation.

Proof: First note that any unbounded fan-in circuit can be efficiently converted into a circuit with only unbounded fan-in NAND gates (allowing also unary NAND gates), with only a constant factor blowup in depth. So, as a first step, we observe that \mathcal{C} is closed under unbounded NAND: for any promise problem Π , $\text{NAND}(\Pi) \stackrel{\text{def}}{=} \overline{\text{AND}(\Pi)} \in \mathcal{C}$, by closure under unbounded AND and complementation. To generalize this to constant depth circuits with unbounded fan-in NAND gates, we first need a definition.

Definition 5.8 For any promise problem Π , and for all natural numbers $d \geq 0$ we define $\text{Depth}^d(\Pi)$ to be the promise problem whose instances are tuples $(C, (x_1, x_2, \dots, x_k))$, where C is a circuit of depth at most d (using unbounded fan-in NAND gates only). The YES instances are those such that for all valid settings of b_1, b_2, \dots, b_k , $C(b_1, b_2, \dots, b_k) = 1$; whereas the NO instances are those tuples such that for all valid settings of b_1, b_2, \dots, b_k , $C(b_1, b_2, \dots, b_k) = 0$. Here, a setting for b_i is considered valid when $b_i = 1$ if $x_i \in \Pi_{\text{YES}}$ and $b_i = 0$ if $x_i \in \Pi_{\text{NO}}$ (and b_i is unrestricted when x_i violates the promise).

Using the fact that every \mathcal{AC}^0 circuit can be efficiently transformed into one with only NAND gates, we see that $\Pi \leq_{\mathcal{AC}^0\text{-tt}} \Gamma$ means that there exists some d such that $\Pi \leq_{\text{Karp}} \text{Depth}^d(\Gamma)$ under a Karp reduction. Hence if we can show that for all $d \geq 0$ and promise problems Π , $\text{Depth}^d(\Pi) \in \mathcal{C}$, the lemma will be established. We will prove this by induction on d .

First, observe that a depth 0 circuit is simply a variable (negations of variables are achieved with one unary NAND gate, so count as depth 1). Hence, $\text{Depth}^0(\Pi) \leq_{\text{Karp}} \Pi \in \mathcal{C}$. Now assume that $\text{Depth}^d(\Pi) \in \mathcal{C}$. Observe that a depth $d + 1$ circuit is simply a NAND of some number of depth d circuits. Using this observation, we will argue that that

$$\text{Depth}^{d+1}(\Pi) \leq_{\text{Karp}} \text{DisjUn}(\text{Depth}^d(\Pi), \text{NAND}(\text{Depth}^d(\Pi))).$$

By the hypothesized closure properties of \mathcal{C} , this implies that $\text{Depth}^{d+1}(\Pi) \in \mathcal{C}$. The reduction works as follows. The input to the reduction is a tuple (C, \bar{x}) where $\bar{x} = (x_1, x_2, \dots, x_k)$. If C is actually a depth d circuit, then it simply outputs $(0, (C, \bar{x}))$. If not, then it extracts from C the circuits C_1, C_2, \dots, C_s that provide input to the topmost NAND gate. Then the reduction outputs $(1, ((C_1, \bar{x}), (C_2, \bar{x}), \dots, (C_s, \bar{x})))$. It is clear that map gives a Karp reduction from $\text{Depth}^{d+1}(\Pi)$ to $\text{DisjUn}(\text{Depth}^d(\Pi), \text{NAND}(\text{Depth}^d(\Pi)))$, completing the induction step and the proof. ■

Which of the conditions of Lemma 5.7 does \mathcal{NISZK} satisfy? We argue that Conditions 1, 2, and 3 are satisfied by \mathcal{NISZK} :

Lemma 5.9 \mathcal{NISZK} is closed under Karp reductions.

Proof: Suppose $\Gamma \in \mathcal{NISZK}$, and $\Pi \leq_{\text{Karp}} \Gamma$. Since EA is complete for \mathcal{NISZK} , we have $\Gamma \leq_{\text{Karp}} \text{EA}$. By composing reductions, we see that $\Pi \leq_{\text{Karp}} \text{EA}$. By Proposition 3.3, $\Pi \in \mathcal{NISZK}$. ■

Lemma 5.10 \mathcal{NISZK} is closed under unbounded AND.

Proof: First, we argue that $\text{AND}(\text{EA}) \in \mathcal{NISZK}$ by describing a \mathcal{NISZK} proof system for $\text{AND}(\text{EA})$: Let $((X_1, k_1), \dots, (X_m, k_m))$ be an instance of $\text{AND}(\text{EA})$, and say ℓ is the total length of the instance. Artificially pad each circuit X_i to be of description size ℓ (by adding unused gates) and let Y_i be the resulting circuit. Now execute the \mathcal{NISZK} proof system for EA given by Lemma 3.1 on each pair (Y_i, k_i) in parallel, and have the $\text{AND}(\text{EA})$ -verifier accept if the EA-verifier would have accepted on each pair.

If every pair (X_i, k_i) is a YES instance of EA, the $\text{AND}(\text{EA})$ verifier will accept with probability at least $1 - m \cdot 2^{-\ell} = 1 - 2^{-\Omega(\ell)}$, as the completeness error of the EA proof system is at most $2^{-\ell}$. Similarly, running the simulator for the EA proof system m times independently will give a simulation for the $\text{AND}(\text{EA})$ proof system with simulator deviation at most $m \cdot 2^{-\ell} = 2^{-\Omega(\ell)}$. Finally, if just one pair (X_i, k_i) is a NO instance of EA (even if the others violate the promise), the verifier will accept with probability at most $2^{-\ell}$ in the i 'th execution of the EA protocol, and so the $\text{AND}(\text{EA})$ verifier will accept with probability at most $2^{-\ell}$.

This shows that $\text{AND}(\text{EA}) \in \mathcal{NLSZK}$. Now let Π be any promise problem in \mathcal{NLSZK} . Since EA is complete for \mathcal{NLSZK} , there is a Karp reduction f from Π to EA. This induces a Karp reduction from $\text{AND}(\Pi)$ to $\text{AND}(\text{EA})$ in the obvious way (i.e. $(x_1, \dots, x_k) \mapsto (f(x_1), \dots, f(x_k))$). As $\text{AND}(\text{EA})$ is in \mathcal{NLSZK} and \mathcal{NLSZK} is closed under Karp reductions, $\text{AND}(\Pi) \in \mathcal{NLSZK}$. ■

Lemma 5.11 *\mathcal{NLSZK} is closed under disjoint union.*

Proof: For any two promise problem Π and Γ in \mathcal{NLSZK} , the Karp reductions f_0 from Π to EA and f_1 from Γ to EA induce a Karp reduction from $\text{DisjUn}(\Pi, \Gamma)$ to EA given by $(\sigma, x) \mapsto f_\sigma(x)$. By Proposition 3.3, $\text{DisjUn}(\Pi, \Gamma) \in \mathcal{NLSZK}$. ■

Combining everything, we can give a condition under which $\mathcal{SZK} = \mathcal{NLSZK}$.

Proposition 5.12 *If \mathcal{NLSZK} is closed under complementation, then $\mathcal{SZK} = \mathcal{NLSZK}$.*

Proof: Suppose \mathcal{NLSZK} is closed under complementation. Combining this with Lemmas 5.7, 5.9, 5.10, and 5.11, it follows that \mathcal{NLSZK} is closed under \mathcal{AC}^0 truth-table reductions. Applying Proposition 5.4 ($\text{ED} \leq_{\mathcal{AC}^0\text{-tt}} \text{EA}$) and Lemma 3.1 ($\text{EA} \in \mathcal{NLSZK}$), we conclude that $\text{ED} \in \mathcal{NLSZK}$. Since ED is complete for \mathcal{SZK} [GV99] and \mathcal{NLSZK} is closed under Karp reductions (Lemma 5.9), we have $\mathcal{SZK} \subset \mathcal{NLSZK}$. As $\mathcal{NLSZK} \subset \mathcal{SZK}$ is true from the definition of \mathcal{NLSZK} , we conclude that $\mathcal{NLSZK} = \mathcal{SZK}$. ■

Finally, we deduce Theorem 1.6, which gives a number of conditions equivalent to $\mathcal{NLSZK} = \mathcal{SZK}$.

Proof of Theorem 1.6:

1 \Rightarrow 3. This follows from the result of [SV99] that \mathcal{SZK} is closed under \mathcal{NC}^1 truth-table reductions.

3 \Rightarrow 2 \Rightarrow 1. The first is trivial and the second is Proposition 5.12.

1 \Leftrightarrow 4. This follows from Theorem 1.3 (which asserts that EA and SDU are complete for \mathcal{NLSZK}), the fact that ED and SD are complete for \mathcal{SZK} [SV97, GV99], and Lemma 5.9 (that \mathcal{NLSZK} is closed under Karp reductions).

2 \Leftrightarrow 5. This follows from Theorem 1.3 (that EA and SDU are complete for \mathcal{NLSZK}) and Lemma 5.9 (that \mathcal{NLSZK} is closed under Karp reductions).

Acknowledgments

We thank the CRYPTO '99 program committee and reviewers for helpful comments on the presentation.

References

- [ALM⁺92] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and hardness of approximation problems. In *Proceedings of the Thirty Third Annual Symposium on Foundations of Computer Science*, pages 14–23, 1992.

- [AS92] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs. In *Proceedings of the Thirty Third Annual Symposium on Foundations of Computer Science*, pages 2–13, 1992.
- [BDMP91] Manuel Blum, Alfredo De Santis, Silvio Micali, and Giuseppe Persiano. Noninteractive zero-knowledge. *SIAM Journal on Computing*, 20(6):1084–1118, December 1991.
- [BFM88] Manuel Blum, Paul Feldman, and Silvio Micali. Non-interactive zero-knowledge and its applications (extended abstract). In *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing*, pages 103–112, Chicago, Illinois, 2–4 May 1988.
- [BG89] Mihir Bellare and Shafi Goldwasser. New paradigms for digital signatures and message authentication based on non-interactive zero knowledge proofs. In G. Brassard, editor, *Advances in Cryptology—CRYPTO '89*, volume 435 of *Lecture Notes in Computer Science*, pages 194–211. Springer-Verlag, 1990, 20–24 August 1989.
- [BMO90] Mihir Bellare, Silvio Micali, and Rafail Ostrovsky. Perfect zero-knowledge in constant rounds. In *Proceedings of the Twenty Second Annual ACM Symposium on Theory of Computing*, pages 482–493, 1990.
- [BR90] Mihir Bellare and Phillip Rogaway. Non-interactive perfect zero-knowledge. Unpublished manuscript, June 1990.
- [CT91] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley Series in Telecommunications. John Wiley & Sons, Inc., 2nd edition, 1991.
- [Dam94] Ivan Damgård. Interactive hashing can simplify zero-knowledge protocol design. In *Proceedings of Crypto '95, Lecture Notes in Computer Science*, volume 403, pages 100–109. Springer-Verlag, 1994.
- [DDN91] Danny Dolev, Cynthia Dwork, and Moni Naor. Non-malleable cryptography (extended abstract). In *Proceedings of the Twenty Third Annual ACM Symposium on Theory of Computing*, pages 542–552, New Orleans, Louisiana, 6–8 May 1991.
- [DDP94] Alfredo De Santis, Giovanni Di Crescenzo, and Giuseppe Persiano. The knowledge complexity of quadratic residuosity languages. *Theoretical Computer Science*, 132(1–2):291–317, 26 September 1994.
- [DDP97] Alfredo De Santis, Giovanni Di Crescenzo, and Giuseppe Persiano. Randomness-efficient non-interactive zero-knowledge (extended abstract). In Pierpaolo Degano, Robert Gorrieri, and Alberto Marchetti-Spaccamela, editors, *Automata, Languages and Programming, 24th International Colloquium*, volume 1256 of *Lecture Notes in Computer Science*, pages 716–726, Bologna, Italy, 7–11 July 1997. Springer-Verlag.
- [DDPY98] Alfredo De Santis, Giovanni Di Crescenzo, Giuseppe Persiano, and Moti Yung. Image Density is complete for non-interactive-SZK. In *Automata, Languages and Programming, 25th International Colloquium*, Lecture Notes in Computer Science, pages 784–795, Aalborg, Denmark, 13–17 July 1998. Springer-Verlag. See [DDPY99].
- [DDPY99] Alfredo De Santis, Giovanni Di Crescenzo, Giuseppe Persiano, and Moti Yung. Image Density is complete for non-interactive-SZK, May 1999. Preliminary draft of full version.

- [DGOW95] Ivan Damgård, Oded Goldreich, Tatsuaki Okamoto, and Avi Wigderson. Honest verifier vs. dishonest verifier in public coin zero-knowledge proofs. In *Proceedings of Crypto '95, Lecture Notes in Computer Science*, volume 403. Springer-Verlag, 1995.
- [DGW94] Ivan Damgård, Oded Goldreich, and Avi Wigderson. Hashing functions can simplify zero-knowledge protocol design (too). Technical Report RS-94-39, BRICS, November 1994. See Part 1 of [DGOW95].
- [DMP87] Alfredo De Santis, Silvio Micali, and Giuseppe Persiano. Non-interactive zero-knowledge proof systems. In Carl Pomerance, editor, *Advances in Cryptology—CRYPTO '87*, volume 293 of *Lecture Notes in Computer Science*, pages 52–72. Springer-Verlag, 1988, 16–20 August 1987.
- [DMP88] Alfredo De Santis, Silvio Micali, and Giuseppe Persiano. Non-interactive zero-knowledge with preprocessing. In S. Goldwasser, editor, *Advances in Cryptology—CRYPTO '88*, volume 403 of *Lecture Notes in Computer Science*, pages 269–282. Springer-Verlag, 1990, 21–25 August 1988.
- [DNS98] Cynthia Dwork, Moni Naor, and Amit Sahai. Concurrent zero-knowledge. In *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing*, pages 409–418, 1998.
- [DOY97] Giovanni Di Crescenzo, Tatsuaki Okamoto, and Moti Yung. Keeping the SZK-verifier honest unconditionally. In *Advances in Cryptology—CRYPTO '97*, pages 31–45, 1997.
- [FLS90] Uriel Feige, Dror Lapidot, and Adi Shamir. Multiple non-interactive zero knowledge proofs based on a single random string (extended abstract). In *31st Annual Symposium on Foundations of Computer Science*, volume I, pages 308–317, St. Louis, Missouri, 22–24 October 1990. IEEE.
- [For89] Lance Fortnow. The complexity of perfect zero-knowledge. In Silvio Micali, editor, *Advances in Computing Research*, volume 5, pages 327–343. JAC Press, Inc., 1989.
- [GK93] Oded Goldreich and Eyal Kushilevitz. A perfect zero-knowledge proof system for a problem equivalent to the discrete logarithm. *Journal of Cryptology*, 6:97–116, 1993.
- [GM84] Shafi Goldwasser and Silvio Micali. Probabilistic encryption. *Journal of Computer and System Sciences*, 28(2):270–299, 1984.
- [GMR89] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, February 1989.
- [GMW91] Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. *Journal of the Association for Computing Machinery*, 38(1):691–729, 1991.
- [GO94] Oded Goldreich and Yair Oren. Definitions and properties of zero-knowledge proof systems. *Journal of Cryptology*, 7(1):1–32, Winter 1994.
- [GSV98] Oded Goldreich, Amit Sahai, and Salil Vadhan. Honest-verifier statistical zero-knowledge equals general statistical zero-knowledge. In *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing*, pages 399–408, 1998.

- [GSV99] Oded Goldreich, Amit Sahai, and Salil Vadhan. Can statistical zero knowledge be made non-interactive? or On the relationship of SZK and NISZK [extended abstract]. In Michael Wiener, editor, *Advances in Cryptology—CRYPTO '99*, Lecture Notes in Computer Science. Springer-Verlag, 15–19 August 1999.
- [GV99] Oded Goldreich and Salil Vadhan. Comparing entropies in statistical zero-knowledge with applications to the structure of SZK. In *Proceedings of the Fourteenth Annual IEEE Conference on Computational Complexity*, pages 54–73, Atlanta, GA, May 1999. IEEE Computer Society Press.
- [ILL89] Russell Impagliazzo, Leonid A. Levin, and Michael Luby. Pseudo-random generation from one-way functions (extended abstracts). In *Proceedings of the Twenty-First Annual ACM Symposium on Theory of Computing*, pages 12–24, Seattle, Washington, 15–17 May 1989.
- [KP98] Joe Kilian and Erez Petrank. An efficient noninteractive zero-knowledge proof system for NP with general assumptions. *Journal of Cryptology*, 11(1):1–27, Winter 1998.
- [LFKN90] Carsten Lund, Lance Fortnow, Howard Karloff, and Noam Nisan. Algebraic methods for interactive proofs. In *Proceedings of the Thirty First Annual Symposium on Foundations of Computer Science*, pages 1–10, 1990.
- [LLS75] R. E. Ladner, N. A. Lynch, and A. L. Selman. A comparison of polynomial time reducibilities. *Theoretical Computer Science*, 1(2):103–123, December 1975.
- [NY90] Moni Naor and Moti Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *Proceedings of the Twenty Second Annual ACM Symposium on Theory of Computing*, pages 427–437, Baltimore, Maryland, 14–16 May 1990.
- [Oka96] Tatsuaki Okamoto. On relationships between statistical zero-knowledge proofs. In *Proceedings of the Twenty Eighth Annual ACM Symposium on the Theory of Computing*, 1996. See also preprint of full version, Oct. 1997.
- [Ost91] Rafail Ostrovsky. One-way functions, hard on average problems, and statistical zero-knowledge proofs. In *Proceedings of the Thirty Second Annual Symposium on Foundations of Computer Science*, pages 133–138, 1991.
- [OW93] Rafail Ostrovsky and Avi Wigderson. One-way functions are essential for non-trivial zero-knowledge. In *Proceedings of the Second Israel Symposium on Theory of Computing and Systems*, 1993.
- [Sha90] Adi Shamir. $IP=PSPACE$. In *Proceedings of the Thirty First Annual Symposium on Foundations of Computer Science*, pages 11–15, 1990.
- [SV97] Amit Sahai and Salil Vadhan. A complete promise problem for statistical zero-knowledge. In *Proceedings of the Thirty Eighth Annual Symposium on Foundations of Computer Science*, pages 448–457, 1997.
- [SV99] Amit Sahai and Salil Vadhan. Manipulating statistical difference. In Panos Pardalos, Sanguthevar Rajasekaran, and José Rolim, editors, *Randomization Methods in Algorithm Design (DIMACS Workshop, December 1997)*, volume 43 of *DIMACS Series*

in *Discrete Mathematics and Theoretical Computer Science*, pages 251–270. American Mathematical Society, 1999.

[Yao82] Andrew C. Yao. Theory and application of trapdoor functions. In *Proceedings of the Twenty Third Annual Symposium on Foundations of Computer Science*, pages 80–91, 1982.

A Definitions

Following [GK93], we extend the standard definition of interactive proof systems to promise problems $\Pi = (\Pi_{\text{YES}}, \Pi_{\text{NO}})$. That is, we require the completeness condition to hold for YES instances (i.e., $x \in \Pi_{\text{YES}}$), require the soundness condition to hold for NO instances (i.e., $x \in \Pi_{\text{NO}}$), and do not require anything for inputs which violate the promise (i.e., $x \notin \Pi_{\text{YES}} \cup \Pi_{\text{NO}}$).

We are mainly interested in such proof systems which are statistical zero knowledge:

Definition A.1 (Statistical Zero Knowledge – *SZK*): *Let (P, V) be an interactive proof system for a promise problem $\Pi = (\Pi_{\text{YES}}, \Pi_{\text{NO}})$.*

- *We denote by $\langle P, V \rangle(x)$ the view of the verifier V while interacting with P on common input x ; this consists of the common input, V 's internal coin tosses, and all messages it has received.*
- *(P, V) is said to be (general) statistical zero knowledge if, for every probabilistic polynomial-time V^* , there exists a probabilistic polynomial-time machine (called a simulator), S , and a negligible function $\mu : \mathbb{N} \mapsto [0, 1]$ (called the simulator deviation) so that for every $x \in \Pi_{\text{YES}}$ the statistical difference between $S(x)$ and $\langle P, V^* \rangle(x)$ is at most $\mu(|x|)$.*
- *SZK denotes the class of promise problems having statistical zero-knowledge interactive proof systems.*

Honest-verifier statistical zero-knowledge proof systems are such where the zero-knowledge requirement is only required to hold for the prescribed/honest verifier V , rather than for every polynomial-time computable V^* . Every honest-verifier statistical zero-knowledge proof system can be transformed into a general statistical zero-knowledge proof system (actually meeting an even stronger zero-knowledge requirement) [GSV98].

B Statistical Inequalities

Fact B.1 *For any two random variables, X and Y , ranging over a domain D it holds that*

$$|\mathbb{H}(X) - \mathbb{H}(Y)| \leq \log(|D| - 1) \cdot \delta + \mathbb{H}_2(\delta)$$

where $\delta \stackrel{\text{def}}{=} \Delta(X, Y)$.

This fact can be inferred from Fano's Inequality (cf., [CT91, Thm. 2.11.1]), and a slightly weaker bound can be found in [CT91, Thm. 16.3.2]. A more direct proof follows.

Proof: Assume $\delta > 0$ or else the claim is obvious. Let $p(x) \stackrel{\text{def}}{=} \Pr[X = x]$ and $q(x) \stackrel{\text{def}}{=} \Pr[Y = x]$. Define $m(x) \stackrel{\text{def}}{=} \min\{p(x), q(x)\}$. Then $\sum_{x \in D} m(x) = 1 - \delta$. Define random variables Z' , X' and

Y' so that

$$\begin{aligned}\Pr [Z' = x] &= m'(x) \stackrel{\text{def}}{=} \frac{1}{1-\delta} \cdot m(x) \\ \Pr [X' = x] &= p'(x) \stackrel{\text{def}}{=} \frac{1}{\delta} \cdot (p(x) - m(x)) \\ \Pr [Y' = x] &= q'(x) \stackrel{\text{def}}{=} \frac{1}{\delta} \cdot (q(x) - m(x))\end{aligned}$$

Think of X (resp., Y) as being generated by picking Z' with probability $1 - \delta$ and X' (resp., Y') otherwise. Then

$$\begin{aligned}\mathsf{H}(X) &\leq (1 - \delta) \cdot \mathsf{H}(Z') + \delta \cdot \mathsf{H}(X') + \mathsf{H}_2(\delta) \\ \mathsf{H}(Y) &\geq (1 - \delta) \cdot \mathsf{H}(Z')\end{aligned}$$

Observing that $\Pr [X' = x] = 0$ on at least one $x \in D$, it follows that $\mathsf{H}(X') \leq \log(|D| - 1)$, and the fact follows. ■

Comment: The above bound is tight. Let $e \in D$ and consider X which is identically e , and Y which with probability $1 - \delta$ equals e and otherwise is uniform over $D \setminus \{e\}$. Clearly, $\Delta(X, Y) = \delta$ and $\mathsf{H}(Y) - \mathsf{H}(X) = \delta \log(|D| - 1) + \mathsf{H}_2(\delta) - 0$.

C Proof of the Flattening Lemma

For every x in the support of X , we let $w(x) = -\log \Pr [X = x]$. Then w maps the support of X , denoted D , to $[0, m]$. Let X_1, \dots, X_k be identical and independent copies of X . The lemma asserts that for every t

$$\Pr \left[\left| \sum_{i=1}^k w(X_i) - k \cdot \mathsf{H}(X) \right| > t \cdot m \sqrt{k} \right] < 2^{-t^2+1}$$

Observe that $E(w(X_i)) = \sum_x \Pr [X = x] w(x) = \mathsf{H}(X)$, for every i . Thus, the lemma follows by a straightforward application of Hoeffding Inequality: Specifically, define random variables $\xi_i = w(X_i)$, let $\mu = E(\xi_i)$ and $\delta = tm/\sqrt{k}$, and use

$$\begin{aligned}\Pr \left[\left| \frac{\sum_{i=1}^k \xi_i}{k} - \mu \right| > \delta \right] &< 2 \cdot \exp \left(-\frac{2\delta^2}{m^2} \cdot k \right) \\ &= 2 \cdot \exp (-2t^2)\end{aligned}$$

The lemma follows. ■