

A Complete Problem for Statistical Zero Knowledge

AMIT SAHAI

Princeton University, Princeton, New Jersey

AND

SALIL VADHAN

Harvard University, Cambridge, Massachusetts

Abstract. We present the first complete problem for SZK, the class of promise problems possessing statistical zero-knowledge proofs (against an honest verifier). The problem, called STATISTICAL DIFFERENCE, is to decide whether two efficiently samplable distributions are either statistically close or far apart. This gives a new characterization of SZK that makes *no reference to interaction or zero knowledge*.

We propose the use of complete problems to unify and extend the study of statistical zero knowledge. To this end, we examine several consequences of our Completeness Theorem and its proof, such as:

—A way to make every (honest-verifier) statistical zero-knowledge proof very communication efficient, with the prover sending only one bit to the verifier (to achieve soundness error $1/2$).

Preliminary versions of this work appeared as SAHAI, A., AND VADHAN, S. P. 1997. A complete promise problem for statistical zero-knowledge, In *Proceedings of the 38th Annual Symposium on Foundations of Computer Science* (Miami Beach, Fla., Oct. 20–22). IEEE Computer Society Press, Los Alamitos, Calif., pp. 448–457, and SAHAI, A., AND VADHAN, S. P. 1999. Manipulating statistical difference. In *Randomization Methods in Algorithm Design (DIMACS Workshop, December 1997)*. Panos Pardalos, Sanguthevar Rajasekaran, and José Rolim, Eds. DIMACS Series in Discrete Mathematics and Theoretical Computer Science, vol. 43. American Mathematical Society, Providence R.I. This work was done when A. Sahai was at the MIT Laboratory for Computer Science, supported by a DOD NDSEG Graduate Fellowship and partially by DARPA grant DABT63-96-C-0018. He was partially supported by an Alfred P. Sloan Foundation Research Fellowship when preparing this journal article.

This work was done when S. Vadhan was in the MIT Department of Mathematics, supported by a DOD NDSEG Graduate Fellowship and partially by DARPA grant DABT63-96-C-0018. He was supported by an NSF Mathematical Sciences Postdoctoral Research Fellowship and NSF grant CCR-0205423 when preparing this journal article.

Authors' addresses: A. Sahai, Department of Computer Science, Princeton University, Princeton, NJ 08544, e-mail: sahai@cs.princeton.edu, URL: <http://www.cs.princeton.edu/~sahai>; S. Vadhan, Division of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138, e-mail: salil@eecs.harvard.edu, URL: <http://eecs.harvard.edu/~salil>.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 1515 Broadway, New York, NY 10036 USA, fax: +1 (212) 869-0481, or permissions@acm.org.

© 2003 ACM 0004-5411/03/0300-0196 \$5.00

- Simpler proofs of many of the previously known results about statistical zero knowledge, such as the Fortnow and Aiello–Håstad upper bounds on the complexity of SZK and Okamoto’s result that SZK is closed under complement.
- Strong closure properties of SZK that amount to constructing statistical zero-knowledge proofs for complex assertions built out of simpler assertions already shown to be in SZK.
- New results about the various measures of “knowledge complexity,” including a collapse in the hierarchy corresponding to knowledge complexity in the “hint” sense.
- Algorithms for manipulating the statistical difference between efficiently samplable distributions, including transformations that “polarize” and “reverse” the statistical relationship between a pair of distributions.

Categories and Subject Descriptors: F.1.2 [Computation by Abstract Devices]: Modes of Computation—*interactive and reactive computation*; F.1.3 [Computation by Abstract Devices]: Complexity Measures Classes

General Terms: Security, Theory

Additional Key Words and Phrases: Knowledge complexity, proof systems, statistical difference, zero knowledge

1. Introduction

A revolution in theoretical computer science occurred when it was discovered that NP has complete problems [Cook 1971; Levin 1973; Karp 1972]. Most often, these theorems and other completeness results are viewed as negative statements, as they provide evidence of a problem’s intractability. These same results, viewed as positive statements, enable one to study an entire class of problems by focusing on a single problem. For example, all languages in NP were shown to have computational zero-knowledge proofs when such a proof was exhibited for GRAPH 3-COLORABILITY [Goldreich et al. 1991]. Similarly, the result that $IP = PSPACE$ was shown by giving an interactive proof for QUANTIFIED BOOLEAN FORMULA, which is complete for PSPACE [Lund et al. 1992; Shamir 1992]. More recently, the celebrated PCP theorem characterizing NP was proven by designing efficient probabilistically checkable proofs for a specific NP-complete language [Arora et al. 1998; Arora and Safra 1998].

In this article, we present a complete problem for SZK, the class of promise problems¹ possessing statistical zero-knowledge proofs (against an honest verifier). This problem provides a new and simple characterization of SZK—one which makes no reference to interaction or zero knowledge. We propose the use of complete problems as a tool to unify and extend the study of statistical zero knowledge. To this end, we use our complete problem to both establish a number of new results about SZK and easily deduce nearly all previous results about SZK.

¹ A *promise problem* is a decision problem given by a pair of disjoint sets of strings, corresponding to YES and NO instances. In contrast to languages, there may be strings that are neither YES instances nor NO instances. A formal definition of promise problems is given in Section 2.1.

Although our complete problem is not a language, it still provides a meaningful characterization of the class of languages possessing statistical zero-knowledge proofs. Moreover, essentially all of the applications of our Completeness Theorem to prove results about the promise class SZK also imply the analogous results for the language class. Thus, throughout the article, all discussion refers to promise problems rather than languages (except where otherwise noted). Section 2.1 contains further elaborates on the issue of promise problems vs. languages.

1.1. STATISTICAL ZERO KNOWLEDGE. Zero knowledge was introduced in the seminal paper of Goldwasser et al. [1989] within the context of their new notion of interactive proof systems. Informally, an *interactive proof* is a protocol in which a computationally unbounded prover P attempts to convince a probabilistic polynomial-time verifier V of an assertion, namely that a string x is a YES instance of some (promise) problem. The *zero knowledge* property requires that, during this process, the verifier learns nothing beyond the validity of the assertion being proven. To formalize this seemingly impossible notion, two probability distributions are considered:

- (1) The interaction of P and V from V 's point of view.
- (2) The output of a probabilistic polynomial-time machine not interacting with anyone, called the *simulator*, on input x .

An interactive proof system (P, V) is said to be *zero knowledge* if, for every YES instance x , the two distributions above are “alike.” Intuitively, the verifier gains no knowledge by interacting with the prover except that x is a YES instance, since it could have run the simulator instead. The specific variants of zero knowledge differ by the interpretation given to “alike.” The most strict interpretation, leading to *perfect zero knowledge*, requires that the distributions be identical. A slightly relaxed interpretation, leading to *statistical zero knowledge* (sometimes called *almost perfect zero knowledge*), requires that the distributions have negligible statistical difference from one another. The most liberal interpretation, leading to *computational zero knowledge*, requires that samples from the two distributions be indistinguishable by any polynomial-time machine.

In this work, we focus on the class of problems possessing *statistical zero-knowledge* proof systems, which we denote **SZK**. We remark that we concentrate on zero-knowledge proofs against an *honest verifier*, that is, the verifier that follows the specified protocol. In cryptographic applications, one usually wants the zero-knowledge condition to hold for all (even cheating) verifier strategies. However, subsequent to this work, it has been shown that one can transform any proof system that is statistical zero knowledge against an honest verifier into one that is statistical zero knowledge against all verifiers [Goldreich et al. 1998], so restricting our attention to honest verifiers causes no loss of generality.

One reason for interest in **SZK** is that it contains a number of important computational problems. These include problems not known to be in **NP**, such as GRAPH NONISOMORPHISM [Goldreich et al. 1991] and PERMUTATION GROUP NONISOMORPHISM [Kannan 1989]. It also contains problems with cryptographic application and significance that are believed to be hard on average, such as QUADRATIC RESIDUOSITY (and its complement) [Goldwasser et al. 1989], a problem equivalent to the DISCRETE LOGARITHM problem [Goldreich and Kushilevitz 1993], and approximate versions of the SHORTEST VECTOR and CLOSEST VECTOR problems in lattices [Goldreich and Goldwasser 2000]. At the same time, the statistical zero knowledge property has several strong consequences. Unlike a computational zero-knowledge protocol, a statistical zero-knowledge protocol typically remains zero knowledge even against a computationally unbounded verifier.² In addition, a problem

² A rare exception are the results of Bellare et al. [1990], which yield protocols that are only zero knowledge for polynomial-time verifiers.

that has a statistical zero-knowledge proof must lie low in the polynomial-time hierarchy. In fact, such a problem cannot be NP-complete unless the polynomial-time hierarchy collapses [Fortnow 1989; Aiello and Håstad 1991; Boppana et al. 1987]. Because SZK contains problems believed to be hard yet cannot contain NP-complete problems, it holds an intriguing position in complexity theory.

1.2. THE COMPLETE PROBLEM. The promise problem we show to be complete for SZK is STATISTICAL DIFFERENCE. An instance of STATISTICAL DIFFERENCE consists of a pair of probability distributions, specified by circuits that sample from them. Roughly speaking, the problem is to decide whether the distributions defined by the two circuits are statistically “close” or “far apart”. (The gap between “close” and “far apart” is what makes it a promise problem and not just a language.) Our main theorem is that STATISTICAL DIFFERENCE is complete for SZK. This Completeness Theorem gives a new characterization of SZK. Informally, it says that the assertions that can be proven in statistical zero knowledge are exactly those that can be cast as deciding whether a pair of efficiently samplable distributions are statistically close or far apart.

The starting point for our proof of the Completeness Theorem is a powerful theorem of Okamoto [2000], which states that all languages in SZK have *public-coin* (also known as Arthur–Merlin [Babai and Moran 1988]) statistical zero-knowledge proofs. Using the approach pioneered by Fortnow and others [Fortnow 1989; Aiello and Håstad 1991; Ostrovsky 1991], we analyze the simulator of such a proof system and show that statistical properties of the simulator’s output distribution can be used to distinguish between YES and NO instances of the problem in consideration. Our key new observation is that, for a *public-coin* proof system, these statistical properties can be captured by the statistical difference between efficiently samplable distributions. We thereby conclude that every problem in SZK reduces to STATISTICAL DIFFERENCE.

To show that STATISTICAL DIFFERENCE is in SZK, we exhibit a simple 2-message proof system for it, generalizing the well-known proof systems for QUADRATIC NONRESIDUOSITY [Goldwasser et al. 1989], and GRAPH NONISOMORPHISM [Goldreich et al. 1991]. One ingredient in our proof system is a new “Polarization Lemma” for statistical difference, which may be of independent interest. Roughly speaking, this lemma gives an efficient transformation that takes as input a pair of probability distributions (specified by circuits that sample from them) and produces a new pair of distributions such that if the original pair is statistically close (respectively, far apart), the new pair is statistically much closer (respectively, much further apart).

1.3. CONSEQUENCES. We propose using complete problems, such as STATISTICAL DIFFERENCE, to unify and extend the study of SZK. We also use the connection between SZK and statistical properties of samplable distributions to establish new techniques for manipulating such distributions. The results we obtain along these lines are summarized below.

The Relationship between SZK and BPP. Our complete problem illustrates that statistical zero knowledge is a natural generalization of BPP. In the definition of STATISTICAL DIFFERENCE, the circuits can output strings of any length. If we restrict the circuits to have output of logarithmic length, the resulting problem is easily shown to be complete for BPP.

Efficient SZK Proof Systems. The zero-knowledge proof system we exhibit for STATISTICAL DIFFERENCE has many attractive properties (which we describe shortly); by the Completeness Theorem, it follows that every problem in SZK also has a proof system with such properties. First, the protocol is very communication efficient—only two messages are exchanged between the prover and verifier, and the prover only sends *one bit* to the verifier (to achieve soundness error $1/2$). In addition, we will show that when the input is a YES instance, the verifier’s view of the interaction can be simulated by a polynomial-time simulator with *exponentially small* statistical deviation.

Closure Properties. Using the complete problem, we demonstrate that SZK has some very strong closure properties. These can be informally described as asserting the existence of statistical zero-knowledge proofs for complex assertions built out simpler assertions already known to be in SZK. These complex assertions take the form of arbitrary propositional formulas whose atoms are statements about membership in some problem in SZK, and the statistical zero-knowledge proofs we exhibit have complexity that is polynomial in the size of these formulas. These results strengthen earlier ones of De Santis et al. [1994] and Damgård and Cramer [1996], which held for monotone formulas and various subclasses of SZK, such as random self-reducible problems.

By the Completeness Theorem, the closure properties we establish are equivalent to the existence of efficient transformations that manipulate the statistical difference between samplable distributions in various ways. Indeed, it is by exhibiting such transformations that we prove the closure properties of SZK. The transformations we give (and their application to closure properties) are inspired by the techniques of De Santis et al. [1994].

Simpler Proofs of Previous Results. Many of the previous results about SZK can be deduced as immediate corollaries of our Completeness Theorem and its proof. For example, the result of Okamoto [2000] that SZK is closed under complement follows directly from our proof of the Completeness Theorem. Then, using the fact that our proof system for STATISTICAL DIFFERENCE is a constant-round one, we deduce that $\text{SZK} \subset \text{AM} \cap \text{co-AM}$, as originally proven by Fortnow [1989] and Aiello and Håstad [1991]. In addition, the result of Ostrovsky [1991] that one-way functions exist if SZK contains a hard-on-average problem follows immediately by combining our Completeness Theorem with a result of Goldreich [1990] on computational indistinguishability.

Knowledge Complexity. In addition to introducing zero-knowledge proofs, the conference version of the paper of Goldwasser et al. [1989] proposed a more general idea of measuring the amount of knowledge leaked in an interactive proof. Goldreich and Petrank [1999] suggested several definitions of *knowledge complexity* to accomplish this, and relationships between these various types of knowledge complexity were explored by Goldreich and Petrank [1999], Bellare and Petrank [1992], Goldreich et al. [1998], Aiello et al. [1995], and Petrank and Tardos [1996]. Loosely speaking, the definitions of (statistical) knowledge complexity measure the “amount of help” a verifier needs to generate a distribution that is statistically close to its real interaction with the prover. There are several ways of formalizing the “amount of help” the verifier needs and each leads to a different notion of knowledge complexity.

Our work on SZK turns out to have consequences for (nonzero) knowledge complexity as well. First, we show that for the weakest of the various measures of knowledge complexity, namely statistical knowledge complexity in the “hint sense,” the corresponding hierarchy collapses by logarithmic additive factors at all levels, and in particular, knowledge complexity $\log n$ equals statistical zero knowledge. No collapse was previously known for any of the variants of knowledge complexity suggested by Goldreich and Petrank [1999]. Our results are obtained by combining our results on SZK with a general lemma relating knowledge complexity in the hint sense to zero knowledge *for promise problems*.

As with zero knowledge, *perfect* knowledge complexity can also be defined. This measures the number of bits of help the verifier needs to simulate the interaction *exactly*, rather than statistically closely. Using our complete problem for SZK, we improve some results of Aiello et al. [1995] on the perfect knowledge complexity of statistical zero knowledge.

Reversing Statistical Difference. One interesting result that follows from the completeness of STATISTICAL DIFFERENCE and the closure of SZK under complement is the existence of an efficient mapping which “reverses” statistical difference. That is, for every pair of efficiently samplable distributions, we can construct another pair of efficiently samplable distributions such that when the former are statistically close, the latter are statistically far apart, and when the former are far apart, the latter are close.

This motivated us to search for a more explicit description of such a transformation. By extracting ideas from the work of Okamoto [2000] and our proof of the Completeness Theorem, we have obtained such a description (which we give in Section 3.4).

Weak SZK and Expected Polynomial-Time Simulators. The original definition of SZK by Goldwasser et al. [1989] allows the simulator to run in *expected* polynomial time, whereas we insist on strict polynomial time, following Goldreich [2001]. Actually, our proof of the Completeness Theorem shows that the two definitions are equivalent for *public-coin* proof systems. That is, if a problem possesses a public-coin SZK proof system with an expected polynomial-time simulator, then it also possesses an SZK proof system with a strict polynomial-time simulator (which can be made public coin by Okamoto [2000]). In fact, the equivalence extends to an even weaker definition of SZK, in which it is only required that for every polynomial $p(n)$, there exists a simulator achieving simulator deviation $1/p(n)$.

Perfect and Computational Zero Knowledge. Our techniques can also be used to analyze public-coin perfect and computational zero-knowledge proofs. Although we do not obtain complete problems in these cases, we do obtain some novel insights into the corresponding complexity classes. Specifically, in Section 3.2, we show that every problem possessing a public-coin perfect zero-knowledge proof (essentially) reduces to a restricted version of STATISTICAL DIFFERENCE. We also show that for any problem possessing a public-coin computational zero-knowledge proof, there exist ensembles of samplable distributions indexed by instances of the problem such that on YES instances, the distributions are computationally indistinguishable and on NO instances, the distributions are statistically far apart.

Cheating-Verifier Zero Knowledge. Although, in this article, we primarily focus on honest-verifier statistical zero knowledge, there have been a number of

works examining “cheating-verifier” statistical zero knowledge, and in particular relating the honest and cheating-verifier definitions. Some of these works exhibited transformations from honest-verifier SZK proofs to cheating-verifier ones under (successively weaker) complexity assumptions [Bellare et al. 1990; Ostrovsky et al. 1993; Okamoto 2000], and others gave unconditional transformations for restricted subclasses of SZK [Damgård 1993; Damgård et al. 1995]. Finally, subsequent to our article, it was proven by Goldreich et al. [1998] that honest-verifier and cheating-verifier SZK are equal, unconditionally and with no restrictions.

Following the paradigm advocated by Bellare et al. [1990], we use the above transformations to translate our results about honest-verifier SZK, namely the Completeness Theorem and its corollaries, to the cheating-verifier class. In Section 3.4, we precisely state the results thereby obtained for cheating-verifier statistical zero knowledge.

1.4. SUBSEQUENT WORK. Subsequent to the conference version of this article [Sahai and Vadhan 1997], there have been a number of other works improving our understanding of SZK, many of which make use of the complete problem methodology advocated here. As mentioned above, Goldreich et al. [1998] show that honest-verifier statistical zero knowledge equals cheating-verifier statistical zero knowledge. Goldreich and Vadhan [1999] use the complete problem methodology to give a simpler proof of Okamoto’s theorem that private-coin SZK equals public-coin SZK (on which our work relies). In the process, they exhibit another complete problem for SZK, called ENTROPY DIFFERENCE, which amounts to deciding which of two given distributions (specified by circuits that sample from them) has noticeably higher entropy than the other. Di Crescenzo et al. [2000] consider two variants of (honest-verifier) statistical zero-knowledge proofs, namely “proofs of decision power” and “proofs of decision,” and exhibit such proof systems for all of SZK. Their construction makes use of the complete problems for SZK given here and in [Goldreich and Vadhan 1999] and special properties of their proof systems. Goldreich et al. [2001] study the complexity of interactive proofs with low prover-to-verifier communication. Using our complete problem, they show that the class of problems with interactive proofs in which the prover sends only one bit to the verifier is exactly SZK (modulo some constraints on the completeness and soundness probabilities).

De Santis et al. [1998] extend the use of complete problems to study “non-interactive” statistical zero knowledge; they exhibit a complete problem for the corresponding complexity class NISZK and use it to prove some general results about the class. Goldreich et al. [1999], exhibit two more complete problems for NISZK. These problems are natural restrictions of the complete problems for SZK given here and in Goldreich et al. [1999] and thus they are able to use the complete problems to relate SZK and NISZK. Gutfreund and Ben-Or [2000] examine weaker models of noninteractive zero knowledge proofs, and, using our complete problem and reversal mapping, show that every problem in SZK has a noninteractive statistical zero-knowledge proof in one of their models.

Finally, Vadhan [2000] examines the blow-up in the prover’s complexity incurred by transformations from private-coin proof systems to public-coin proof systems, such as those in Goldwasser and Sipser [1989] and Okamoto [2000], and shows that this inefficiency is inherent in the fact that the transformations use the original prover and verifier strategies as “black boxes”. In fact, it is shown that any black-box

transformation which preserves the prover's complexity must fail on our proof system for STATISTICAL DIFFERENCE.

Unified presentations of many of the above results, together with the results in this article, can be found in the Ph.D. dissertations of the authors [Vadhan 1999; Sahai 2000].

2. Preliminaries

2.1. PROMISE PROBLEMS AND COMPLETENESS. The problem we prove to be complete for SZK is not a language, but rather a *promise problem* [Even et al. 1984]. Formally, a promise problem Π consists of two disjoint sets of strings Π_Y and Π_N , where Π_Y is the set of YES *instances* and Π_N is the set of NO *instances*. A promise problem Π is associated with the following computational problem: Given an input which is “promised” to lie in $\Pi_Y \cup \Pi_N$, decide whether it comes from Π_Y or Π_N . The *complement* of Π is the promise problem $\bar{\Pi}$, where $\bar{\Pi}_Y = \Pi_N$ and $\bar{\Pi}_N = \Pi_Y$. Note that languages are a special case of promise problems.

We say that promise problem Π *reduces* to promise problem Γ if there is a polynomial-time computable function f such that

$$\begin{aligned} x \in \Pi_Y &\Rightarrow f(x) \in \Gamma_Y \\ x \in \Pi_N &\Rightarrow f(x) \in \Gamma_N. \end{aligned}$$

That is, we work with polynomial-time many-one (or Karp) reductions, unless otherwise specified. We say that \mathbf{C} is *closed under reductions* if $[\Pi \text{ reduces to } \Gamma \text{ and } \Gamma \in \mathbf{C}] \Rightarrow \Pi \in \mathbf{C}$.

If \mathbf{C} is a class of promise problems, we say that promise problem Π is *complete* for \mathbf{C} if $\Pi \in \mathbf{C}$ and every promise problem in \mathbf{C} reduces to Π .

If \mathbf{C} is a class of promise problems, the corresponding *language class* $\mathbf{C}_{\text{lang}} \subset \mathbf{C}$ is the class of languages in \mathbf{C} . It should be noted that \mathbf{C} and \mathbf{C}_{lang} do not always have the same complexity-theoretic properties, particular when allowing reductions that can violate the promise. For example, the promise-problem version of $\text{NP} \cap \text{co-NP}$ contains a problem that is NP-hard with respect to Cook reductions, whereas this does not hold for the language class $\text{NP} \cap \text{co-NP}$ assuming $\text{NP} \neq \text{co-NP}$ [Even et al. 1984]. (See also Goldreich and Goldwasser [2000]).

Nevertheless, in this article, the study of a promise class \mathbf{C} (namely, $\mathbf{C} = \text{SZK}$) proves to be very useful in understanding the corresponding language class. There are two main reasons for this. First, if \mathbf{C} is closed under reductions (as we will prove for $\mathbf{C} = \text{SZK}$) and Π is complete for \mathbf{C} , then Π also meaningfully characterizes \mathbf{C}_{lang} in that $\mathbf{C}_{\text{lang}} = \{L : L \text{ reduces to } \Pi\}$. Second, many of our results are of the form “For every problem $\Pi \in \mathbf{C}$, . . .”. Clearly, all such results will also apply to \mathbf{C}_{lang} because $\mathbf{C}_{\text{lang}} \subset \mathbf{C}$.

2.2. PROBABILITY DISTRIBUTIONS. If X is a probability distribution (or random variable), we write $x \leftarrow X$ to indicate that x is a sample taken from X . If S is a set, we write $x \in_R S$ to indicate that x is uniformly selected from S .

In this article, we will consider probability distributions defined both by circuits and by probabilistic algorithms (i.e., Turing machines). If A is a probabilistic algorithm, we use $A(x)$ to denote the output distribution of A on input x . A *PPT* algorithm (for “probabilistic polynomial time”) is a probabilistic algorithm that runs in strict polynomial time. If C is a circuit mapping m -bit strings to n -bit strings,

then choosing an input u uniformly at random from $\{0, 1\}^m$ defines a probability distribution on $\{0, 1\}^n$ given by $C(u)$. For notational convenience, we also denote this probability distribution by C . These definitions capture the idea of an “(efficiently) samplable” distribution, as to sample from the distribution one need only run the algorithm or evaluate the circuit.

2.3. THE STATISTICAL DIFFERENCE METRIC. For probability distributions (or random variables) X and Y on a discrete set D , the *statistical difference* between X and Y is defined to be

$$\|X - Y\| = \max_{S \subseteq D} |\Pr[X \in S] - \Pr[Y \in S]|. \quad (1)$$

This is often also called the *variation distance* between X and Y . Removing the absolute values in (1) does not change the definition because replacing S by its complement changes the sign (but not magnitude) of $\Pr[X \in S] - \Pr[Y \in S]$. The maximum in (1) can be achieved by taking $S = \{x : \Pr[X = x] > \Pr[Y = x]\}$ (or its complement); this can be seen directly or in the proof of Fact 2.1 below.

There is an equivalent formulation of statistical difference in terms of the ℓ_1 norm $|\cdot|_1$ that will sometimes be more convenient for us. To every probability distribution X on a discrete set D , the *mass function* of X is a vector in \mathbb{R}^D whose x th coordinate is $\Pr[X = x]$. For the sake of elegance, we also denote this vector by X . With this notation, we can state the following well-known fact.

$$\text{FACT 2.1. } \|X - Y\| = \frac{1}{2} |X - Y|_1.$$

The proof of this fact and others in this section are deferred to Appendix 6. It is immediate from this characterization of statistical difference that it is a metric (as long as we identify random variables that are identically distributed). In particular, it satisfies the Triangle Inequality.

FACT 2.2 (TRIANGLE INEQUALITY). *For any probability distributions X , Y , and Z ,*

$$\|X - Y\| \leq \|X - Z\| + \|Z - Y\|.$$

Recall that for any two vectors $v \in \mathbb{R}^m$ and $w \in \mathbb{R}^n$, their *tensor product* $v \otimes w$ is the vector in \mathbb{R}^{nm} , whose (i, j) th component is $v_i w_j$. Now, if we have a pair of random variables (X, Y) (on the same probability space) taking values in $D \times E$, then X is independent from Y iff the mass function of (X, Y) is the tensor product of the mass functions of X and Y (which are elements of \mathbb{R}^D and \mathbb{R}^E , respectively). For this reason, if we have random variables X and Y taking values in sets D and E , respectively, we write $X \otimes Y$ for the random variable taking values in $D \times E$ which consists of independent samples of X and Y . Similarly, $\otimes^k X$ denotes the random variable taking values in D^k consisting of k independent copies of X , that is, $X \otimes X \otimes \cdots \otimes X$.

Now, for any two vectors v and w , $|v \otimes w|_1 = |v|_1 \cdot |w|_1$. In addition, for any mass function X , $|X|_1 = 1$. These facts enable one to show that the statistical difference behaves well with respect to independent random variables:

FACT 2.3. *Suppose X_1 and X_2 are independent random variables on one probability space and Y_1 and Y_2 are independent random variables on another*

probability space. Then,

$$\|(X_1, X_2) - (Y_1, Y_2)\| \leq \|X_1 - Y_1\| + \|X_2 - Y_2\|.$$

One basic fact about statistical difference is that it cannot be created out of nothing. That is, for any procedure A , even if it is randomized, the statistical difference between $A(X)$ and $A(Y)$ is no greater than the statistical difference between X and Y . Formally, if D is any set, a *randomized procedure* on D is a pair $A = (f, R)$, where R is a probability distribution on some set E and f is a function from $D \times E$ to any set F . Think of the distribution R as providing a “random seed” to the procedure A . If X is a probability distribution on D , then $A(X)$ denotes the probability distribution on F obtained by sampling $X \otimes R$ and applying f to the result. Note that applying a *function* is a special case of applying a randomized procedure.

FACT 2.4. *If X and Y are random variables and A is any randomized procedure, then*

$$\|A(X) - A(Y)\| \leq \|X - Y\|.$$

The next fact is useful when arguing that the statistical difference between distributions is small.

FACT 2.5. *Suppose $X = (X_1, X_2)$ and $Y = (Y_1, Y_2)$ are probability distributions on a set $D \times E$ such that*

- (1) X_1 and Y_1 are identically distributed, and
- (2) With probability greater than $(1 - \epsilon)$ over $x \leftarrow X_1$ (equivalently, $x \leftarrow Y_1$),

$$\|X_2|_{X_1=x} - Y_2|_{Y_1=x}\| < \delta,$$

(where $B|_{A=a}$ denotes the conditional distribution of B given that $A = a$ for jointly distributed random variables A and B).

Then $\|X - Y\| < \epsilon + \delta$.

The next fact says that if two distributions have small statistical difference, then their mass functions must be close at most points.

FACT 2.6. *If X and Y are any two distributions such that $\|X - Y\| < \epsilon$, then with probability $> 1 - 2\sqrt{\epsilon}$ over $x \leftarrow X$,*

$$(1 - \sqrt{\epsilon}) \Pr[X = x] < \Pr[Y = x] < (1 + \sqrt{\epsilon}) \Pr[X = x].$$

2.4. ZERO-KNOWLEDGE PROOFS. Before defining zero knowledge, we need to introduce some more terminology. Recall that a *PPT* algorithm is a probabilistic algorithm which runs in *strict* polynomial time. A function $f(n)$ is *negligible* if for all polynomials $p(n)$, $f(n) \leq 1/p(n)$ for all sufficiently large n .

We follow Goldwasser et al. [1989] and Goldreich [2001] in defining interactive proofs and zero-knowledge. The original definitions in Goldwasser et al. [1989] were given for languages. We generalize these definitions to promise problems in the natural way, as previously done in Goldreich and Kushilevitz [1993]. That is, conditions previously required for inputs in the language are now required for YES instances of a promise problem and conditions previously required for inputs not in the language are now required for NO instances.

Informally, an interactive proof is a protocol in which a computationally unbounded prover attempts to convince a polynomial-time verifier V that an assertion

is true, that is, that a string x is a YES instance of a promise problem. More formally, an interactive protocol (P, V) between a computationally unbounded prover P and a PPT verifier V is said to be an *interactive proof system* for a promise problem Π with *completeness error* $c(n)$ and *soundness error* $s(n)$ if the following conditions hold:

- (1) If $x \in \Pi_Y$, then $\Pr[(P, V)(x) = \text{accept}] \geq 1 - c(|x|)$.
- (2) If $x \in \Pi_N$, then for all P^* , $\Pr[(P^*, V)(x) = \text{accept}] \leq s(|x|)$.

We always require that $1 - c(n) > s(n) + 1/\text{poly}(n)$ and that both can be computed in time $\text{poly}(n)$; under this assumption, parallel repetition can be used to obtain a new interactive proof for Π with completeness error and soundness error 2^{-n^k} , for any constant k . We say that (P, V) *exchanges* $m(n)$ *messages* if the prover and verifier exchange at most $m(n)$ messages on any input of length n . An interactive proof system is said to be *public coin* if, on every input, the verifier's random coins r can be written as a concatenation of strings $r_1 r_2 \cdots r_l$ such that the i th message sent from the verifier to the prover is simply r_i .

Roughly speaking, an interactive proof is said to be zero knowledge if, when the input is a YES instance, the verifier can simulate its view of the interaction on its own. To formalize this, let (P, V) be an interactive proof system for a promise problem Π . Let $\text{View}_{P,V}(x)$ be a random variable describing the random coins of V and the messages exchanged between P and V during their interaction on input x . (P, V) is said to be a *statistical zero-knowledge* proof system against the *honest verifier* if there exists a PPT simulator S and a negligible function α (called the *simulator deviation*) such that

$$\text{if } x \in \Pi_Y, \text{ then } \|S(x) - \text{View}_{P,V}(x)\| \leq \alpha(|x|). \quad (2)$$

A *perfect zero-knowledge* proof system is defined in the same way, except that (2) is replaced by $\|S(x) - \text{View}_{P,V}(x)\| = 0$, where S is allowed to output 'fail' with probability at most $1/2$ and $S(x)$ denotes the conditional distribution of S given that the output is not fail.³ A *computational zero-knowledge* proof system replaces (2) with the requirement that $\{S(x)\}_{x \in \Pi_Y}$ and $\{\text{View}_{P,V}(x)\}_{x \in \Pi_Y}$ are *computationally indistinguishable* [Goldwasser and Micali 1984; Yao 1982] ensembles of distributions. That is, for every *nonuniform* polynomial-time algorithm D , there is a negligible function α such that $|\Pr[D(x, S(x)) = 1] - \Pr[D(x, \text{View}_{P,V}(x))]| \leq \alpha(|x|)$ for all $x \in \Pi_Y$.

We let SZK (respectively, PZK, CZK) denote the class of promise problems with statistical (respectively perfect, computational) zero-knowledge proof systems against the honest verifier.

Remarks on the Definitions

- (1) (*Honest verifiers*). We only require that the zero-knowledge condition to hold against the honest verifier, that is, the verifier that follows the protocol as specified. The usual definition requires the zero-knowledge property to hold against any polynomial-time verifier strategy. However, subsequent to this work, it has

³ A failure probability can also be allowed in the definition of statistical zero-knowledge, but this can easily be reduced to an 2^{-n^k} for any constant k by repeated trials and absorbed into the simulator deviation.

been shown that any proof system that is statistical zero knowledge against the honest verifier can be transformed into one that is zero knowledge against cheating verifiers [Goldreich et al. 1998]. Via this transformation, many of our results directly translate to the class of promise problems possessing statistical zero-knowledge proofs against cheating verifiers. This is discussed in detail in Section 3.4.

- (2) (*Error probabilities*). The completeness and soundness error probabilities can be made exponentially small without increasing the number of rounds, because zero-knowledge *against an honest verifier* is preserved under parallel repetition.
- (3) (*Strict polynomial-time simulation*). Following Goldreich [2001], we work with the variant of zero knowledge in which the simulator is required to run in *strict* polynomial time, with some probability of failure in the perfect case. The original definition by Goldwasser et al. [1989] allows the simulator to run in expected polynomial time, but with zero probability of failure. Our choice is not very restrictive, because we are only discussing honest-verifier statistical zero-knowledge and we do not know of any problems that require an expected polynomial-time simulator for the honest verifier. In addition, as shown in Section 3.1, our techniques can be used to prove that expected polynomial-time simulators and strict polynomial-time simulators are actually *equivalent* for public-coin statistical zero-knowledge proofs against an honest verifier.
- (4) (*Promise problems vs. languages*). Our definitions above generalize the original definitions of Goldwasser et al. [1989] from languages to promise problems, and we focus on the “promise class” **SZK** rather than the class of languages possessing statistical zero-knowledge proofs. As this was discussed in Section 2.1, here we simply reiterate the main justifications for this extension. First, for essentially all of our results, the fact that we prove them for the promise class only makes them stronger, by virtue of the fact that the promise class contains the language class. Second, several of the most important natural problems known to be in **SZK**, such as those in Goldreich and Kushilevitz [1993] and Goldreich and Goldwasser [2000], are not languages, but promise problems, so it may actually be preferable to study the promise class.

Our only result that requires new interpretation for the language class is the Completeness Theorem (Thm. 3.1 below). As the complete problem is a promise problem, it is not complete for the language class in the usual sense. Nevertheless, it still gives a characterization of the language class, in that a language has a statistical zero-knowledge proof *if and only if* it reduces to the complete problem. (This requires that **SZK** is closed under reductions (Corollary 4.3), which we will prove using the Completeness Theorem.)

We refer the reader to Goldreich [2001] for a more detailed discussion of the definitional issues with zero-knowledge proofs and a wider perspective on the subject.

3. The Completeness Theorem

3.1. THE COMPLETE PROBLEM. The main aim of this article is to demonstrate that **SZK** consists exactly of the problems that involve deciding whether two

efficiently samplable distributions are either far apart or close together. This can be formally captured by the following promise problem STATISTICAL DIFFERENCE (abbreviated SD):

$$\begin{aligned} \text{SD}_Y &= \left\{ (C_0, C_1) : \|C_0 - C_1\| > \frac{2}{3} \right\} \\ \text{SD}_N &= \left\{ (C_0, C_1) : \|C_0 - C_1\| < \frac{1}{3} \right\}. \end{aligned}$$

In the above definition, C_0 and C_1 are circuits; these define probability distributions as discussed in Section 2. The thresholds of $1/3$ and $2/3$ in this definition are not completely arbitrary; it is important for the Polarization Lemma of Section 3.2 that $(2/3)^2 > 1/3$.

We can now state the main theorem of the article.

THEOREM 3.1 (COMPLETENESS THEOREM).

STATISTICAL DIFFERENCE is complete for SZK.

The most striking thing about Theorem 3.1 is that it characterizes statistical zero knowledge *with no reference to interaction or zero knowledge*. Future investigation of the class SZK can focus on the single problem SD, instead of dealing with arbitrarily complicated protocols, problems, and simulators.⁴

We emphasize that the novelty of this result lies in the specific complete problem we present and not merely the *existence* of a complete promise problem. For example, it is fairly straightforward to construct a complete promise problem for PZK involving descriptions of Turing machines for the verifier and simulator. (See Appendix 6.) However, in contrast to SD, a complete problem constructed in this manner is essentially restatement of the definition of the class and therefore does not simplify the study of the class at all.

The proof of Theorem 3.1 comes in Sections 3.3 and 3.4 via two lemmas and a theorem of Okamoto [2000]. But first, we observe that a statement analogous to Theorem 3.1 can be made for BPP, if we generalize BPP to promise problems in the obvious way.

PROPOSITION 3.2. *If SD' is the promise problem obtained by modifying the definition of SD so that C_0 and C_1 only have 1 bit of output, then SD' is complete for BPP.*

PROOF. To see that SD' is in BPP, first observe that for circuits C_0 and C_1 (or any random variables) that always output either 0 or 1,

$$\|C_0 - C_1\| = |\Pr[C_0 = 1] - \Pr[C_1 = 1]|.$$

Thus, an estimate on $\|C_0 - C_1\|$ that is correct within an additive factor of $1/3$ can be obtained by sampling C_0 and C_1 polynomially many times and counting the number of ones that occur for each. This is sufficient to decide SD' .

Now we show that every promise problem Π in BPP reduces to SD' . Let A be the PPT machine that outputs 1 with probability greater than $2/3$ when $x \in \Pi_Y$, but

⁴ It should be noted that completeness is most meaningful for classes that are closed under reductions, which is not a priori clear for SZK. Later, we prove that SZK is indeed closed under reductions, as a corollary of the Completeness Theorem itself (Corollary 4.3).

outputs 1 with probability less than $1/3$ when $x \in \Pi_N$. Let $p(n)$ be a polynomial bound on the running time of A . Given an input x , we can, by standard techniques,⁵ produce in polynomial time a circuit C_x describing the computation of A on x for $p(|x|)$ steps. The input to C_x is the first $p(|x|)$ bits on the random tape of A and the output is the first bit on the output tape. Let D be a circuit that always outputs 0. Then $\|C_x - D\| = \Pr[A(x) = 1]$, so $x \mapsto (C_x, D)$ is a polynomial-time reduction from Π to SD' . \square

Proposition 3.2 remains true even if we allow C_0 and C_1 to output strings of logarithmic length. Other classes such as \mathbf{P} and $\mathbf{co-RP}$ can be obtained by modifying the definition of SD in a similar fashion (and changing the thresholds). This demonstrates that \mathbf{SZK} is a natural generalization of these well-known classes.

3.2. A POLARIZATION LEMMA. In this section, we exhibit a transformation that “polarizes” the statistical relationship between two distributions. That is, pairs of distributions that are statistically close become much closer and pairs of distributions that are statistically far apart become much further apart.

LEMMA 3.3 (POLARIZATION LEMMA).⁶ *There is a polynomial-time-computable function that takes a triple $(C_0, C_1, 1^k)$, where C_0 and C_1 are circuits, and outputs a pair of circuits (D_0, D_1) such that*

$$\begin{aligned} \|C_0 - C_1\| < \frac{1}{3} &\Rightarrow \|D_0 - D_1\| < 2^{-k} \\ \|C_0 - C_1\| > \frac{2}{3} &\Rightarrow \|D_0 - D_1\| > 1 - 2^{-k}. \end{aligned}$$

The usefulness of the Polarization Lemma comes from the fact that the two distributions it produces can be treated almost as if they were identically distributed or disjoint (i.e., statistical difference 0 and 1, respectively). Indeed, it will be essential in proving that SD (with thresholds of $2/3$ and $1/3$, as we’ve defined it) is in \mathbf{SZK} and we will make further use of it in deriving consequences of Theorem 3.1.

Superficially, it may seem that a Chernoff-bound argument is all that is needed to prove Lemma 3.3. However, Chernoff bounds are primarily useful for distinguishing between two events. This corresponds to *increasing* statistical difference, as formalized in the following “direct product” lemma:

LEMMA 3.4 (DIRECT PRODUCT LEMMA). *Let X and Y be distributions such that $\|X - Y\| = \epsilon$. Then for all k ,*

$$k\epsilon \geq \|\otimes^k X - \otimes^k Y\| \geq 1 - 2 \exp(-k\epsilon^2/2)$$

PROOF. The upper bound of $k\epsilon$ follows immediately from Fact 2.3, so we proceed to the proof of the lower bound. Recall, from the definition of statistical difference, that there must exist a set S such that

$$\Pr[X \in S] - \Pr[Y \in S] = \epsilon.$$

⁵ See, for example, [Papadimitriou 1994, Thms. 8.1 and 8.2].

⁶ The Polarization Lemma stated here is called the Amplification Lemma in [Sahai and Vadhan 1997]. We change the name here to stress that the Polarization Lemma does not merely increase statistical difference.

Let $p = \Pr[Y \in S]$, so $\Pr[X \in S] = p + \epsilon$. Hence, in k independent samples of X , the expected number of samples that lie in S is $(p + \epsilon)k$, whereas in k independent samples of Y , the expected number of samples that lie in S is pk . The Chernoff bound⁷ tells us that the probability that *at least* $(p + (\epsilon/2))k$ components of $\otimes^k Y$ lie in S is at most $\exp(-k\epsilon^2/2)$, whereas the probability that *at most* $(p + (\epsilon/2))k$ components of $\otimes^k X$ lie in S is at most $\exp(-k\epsilon^2/2)$. Let S' be the set of all k -tuples that contain more than $(p + (\epsilon/2))k$ components that lie in S . Then,

$$\|\otimes^k X - \otimes^k Y\| \geq \Pr[\otimes^k X \in S'] - \Pr[\otimes^k Y \in S'] \geq 1 - 2\exp(-k\epsilon^2/2). \quad \square$$

Note the gap between the upper and lower bounds in Lemma 3.4; the lower bound says that taking $O(1/\epsilon^2)$ copies is sufficient to increase statistical difference from ϵ to a constant, while the upper bound says that $\Omega(1/\epsilon)$ copies are necessary. This gap is inherent, as the following example illustrates: Taking X and Y to be distributions on $\{0, 1\}$ that are 1 with probability 1 and $1 - \epsilon$, respectively, we see that the statistical difference between $\otimes^k X$ and $\otimes^k Y$ is exactly $1 - (1 - \epsilon)^k$, which is a constant for $k = \Theta(1/\epsilon)$. On the other hand, when X and Y are 1 with probability $(1 + \epsilon)/2$ and $(1 - \epsilon)/2$, respectively, it can be shown that $k = \Theta(1/\epsilon^2)$ copies are necessary to increase the statistical difference to a constant. Furthermore, in this latter example, $\|X \otimes X - Y \otimes Y\| = \epsilon = \|X - Y\|$, so we cannot even hope to show that statistical difference always increases for every $k > 1$ (as pointed out to us by Madhu Sudan).

Notice that the Direct Product Lemma 3.4 is *not* sufficient to prove the Polarization Lemma, because it only increases statistical difference, whereas we would like to increase statistical difference in some cases and decrease it in others. However, it does drive larger values of the statistical difference to 1 more quickly than it drives smaller values to 1, so it is a step in the right direction. The following lemma provides a complementary technique that decreases the statistical difference to 0, with small values going to 0 faster than large values.

LEMMA 3.5 (XOR LEMMA). *There is a polynomial-time computable function that maps a triple $(C_0, C_1, 1^k)$, where C_0 and C_1 are circuits, to a pair of circuits (D_0, D_1) such that $\|D_0 - D_1\| = \|C_0 - C_1\|^k$. Specifically, D_0 and D_1 are defined as follows:*

D_0 : Uniformly select $(b_1, \dots, b_k) \in \{0, 1\}^k$ such that $b_1 \oplus \dots \oplus b_k = 0$, and output a sample of $C_{b_1} \otimes \dots \otimes C_{b_k}$.
 D_1 : Uniformly select $(b_1, \dots, b_k) \in \{0, 1\}^k$ such that $b_1 \oplus \dots \oplus b_k = 1$, and output a sample of $C_{b_1} \otimes \dots \otimes C_{b_k}$.

In order to prove this lemma, we employ a generalization of the technique used in De Santis et al. [1994] to represent the logical AND of statements about GRAPH NONISOMORPHISM. This tool is described in the following Proposition.

⁷ For the formulation of the Chernoff bound we use, see, for example, the formulation of Hoeffding's inequality in Hofri [1995, Sect. 7.2.1].

PROPOSITION 3.6. *Let X_0, X_1, Y_0, Y_1 be any random variables, and define the following pair of random variables:*

Z_0 : *Choose $a, b \in_R \{0, 1\}$ such that $a \oplus b = 0$. Output a sample of $X_a \otimes Y_b$.*

Z_1 : *Choose $a, b \in_R \{0, 1\}$ such that $a \oplus b = 1$. Output a sample of $X_a \otimes Y_b$.*

Then $\|Z_0 - Z_1\| = \|X_0 - X_1\| \cdot \|Y_0 - Y_1\|$.

The statistical difference between X_0 and X_1 (or Y_0 and Y_1) measures the advantage a computationally unbounded party has, over random guessing, in guessing b given a sample from X_b , where b is selected uniformly from $\{0, 1\}$. (This view of statistical difference will become more apparent in the subsequent section.) Intuitively, the above Proposition says that the advantage one has in guessing the XOR of two independent bits is the product of the advantages one has for guessing each individual bit.

PROOF.

$$\begin{aligned} \|Z_0 - Z_1\| &= \frac{1}{2} |Z_0 - Z_1|_1 \\ &= \frac{1}{2} \left| \left(\frac{1}{2} X_0 \otimes Y_0 + \frac{1}{2} X_1 \otimes Y_1 \right) - \left(\frac{1}{2} X_1 \otimes Y_0 + \frac{1}{2} X_0 \otimes Y_1 \right) \right|_1 \\ &= \frac{1}{4} |(X_0 - X_1) \otimes (Y_0 - Y_1)|_1 \\ &= \left(\frac{1}{2} |X_0 - X_1|_1 \right) \cdot \left(\frac{1}{2} |Y_0 - Y_1|_1 \right) \\ &= \|X_0 - X_1\| \cdot \|Y_0 - Y_1\| \end{aligned}$$

Recall that the penultimate equality above follows because $|v \otimes w| = |v| \cdot |w|$. \square

Proposition 3.6 and an induction argument establish Lemma 3.5. Yao's XOR Lemma [1982] (cf., Goldreich et al. [1995]) can be seen as an analogue of Lemma 3.5 in the computational setting, where the analysis is much more difficult.⁸

Now we combine the Direct Product and XOR constructions of Lemmas 3.4 and 3.5 to prove Lemma 3.3. The Direct Product Lemma gives a way to increase statistical difference with large values going to 1 faster than small values. Similarly, the XOR Lemma shows how to decrease statistical difference with small values going to 0 faster than large values. Intuitively, alternating these procedures should "polarize" large and small values of statistical difference, pushing them closer to 1 and 0, respectively. A similar alternation between procedures with complementary effects was used by Ajtai and Ben-Or [1984] to amplify the success probability of randomized constant-depth circuits.

⁸ To see the analogy, recall that Yao's XOR Lemma considers the maximum advantage an *efficient* algorithm has, over random guessing, in computing a bit b from string x when they are selected according to some distribution $(b, x) \leftarrow (B, X)$ (e.g., X is uniform and B is a hardcore bit of $f^{-1}(X)$ for some one-way permutation f). It states that the maximum advantage an efficient algorithm has in computing the XOR $b_1 \oplus \dots \oplus b_k$ from (x_1, \dots, x_k) decreases exponentially with k when the pairs (b_i, x_i) are independently distributed according to (B, X) .

PROOF. Let $\ell = \lceil \log_{4/3} 6k \rceil$. Apply Lemma 3.5 to the triple $(C_0, C_1, 1^\ell)$ to produce (C'_0, C'_1) such that if

$$\|C_0 - C_1\| < \frac{1}{3} \Rightarrow \|C'_0 - C'_1\| < \left(\frac{1}{3}\right)^\ell$$

$$\|C_0 - C_1\| > \frac{2}{3} \Rightarrow \|C'_0 - C'_1\| > \left(\frac{2}{3}\right)^\ell.$$

Let $m = 3^{\ell-1}$. Let $C''_0 = \otimes^m C'_0$ and let $C''_1 = \otimes^m C'_1$. Then, by Fact 2.3 and the Direct Product Lemma,

$$\|C_0 - C_1\| < \frac{1}{3} \Rightarrow \|C''_0 - C''_1\| < \frac{1}{3}$$

$$\|C_0 - C_1\| > \frac{2}{3} \Rightarrow \|C''_0 - C''_1\| > 1 - 2 \exp\left(\frac{-3^{\ell-1}(2/3)^{2\ell}}{2}\right) > 1 - 2 \exp(-k).$$

Finally, apply the transformation of Lemma 3.5 one more time to $(C''_0, C''_1, 1^k)$ to produce (D_0, D_1) such that

$$\|C_0 - C_1\| < \frac{1}{3} \Rightarrow \|D_0 - D_1\| < 3^{-k} < 2^{-k}$$

$$\|C_0 - C_1\| > \frac{2}{3} \Rightarrow \|D_0 - D_1\| > (1 - 2 \exp(-k))^k$$

$$> 1 - 2k \exp(-k) > 1 - 2^{-k}. \quad \square$$

Notice that the above analysis relies on the fact that $(2/3)^2 > (1/3)$, so it will not work if $2/3$ and $1/3$ are replaced by, say, $.51$ and $.49$. We do not know how to prove such a Polarization Lemma for arbitrary constant thresholds. We can however extend it to thresholds α and β , where $\alpha^2 > \beta$, and the running time will be polynomial in $1/(\alpha - \beta)$ and $\exp(1/\delta)$, where δ is defined by $\alpha^{2+\delta} = \beta$, along with the input size. See Sahai and Vadhan [1999] for more details.

3.3. A PROTOCOL FOR STATISTICAL DIFFERENCE. In this section, we show that SD has a simple two-message statistical zero-knowledge proof system, which is a generalization of the standard protocols for QUADRATIC NONRESIDUOSITY Goldwasser et al. [1989], and GRAPH NONISOMORPHISM [Goldreich et al. 1991]. Intuitively, if two distributions are statistically far apart, then, when given a random sample from one of the distributions, a computationally unbounded party should have a good chance of guessing from which distribution it came. However, if the two distributions are statistically very close, even a computationally unbounded party should not have much better than a 50% chance of guessing correctly. This suggests the following two-message protocol for SD (note that this protocol is *not* a public-coin protocol):

Zero-knowledge Proof System for SD

Input: (C_0, C_1) (such that either $\|C_0 - C_1\| > 2/3$ or $\|C_1 - C_1\| < 1/3$)

1. V, P : Compute $(D_0, D_1) = \mathbf{Polarize}(C_0, C_1, 1^n)$, where $n = |(C_0, C_1)|$.
2. V : Flip one random coin $r \leftarrow \{0, 1\}$. Let z be a sample of D_r . Send z to P .
3. P : If $\Pr[D_0 = z] > \Pr[D_1 = z]$, answer 0, otherwise answer 1.
4. V : Accept if P 's answer equals r , reject otherwise.

LEMMA 3.7. *The above is a statistical zero-knowledge proof system for SD, with soundness error $(1/2) + 2^{-n}$, and completeness error and simulator deviation both 2^{-n} . Thus, $SD \in \text{SZK}$.*

PROOF. We observe that the prover strategy given in the protocol is optimal (that is, maximizes the verifier's acceptance probability), and use this to bound both the soundness and completeness error. The simulator deviation will then follow easily.

Indeed the optimality of the given prover strategy follows from a standard argument: Consider any prover P^* . Suppose for some z the prover P^* fails to follow the strategy we present. If $\Pr[D_0 = z] \neq \Pr[D_1 = z]$, this means that with nonzero probability, P^* chooses the distribution in which z is less likely to occur. Then, conditioned on z , the success probability of P^* will certainly be lower than that of the prover in our protocol. If $\Pr[D_0 = z] = \Pr[D_1 = z]$, the prover has no information about r , so no matter what strategy it uses, it has exactly even odds of guessing correctly. Since these observations hold for all z , the given prover is optimal.

We now analyze the probability of success of the optimal prover. Recall that $\|D_0 - D_1\| = \Pr[D_0 \in S] - \Pr[D_1 \in S]$ for $S = \{z : \Pr[D_0 = z] > \Pr[D_1 = z]\}$. The probability that the optimal prover guesses correctly is exactly

$$\begin{aligned} \frac{1}{2} \Pr[D_0 \in S] + \frac{1}{2} \Pr[D_1 \notin S] &= \frac{1}{2} (\Pr[D_0 \in S] + 1 - \Pr[D_1 \in S]) \\ &= \frac{1 + \|D_0 - D_1\|}{2}. \end{aligned}$$

By Lemma 3.3, $\|D_0 - D_1\| > 1 - 2^{-n}$ when (C_0, C_1) is a YES instance of SD, and $\|D_0 - D_1\| < 2^{-n}$ when (C_0, C_1) is a NO instance. Hence, the probability that the prover convinces the verifier to accept is greater than $(1 + 1 - 2^{-n})/2 > 1 - 2^{-n}$ for YES instances, and less than $(1 + 2^{-n})/2 < 1/2 + 2^{-n}$ for NO instances. This immediately gives the completeness error; the soundness error also follows because we considered the optimal prover strategy.

Now, notice that when the prover answers correctly, all the verifier receives from the prover is the value of r , which the verifier already knew. Thus, since we have shown that the prover is answering correctly with all but exponentially small probability, intuitively the verifier learns nothing. To turn this intuition into a proof of statistical zero knowledge, we consider the following probabilistic polynomial-time simulator: On input (C_0, C_1) , the simulator first computes $(D_0, D_1) = \text{Polarize}(C_0, C_1, 1^n)$, where $n = |(C_0, C_1)|$. The simulator then flips one random coin $r \in \{0, 1\}$. If $r = 0$, it samples z from D_0 ; otherwise, it samples z from D_1 . The simulator then outputs a conversation in which the verifier sends z to the prover, and the prover responds with r . The simulator also outputs the random coins it used to generate r and z as the coins of the verifier. Thus, the simulator presented here always outputs conversations in which the prover responds correctly. Except for the prover's response, all other components of the simulator's output distribution are distributed identically to the verifier's view of the real interaction. Hence, the simulator deviation is bounded by the probability that the prover responds incorrectly in the real interaction, which we have already argued is at most 2^{-n} in the case of YES instances. \square

Note that the above proof system remains complete and sound even without polarization, but for the zero-knowledge property, we need to make the statistical difference very close to 1 on YES instances.

By using a security parameter k rather than n in the call to **Polarize**, both the completeness error and simulator deviation can be reduced to 2^{-k} . Thus, even very short assertions about SD can be proven with very high security. Contrast this with the original definition of SZK [Goldwasser et al. 1989], which only requires that the simulator deviation vanish as an *negligible* function of the *input length*. This property has obvious cryptographic significance, so we formulate it more precisely in Section 3.1.

3.4. **SZK-HARDNESS OF SD.** The other main lemma we prove to show that SD is complete for SZK follows:

LEMMA 3.8. *Suppose promise problem Π has a public coin statistical zero-knowledge proof system. Then there exist PPTs A and B and a negligible function α such that*

$$\begin{aligned} x \in \Pi_Y &\Rightarrow \|A(x) - B(x)\| \leq \alpha(|x|), \text{ and} \\ x \in \Pi_N &\Rightarrow \|A(x) - B(x)\| \geq 1 - 2^{-\Omega(|x|)}. \end{aligned}$$

We defer the proof of this lemma to Section 3.5, and first observe how it gives a reduction to SD for problems with public-coin statistical zero-knowledge proofs.

COROLLARY 3.9. *Suppose promise problem Π has a public-coin statistical zero-knowledge proof system. Then Π reduces to $\overline{\text{SD}}$. (Equivalently, $\overline{\Pi}$ reduces to SD.)*

PROOF. First, apply Lemma 3.8 to obtain A and B , with $p(|x|)$ being a polynomial bound on the running times of $A(x)$ and $B(x)$. Given a string x , we can, by standard techniques,⁹ produce in polynomial-time circuits C_0 and C_1 simulating the computation of A and B , respectively, on x for $p(|x|)$ steps. The inputs to C_0 and C_1 are the first $p(|x|)$ bits on the random tapes of A and B and the outputs are the first $p(|x|)$ positions on the output tapes. Then $\|C_0 - C_1\| = \|A(x) - B(x)\|$, which is at most $\alpha(|x|) < 1/3$ if $x \in \Pi_Y$ and at least $1 - 2^{-|x|} > 2/3$ if $x \in \Pi_N$ (for all sufficiently long x). So $x \mapsto (C_0, C_1)$ is a reduction from Π to $\overline{\text{SD}}$ (for all but finitely many x). \square

The final ingredient in the proof of Theorem 3.1 is a theorem of Okamoto [2000], which we state in terms of promise problems.¹⁰

THEOREM 3.10 [OKAMOTO 2000, THM. 1]. *If a promise problem Π has a statistical zero-knowledge proof system, then Π has a public-coin statistical zero-knowledge proof system.*

Now it will be easy to show that SD is complete for SZK.

⁹ See, for example, Papadimitriou [1994, Thms. 8.1 and 8.2].

¹⁰ Okamoto stated his result in terms of languages, but the proof readily extends to promise problems (cf., Goldreich and Vadhan [1999]).

PROOF OF THEOREM 3.1. Lemma 3.7 tells us that $\text{SD} \in \text{SZK}$, so we only need to show that every problem in SZK reduces to SD . Corollary 3.9 and Theorem 3.10 imply that every problem $\Pi \in \text{SZK}$ reduces to $\overline{\text{SD}}$. In particular, SD reduces to $\overline{\text{SD}}$, or, equivalently, $\overline{\text{SD}}$ reduces to SD . Composing reductions, it follows that every problem $\Pi \in \text{SZK}$ reduces to SD . \square

3.5. PROOF OF LEMMA 3.8

Intuition. Recall that we wish to construct a pair of probabilistic polynomial-time machines A and B such that if $x \in \Pi_Y$, the distributions $A(x)$ and $B(x)$ are statistically very close, but when $x \in \Pi_N$, $A(x)$ and $B(x)$ are far apart. We are given that Π has a *public-coin* statistical zero-knowledge proof system. A natural place to search for the desired distributions is in the output of the simulator for this proof system. We wish to find properties of the simulator's output that (1) distinguish the case $x \in \Pi_Y$ from $x \in \Pi_N$ and (2) are captured by the statistical difference between samplable distributions. Following Aiello and Håstad [1991], we think of the simulator as describing the moves of a *virtual prover* and a *virtual verifier*.

In the case that $x \in \Pi_Y$, we have strong guarantees on the simulator's output. Namely, its output distribution is statistically very close to the real interaction. In particular, it outputs accepting conversations with high probability and the virtual prover and verifier "behave" similarly to the real prover and verifier.

When $x \in \Pi_N$, there are two cases. If the simulator outputs accepting conversations with low probability, this easily distinguishes it from the simulator output when $x \in \Pi_Y$. However, it is possible that the simulator will output accepting conversations with high probability even when $x \in \Pi_N$. This means that the virtual prover is doing quite well in fooling the virtual verifier. This naturally suggests a strategy for a real prover—imitate the virtual prover's behavior. Such a prover, called a *simulation-based prover*, was introduced by Fortnow [1989] and is a crucial construct in our proof.

The soundness of the proof system tells us that the simulation-based prover cannot hope to convince the *real* verifier with high probability. There must be a reason for this discrepancy between the success rates of the virtual prover and the simulation-based prover. Since the virtual prover behaves exactly like the simulation-based prover, it must be that the virtual verifier does not behave like the real verifier. Note that, in a *public-coin* proof system, the behavior of the real verifier is extremely simple: it chooses each of its messages independently and uniformly at random. If the virtual verifier does not behave as the real verifier, then it must be that either the virtual verifier's messages are *far from uniform*, or that they are *dependent on previous messages*.

We therefore exhibit two efficiently samplable distributions, one describing the messages of the real verifier and the other describing the messages of the virtual verifier. We show that if $x \in \Pi_Y$, these two distributions must be nearly identical; whereas if $x \in \Pi_N$, they must be far apart.

The basic approach described above is quite similar to the approaches developed in previous work on SZK , such as that of Fortnow [1989], Aiello and Håstad [1991], and Ostrovsky [1991]. However, by focusing on public-coin proofs (thanks to Theorem 3.10 [Okamoto 2000]), we are able to carry out a cleaner analysis and reach a novel conclusion (namely, the Completeness Theorem).

Notation. Let (P, V) be a public-coin interactive proof system for a promise problem Π , which is (honest-verifier) statistical zero knowledge, and let S be a simulator for this proof system. Without loss of generality, we may assume that the interaction of P and V on input x always has $2r(|x|)$ exchanged messages, with V sending the first message and each message consisting of exactly $q(|x|)$ bits, for some polynomials q and r . Moreover, it may be assumed that S 's output always consists of $2r(|x|)$ strings of length $q(|x|)$. The output of S and the conversation between P and V on input x will be written in the form $S(x) = (c_1, p_1, \dots, c_r, p_r)_S$ and $(P, V)(x) = (c_1, p_1, \dots, c_r, p_r)_{(P, V)}$, respectively, where c_1, \dots, c_r represent the messages (equivalently coin tosses, since we are in the public-coin setting) of V , p_1, \dots, p_r represent the prover messages, and $r = r(|x|)$. (Dependence on x will often be omitted in this manner for notational convenience.) We use notation such as $(c_i)_S$ for the random variable obtained by running S once and taking the c_i -component of its output. More generally, partial conversation transcripts will be written like $(c_1, p_1, c_2, p_2)_S$. We call a conversation transcript $(c_1, p_1, \dots, c_r, p_r)$ that would make V accept (respectively, reject) an *accepting conversation* (respectively, *rejecting conversation*). We denote by $U(n)$ the uniform distribution on strings of length n .

The Proof. In order to formalize the above intuition, a definition of the simulation-based prover needs to be given. This is the prover P^* that imitates the virtual prover, that is, P^* does the following to compute its next message when the current conversation transcript is (c_1, p_1, \dots, c_i) :

If $S(x)$ outputs conversations that begin with (c_1, p_1, \dots, c_i) with probability 0, then output 0.

Else output $y \in \{0, 1\}^{q(|x|)}$ with probability

$$p_y = \Pr[S(x) \text{ begins with } (c_1, p_1, \dots, c_i, y) | S(x) \text{ begins with } (c_1, p_1, \dots, c_i)].$$

In order to analyze the success probability of P^* , we first compare the output of S to the actual conversations between P^* and V . For $i = 1, \dots, r$, consider the distributions $A_i = (c_1, p_1, \dots, c_{i-1}, p_{i-1}, c_i)_S$ and $B_i = (c_1, p_1, \dots, c_{i-1}, p_{i-1})_S \otimes U(q(|x|))$, and let $\epsilon_i = \|A_i - B_i\|$. The last component of A_i is a move of the virtual verifier whereas the last component of B_i is chosen uniformly and independently of the history, just like a move of the real verifier. Thus, the ϵ_i 's measure the similarity between the behavior of the virtual verifier and the real verifier.

The following claim formalizes our intuition that, if the virtual verifier and real verifier have similar behavior, then the interaction between the simulation-based prover P^* and the real verifier V is similar to the interaction between the virtual prover and virtual verifier (as described by the simulator).

$$\text{CLAIM 3.11. } \|S(x) - (P^*, V)(x)\| \leq \sum_{i=1}^r \epsilon_i.$$

PROOF OF CLAIM. Let $C_i^S = (c_1, p_1, \dots, c_i)_S$ be the random variable of partial simulator transcripts ending with the i th coins of the virtual verifier. Let $P_i^S = (c_1, p_1, \dots, c_i, p_i)_S$ be the random variable of partial transcripts ending with the i th virtual prover response. Similarly define C_i^* and P_i^* as partial conversation transcripts of (P^*, V) . The aim is to show that at round k , the statistical difference

TABLE I. THE COMPONENTS OF A AND B

Algorithm A		Algorithm B	
$A_0(x)$	Run $S(x)$ for $ x $ repetitions. Output '1' if the majority are accepting conversations and '0' otherwise.	$B_0(x)$	Output 1.
$A_i(x)$	Run $S(x)$ to output $(c_1, p_1, \dots, c_i)_{S(x)}$.	$B_i(x)$	Run $S(x)$ and flip $q(x)$ more coins to output $(c_1, p_1, \dots, c_{i-1}, p_{i-1})_{S(x)} \otimes U(q(x))$.

grows by at most ϵ_k . Formally, it will be shown by induction on k that

$$\|P_k^S - P_k^*\| \leq \sum_{i=1}^k \epsilon_i$$

The case $k = 0$ is trivial. For general k , first note that since P^* gives a response chosen according to the same distribution as the virtual prover, adding these responses to the conversations cannot increase the statistical difference (by Fact 2.4). That is,

$$\|P_{k+1}^S - P_{k+1}^*\| = \|C_{k+1}^S - C_{k+1}^*\|.$$

The idea now is to extract the parts of $\|C_{k+1}^S - C_{k+1}^*\|$ corresponding to ϵ_{k+1} and observe that what is left is simply the error from the previous round. Note that $C_{k+1}^* = P_k^* \otimes U(q(|x|))$, since the real verifier's coins are always uniform and independent from what came before.

Then, applying Fact 2.3 and the Triangle Inequality,

$$\begin{aligned} \|C_{k+1}^S - C_{k+1}^*\| &\leq \|C_{k+1}^S - P_k^S \otimes U(q(|x|))\| + \\ &\quad \|P_k^S \otimes U(q(|x|)) - P_k^* \otimes U(q(|x|))\| \\ &\leq \epsilon_{k+1} + \|P_k^S - P_k^*\| + \|U(q(|x|)) - U(q(|x|))\| \\ &\leq \epsilon_{k+1} + \sum_{i=1}^k \epsilon_i. \end{aligned}$$

This completes the induction. Since $P_r^S = S(x)$ and $P_r^* = (P^*, V)(x)$, the Claim is proved. \square

We are now ready to construct the distributions we seek. The two distributions A and B each consist of $r + 1$ components, shown in Table I. A is the algorithm whose output on input x is $(A_0(x), A_1(x), \dots, A_r(x))$, all run independently, and B is the algorithm whose output is $(B_0(x), B_1(x), \dots, B_r(x))$, all run independently. Recall that, for $i \geq 1$, ϵ_i is the statistical difference between A_i and B_i .

We show that the statistical difference between A and B is negligible if $x \in \Pi_Y$ and is noticeable if $x \in \Pi_N$. Amplifying this gap by repetition will yield Lemma 3.8.

CLAIM 3.12. *There exists a negligible function α such that if $x \in \Pi_Y$, then $\|A(x) - B(x)\| \leq \alpha(|x|)$.*

PROOF OF CLAIM. The statistical difference between $A(x)$ and $B(x)$ is bounded above by the sum of the statistical differences between $A_i(x)$ and $B_i(x)$ over $i = 0, \dots, r(|x|)$ (by Fact 2.3). First, let's examine $i = 0$. Since $S(x)$ outputs a conversation that makes V accept with probability at least $2/3 - \text{neg}(|x|)$, the

Chernoff bound implies that $\Pr[A_0(x) = 1] = 1 - 2^{-\Omega(|x|)}$, so the statistical difference between A_0 and B_0 is negligible. For $i \geq 1$, recall that in the real conversations of P and V , the verifier's coins are truly uniform and independent from prior rounds, so $\|A_i(x) - B_i(x)\|$ should essentially be bounded by the statistical difference between the simulator's output and the real interaction. This is in fact true, as (omitting x from the notation):

$$\begin{aligned} \|A_i - B_i\| &\leq \|A_i - (c_1, p_1, \dots, c_i)_{P,V}\| + \|(c_1, p_1, \dots, c_i)_{P,V} - B_i\| \\ &\leq \|S - (P, V)\| + \|S - (P, V)\|. \end{aligned}$$

Thus,

$$\|A(x) - B(x)\| \leq 2^{-\Omega(|x|)} + 2r(|x|) \cdot \|S(x) - (P, V)(x)\|,$$

which is negligible since $\|S(x) - (P, V)(x)\|$ is negligible and $r(x)$ is polynomial. \square

CLAIM 3.13. *If $x \in \Pi_N$, then $\|A(x) - B(x)\| \geq 1/12r(|x|)$.*

PROOF OF CLAIM. It suffices to show that for some i , $\epsilon_i = \|A_i(x) - B_i(x)\| > 1/12r(|x|)$ (by Fact 2.4). We deal with two cases depending on the probability that S outputs an accepting conversation.

Case 1. $\Pr[S(x) \text{ accepts}] \leq 5/12$. Then, by the Chernoff bound, $\Pr[A_0(x) = 1] \leq 2^{-\Omega(|x|)}$, so the statistical difference between $A_0(x)$ and $B_0(x)$ is at least $1 - 2^{-\Omega(|x|)} > 1/12r(|x|)$.

Case 2. $\Pr[S(x) \text{ accepts}] > 5/12$. Then, since $\Pr[(P^*, V)(x) \text{ accepts}]$ is at most $1/3$, we must have

$$\sum_{i=0}^r \epsilon_i \geq \|S(x) - (P^*, V)(x)\| > \frac{5}{12} - \frac{1}{3} = \frac{1}{12}.$$

Thus, at least one ϵ_i must be greater than $1/12r(|x|)$. \square

Now consider the samplable distributions $\hat{A}(x) = \otimes^{s(|x|)} A(x)$ and $\hat{B}(x) = \otimes^{s(|x|)} B(x)$, where $s(n) = n \cdot r(n)^2$. If $x \in \Pi_Y$, $\|\hat{A}(x) - \hat{B}(x)\| \leq s(|x|) \cdot \|A(x) - B(x)\|$, which is still negligible. If $x \in \Pi_N$, then, by the Direct Product Lemma (Lemma 3.4), $\|\hat{A}(x) - \hat{B}(x)\| \geq 1 - 2^{-\Omega(|x|)}$. This completes the proof of Lemma 3.8. \square

We illustrate the constructions in this lemma and the statistical zero-knowledge proof system for STATISTICAL DIFFERENCE for the specific example of GRAPH ISOMORPHISM in Appendix C.

4. Applications

4.1. EFFICIENT STATISTICAL ZERO-KNOWLEDGE PROOFS. The proof system for STATISTICAL DIFFERENCE given in Section 3.3 has a number of desirable features. It is very efficient in terms of communication and interaction, and the simulator deviation can be made exponentially small in a security parameter (that can be varied independently of the input length). By the Completeness Theorem, it follows that every problem in SZK also has a proof system with these properties.

We begin by formalizing one of these properties that was informally discussed in Section 3.3.

Definition 4.1. An interactive protocol (P, V) is called a *security-parameterized statistical zero-knowledge proof system* for a promise problem Π if there exists a PPT simulator S , a negligible function $\alpha(k)$ (called the *simulator deviation*), and completeness and soundness errors $c(k)$ and $s(k)$ such that for all strings x and all $k \in \mathbb{N}$,

- (1) If $x \in \Pi_Y$, then $\Pr[(P, V)(x, 1^k) = \text{accept}] \geq 1 - c(k)$.
- (2) If $x \in \Pi_N$, then for all P^* , $\Pr[(P^*, V)(x, 1^k) = \text{accept}] \leq s(k)$.
- (3) If $x \in \Pi_Y$, then $\|S(x, 1^k) - \text{View}_{P,V}(x, 1^k)\| \leq \alpha(k)$.

As usual, we require that $c(k)$ and $s(k)$ are computable in time $\text{poly}(k)$ and $1 - c(k) > s(k) + 1/\text{poly}(k)$

We now describe the efficient proof systems inherited by all of SZK.

COROLLARY 4.2. *Every problem in SZK possesses a security-parameterized statistical zero-knowledge proof system with the following properties:*

- (1) *Simulator deviation 2^{-k} , completeness error 2^{-k} , and soundness error $1/2 + 2^{-k}$.*
- (2) *The prover and verifier exchange only 2 messages.*
- (3) *The prover sends only 1 bit to the verifier.*
- (4) *The prover is deterministic.*

PROOF. Let Π be any promise problem in SZK. Let f be the reduction from Π to SD guaranteed by the Completeness Theorem. A protocol with the desired properties for Π can be obtained as follows: on input $(x, 1^k)$, execute the proof system for SD, given in Section 3.3, on input $f(x)$ and using k rather than n in the call to **Polarize**. \square

4.2. CLOSURE PROPERTIES. In this section, we prove several closure properties of SZK. The first, closure under reductions, is a direct consequence of the “security parametrization” property shown to hold for SZK in the previous section.

COROLLARY 4.3. *SZK is closed under (Karp) reductions. That is, if $\Pi \in \text{SZK}$ and Γ reduces to Π , then $\Gamma \in \text{SZK}$.*

PROOF. By Corollary 4.2, Π has a security-parameterized statistical zero-knowledge proof. A statistical zero-knowledge proof for Γ can be obtained as follows: On input x , the prover, verifier, and simulator run the security-parameterized proof for Π on input $(f(x), 1^{|x|})$, where f is the reduction from Γ to Π . \square

The security-parametrization property is necessary in the above proof, because an arbitrary reduction f could potentially shrink string lengths dramatically and we want the simulator deviation to be negligible as a function of $|x|$, not $|f(x)|$.

Next, we show how Okamoto’s result that SZK is closed under complement follows immediately from our proof of Completeness Theorem.

COROLLARY 4.4 ([OKAMOTO 2000, THM. 2]). *SZK is closed under complement, even for promise problems.*

PROOF. Let Π be any problem in **SZK**. By Theorem 3.10 and Corollary 3.9, $\overline{\Pi}$ reduces to **SD**, which is in **SZK**. By Corollary 4.3, $\overline{\Pi} \in \mathbf{SZK}$. \square

Before moving on to additional closure properties, we deduce the upper bounds of Fortnow [1989] and Aiello and Håstad [1991] on the complexity of **SZK**.

COROLLARY 4.5 ([FORTNOW 1989; AIELLO AND HÅSTAD 1991]). $\mathbf{SZK} \subset \mathbf{AM} \cap \mathbf{co-AM}$, where **AM** denotes the class of problems possessing constant-message interactive proofs.

PROOF. Immediate from Corollaries 4.2 and 4.4. \square

Above, we have seen that **SZK** satisfies a computational closure property (Corollary 4.3) and a Boolean closure property (Corollary 4.4 [Okamoto 2000]). Now we will exhibit a stronger closure property, which can be viewed as both a computational one and a Boolean one: given an arbitrary Boolean formula whose atoms are statements about membership in *any* problem in **SZK**, one can efficiently construct a statistical zero-knowledge interactive proof for its validity. Note that such a property does not follow immediately from the fact that a class is closed under intersection, union, and complementation, because applying these more than a constant number of times could incur a superpolynomial cost in efficiency, while we ask that the construction can be done efficiently with respect to the size of the formula. We achieve this by applying a construction of De Santis et al. [1994] to **STATISTICAL DIFFERENCE**, and then appealing to the Completeness Theorem.

We begin with some definitions describing precisely what kind of Boolean closure properties we will achieve. (Later, we will see how it can also be interpreted as closure under a certain class of polynomial-time reductions.) In order to deal with instances of promise problems that violate the promise, we will work with an extension of Boolean algebra that includes an additional “ambiguous” value \star .

Definition 4.6. A *partial assignment* to variables v_1, \dots, v_k is a k -tuple $\bar{a} = (a_1, \dots, a_k) \in \{0, 1, \star\}^k$. For a propositional formula (or circuit) ϕ on variables v_1, \dots, v_k , the evaluation $\phi(\bar{a})$ is recursively defined as follows:

$$v_i(\bar{a}) = a_i \quad \left| \quad \begin{array}{l} (\phi \wedge \psi)(\bar{a}) = \begin{cases} 1 & \text{if } \phi(\bar{a}) = 1 \text{ and } \psi(\bar{a}) = 1 \\ 0 & \text{if } \phi(\bar{a}) = 0 \text{ or } \psi(\bar{a}) = 0 \\ \star & \text{otherwise} \end{cases} \\ (\neg\phi)(\bar{a}) = \begin{cases} 1 & \text{if } \phi(\bar{a}) = 0 \\ 0 & \text{if } \phi(\bar{a}) = 1 \\ \star & \text{if } \phi(\bar{a}) = \star \end{cases} \quad (\phi \vee \psi)(\bar{a}) = \begin{cases} 1 & \text{if } \phi(\bar{a}) = 1 \text{ or } \psi(\bar{a}) = 1 \\ 0 & \text{if } \phi(\bar{a}) = 0 \text{ and } \psi(\bar{a}) = 0 \\ \star & \text{otherwise} \end{cases} \end{array} \right.$$

Note that $\phi(\bar{a})$ equals 1 (respectively, 0) for some partial assignment \bar{a} , then $\phi(\bar{a}')$ also equals 1 (respectively, 0) for every Boolean \bar{a}' obtained by replacing every \star in \bar{a} with either a 0 or 1. The converse, however, is not true: The formula $\phi = v \vee \neg v$ evaluates to 1 on every Boolean assignment, yet is not 1 when evaluated at \star . Thus, the “law of excluded middle” $\phi \vee \neg\phi \equiv 1$ no longer holds in this setting. However, other identities in Boolean algebra, such as De Morgan’s laws (e.g., $\neg(\phi \vee \psi) \equiv \neg\phi \wedge \neg\psi$), do remain true.

Definition 4.7. For a promise problem Π , the *characteristic function* of Π is the map $\chi_\Pi : \{0, 1\}^* \rightarrow \{0, 1, \star\}$ given by

$$\chi_\Pi(x) = \begin{cases} 1 & \text{if } x \in \Pi_Y \\ 0 & \text{if } x \in \Pi_N \\ \star & \text{otherwise} \end{cases}$$

Definition 4.8. For any promise problem Π , we define a new promise problem $\Phi(\Pi)$ as follows:

$$\begin{aligned} \Phi(\Pi)_Y &= \{(\phi, x_1, \dots, x_k) : \phi(\chi_\Pi(x_1), \dots, \chi_\Pi(x_k)) = 1\} \\ \Phi(\Pi)_N &= \{(\phi, x_1, \dots, x_k) : \phi(\chi_\Pi(x_1), \dots, \chi_\Pi(x_k)) = 0\}, \end{aligned}$$

where ϕ is a k -ary propositional formula. $\text{Mon}(\Pi)$ is defined analogously, except that only monotone ϕ are considered.¹¹

De Santis et al. [1994] show that $\text{Mon}(L) \in \text{SZK}$ for any language L that is random self-reducible, whose complement is self-reducible, or whose complement has a noninteractive statistical zero-knowledge proof. They also give statistical zero-knowledge proofs for some simple statements involving a random-self-reducible language and its complement. Damgård and Cramer [1996] extend these results by showing that $\text{Mon}(L) \in \text{SZK}$ as long as L or its complement has a 3-message public-coin statistical zero-knowledge proof, and also treat a larger class of monotone functions.

Our result holds for all of SZK and for all Boolean formulas, not just monotone ones:

THEOREM 4.9. *For any promise problem $\Pi \in \text{SZK}$, $\Phi(\Pi) \in \text{SZK}$.*

This theorem can be generalized to work for all Boolean formulas whose atoms are statements about membership in any finite set of languages in SZK , but we omit the notationally cumbersome formal statement since it is immediate from the completeness of $\text{STATISTICAL DIFFERENCE}$.

Our proof of Theorem 4.9 is based heavily on the work of De Santis et al. [1994], which constructs a statistical zero-knowledge proof for $\text{Mon}(L)$ for any random self-reducible language L . Their zero-knowledge proofs are constructed by producing two distributions that are either disjoint or identical, depending on whether or not the formula is true. Hence, their construction can be viewed as a reduction to extreme instances of SD , in which the thresholds are 1 and 0. Here, we begin by applying essentially the same construction to SD , but use the Direct Product, XOR, and the Polarization Lemmas of Section 3.2 to analyze it for all instances of SD (rather than just the extreme ones). This proves that $\text{Mon}(\text{SD})$ is in SZK . Then, using the completeness of SD (Theorem 3.1) and closure under complement (Corollary 4.4 [Okamoto 2000]), we deduce the result for general (i.e., nonmonotone) formulas and every promise problem in SZK .

As stated above, the main step in proving Theorem 4.9 is the following lemma:

LEMMA 4.10. $\text{Mon}(\text{SD}) \in \text{SZK}$.

¹¹ In De Santis et al. [1994], only monotone formulas are treated. What they call $\Phi(L)$ is what we call $\text{Mon}(L)$.

Sample(ψ, b)

If $\psi = v_i$, sample $z \leftarrow D_b^i$.

If $\psi = \tau \vee \mu$,

Sample $z_1 \leftarrow \text{Sample}(\tau, b)$;

Sample $z_2 \leftarrow \text{Sample}(\mu, b)$;

Let $z = (z_1, z_2)$.

If $\psi = \tau \wedge \mu$,

Choose $c, d \in_R \{0, 1\}$ subject to $c \oplus d = b$;

Sample $z_1 \leftarrow \text{Sample}(\tau, c)$;

Sample $z_2 \leftarrow \text{Sample}(\mu, d)$;

Let $z = (z_1, z_2)$.

Output z .

FIG. 1.

PROOF. For intuition, consider two instances of statistical difference (C_0, C_1) and (D_0, D_1) , both of which have statistical difference very close to 1 or very close to 0 (which can be achieved by the Polarization Lemma). Then $(C_0 \otimes D_0, C_1 \otimes D_1)$ will have statistical difference very close to 1 if either of the original statistical differences is very close to 1 and will have statistical difference very close to 0 otherwise. Thus, this Direct Product operation represents OR. Similarly, the XOR operation in Proposition 3.6 represents AND. We will recursively apply these constructions to obtain a reduction from Mon(SD) to SD. By closure under reductions (Corollary 4.3), Lemma 4.10 will follow.

Let $w = (\phi, (C_0^1, C_1^1), \dots, (C_0^k, C_1^k))$ be an instance of Mon(SD) and let $n = |w|$. By applying the Polarization Lemma (Lemma 3.3), we can construct in polynomial time pairs of circuits $(D_0^1, D_1^1), \dots, (D_0^k, D_1^k)$ such that the statistical difference between D_0^i and D_1^i is greater than $1 - 2^{-n}$ if $(C_0^i, C_1^i) \in \text{SD}_Y$ and is less than 2^{-n} if $(C_0^i, C_1^i) \in \text{SD}_N$.

Consider the randomized recursive procedure Sample(ψ, b) in Figure 1 which takes a subformula ψ of $\phi = \phi(v_1, \dots, v_n)$ and a bit $b \in \{0, 1\}$ as input. Executing Sample(ϕ, b) for $b \in \{0, 1\}$ takes time polynomial in n , because the number of recursive calls is equal to the number of subformulas of ϕ . For a subformula ψ of ϕ , define

$$\text{Dif}(\psi) = \|\text{Sample}(\psi, 0) - \text{Sample}(\psi, 1)\|.$$

Then we can prove the following about Dif:

CLAIM 4.11. Let $\bar{a} = (\chi_{\text{SD}}(C_0^1, C_1^1), \dots, \chi_{\text{SD}}(C_0^k, C_1^k))$. For every subformula ψ of ϕ , we have:

$$\begin{aligned} \psi(\bar{a}) = 1 &\Rightarrow \text{Dif}(\psi) > 1 - |\psi|2^{-n} \\ \psi(\bar{a}) = 0 &\Rightarrow \text{Dif}(\psi) < |\psi|2^{-n} \end{aligned}$$

Note that nothing is claimed when $\psi(\bar{a}) = \star$.

PROOF OF CLAIM. The proof of the claim is by induction on subformulas ψ of ϕ . It holds for atomic subformulas (i.e., the variables v_i) by the properties of the D_b^i 's.

Case I. $\psi = \tau \vee \mu$.] If $\psi(\bar{a}) = 1$, then either $\tau(\bar{a}) = 1$ or $\mu(\bar{a}) = 1$. Without loss of generality, say $\tau(\bar{a}) = 1$. Then, by Fact 2.4 and induction,

$$\text{Dif}(\psi) \geq \text{Dif}(\tau) > 1 - |\tau|2^{-n} > 1 - |\psi|2^{-n}.$$

If $\psi(\bar{a}) = 0$, then $\tau(\bar{a}) = \mu(\bar{a}) = 0$. By Fact 2.3 and induction,

$$\text{Dif}(\psi) \leq \text{Dif}(\tau) + \text{Dif}(\mu) < |\tau|2^{-n} + |\mu|2^{-n} \leq |\psi|2^{-n}.$$

Case II. $\psi = \tau \wedge \mu$.] By Proposition 3.6, $\text{Dif}(\psi) = \text{Dif}(\tau) \cdot \text{Dif}(\mu)$. If $\psi(\bar{a}) = 1$, then, by induction,

$$\text{Dif}(\psi) \geq (1 - |\tau|2^{-n})(1 - |\mu|2^{-n}) > 1 - (|\tau| + |\mu|)2^{-n} \geq 1 - |\psi|2^{-n}.$$

If $\psi(\bar{a}) = 0$, then, without loss of generality, say $\tau(\bar{a}) = 0$. By induction,

$$\text{Dif}(\psi) \leq \text{Dif}(\tau) < |\tau|2^{-n} \leq |\psi|2^{-n}. \quad \square$$

Now, let A and B be the circuits which sample from the distributions $\text{Sample}(\phi, 0)$ and $\text{Sample}(\phi, 1)$, respectively. (The random bits each procedure uses are the inputs to the circuits.) By the above claim, $\|A - B\| > 1 - n2^{-n} > 2/3$ if $\phi(\bar{a}) = 1$, and $\|A - B\| < n2^{-n} < 1/3$ if $\phi(\bar{a}) = 0$. In other words, the construction of A and B from w is a reduction from $\text{Mon}(\text{SD})$ to SD . This reduction can be computed in polynomial time because Sample runs in polynomial time. Thus, by Corollary 4.3, $\text{Mon}(\text{SD}) \in \text{SZK}$. \square

Now it is straightforward to deduce Theorem 4.9.

PROOF. Let Π be any promise problem in SZK . By closure under complement (Corollary 4.4 [Okamoto 2000]) and the completeness of SD (Theorem 3.1), both Π and $\bar{\Pi}$ reduce to SD . Let f and g be these reductions, respectively. Now, let (ϕ, x_1, \dots, x_k) be any instance of $\Phi(\Pi)$, where $\phi = \phi(v_1, \dots, v_k)$. Use De Morgan's laws to propagate all negations of ϕ to its variables. Now replace all occurrences of the literal $\neg v_i$ with a new variable w_i . Let $\psi(v_1, \dots, v_k, w_1, \dots, w_k)$ be the resulting (monotone) formula. It is clear that

$$(\phi, x_1, \dots, x_k) \mapsto (\psi, f(x_1), \dots, f(x_k), g(x_1), \dots, g(x_k))$$

is a reduction from $\Phi(\Pi)$ to $\text{Mon}(\text{SD})$. Since $\text{Mon}(\text{SD}) \in \text{SZK}$ (Lemma 4.10) and SZK is closed under reductions (Corollary 4.3), Theorem 4.9 follows. \square

Theorem 4.9 can be also viewed as demonstrating that SZK is closed under a type of polynomial-time reducibility, which is formalized by the following two definitions.

Definition 4.12 (Truth-Table Reduction [Ladner et al. 1975]). We say a promise problem Π *truth-table reduces* to a promise problem Γ if there exists a (deterministic) polynomial-time computable function f , which on input x produces a tuple (y_1, \dots, y_k) and a Boolean circuit C (with k input gates) such that

$$\begin{aligned} x \in \Pi_Y &\Rightarrow C(\chi_\Gamma(y_1), \dots, \chi_\Gamma(y_k)) = 1 \\ x \in \Pi_N &\Rightarrow C(\chi_\Gamma(y_1), \dots, \chi_\Gamma(y_k)) = 0. \end{aligned}$$

In other words, a truth-table reduction for promise problems is a nonadaptive Cook reduction that is allowed to make queries that violate the promise, but still must have an unambiguous output (in the strong sense formalized by Definition 4.6).

We further consider the case where we restrict the complexity of computing the output of the reduction from the queries:

Definition 4.13 (NC¹ Truth-Table Reductions). A truth-table reduction f between promise problems is an NC¹ truth-table reduction if the circuit C produced by the reduction on input x has depth bounded by $c_f \log |x|$, where c_f is a constant independent of x .

With these definitions, we can restate Theorem 4.9 as follows:

COROLLARY 4.14. *SZK is closed under NC¹ truth-table reductions.*

PROOF. Any circuit of size s and depth d can be efficiently “unrolled” into a formula of size $2^d \cdot s$. Hence, an NC¹ truth-table reduction from Γ to Π gives rise to a Karp reduction from Γ to $\Phi(\Pi)$. Since SZK is closed under $\Phi(\cdot)$ and Karp reductions, it is also closed under NC¹ truth-table reductions. \square

It would be interesting to prove that SZK is closed under general truth-table reductions (or, even better, adaptive Cook reductions), or give evidence that this is not the case.

4.3. KNOWLEDGE COMPLEXITY. Knowledge complexity [Goldwasser et al. 1989; Goldreich and Petrank 1999] is a generalization of zero knowledge which attempts to quantify how much a verifier learns from an interactive proof. A number of different measures have been proposed to accomplish this, most of which are based on the intuition that a verifier gains at most k bits of “knowledge” from an interaction if it can simulate the interaction with at most k bits of “help”. Below we give terse definitions of the variants we consider. The first three definitions come from Goldreich and Petrank [1999], and the last comes from [Aiello et al. 1995]. Let (P, V) be an interactive proof system for a promise problem Π . Then the knowledge complexity of (P, V) in various senses is defined as follows:

- Hint Sense.* We say that (P, V) has perfect (respectively, statistical) knowledge complexity $k(n)$ in the *hint* sense if there exists a PPT simulator S and a hint function $h : \Pi_Y \rightarrow \{0, 1\}^*$ such that for all $x \in \Pi_Y$, $|h(x)| = k(|x|)$ and $\|S(x, h(x)) - \text{View}_{P,V}(x)\|$ is 0 (respectively, is bounded by a negligible function of $|x|$.)
- Strict Oracle Sense.* (P, V) is said to have perfect (respectively, statistical) knowledge complexity $k(n)$ in the *strict oracle* sense if there exists a PPT oracle-machine S and an oracle \mathcal{O} such that on every input $x \in \Pi_Y$, S queries \mathcal{O} at most $k(|x|)$ times and $\|S^{\mathcal{O}}(x) - \text{View}_{P,V}(x)\|$ is 0 (respectively, is bounded by a negligible function of $|x|$.)
- Oracle Sense.* (P, V) is said to have perfect (respectively, statistical) knowledge complexity $k(n)$ in the *oracle* sense if there exists a PPT oracle-machine S and an oracle \mathcal{O} such that on every input $x \in \Pi_Y$, S queries \mathcal{O} at most $k(|x|)$ times, S outputs ‘fail’ with probability at most $1/2$, and $\|S^{\mathcal{O}}(x) - \text{View}_{P,V}(x)\|$ is 0 (respectively, is bounded by a negligible function of $|x|$), where $S^{\mathcal{O}}(x)$ denotes the output distribution of S conditioned on nonfailure.
- Average Oracle Sense.* (P, V) has perfect (respectively, statistical) knowledge complexity $k(n)$ in the *average oracle* sense if there exists a PPT oracle-machine S and an oracle \mathcal{O} such that for every input $x \in \Pi_Y$, the average number of queries

S makes to \mathcal{O} is at most $k(|x|)$ and $\|S^{\mathcal{O}}(x) - \text{View}_{P,V}(x)\|$ is 0 (respectively, is bounded by a negligible function of $|x|$.)

—*Entropy Sense.* (P, V) has perfect (respectively, statistical) knowledge complexity $k(n)$ in the *entropy* sense if there exists a PPT oracle-machine S , an oracle \mathcal{O} , and a PPT oracle-simulator A such that for all $x \in \Pi_Y$, $E_R[\log P_x(R)^{-1}] \leq k(|x|)$, where $P_x(R) = \Pr_{\rho}[A(x, R; \rho) = S^{\mathcal{O}}(x; R)]$ and $\|S^{\mathcal{O}}(x) - \text{View}_{P,V}(x)\|$ is 0 (respectively, is bounded by a negligible function of $|x|$). Here, the notation $M(y; r)$ denotes the output of PPT M on input y and random coins r .

The *knowledge complexity* (in some specified sense) of a promise problem Π is $k(n)$ if there exists an interactive proof system (P, V) for Π achieving negligible error probability in both the completeness and soundness conditions such that the knowledge complexity of (P, V) is $k(n)$. The class of languages possessing perfect knowledge complexity $k(n)$ in the hint, strict oracle, average oracle, and entropy senses are denoted by $\text{PKC}_{\text{hint}}(k(n))$, $\text{PKC}_{\text{strict}}(k(n))$, $\text{PKC}_{\text{avg}}(k(n))$, and $\text{PKC}_{\text{ent}}(k(n))$, respectively. Statistical knowledge complexity is denoted by $\text{SKC}(k(n))$ with the appropriate subscript.

4.3.1. *A Collapse for the Hint Sense.* Our first result about knowledge complexity is that the SKC_{hint} hierarchy collapses by logarithmic additive factors. Previously, Goldreich and Petrank [1999] have shown that $\text{SKC}_{\text{hint}}(\text{poly}(n)) \subset \text{AM}$ and $\text{SKC}_{\text{hint}}(O(\log(n))) \subset \text{co-AM}$; the second of these results can be derived immediately from our result and the result $\text{SZK} \subset \text{co-AM}$ (Corollary 4.5 [Fortnow 1989; Aiello and Håstad 1991]).

THEOREM 4.15. *For any polynomially bounded function $k(n)$,*

$$\text{SKC}_{\text{hint}}(k(n) + \log n) = \text{SKC}_{\text{hint}}(k(n)).$$

For intuition, consider the case that $k(n) = 0$. Loosely speaking, if the verifier is given the hint along with the input (with the “promise” that the hint is correct), then the original proof system becomes *zero* knowledge, so we can apply the results of the previous section. By the Boolean closure properties established in Theorem 4.9, we can take the “union over all possible hints” (there are only polynomially many of them) without leaving SZK .

In order to turn this intuition into a proof, we first show that knowledge complexity in the hint sense can be characterized in terms of zero-knowledge promise problems, so that questions about the SKC_{hint} hierarchy are reduced to questions about statistical zero knowledge. This equivalence is obtained by providing the hint along with the input and “promising” that the hint is correct.

LEMMA 4.16. *Let $k(n)$ be any polynomially bounded function. Then $\Pi \in \text{SKC}_{\text{hint}}(k(n))$ (respectively, $\text{PKC}_{\text{hint}}(k(n))$) iff there exists a promise problem $\Gamma \in \text{SZK}$ (respectively, PZK) such that*

- (1) $x \in \Pi_Y \Rightarrow$ *there exists a such that $|a| = k(|x|)$ and $(x, a) \in \Gamma_Y$, and*
- (2) $x \in \Pi_N \Rightarrow$ *for all a , $(x, a) \in \Gamma_N$.*

PROOF. We only give the proof for statistical knowledge complexity and zero knowledge; the perfect case is identical.

\Rightarrow Let Π be a promise problem in $\text{SKC}_{\text{hint}}(k(n))$ and let $h : \Pi_Y \rightarrow \{0, 1\}^*$ be a hint function corresponding to an appropriate interactive proof system and

simulator for Π . Consider the following promise problem Γ :

$$\begin{aligned}\Gamma_Y &= \{(x, h(x)) : x \in \Pi_Y\} \\ \Gamma_N &= \{(x, a) : x \in \Pi_N\}\end{aligned}$$

By using the protocol and simulator for Π , we see that $\Gamma \in \mathbf{SZK}$ (the verifier and prover for Γ should ignore the second component, whereas the simulator uses it as a hint.) It is clear that Γ satisfies the other conditions of Lemma 4.16.

\Leftarrow Let $\Gamma \in \mathbf{SZK}$ be the promise problem satisfying the stated conditions. Let $h : \Pi_Y \rightarrow \{0, 1\}^*$ be any function such that for all $x \in \Pi_Y$,

- (1) $|h(x)| = k(|x|)$,
- (2) $(x, h(x)) \in \Gamma_Y$.

(Such a function is guaranteed by Condition 1.) We now give a proof system for Π of knowledge complexity $k(n)$. On input x , the prover gives the verifier $h(x)$ in the first step, and then they execute the protocol for Γ on $(x, h(x))$. The completeness and soundness of this protocol follow from the properties of the Γ proof system. This proof system is easily seen to have knowledge complexity $k(n)$ in the hint sense, using the hint $h(x)$ with the zero-knowledge simulator for Γ . \square

We now prove Theorem 4.15.

PROOF. Let Π be a problem in $\mathbf{SKC}_{\text{hint}}(k(n) + \log n)$ and let Γ be the promise problem guaranteed by Lemma 4.16. By Theorem 4.9, $\Phi(\Gamma) \in \mathbf{SZK}$. Now consider a different, but related promise problem Γ' , defined by

$$\begin{aligned}\Gamma'_Y &= \{(x, a) : \text{there exists } b \text{ such that } |b| = \log |x| \text{ and } (x, ab) \in \Gamma_Y\} \\ \Gamma'_N &= \{(x, a) : \text{for all } b, (x, ab) \in \Gamma_N\} = \{(x, a) : x \in \Pi_N\}.\end{aligned}$$

For any string x , let b_1, \dots, b_n be all strings of length $\log |x|$, and let C be the circuit of depth $O(\log |x|)$ computing the function $\phi(v_1, \dots, v_n) = \bigvee_i v_i$. The relationship between Γ and Γ' above implies that

$$(x, a) \mapsto (\phi, (x, ab_1), \dots, (x, ab_n))$$

is an \mathbf{NC}^1 truth-table reduction from Γ' to Γ . Since \mathbf{SZK} is closed under such reductions (Corollary 4.14), we conclude that $\Gamma' \in \mathbf{SZK}$.

Now, $x \in \Pi_Y$, then there exists an a of length $k(|x|) + \log(|x|)$ such that $(x, a) \in \Gamma_Y$. Taking a' to be the first $k(|x|)$ bits of a , we see that there exists an a' of length $k(|x|)$ such that $(x, a') \in \Gamma'_Y$. Moreover, if $x \in \Pi_N$, then for all a , $(x, a) \in \Gamma'_N$. Thus, by Lemma 4.16, we conclude that $\Pi \in \mathbf{SKC}_{\text{hint}}(k(n))$. \square

4.3.2. The Perfect Knowledge Complexity of \mathbf{SZK} . The next theorem establishes tighter bounds on the perfect knowledge complexity of \mathbf{SZK} . Aiello et al. [1995] have previously demonstrated that every language in \mathbf{SZK} has perfect knowledge complexity $n^{-\omega(1)}$ (respectively, $1 + n^{-\omega(1)}$) in the entropy (respectively, average oracle) sense. Our results improve on these bounds, although the results of Aiello et al. [1995] also apply to cheating-verifier classes and ours do not. Goldreich et al. [1998] show that \mathbf{SZK} has logarithmic perfect knowledge complexity in the oracle sense, so our results are incomparable to theirs. Our result for the strict oracle sense is the first that we know of.

THEOREM 4.17¹²

- (1) For every polynomial-time computable $m(n) = \omega(\log n)$, $\text{SZK} \subset \text{PKC}_{\text{strict}}(m(n))$.
- (2) $\text{SZK} \subset \text{PKC}_{\text{avg}}(1 + 2^{-n})$.
- (3) $\text{SZK} = \text{PKC}_{\text{ent}}(2^{-n})$.

Corollary 4.2 tells us that every problem in SZK has a simple two-message proof system like the SD proof system of Section 3.3. Thus, in order to measure the perfect knowledge complexity of SZK and prove Theorem 4.17, it suffices to analyze this protocol. Intuitively, since the prover is only sending the verifier one bit and this bit is almost always a value the verifier knows, the knowledge complexity of this protocol should be extremely small. However, this argument does not suffice, because the knowledge complexity of a problem Π is determined only by proof systems for Π that achieve *negligible* error probability in both the completeness and soundness conditions. We can overcome this difficulty by performing $\omega(\log n)$ parallel repetitions.

PROOF. Let Π be any problem in SZK and let (P, V) be the proof system for Π constructed in Corollary 4.2 (from the SD proof system of Section 3.3) with the security parameter set to $k = 4n$ (so the completeness error is 2^{-4n}). Let $m = m(n)$ be any function computable in time $\text{poly}(n)$ such that $\omega(\log n) \leq m \leq n$. Consider the proof system (P', V') obtained by m parallel repetitions of (P, V) ; this has negligible completeness and soundness errors. We now analyze its perfect knowledge complexity.

- (1) The prover sends at most m bits to the verifier on inputs of length n , so the perfect knowledge complexity of this protocol in the strict oracle sense is bounded by m .
- (2) A perfect simulator for (P', V') can be obtained as follows: On input x of length n , the simulator runs $V(x)$ for m times independently and queries the oracle *once* to find out if any of these runs would result in an incorrect prover response. If the oracle replies yes, the simulator queries the oracle m more times to find out which runs would result in an incorrect response. The simulator then outputs the random coins used for running V and the appropriate prover responses.

In each subprotocol, the prover gives an incorrect response with probability at most 2^{-4n} . Thus, the simulator has to query the oracle for more than one bit with probability at most $m2^{-4n}$. Thus, on average, the simulator queries the oracle for at most $1 + m^22^{-4n} < 1 + 2^{-n}$ bits.

- (3) Let S be the simulator for (P', V') that simply simulates V' and queries the oracle \mathcal{O} for all prover responses. One possible oracle simulator would assume that the prover is correct in all subprotocols. Unfortunately, this gives $1/P_x(R) = \infty$ for some R and yields infinite knowledge complexity. Thus, we instead have our oracle simulator A assume that the prover is right in each subprotocol independently with probability $1 - \delta$, where $\delta = 2^{-2n}$. Thus, $P_x(R) = (1 - \delta)^k \delta^{m-k}$,

¹² The $2^{-\Omega(n)}$ in these results can be improved to $2^{-\Omega(n^k)}$ for any constant k by polarizing with security parameter n^k instead of n in the SD proof system of Section 3.3.

if R is a set of random coins for V' (equivalently S , since S mimics V') that would elicit a correct prover response in exactly k of the subprotocols. Let ϵ be the probability that the prover is incorrect in an individual subprotocol. Then, $\epsilon \leq \delta^2$, and we have

$$\begin{aligned}
\mathbb{E}_R \left[\log \frac{1}{P_x(R)} \right] &= \sum_{k=0}^m \binom{m}{k} \epsilon^{m-k} (1-\epsilon)^k \log \left(\frac{1}{(1-\delta)^k \delta^{m-k}} \right) \\
&= \left(\log \frac{1}{\delta^m} \right) \sum_{k=0}^m \binom{m}{k} \epsilon^{m-k} (1-\epsilon)^k \\
&\quad + \left(\log \frac{\delta}{1-\delta} \right) \sum_{k=0}^m k \binom{m}{k} \epsilon^{m-k} (1-\epsilon)^k \\
&= \log \frac{1}{\delta^m} + m(1-\epsilon) \left(\log \frac{\delta}{1-\delta} \right) \\
&= m \left(\log \frac{1}{1-\delta} + \epsilon \log \frac{1-\delta}{\delta} \right) \\
&\leq m \left(\log \frac{1}{1-\delta} + \delta^2 \log \frac{1}{\delta} \right) \\
&\leq 2m\delta < 2^{-n}
\end{aligned}$$

for sufficiently large n . Note that the third equality above follows from the identity $k \binom{m}{k} = m \binom{m-1}{k-1}$.

The opposite inclusion follows from the result of Aiello et al. [1995] that $\text{PKC}_{\text{ent}}(\text{neg}(n)) \subset \text{SZK}$ for any negligible function $\text{neg}(n)$. \square

4.4. REVERSING STATISTICAL DIFFERENCE. By the completeness of SD (Theorem 3.1) and SZK's closure under complement (Corollary 4.4), it follows that $\overline{\text{SD}}$ reduces to SD. This is equivalent to the following surprising result:

COROLLARY 4.18 (REVERSAL MAPPING). *There is a polynomial-time computable function that maps pairs of circuits (C_0, C_1) to pairs of circuits (D_0, D_1) such that*

$$\begin{aligned}
\|C_0 - C_1\| < \frac{1}{3} &\Rightarrow \|D_0 - D_1\| > \frac{2}{3} \\
\|C_0 - C_1\| > \frac{2}{3} &\Rightarrow \|D_0 - D_1\| < \frac{1}{3}.
\end{aligned}$$

That is, SD reduces to $\overline{\text{SD}}$.

This corollary motivated our search for a more explicit description of such a mapping. By extracting ideas used in the transformations of statistical zero-knowledge proofs given in Okamoto [2000] and Sahai and Vadhan [1997], we obtained the description of this transformation given below.

The Construction. Let (C_0, C_1) be any pair of circuits and let $n = |(C_0, C_1)|$. By the Polarization Lemma (Lemma 3.3), we can produce in polynomial time

a pair of circuits $(\widetilde{C}_0, \widetilde{C}_1)$ such that

$$\|C_0 - C_1\| < \frac{1}{3} \Rightarrow \|\widetilde{C}_0 - \widetilde{C}_1\| > 1 - 2^{-n}$$

$$\|C_0 - C_1\| > \frac{2}{3} \Rightarrow \|\widetilde{C}_0 - \widetilde{C}_1\| < 2^{-n}$$

Let $q = \text{poly}(n)$ be the number of input gates of \widetilde{C}_0 and \widetilde{C}_1 (without loss of generality, we may assume they have the same number) and let $\ell = \text{poly}(n)$ be the number of output gates. For notational convenience, let $R = \{0, 1\}^q$ and $L = \{0, 1\}^\ell$. Let $m = n^3 q^2$ and define a new distribution $\vec{C}: \{0, 1\}^m \times R^m \rightarrow L^m$ as follows:

$$\vec{C}(\vec{b}, \vec{r}) = (\widetilde{C}_{b_1}(r_1), \dots, \widetilde{C}_{b_m}(r_m)).$$

We use the notation $\vec{z} \leftarrow \vec{C}$ to denote \vec{z} chosen according to \vec{C} , that is, select \vec{b} and \vec{r} uniformly and let $\vec{z} = \vec{C}(\vec{b}, \vec{r})$.

Let \mathcal{H} be a 2-universal family of hash functions from $\{0, 1\}^m \times R^m \times L^m$ to $S = \{0, 1\}^{(q+1)m - 2\Delta - n}$, where $\Delta = \sqrt{nmq^2} = m/n$. We can now describe the new distributions:

D_0 : Choose $(\vec{b}, \vec{r}) \in_R \{0, 1\}^m \times R^m$, $\vec{y} \leftarrow \vec{C}$, and $h \in_R \mathcal{H}$.
 Output $(\vec{C}(\vec{b}, \vec{r}), \vec{b}, h, h(\vec{b}, \vec{r}, \vec{y}))$.

D_1 : Choose $(\vec{b}, \vec{r}) \in_R \{0, 1\}^m \times R^m$, $h \in_R \mathcal{H}$, and $s \in_R S$.
 Output $(\vec{C}(\vec{b}, \vec{r}), \vec{b}, h, s)$.

The important things to note about these distributions are that \vec{b} is part of the output, and that D_0 and D_1 only differ in the last component, where D_0 has the value of the hash function and D_1 has a truly random element of S . Also note that the size of S is chosen to be $|\{0, 1\}^m \times R^m|/2^{2\Delta+n}$, which is essentially $|\{0, 1\}^m \times R^m|$, scaled down by a “slackness” factor of $2^{2\Delta+n}$. The introduction of the sample \vec{y} in D_0 may at first seem superfluous; we explain below.

Intuition. For intuition, consider the case that \vec{C} is a *flat* distribution; that is, for every $\vec{z} \in \text{range}(\vec{C})$, the size of the preimage set $|\{(\vec{b}, \vec{r}): \vec{C}(\vec{b}, \vec{r}) = \vec{z}\}|$ is the same value N . Then the range of \vec{C} has size $|\{0, 1\}^m \times R^m|/N = 2^{(q+1)m}/N$. So, in D_0 , conditioned on a value for $\vec{C}(\vec{b}, \vec{r})$, the triple $(\vec{b}, \vec{r}, \vec{y})$ is selected uniformly from a set of size $2^{(q+1)m}$. Since this is much greater than $|S|$, the Leftover Hash Lemma of Håstad et al. [1999] implies that conditioned on any value for the first component of D_0 , the last two components $(h, h(\vec{b}, \vec{r}, \vec{y}))$ are distributed close to the uniform distribution on $\mathcal{H} \times S$, which is the distribution that D_1 has in its last two components.¹³ Thus, if their second components were missing, D_0 and D_1

¹³ Here we see the importance of \vec{y} : Without \vec{y} , conditioned on some value of $\vec{C}(\vec{b}, \vec{r})$, the pair (\vec{b}, \vec{r}) would be selected uniformly from a space of size N . If we were only hashing this pair, for the distribution $h(\vec{b}, \vec{r})$ to be uniform by the Leftover Hash Lemma, S would have had to be chosen so that

would be statistically close. Now, consider the case that $\|\widetilde{C}_0 - \widetilde{C}_1\| \approx 1$. Then \vec{b} is essentially “determined” by $\vec{C}(\vec{b}, \vec{r})$. So the presence of \vec{b} can be ignored, and the above argument says that D_0 and D_1 are statistically very close. Now, consider the case that $\|\widetilde{C}_0 - \widetilde{C}_1\| \approx 0$. Then \vec{b} is essentially “unrestricted” by $\vec{C}(\vec{b}, \vec{r})$. Since there are 2^m choices for \vec{b} , conditioning on \vec{b} in addition to $\vec{C}(\vec{b}, \vec{r})$, cuts the number of triples $(\vec{b}, \vec{r}, \vec{y})$ down from $2^{m(q+1)}$ to roughly $2^{m(q+1)}/2^m$. Since $2^{m(q+1)}/2^m$ is much smaller than $|S|$, $h(\vec{b}, \vec{r}, \vec{y})$ will cover only a small fraction of $|S|$ and thus will be far from uniform (conditioned on values for $\vec{C}(\vec{b}, \vec{r})$, \vec{b} , and h).

Direct Proof of Corollary 4.18. First, we will argue that \vec{C} is close to being flat (in a particular sense), so that we can apply arguments like those given above. The fact that \vec{C} is close to flat will follow from the fact that it is composed of many independent, identically distributed random variables. This is a form of the Asymptotic Equipartition Property commonly used in Information Theory (cf., Cover and Thomas [1991]).

For $\vec{z} \in L^m$, we say the *weight* of \vec{z} is the logarithm of the size of the preimage set of \vec{z} . Formally, let $\text{wt}(\vec{z}) = \log_2 |\{(\vec{b}, \vec{r}) : \vec{C}(\vec{b}, \vec{r}) = \vec{z}\}|$. Let w be the expected weight of an image, that is, $w = E_{\vec{z} \leftarrow \vec{C}}[\text{wt}(\vec{z})]$. Then, we can show the following:

LEMMA 4.19. $\Pr_{\vec{z} \leftarrow \vec{C}} [|\text{wt}(\vec{z}) - w| > \Delta] < 2^{-\Omega(n)}$.

PROOF. For $z \in L$, let $\text{wt}_0(z) = \log_2 |\{(b, r) : \widetilde{C}_b(r) = z\}|$. Then, for $\vec{z} \in L^m$, $\text{wt}(\vec{z}) = \text{wt}_0(z_1) + \dots + \text{wt}_0(z_m)$. Observe that, when \vec{z} is selected according to \vec{C} , z_1, \dots, z_m are independent and identically distributed. Moreover, for any $z \in L$, $0 \leq \text{wt}_0(z) \leq q$. So, by the Hoeffding inequality [Hofri 1995, Sect. 7.2.1], we have

$$\Pr_{\vec{z} \leftarrow \vec{C}} [|\text{wt}(\vec{z}) - w| > \Delta] < 2 \exp(-2\Delta^2/mq^2) = 2 \exp(-2n). \quad \square$$

It will be convenient to eliminate those $\vec{z} \in L^m$ that have weight far above or below the mean. Let $T = \{(\vec{b}, \vec{r}) : |\text{wt}(\vec{C}(\vec{b}, \vec{r})) - w| \leq \Delta\}$, which we will call the set of *typical* pairs (\vec{b}, \vec{r}) . The above Lemma says that $|T| \geq (1 - 2^{-\Omega(n)})|\{0, 1\}^m \times R^m|$. Thus, $\|T - \{0, 1\}^m \times R^m\| \leq 2^{-\Omega(n)}$, where for simplicity of notation, we let the name of a set also refer to the uniform distribution on the same set. Define \vec{C}' to be the distribution obtained by selecting $(\vec{b}, \vec{r}) \leftarrow T$ and outputting $\vec{C}(\vec{b}, \vec{r})$. Then, since \vec{C} is a function, Fact 2.4 tells us that $\|\vec{C} - \vec{C}'\| = 2^{-\Omega(n)}$. Similarly, we define

$|S| \ll N$. The value of N , however, depends on the inner workings of the circuit C , and is in general unknown. By including \vec{y} , which comes uniformly from a space of size $2^{(q+1)m}/N$, we balance the arguments to h so that they come from a space of size $2^{(q+1)m}$, a known quantity. This use of “dummy” samples to form a space whose size is known is the “complementary usage of messages” technique of Okamoto [2000].

variants of D_0 and D_1 that sample from T instead of $\{0, 1\}^m \times R^m$:

$$\begin{array}{l} D'_0: \text{ Let } (\vec{b}, \vec{r}) \in_R T, \vec{y} \leftarrow \vec{C}', \text{ and } h \in_R \mathcal{H}. \quad \text{Output } (\vec{C}(\vec{b}, \vec{r}), \vec{b}, h, h(\vec{b}, \vec{r}, \vec{y})). \\ D'_1: \text{ Let } (\vec{b}, \vec{r}) \in_R T, h \in_R \mathcal{H}, \text{ and } s \in_R S. \quad \text{Output } (\vec{C}(\vec{b}, \vec{r}), \vec{b}, h, s). \end{array}$$

Since D'_0 (respectively, D'_1) is a randomized procedure applied to two (respectively, one) independent samplings from T , Fact 2.4 tells us that $\|D_0 - D'_0\| = 2^{-\Omega(n)}$ (respectively, $\|D_1 - D'_1\| = 2^{-\Omega(n)}$). Hence, it suffices to prove that these modified distributions have the properties we want in each case. For the case when C_0 and C_1 are statistically far, we prove the following claim:

CLAIM 4.20. *If $\|\vec{C}_0 - \vec{C}_1\| > 1 - 2^{-n}$, then $\|D'_0 - D'_1\| < 2^{-\Omega(n)}$.*

PROOF OF CLAIM. First, we formalize the idea that \vec{b} is “determined” by \vec{C} . Define $f : L \rightarrow \{0, 1\}$ by

$$f(z) = \begin{cases} 0 & \text{if } \Pr[\vec{C}_0 = z] > \Pr[\vec{C}_1 = z] \\ 1 & \text{otherwise} \end{cases}$$

In other words, f is exactly the prover strategy from the proof system for STATISTICAL DIFFERENCE given in Section 3.3. The completeness of that proof system (Lemma 3.7) says that $\Pr_{b,r}[f(\vec{C}_b(r)) = b] > 1 - 2^{-n}$. Now define $\vec{f} : L^m \rightarrow \{0, 1\}^m$ by $\vec{f}(\vec{z}) = (f(z_1), \dots, f(z_m))$. Then

$$\Pr_{\vec{b}, \vec{r}}[\vec{f}(\vec{C}(\vec{b}, \vec{r})) = \vec{b}] > (1 - 2^{-n})^m = 1 - 2^{-\Omega(n)}.$$

Since T is a $1 - 2^{-\Omega(n)}$ fraction of $\{0, 1\}^m \times R^m$, the same is true when (\vec{b}, \vec{r}) is selected uniformly from T . Thus, if we define:

$$\begin{array}{l} D''_0: \text{ Let } (\vec{b}, \vec{r}) \in_R T, \vec{y} \leftarrow \vec{C}', \text{ and } h \in_R \mathcal{H}. \\ \quad \text{Output } (\vec{C}(\vec{b}, \vec{r}), \vec{f}(\vec{C}(\vec{b}, \vec{r})), h, h(\vec{b}, \vec{r}, \vec{y})). \\ D''_1: \text{ Let } (\vec{b}, \vec{r}) \in_R T, h \in_R \mathcal{H}, \text{ and } s \in_R S. \text{ Output } (\vec{C}(\vec{b}, \vec{r}), \vec{f}(\vec{C}(\vec{b}, \vec{r})), h, s). \end{array}$$

Then, $\|D'_0 - D''_0\| = 2^{-\Omega(n)}$ and $\|D'_1 - D''_1\| = 2^{-\Omega(n)}$. So it suffices to show that $\|D''_0 - D''_1\| = 2^{-\Omega(n)}$. Since the first components of D''_0 and D''_1 are identically distributed and the second components are determined by the first ones, it suffices to show (by Fact 2.5) that, conditioned on any value for the first coordinate, the third and fourth components have statistical difference $2^{-\Omega(n)}$. This will follow from the Leftover Hash Lemma [Håstad et al. 1999]:

LEMMA 4.21 (LEFTOVER HASH LEMMA [HÅSTAD ET AL. 1999]). *Let \mathcal{H} be a family of 2-universal hash functions from D to S . Let X be a probability distribution on D such that for all $x \in D$, $\Pr[X = x] \leq \epsilon/|S|$. Then the following two distributions have statistical difference at most $\epsilon^{1/3}$.*

- (1) Choose $x \leftarrow X$, $h \in_R \mathcal{H}$. Output $(h, h(x))$.
- (2) Choose $h \in_R \mathcal{H}$, $s \in_R S$. Output (h, s) .

By the above argument and the Leftover Hash Lemma, it suffices to show that, conditioned on any value \vec{z} for $\vec{C}'(\vec{b}, \vec{r})$, no triple $(\vec{b}, \vec{r}, \vec{y})$ has probability more than

$2^{-\Omega(n)}/|S|$. The pair (\vec{b}, \vec{r}) comes uniformly from a set of size $2^{\text{wt}(\vec{z})} \geq 2^{w-\Delta}$, and \vec{y} is selected independently according to \vec{C}' , so the probability of any triple $(\vec{b}, \vec{r}, \vec{y})$ is at most

$$\left(\frac{1}{2^{w-\Delta}}\right) \left(\frac{2^{w+\Delta}}{|T|}\right) \leq \frac{2^{2\Delta}}{(1-2^{-\Omega(n)})2^{(q+1)m}} = \frac{2^{-\Omega(n)}}{|S|}.$$

Thus, $\|D_0'' - D_1''\| \leq 2^{-\Omega(n)}$, and the claim is established. \square

Now we treat the other case, when C_0 and C_1 are statistically close.

CLAIM 4.22. *If $\|\widetilde{C}_0 - \widetilde{C}_1\| < 2^{-n}$, then $\|D'_0 - D'_1\| > 1 - 2^{-\Omega(n)}$.*

PROOF OF CLAIM. First, we formalize the idea that \vec{b} is almost completely “undetermined” by $\vec{C}(\vec{b}, \vec{r})$. Since $\|\widetilde{C}_0 - \widetilde{C}_1\| < 2^{-n}$, it follows from Fact 2.6 that, with probability $1 - 2^{-\Omega(n)}$ over $z \leftarrow \widetilde{C}_0$,

$$(1 - 2^{-\Omega(n)}) \Pr[\widetilde{C}_1 = z] \leq \Pr[\widetilde{C}_0 = z] \leq (1 + 2^{-\Omega(n)}) \Pr[\widetilde{C}_1 = z].$$

In other words,

$$1 - 2^{-\Omega(n)} \leq \frac{|\{r : \widetilde{C}_0(r) = z\}|}{|\{r : \widetilde{C}_1(r) = z\}|} \leq 1 + 2^{-\Omega(n)}.$$

The same is true with probability $1 - 2^{-\Omega(n)}$ when the roles of \widetilde{C}_0 and \widetilde{C}_1 are reversed. Thus, with probability $1 - m2^{-\Omega(n)} = 1 - 2^{-\Omega(n)}$ over $\vec{z} \leftarrow \vec{C}$, we have for every pair $\vec{b}, \vec{c} \in \{0, 1\}^m$,

$$1 - 2^{-\Omega(n)} = (1 - 2^{\Omega(n)})^m \leq \frac{|\{\vec{r} : \vec{C}(\vec{b}, \vec{r}) = \vec{z}\}|}{|\{\vec{r} : \vec{C}(\vec{c}, \vec{r}) = \vec{z}\}|} \leq (1 + 2^{-\Omega(n)})^m = 1 + 2^{-\Omega(n)}.$$

Since there are 2^m choices for \vec{c} , this, combined with Lemma 4.19, implies that, with probability $1 - 2^{-\Omega(n)}$ over $\vec{z} \leftarrow \vec{C}$, the following holds for every $\vec{b} \in \{0, 1\}^m$:

$$|\{\vec{r} : \vec{C}(\vec{b}, \vec{r}) = \vec{z}\}| \leq (1 + 2^{-\Omega(n)}) \cdot \frac{2^{\text{wt}(\vec{z})}}{2^m} \leq (1 + 2^{-\Omega(n)}) \cdot 2^{w+\Delta-m}.$$

Since this is true with probability $1 - 2^{-\Omega(n)}$ for \vec{z} selected according to \vec{C} , it is also true with probability $1 - 2^{-\Omega(n)}$ for \vec{z} selected according to \vec{C}' . Fix any such \vec{z} and fix any $\vec{b} \in \{0, 1\}^m$ and $h \in \mathcal{H}$. Then, in D'_0 , conditioned on $\vec{C}'(\vec{b}, \vec{r}) = \vec{z}$, \vec{b} , and h , there are at most

$$\begin{aligned} (1 + 2^{-\Omega(n)}) \cdot 2^{w+\Delta-m} \left(\frac{|T|}{2^{w-\Delta}}\right) &\leq (1 + 2^{-\Omega(n)}) \cdot 2^{2\Delta-m} \cdot 2^{m(q+1)} \\ &= (1 + 2^{-\Omega(n)}) \cdot 2^{4\Delta+n-m} \cdot |S| \\ &= 2^{-\Omega(m)} \cdot |S| \end{aligned}$$

possible values for (\vec{r}, \vec{y}) . Thus, with probability at least $1 - 2^{-\Omega(n)}$, conditioned on values for the first three components of D'_0 , the fourth component $h(\vec{b}, \vec{r}, \vec{y})$ can cover at most a $2^{-\Omega(m)} \leq 2^{-\Omega(n)}$ fraction of S . In contrast, conditioned on values for the first three components of D'_1 , the fourth component is uniformly distributed on S . Therefore, $\|D'_0 - D'_1\| \geq 1 - 2^{-\Omega(n)}$. \square

In Vadhan [1999], it is shown that this Reversal Mapping can be better understood as a composition of two reductions, going the two directions between STATISTICAL DIFFERENCE and ENTROPY DIFFERENCE (the complete problem for SZK given in Goldreich and Vadhan [1999], which trivially reduces to its complement).

5. Other Forms of Zero Knowledge

3.1. WEAK-SZK AND EXPECTED POLYNOMIAL-TIME SIMULATORS. Recall that, in this article, we defined statistical zero-knowledge with respect to *strict* polynomial-time simulators. As noted in Section 2, the original definition of statistical zero-knowledge permits *expected* polynomial-time simulators, but only allowing strict polynomial-time simulators is not very restrictive when discussing honest-verifier proofs, as we are.

However, our techniques do say something about expected polynomial-time simulators, and, in particular, show that expected polynomial-time simulators are no more powerful than strict ones for public-coin statistical zero-knowledge. This is the first general equivalence between strict and expected polynomial-time simulators for statistical zero knowledge that we know of.

Indeed, we are able to generalize further to an even weaker notion, that of *weak* statistical zero knowledge (as previously considered in Di Crescenzo et al. [1997], where it was referred to as “nonuniform simulation”):

Definition 5.1. An interactive proof system (P, V) for a promise problem Π is *weak statistical zero knowledge* if for all polynomials p , there exists an efficient probabilistic (strict) polynomial-time algorithm S_p such that

$$\|S_p(x) - (P, V)(x)\| \leq \frac{1}{p}(|x|),$$

for all sufficiently long $x \in \Pi_Y$.

We denote by *weak-SZK* the class of promise problems admitting weak statistical zero-knowledge proofs, and by *public-coin weak-SZK* the class corresponding to such proofs that are also public coin. Note that any proof system admitting an expected polynomial-time simulator (in the usual sense) certainly also satisfies the requirements of weak statistical zero-knowledge. We show that in fact any public-coin weak statistical zero-knowledge proof system can be transformed into a statistical zero-knowledge proof system with a strict polynomial-time simulator achieving negligible (in fact, exponentially small) simulator deviation. In other words, *public-coin weak-SZK = SZK*.

PROPOSITION 5.2. *public-coin weak-SZK = SZK = public-coin SZK.*

The only obstacle in generalizing Proposition 5.2 to all weak statistical zero-knowledge proofs (instead of just public-coin ones) is that Okamoto’s [2000] private to public-coin transformation is only given for strict polynomial-time simulators achieving negligible simulator deviation. In fact, this generalization was accomplished in work (subsequent to ours) by Goldreich and Vadhan [1999].

In order to establish Proposition 5.2, it suffices to show that every problem in *public-coin weak-SZK* reduces to SD, as the proposition follows by closure under reductions (Corollary 4.3) and Okamoto’s theorem that *SZK = public-coin SZK*

(Theorem 3.10). Therefore, we need only establish the following generalization of Lemma 3.8:

LEMMA 5.3. *Suppose promise problem Π has a public-coin weak statistical zero-knowledge proof. Then there exist probabilistic (strict) polynomial-time machines A and B such that*

$$\begin{aligned} x \in \Pi_Y &\Rightarrow \|A(x) - B(x)\| < \frac{1}{3}, \text{ and} \\ x \in \Pi_N &\Rightarrow \|A(x) - B(x)\| > \frac{2}{3}. \end{aligned}$$

PROOF. The proof is identical to the proof of Lemma 3.8, except that wherever the simulator S is used in that proof, we replace it with S_p , a simulator with deviation $1/p(n)$, where $p(n) = 7n \cdot r(n)^3$. Then we replace Claim 3.12 with the following:

CLAIM 5.4. *If $x \in \Pi_Y$, then $\|A(x) - B(x)\| \leq 1/(3|x| \cdot r(|x|)^2)$.*

PROOF OF CLAIM. The proof is identical to the proof of Claim 3.12, except that now we have

$$\|A(x) - B(x)\| \leq 2^{-\Omega(|x|)} + 2r(|x|) \cdot \|S_p(x) - (P, V)(x)\| < \frac{1}{3|x| \cdot r(|x|)^2}. \quad \square$$

On the other hand, Claim 3.13 remains true, that is, $x \in \Pi_N$ implies $\|A(x) - B(x)\| \geq 1/12r(n)$. Then, as in the original proof, we consider the samplable distributions $\hat{A}(x) = \otimes^{s(|x|)} A(x)$ and $\hat{B}(x) = \otimes^{s(|x|)} B(x)$, where $s(n) = n \cdot r(n)^2$. If $x \in \Pi_Y$, $\|\hat{A}(x) - \hat{B}(x)\| \leq s(|x|)\|A(x) - B(x)\| < 1/3$, as desired. If $x \in \Pi_N$, then by the Direct Product Lemma (Lemma 3.4), $\|\hat{A}(x) - \hat{B}(x)\| \geq 1 - 2^{-\Omega(|x|)}$. \square

3.2. PERFECT AND COMPUTATIONAL ZERO KNOWLEDGE. Although the focus of this article is statistical zero knowledge, some of the techniques also apply to perfect and computational zero knowledge. In particular, for public-coin proof systems, we obtain variants of Lemma 3.8 for both perfect and computational zero knowledge. In addition, a restricted version of STATISTICAL DIFFERENCE can be shown to have perfect zero-knowledge proof.

First, we define some variants of SD. For any two constants α and β with $\alpha > \beta$, define:

$$\begin{aligned} \text{SD}_Y^{\alpha, \beta} &= \{(C_0, C_1) : \|C_0 - C_1\| \geq \alpha\} \\ \text{SD}_N^{\alpha, \beta} &= \{(C_0, C_1) : \|C_0 - C_1\| \leq \beta\}. \end{aligned}$$

$\text{SD}^{\alpha, \beta}$ is interreducible with SD and hence complete for SZK whenever $1 > \alpha^2 > \beta > 0$, because the Polarization Lemma generalizes to such thresholds. (See discussion at the end of Section 3.2.)

We can *almost* show that every problem that has a public-coin perfect zero-knowledge proof reduces to $\overline{\text{SD}^{1/2, 0}}$. The caveats are that either the original proof system must have perfect completeness, or we obtain distributions that are samplable in *expected* polynomial time rather than circuits.

PROPOSITION 5.5. *Every promise problem having a public-coin perfect zero-knowledge proof with perfect completeness reduces to $\overline{\text{SD}^{1/2, 0}}$.*

PROOF. It suffices to show that the distributions $A(x)$ and $B(x)$ constructed in the proof of Lemma 3.8 have statistical difference 0 on YES instances, when the original proof system has perfect completeness and the simulator deviation is 0. Indeed, for $i \geq 1$, the distributions $A_i(x)$ and $B_i(x)$ are identical if the simulator deviation is 0, and the distributions $A_0(x)$ and $B_0(x)$ are identical under the additional assumption that the proof system has perfect completeness. \square

We remark that Proposition 5.5 holds more generally for problems possessing public-coin perfect zero-knowledge proof systems for which the probability that $(P, V)(x)$ accepts can be computed in polynomial time for YES instances x . This is the case since B_0 in the proof above can be changed to output 1 with the appropriate probability.

PROPOSITION 5.6. *Suppose promise problem Π has a public-coin perfect zero-knowledge proof. Then there exist probabilistic expected polynomial time machines A and B such that*

$$\begin{aligned} x \in \Pi_Y &\Rightarrow \|A(x) - B(x)\| = 0, \text{ and} \\ x \in \Pi_N &\Rightarrow \|A(x) - B(x)\| \geq 1 - 2^{-\Omega(|x|)}. \end{aligned}$$

PROOF. The proof is nearly identical to that of Proposition 5.5, except that we must modify $A_0(x)$ and $B_0(x)$ to have statistical difference 0 (at the price of $B_0(x)$ becoming expected polynomial time). Let $c(n)$ be a polynomial bound on the number of random coins S uses on inputs of length n . Then we define A_0 and B_0 as follows (in both descriptions, $n = |x|$):

$A_0(x)$: Run $S(x)$ for $n \cdot c(n)$ repetitions. Output “1” if the majority are accepting conversations and ‘0’ otherwise.

$B_0(x)$: With probability $1 - 2^{-c(n)}$, output “1”. Otherwise, calculate the probability σ that $S(x)$ outputs an accepting conversation (by exhaustive search over all $2^{c(n)}$ random seeds). Now calculate

$$\tau = \sum_{i=0}^{\lfloor \frac{nc(n)}{2} \rfloor} \binom{nc(n)}{i} \sigma^i (1 - \sigma)^{nc(n)-i}.$$

If $\tau > 2^{-c(n)}$, output “1.” Otherwise, output “0” with probability $\tau/2^{-c(n)}$, and “1” otherwise.

Note that $B_0(x)$ runs in expected polynomial time, since, with probability $2^{-c(n)}$, it runs in time $\text{poly}(n) \cdot 2^{c(n)}$ and otherwise it runs in time $\text{poly}(n)$. Also observe that τ is exactly the probability that $A_0(x)$ outputs ‘0’.

Now we argue that, when $x \in \Pi_Y$, $A_0(x)$ and $B_0(x)$ have statistical difference 0, that is, output “1” with the same probability. Since $S(x)$ outputs a conversation that makes V accept with probability at least $2/3 - \text{neg}(n)$, the Chernoff bound implies that $\Pr[A_0(x) = 1] = 1 - 2^{-\Omega(nc(n))}$. This means that τ will always be less than $2^{-c(n)}$ (for sufficiently large n), so B_0 will output “0” with probability $2^{-c(n)} \cdot (\tau/2^{-c(n)}) = \tau$, which is the probability that A_0 outputs “0”. \square

Now, if we could show that $\text{SD}^{1/2,0}$ (or its complement) has a perfect zero-knowledge proof system, we would have something like a completeness result for

PZK. Although we do not know how to do this, we can instead show that $SD^{1,1/2} \in \text{PZK}$. Indeed, consider the protocol of Section 3.3 with the modification that the two parties use the XOR Lemma (Lemma 3.5) instead of the Polarization Lemma. Then the proof of Lemma 3.7 tells us that this protocol, when used for $SD^{1,1/2}$, has completeness error 0, simulator deviation 0, and soundness error $1/2 + 2^{-n}$. Thus, we have:

PROPOSITION 5.7. $SD^{1,1/2} \in \text{PZK}$.

For *computational* zero knowledge, the techniques of Lemma 3.8 give us something significantly weaker:

PROPOSITION 5.8. *Suppose promise problem Π has a public-coin computational zero-knowledge proof. Then there exist probabilistic polynomial-time machines A and B such that*

- (1) $x \in \Pi_N \Rightarrow \|A(x) - B(x)\| \geq 1 - 2^{-\Omega(|x|)}$, and
- (2) $\{A(x)\}_{x \in \Pi_Y}$ and $\{B(x)\}_{x \in \Pi_Y}$ are computationally indistinguishable ensembles of probability distributions.

Note that, in contrast to perfect and statistical zero knowledge, the conditions given in Proposition 5.8 do not give a way to distinguish YES and NO instances; it is possible for $A(x)$ and $B(x)$ to have statistical difference greater than $1 - 2^{-\Omega(|x|)}$ even for $x \in \Pi_Y$. Despite this, it still suffices to establish Theorem 5.13, which we present in Section 3.3. We also remark that Proposition 5.8 holds even when the simulator for the proof system runs in expected polynomial time, except that A and B will also run in expected polynomial time.

PROOF. The proof follows Lemma 3.8 exactly, except for Claim 3.12, which should be replaced with the following:

CLAIM 5.9. $\{A(x)\}_{x \in \Pi_Y}$ and $\{B(x)\}_{x \in \Pi_Y}$ are computationally indistinguishable ensembles of probability distributions.

We omit x from the notation for readability; below all probability distributions actually refer to *ensembles* indexed by $x \in \Pi_Y$. The proof in Claim 3.12 that A_0 and B_0 have exponentially small statistical difference still holds. Let the distributions A' and B' be obtained from A and B by removing the 0th components of A and B , respectively. Since A_0 and B_0 are independent of A' and B' , it suffices to show that A' and B' are computationally indistinguishable. To prove this, we first note that a hybrid argument shows that the distributions $\otimes^r(P, V)$ and $\otimes^r S$ are computationally indistinguishable, since (P, V) and S are computationally indistinguishable. Note that this step uses the fact that our definition of computational indistinguishability is with respect to nonuniform distinguishers, because (P, V) is not a samplable distribution (cf., [Goldreich 2001, Ch. 3, Exercise 9]).

Now we introduce a new distribution C . Define $C_i = (c_1, p_1, \dots, c_i)_{(P, V)}$ for $1 \leq i \leq r$, and let $C = C_1 \otimes \dots \otimes C_r$. Then, C and A' are computationally indistinguishable since a distinguisher D between them could be used to make a distinguisher D' between $\otimes^r(P, V)$ and $\otimes^r S$: Given a sequence of r transcripts (t_1, \dots, t_r) , D' truncates $t_i = (c_1, p_1, \dots, c_r, p_r)$ to produce $t'_i = (c_1, p_1, \dots, c_i)$ and feeds (t'_1, \dots, t'_r) to D . When fed with $\otimes^r S$, D' gives D a sample of A' , and when fed with $\otimes^r(P, V)$, D' gives D a sample of C .

Similarly, C and B' are also computationally indistinguishable because a distinguisher between them could be to make a distinguisher D' between $\otimes^r(P, V)$ and $\otimes^r S$: Given a sequence of r transcripts (t_1, \dots, t_r) , D' truncates $t_i = (c_1, p_1, \dots, c_r, p_r)$ and selects u_i according to the uniform distribution on strings of length $r(|x|)$ to produce $t'_i = (c_1, p_1, \dots, p_{i-1}, u_i)$ and feeds (t'_1, \dots, t'_r) to D . When fed with $\otimes^r S$, D' gives D a sample of B' , and when fed with $\otimes^r(P, V)$, D' gives D a sample of C .

Now, because both A' and B' are computationally indistinguishable from C , they must be computationally indistinguishable from each other, completing the proof. \square

3.3. HARD-ON-AVERAGE PROBLEMS AND ONE-WAY FUNCTIONS. Most, if not all, of cryptography relies on the existence of computational problems that are hard-on-average. However, the mere existence of a hard-on-average problem, even in NP, is not known to imply even the most basic cryptographic primitive, namely a one-way function. Ostrovsky [1991], however, showed that the existence of a hard-on-average problem in SZK *does* imply the existence of one-way functions. This result was subsequently generalized to CZK by Ostrovsky and Wigderson [1993].

In this section, we show how Ostrovsky's result follows readily from our Completeness Theorem and a result of Goldreich [1990] on computational indistinguishability. Using the generalization of our techniques to CZK described in the previous section, we also obtain a simpler proof of the Ostrovsky–Wigderson Theorem restricted to public-coin proof systems.

In order to state these theorems precisely, we need to define what we mean for a problem Π to be “hard.” Informally, we require that membership in Π is (very) hard to decide under some samplable distribution of instances.

Definition 5.10. An ensemble of distributions $\{D_n\}_{n \in \mathbb{N}}$ is said to be *samplable* if there is a probabilistic polynomial-time algorithm that, on input 1^n outputs a string distributed according to D_n .

Definition 5.11. A promise problem Π is *hard-on-average* if there exists a samplable ensemble of distributions $\{D_n\}_{n \in \mathbb{N}}$ such that the following holds: For every nonuniform probabilistic polynomial-time algorithm M , there exists a negligible function $\mu : \mathbb{N} \rightarrow [0, 1]$ such that

$$\Pr[M(x) \text{ correctly decides whether } x \text{ is a YES or NO instance of } \Pi]$$

$$\leq \frac{1}{2} + \mu(n) \quad \forall n \in \mathbb{N},$$

where the probability is taken over $x \leftarrow D_n$ and the coins of M . (If x violates the promise, then M is considered to be correct no matter what it outputs.)

In this section, we give new proofs of the following results.

THEOREM 5.12 ([OSTROVSKY 1991]). *If there is a hard-on-average promise problem in SZK, then one-way functions exist.*

THEOREM 5.13 ([OSTROVSKY AND WIGDERSON 1993] FOR PUBLIC-COIN PROOFS). *If a hard-on-average promise problem possesses a public-coin computational zero-knowledge proof system, then one-way functions exist.*

We only prove Theorem 5.13 as Theorem 5.12 then follows via Theorem 3.10 [Okamoto 2000]. Our proof makes use of Proposition 5.8 in conjunction with the following result of Goldreich [1990]:

PROPOSITION 5.14 ([GOLDREICH 1990]). *Suppose there exist two samplable ensembles of distributions, $\{A_n\}_{n \in \mathbb{N}}$ and $\{B_n\}_{n \in \mathbb{N}}$, such that*

- (1) $\{A_n\}$ and $\{B_n\}$ are computationally indistinguishable.
- (2) There is a polynomial $p : \mathbb{N} \rightarrow \mathbb{N}$ such that for all n , $\|A_n - B_n\| \geq 1/p(n)$.

Then one-way functions exist.

PROOF OF THEOREM 5.12. Suppose Π is a hard-on-average problem with a public-coin computational zero-knowledge proof and let $\{D_n\}$ be the ensemble of distributions under which Π is hard. By Proposition 5.8, there are probabilistic polynomial-time algorithms A and B such that

- (1) $x \in \Pi_N \Rightarrow \|A(x) - B(x)\| \geq 1 - 2^{-\Omega(|x|)}$, and
- (2) $\{A(x)\}_{x \in \Pi_Y}$ and $\{B(x)\}_{x \in \Pi_Y}$ are computationally indistinguishable.

(Note that if $\Pi \in \text{SZK}$, the Completeness Theorem and Polarization Lemma yield such A and B with the computational indistinguishability replaced by statistical difference $2^{-|x|}$.)

We show that the following ensembles $\{A_n\}$ and $\{B_n\}$ meet the requirements of Proposition 5.14:

A_n : Sample x according to D_n . Sample z from $A(x)$. Output (x, z) .

B_n : Sample x according to D_n . Sample z from $B(x)$. Output (x, z) .

The statistical farness of these ensembles will follow from the farness of $A(x)$ and $B(x)$ on NO instances. The computational indistinguishability will follow from the computational indistinguishability of $A(x)$ and $B(x)$ on YES instances, together with the fact that it is hard to distinguish YES instances of Π from NO instances.

To formalize this intuition, we make some observations which follow from the fact that Π is hard-on-average (where here and throughout this proof, we write $\text{neg}(n)$ to denote negligible functions):

- (1) $\Pr[D_n \notin \Pi_Y \cup \Pi_N] = \text{neg}(n)$.
- (2) $|\Pr[D_n \in \Pi_Y] - \frac{1}{2}| = \text{neg}(n)$ and $|\Pr[D_n \in \Pi_Y] - \frac{1}{2}| = \text{neg}(n)$.
- (3) The ensembles $\{D_n^Y\}_{n \in \mathbb{N}}$ and $\{D_n^N\}_{n \in \mathbb{N}}$ obtained by conditioning D_n on being a YES or NO instance, respectively, are computationally indistinguishable.

Items (1) and (2) hold because otherwise the trivial algorithm that always outputs YES or the one that always outputs NO would decide Π correctly with nonnegligible advantage. Item (3) holds because a distinguisher between $\{D_n^Y\}$ and $\{D_n^N\}$ could be used to decide Π with nonnegligible advantage.

CLAIM 5.15. $\|A_n - B_n\| \geq 1/2 - \text{neg}(n)$.

PROOF OF CLAIM. Since D_n must produce a NO instance of Π with probability at least $1/2 - \text{neg}(n)$, $\|A_n - B_n\| \geq (1/2 - \text{neg}(n)) \cdot (1 - 2^{-\Omega(n)}) = 1/2 - \text{neg}(n)$. \square

CLAIM 5.16. $\{A_n\}_{n \in \mathbb{N}}$ and $\{B_n\}_{n \in \mathbb{N}}$ are computationally indistinguishable.

PROOF OF CLAIM. Let M be any probabilistic polynomial-time algorithm. From the fact that $A(x)$ and $B(x)$ are computationally indistinguishable for YES instances, it follows that

$$\left| \Pr[M(x, A(x)) = 1 | x \in \Pi_Y] - \Pr[M(x, B(x)) = 1 | x \in \Pi_Y] \right| = \text{neg}(n), \quad (3)$$

where these probabilities (and all those to follow) are taken over $x \leftarrow D_n$ and the coins of all algorithms (M , A , and B). By the computational indistinguishability of $\{D_n^Y\}$ and $\{D_n^N\}$, we also have

$$\left| \Pr[M(x, A(x)) = 1 | x \in \Pi_Y] - \Pr[M(x, A(x)) = 1 | x \in \Pi_N] \right| = \text{neg}(n)$$

$$\left| \Pr[M(x, B(x)) = 1 | x \in \Pi_Y] - \Pr[M(x, B(x)) = 1 | x \in \Pi_N] \right| = \text{neg}(n).$$

Combining these with Equation 3, we see that all four conditional probabilities differ only by negligible amounts. Therefore,

$$\begin{aligned} & \Pr[M(x, A(x)) = 1] - \Pr[M(x, B(x)) = 1] \\ & \leq \left| \Pr[M(x, A(x)) = 1 | x \in \Pi_Y] - \Pr[M(x, B(x)) = 1 | x \in \Pi_Y] \right| \\ & \quad + \left| \Pr[M(x, A(x)) = 1 | x \in \Pi_N] - \Pr[M(x, B(x)) = 1 | x \in \Pi_N] \right| \\ & \quad + 2 \Pr[x \notin \Pi_Y \cup \Pi_N] \\ & = \text{neg}(n). \end{aligned}$$

This establishes the computational indistinguishability of $\{A_n\}$ and $\{B_n\}$. \square

Given these claims, the result now follows from Proposition 5.14.

3.4. EXTENSIONS TO CHEATING-VERIFIER ZERO KNOWLEDGE. The focus of study in this article has been the class of languages (or promise problems) possessing statistical zero-knowledge proofs *against an honest verifier*. However, in cryptographic applications, one usually wants the zero-knowledge condition to hold even against cheating verifier strategies that deviate arbitrarily from the specified protocol. There have been a number of results showing how to transform proof system that are statistical zero knowledge against the honest-verifier into ones that are statistical zero knowledge against cheating verifier strategies [Bellare et al. 1990; Ostrovsky et al. 1993; Damgård 1993; Damgård et al. 1995; Okamoto 2000; Goldreich et al. 1998]. As advocated by Bellare et al. [1990], one can use such transformations to translate results like ours about honest-verifier statistical zero knowledge to the cheating-verifier definition. In this section, we discuss which of our results apply to the cheating-verifier class and the appropriate formulations in each case.

Definition 5.17. An interactive protocol between a computationally unbounded prover P and a PPT verifier V is said to be a *cheating-verifier statistical zero-knowledge* proof system for a promise problem Π if the following holds.

- (1) If $x \in \Pi_Y$, then $\Pr[(P, V)(x) = \text{accept}] \geq 1 - c(|x|)$.
- (2) If $x \in \Pi_N$, then for all P^* , $\Pr[(P^*, V)(x) = \text{accept}] \leq s(|x|)$.
- (3) For all PPT V^* , there exists a PPT S and negligible function $\alpha(\cdot)$ such that for all $x \in \Pi_Y$, $\|S(x) - \text{View}_{P, V^*}(x)\| \leq \alpha(|x|)$.

As usual, $c(\cdot)$ is called the *completeness error*, $s(\cdot)$ the *soundness error*, and $\alpha(\cdot)$ the *simulator deviation* for V^* . **cheating-ver SZK** denotes the class of promise problems possessing cheating-verifier statistical zero-knowledge proofs.

Bellare et al. [1990] gave the first evidence that cheating-ver SZK = SZK, by proving it under the DISCRETE LOGARITHM assumption. Following sequence of subsequent works [Ostrovsky et al. 1993; Damgård 1993; Damgård et al. 1995], the complexity assumption was recently completely removed:

THEOREM 5.18 ([OKAMOTO 2000; GOLDREICH ET AL. 1998]).
cheating-ver SZK = SZK. Moreover, there is a transformation which converts an honest-verifier statistical zero-knowledge proof (P, V) for a promise problem Π into cheating-verifier statistical zero-knowledge proof (P', V') for Π such that:

- (1) (P', V') is public coin.
- (2) (P', V') has perfect completeness and soundness error 2^{-n} .
- (3) (P', V') has a universal black-box simulator S that works for all (even computationally unbounded) verifier strategies V^* .¹⁴
- (4) If (P, V) has simulator deviation $\alpha(n)$, then (P', V') has simulator deviation $\text{poly}(n) \cdot \alpha(n) + 2^{-n}$ for every V^* .
- (5) If (P, V) is security-parameterized, then so is (P', V') .¹⁵

We now use this result to translate our results about honest-verifier statistical zero knowledge to cheating-verifier statistical zero knowledge.

3.4.1. Class Properties. Since Theorem 5.18 asserts an equality of classes, all of our results about the class SZK also hold for cheating-ver SZK:

COROLLARY 5.19 (PROPERTIES OF THE CLASS cheating-ver SZK).

- (1) STATISTICAL DIFFERENCE is complete for cheating-ver SZK.
- (2) cheating-ver SZK is closed under Karp reductions, complement [Okamoto 2000], $\Phi(\cdot)$, and NC^1 truth-table reductions.
- (3) If there is a hard-on-average problem in cheating-ver SZK, then one-way functions exist [Ostrovsky 1991].

PROOF. Combine Theorem 5.18 with Theorem 3.1, Corollary 4.3, Corollary 4.4, Theorem 4.9, Corollary 4.14, and Theorem 5.12. \square

3.4.2. Protocol Properties. Unfortunately, Theorem 5.18 does not guarantee that every property satisfied by the honest-verifier proof system (P, V) is also satisfied by the cheating-verifier proof system (P', V') . Thus, although we have shown in Corollary 4.2 that every problem in SZK has a very efficient honest-verifier proof, it does not follow that it also has a cheating-verifier proof of similar efficiency. In particular, Theorem 5.18 does not preserve (even up to constant factors) the message complexity, communication complexity, or deterministic prover of Corollary 4.2.¹⁶ The only aspects of Corollary 4.2 that are maintained involve the error probabilities:

¹⁴ See Goldreich et al. [1998] for a definition.

¹⁵ This is not stated explicitly in Okamoto [2000] and Goldreich et al. [1998], but can be achieved replacing the input length n with $\max\{k, n\}$, where k is the security parameter, in their constructions.

¹⁶ In fact it is necessary that some of these properties are not maintained: only problems in BPP have cheating-verifier statistical zero-knowledge proofs with deterministic provers [Goldreich and Oren 1994], and only problems in BPP have cheating-verifier statistical zero-knowledge proofs that are constant round, public coin and have universal black-box simulators [Goldreich and Krawczyk 1996].

COROLLARY 5.20. *Every problem in SZK has a security-parameterized cheating-verifier statistical zero-knowledge proof with perfect completeness, soundness error 2^{-k} , and simulator deviation 2^{-k} .*

PROOF. Apply Theorem 5.18 to Corollary 4.2. \square

We note that the transformation of Bellare et al. [1990] *does* preserve the message complexity of the proof system up to a constant factor, but requires the assumption that the DISCRETE LOGARITHM problem is hard. Thus, under this assumption, their transformation can be combined with Okamoto's [2000] result that every problem in SZK has a constant-message honest-verifier proof system (or Corollary 4.2) and conclude that every problem in SZK has a constant-message cheating-verifier proof system.

3.4.3. Knowledge Complexity. Since Theorem 5.18 only refers to statistical zero-knowledge, we cannot immediately apply it our results about (nonzero) knowledge complexity. Below, we show how this nevertheless can be done for our results about knowledge complexity in the *hint* sense.

First, for all the variants of knowledge complexity discussed in Section 4.3, we can define cheating-verifier knowledge complexity analogously to Definition 5.17. We denote the cheating-verifier variant of a class \mathbf{C} with **cheating-ver C**. We begin by showing that honest-verifier and cheating-verifier statistical knowledge complexity in the *hint* sense coincide. To prove this, we observe one direction of the characterization of knowledge complexity in the hint sense given by Lemma 4.16 also holds for the cheating-verifier classes:

LEMMA 5.21. *Let Π be any promise problem and let $k(n)$ be any polynomially bounded function. Suppose there exists a promise problem $\Gamma \in$ **cheating-ver SZK** (respectively, **cheating-ver PZK**) such that*

- (1) $x \in \Pi_Y \Rightarrow$ *there exists a such that $|a| = k(|x|)$ and $(x, a) \in \Gamma_Y$, and*
- (2) $x \in \Pi_N \Rightarrow$ *for all a , $(x, a) \in \Gamma_N$.*

*Then $\Pi \in$ **cheating-ver SKC**_{hint}($k(n)$) (respectively, **cheating-ver PKC**_{hint}($k(n)$)).*

The proof of Lemma 5.21 is the same as the corresponding direction of Lemma 4.16. The reason the other direction of Lemma 4.16 does not immediately apply to the cheating-verifier case is that the hint function may be different for each verifier. However, we can deduce that direction from the following:

PROPOSITION 5.22. *For every polynomially-bounded function $k(n)$,*

$$\mathbf{SKC}_{\text{hint}}(k(n)) = \mathbf{cheating-ver SKC}_{\text{hint}}(k(n)).$$

PROOF. Clearly, **cheating-ver SKC**_{hint}($k(n)$) \subset **SKC**_{hint}($k(n)$). Now suppose Π is any language in **SKC**_{hint}($k(n)$), and let $\Gamma \in$ **SZK** be the promise problem guaranteed by Lemma 4.16. Then, by Theorem 5.18, $\Gamma \in$ **cheating-ver SZK**. Applying Lemma 5.21, we see that $\Pi \in$ **cheating-ver SKC**_{hint}($k(n)$). \square

Observe that we have actually proved something stronger: if $\Pi \in$ **SKC**_{hint}($k(n)$), then there is a proof system for Π with cheating-verifier statistical knowledge complexity $k(n)$ for which the *same hint function* can be used for every verifier.

Unfortunately, analogues of Proposition 5.22 do not appear to follow immediately from the fact that $\text{SZK} = \text{cheating-ver SZK}$.

Given Proposition 5.22, it follows immediately that the collapse in the SKC_{hint} hierarchy (Theorem 4.15) also holds for the cheating-verifier classes:

PROPOSITION 5.23. *For any polynomially bounded function $k(n)$,*

$$\text{cheating-ver SKC}_{\text{hint}}(k(n) + \log n) = \text{cheating-ver SKC}_{\text{hint}}(k(n)).$$

In contrast, we do not know whether our results on the perfect knowledge complexity of SZK hold for the analogous cheating-verifier classes. To apply the same approach, one would have to analyze the (cheating-verifier) perfect knowledge complexity of the protocols obtained by performing the transformations of Okamoto [2000] and Goldreich et al. [1998] on the protocol for SD . These transformations could conceivably increase the perfect knowledge complexity dramatically.

6. Open Problems

We recall some of the open problems we mentioned throughout the article, along with some additional research problems raised by this work.

- (1) Does $\text{STATISTICAL DIFFERENCE}$ remain complete for SZK and the Polarization Lemma (Lemma 3.3) still hold when the thresholds are α, β such that $\alpha^2 < \beta < \alpha$?
- (2) Is SZK closed under general Cook reductions (adaptive or nonadaptive)? (Recall that in Corollary 4.14, we showed that SZK is closed under NC_1 truth-table reductions.)
- (3) Do the other forms of statistical knowledge complexity collapse like the SKC_{hint} hierarchy (cf., Theorem 4.15)?
- (4) Find natural complete problems for PZK or CZK . (The results of Section 3.2 are efforts in this direction.)
- (5) Find additional natural complete problems for SZK , for example, combinatorial or number-theoretic problems. While we have used SZK -completeness mainly as a positive tool, it could also provide strong evidence of intractability, as SZK contains many problems believed to be hard.
- (6) Does $\text{SZK} = \text{PZK}$? It was this question, posed to us by Shafi Goldwasser, that started us on this research, and unfortunately the answer remains a mystery. However, our Completeness Theorem does imply that $\text{SZK} = \text{PZK}$ if and only if $\text{STATISTICAL DIFFERENCE}$ is in PZK .

Appendixes

A. The Statistical Difference Metric

PROOF OF FACT 2.1. For any set $S \subset D$,

$$\begin{aligned} & 2 |\Pr[X \in S] - \Pr[Y \in S]| \\ &= |\Pr[X \in S] - \Pr[Y \in S]| + |\Pr[X \notin S] - \Pr[Y \notin S]| \end{aligned}$$

$$\begin{aligned}
 &= \left| \sum_{x \in S} (\Pr[X = x] - \Pr[Y = x]) \right| + \left| \sum_{x \notin S} (\Pr[X = x] - \Pr[Y = x]) \right| \\
 &\leq \sum_{x \in S} |\Pr[X = x] - \Pr[Y = x]| + \sum_{x \notin S} |\Pr[X = x] - \Pr[Y = x]| \\
 &= \|X - Y\|_1.
 \end{aligned}$$

Equality is achieved by taking $S = \{x : \Pr[X = x] > \Pr[Y = x]\}$. \square

PROOF OF FACT 2.3

$$\begin{aligned}
 \|(X_1, X_2) - (Y_1, Y_2)\| &\leq \|(X_1, X_2) - (Y_1, X_2)\| + \|(Y_1, X_2) - (Y_1, Y_2)\| \\
 &= \frac{1}{2} |X_1 \otimes X_2 - Y_1 \otimes X_2|_1 + \frac{1}{2} |Y_1 \otimes X_2 - Y_1 \otimes Y_2|_1 \\
 &= \frac{1}{2} |(X_1 - Y_1) \otimes X_2|_1 + \frac{1}{2} |Y_1 \otimes (X_2 - Y_2)|_1 \\
 &= \frac{1}{2} |X_1 - Y_1|_1 \cdot |X_2|_1 + \frac{1}{2} |Y_1|_1 \cdot |X_2 - Y_2|_1 \\
 &= \|X_1 - Y_1\| + \|X_2 - Y_2\|. \quad \square
 \end{aligned}$$

PROOF OF FACT 2.4. Let $A = (f, R)$ be any randomized procedure. Then, for any set $S \subset F$,

$$\begin{aligned}
 &|\Pr[A(X) \in S] - \Pr[A(Y) \in S]| \\
 &= |\Pr[f(X \otimes R) \in S] - \Pr[f(Y \otimes R) \in S]| \\
 &= |\Pr[X \otimes R \in f^{-1}(S)] - \Pr[Y \otimes R \in f^{-1}(S)]| \\
 &\leq \|X \otimes R - Y \otimes R\| \\
 &\leq \|X - Y\| + \|R - R\| \text{ (by Fact 2.3)} \\
 &= \|X - Y\|.
 \end{aligned}$$

Taking the maximum over all sets S completes the proof. \square

PROOF OF FACT 2.5. Let $T \subset D$ be the set of x 's for which $\|X_2|_{X_1=x} - Y_2|_{Y_1=x}\| < \delta$. Now, let S be an arbitrary subset of $D \times E$ and, for every $x \in D$, define $S_x = \{y \in E : (x, y) \in S\}$. Then,

$$\begin{aligned}
 \Pr[X \in S] &\leq \Pr[X_1 \notin T] + \sum_{x \in T} \Pr[X_2 \in S_x | X_1 = x] \cdot \Pr[X_1 = x] \\
 &< \epsilon + \sum_{x \in T} (\Pr[Y_2 \in S_x | Y_1 = x] + \delta) \cdot \Pr[Y_1 = x] \\
 &\leq \epsilon + \delta + \Pr[Y \in S].
 \end{aligned}$$

By symmetry, we also have $\Pr[Y \in S] < \epsilon + \delta + \Pr[X \in S]$. Since S was arbitrary, $\|X - Y\| < \epsilon + \delta$. \square

PROOF OF FACT 2.6. Let $S = \{x : (1 - \sqrt{\epsilon}) \Pr[X = x] \geq \Pr[Y = x]\}$, that is, the set of x 's for which the left-hand inequality in Fact 2.6 is violated. Then,

$$\begin{aligned}
 \Pr[Y \in S] &\leq (1 - \sqrt{\epsilon}) \Pr[X \in S] \\
 &= \Pr[X \in S] - \sqrt{\epsilon} \cdot \Pr[X \in S].
 \end{aligned}$$

Thus, $\sqrt{\epsilon} \cdot \Pr[X \in S] \leq \|X - Y\| < \epsilon$, so we must have $\Pr[X \in S] < \sqrt{\epsilon}$. A similar argument shows that the right-hand inequality in Fact 2.6 is violated with probability less than $\sqrt{\epsilon}$. \square

B. A Generic Complete Problem for PZK

In this appendix, we show how to obtain a complete promise problem for PZK directly from the definition of the class. However, in contrast to STATISTICAL DIFFERENCE, this problem will be essentially a restatement of the definition of the class and therefore of little use.

The complete promise problem for PZK is PZK-GENERIC, which we now define. An instance of PZK-GENERIC is a quadruple $(V, S, x, 1^t)$, where V is a description of an interactive probabilistic Turing machine and S is a description of a (noninteractive) probabilistic Turing machine. A YES instance is such a quadruple for which there exists a prover strategy P such that

- (1) The interaction between P and V on x takes at most t steps (including the computation time for V) and V accepts in this interaction with probability at least $2/3$.
- (2) The running time of S on input x is at most t .
- (3) S outputs `fail` with probability at most $1/2$, and conditioned on not failing, the output distribution of S is *identical* to V 's view of the interaction with P on x .

A NO instance is a quadruple such that for all prover strategies P ,

- (1) The interaction between P and V on x takes at most t steps (including the computation time for V) and V rejects in this interaction with probability at least $2/3$.
- (2) The running time of S on input x is at most t .

PROPOSITION B.1. PZK-GENERIC is complete for PZK.

PROOF. First, we show that every promise problem Π in PZK reduces to PZK-GENERIC. Let (P, V) be the perfect zero-knowledge proof system for Π with simulator S . Let $t(n)$ be a (polynomial) upper bound on both the running time of S and the number of steps of the interaction of P and V on inputs of length n . Then

$$x \mapsto (V, S, x, 1^{t(|x|)})$$

is a polynomial-time reduction from Π to PZK-GENERIC.

Now we argue that PZK-GENERIC \in PZK. Consider the following descriptions of a verifier \bar{V} , a prover \bar{P} , and a simulator \bar{S} :

$\bar{V}(V, S, x, 1^t)$: When interacting with any machine, simulate V on input x .

$\bar{P}(V, S, x, 1^t)$: Exhaustively search for a prover strategy P for which V 's view of $(P, V)(x)$ is identical to the output distribution of $S(x)$ (conditioned on $S(x) \neq \text{fail}$.) If one exists, follow that strategy, otherwise output `fail`.¹⁷

$\bar{S}(V, S, x, 1^t)$: Simulate S on input x .

¹⁷ Alternatively, \bar{P} can act as the *simulation-based prover* (see Section 3.5).

It is easy to see that these definitions provide a perfect zero-knowledge proof system for PZK-GENERIC. \square

The problem with extending this example to SZK is Condition 3 for YES instances. “Identical” needs to be replaced by “negligible statistical difference,” but it is not clear what negligible function to put there. We do not know how to get around this difficulty without using our Completeness Theorem, which implies that every problem in SZK has a statistical zero-knowledge proof with the *same* simulator deviation 2^{-n} (cf., Corollary 4.2).¹⁸

Another observation worth mentioning, pointed out to us by Mihir Bellare, Oded Goldreich, and Madhu Sudan, is that PZK-GENERIC can be modified to obtain complete promise problems for cheating-ver PZK (as long as we restrict to “black-box” simulation) and also the various forms of PKC.

C. An Example for GRAPH ISOMORPHISM

For illustrative purposes, here we explicitly describe what happens when the reduction to and proof system for STATISTICAL DIFFERENCE are applied to the well-known public-coin perfect zero-knowledge proof system for GRAPH ISOMORPHISM [Goldreich et al. 1991]:

Perfect zero-knowledge proof system for GRAPH ISOMORPHISM.

Input: (G_0, G_1) .

1. P sends V a random isomorphic copy H of G_0 .
2. V picks $b \in \{0, 1\}$ at random and sends it to P .
3. P sends V a random isomorphism π between G_b and H , if one exists.
4. V checks that $\pi G_b = H$.

Simulator S , on input (G_0, G_1) :

1. Pick random $b \in \{0, 1\}$ and a random permutation π .
2. Output $(\pi G_b, b, \pi)$.

Notice that the conversations output by S always make V accept.

If the reduction to SD from the proof of Lemma 3.8 is applied to the above protocol, the following distributions are obtained:

$A_0(G_0, G_1)$: Always output 1.

$B_0(G_0, G_1)$: Always output 1.

$A_1(G_0, G_1)$: Output $(\pi G_b, b)$ for a random permutation π and $b \in \{0, 1\}$ chosen at random.

$B_1(G_0, G_1)$: Output $(\pi G_b, c)$ for a random permutation π and b and c chosen uniformly and independently from $\{0, 1\}$.

Thus, $\|A_0(x) - B_0(x)\|$ always equals 0. $\|A_1(x) - B_1(x)\|$ is easily seen to be 0 if $G_0 \cong G_1$ and $1/2$ if $G_0 \not\cong G_1$. For the rest of this section, we ignore A_0 and B_0 since they are irrelevant.

¹⁸ We do not know how to overcome this difficulty by using the result of Bellare [1997], which states that any countable set of negligible functions is “dominated” by a single negligible function. The reason is that there are uncountably many problems in the promise-class SZK.

If we now apply the protocol for SD from Section 3.3 to the distributions A_1 and B_1 (without first applying the Polarization Lemma), we obtain the following proof system (P', V') for GRAPH NONISOMORPHISM:

- (1) V' picks a random bit $d \in \{0, 1\}$. If $d = 0$, V' chooses a random bit $b \in \{0, 1\}$ and a random permutation π and sends $(\pi G_b, b)$ to P' . If $d = 1$, V' chooses random bits $b, c \in \{0, 1\}$ and a random permutation π and sends $(\pi G_b, c)$ to P' .
- (2) P' receives message (H, b) from V' . P' attempts to guess d as follows: If H is isomorphic to G_b , then P' guesses 0, else P' guesses 1.
- (3) V' accepts if the P' guesses d correctly.

Now, if G_0 is not isomorphic to G_1 , then P' will guess correctly with probability $3/4$. However, if G_0 is isomorphic to G_1 , then no prover can guess correctly with probability greater than $1/2$. The above protocol is of the same spirit as the standard GRAPH NONISOMORPHISM protocol [Goldreich et al. 1991]. In both cases, the verifier randomly permutes one of the graphs to obtain a graph H and in order for the prover to succeed with probability greater than $1/2$, the prover needs to be able to identify from which graph H came.

ACKNOWLEDGMENTS. We are grateful to our advisor, Shafi Goldwasser, for getting us started on the topic of statistical zero knowledge and providing direction and advice throughout our work. We are indebted to Oded Goldreich for many enlightening conversations and subsequent collaboration on this topic, and his extensive help with the writing of this article. We also thank Mihir Bellare, Erez Petrank, Tatsuaki Okamoto, Madhu Sudan, Luca Trevisan, Avi Wigderson, and the anonymous conference and journal reviewers for many helpful suggestions, clarifying discussions, and/or encouragement.

REFERENCES

- AIELLO, W., BELLARE, M., AND VENKATESAN, R. 1995. Knowledge on the average—perfect, statistical and logarithmic. In *Proceedings of the 27th Annual ACM Symposium on the Theory of Computing* (Las Vegas, Nev., May 29–June 1). ACM, New York, pp. 469–478.
- AIELLO, W., AND HÅSTAD, J. 1991. Statistical zero-knowledge languages can be recognized in two rounds. *J. Comput. Syst. Sci.* 42, 3 (June), 327–345.
- AJTAI, M., AND BEN-OR, M. 1984. A theorem on probabilistic constant depth computations. In *Proceedings of the 16th Annual ACM Symposium on Theory of Computing* (Washington, D.C.). ACM, New York, pp. 471–474.
- ARORA, S., LUND, C., MOTWANI, R., SUDAN, M., AND SZEGEDY, M. 1998. Proof verification and the hardness of approximation problems. *J. ACM* 45, 3 (May), 501–555.
- ARORA, S., AND SAFRA, S. 1998. Probabilistic checking of proofs: A new characterization of NP. *J. ACM* 45, 1 (Jan.), 70–122.
- BABAI, L., AND MORAN, S. 1988. Arthur–Merlin games: A randomized proof system and a hierarchy of complexity classes. *J. Comput. Syst. Sci.* 36, 254–276.
- BELLARE, M. 1997. A note on negligible functions. Tech. Rep. CS97-529, Dept. Computer Science and Engineering, Univ. California at San Diego, San Diego, Calif., March. Also available from the Theory of Cryptography Library (<http://theory.lcs.mit.edu/~tcrypt01>).
- BELLARE, M., MICALI, S., AND OSTROVSKY, R. 1990. The (true) complexity of statistical zero knowledge. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing* (Baltimore, Md. May 14–16) ACM, New York, pp. 494–502.
- BELLARE, M., AND PETRANK, E. 1992. Making zero-knowledge provers efficient. In *Proceedings of the 24th Annual ACM Symposium on the Theory of Computing* (Victoria, B.C. Canada, May 4–6). ACM, New York, pp. 711–722.

- BOPANA, R. B., HÅSTAD, J., AND ZACHOS, S. 1987. Does co-NP have short interactive proofs? *Inf. Proc. Lett.* 25, 127–132.
- COOK, S. A. 1971. The complexity of theorem-proving procedures. In *Conference Record of the 3rd Annual ACM Symposium on Theory of Computing* (Shaker Heights, Ohio, 3–5). ACM, New York, pp. 151–158.
- COVER, T. M., AND THOMAS, J. A. 1991. *Elements of Information Theory*, 2nd ed. Wiley Series in Telecommunications. Wiley, New York.
- DAMGÅRD, I., AND CRAMER, R. 1996. On monotone function closure of perfect and statistical zero-knowledge. Theory of Cryptography Library: Record 96-03. <http://theory.lcs.mit.edu/~tcryptol>.
- DAMGÅRD, I., GOLDREICH, O., OKAMOTO, T., AND WIGDERSON, A. 1995. Honest verifier vs. dishonest verifier in public coin zero-knowledge proofs. In *Proceedings of Crypto '95*. Lecture Notes in Computer Science, vol. 403. Springer-Verlag, New York.
- DAMGÅRD, I. B. 1993. Interactive hashing can simplify zero-knowledge protocol design without computational assumptions (extended abstract). In *Advances in Cryptology—CRYPTO '93*, Douglas R. Stinson, Ed. Lecture notes in Computer Science, vol. 773. Springer-Verlag, New York, pp. 100–109.
- DE SANTIS, A., DI CRESCENZO, G., PERSIANO, G., AND YUNG, M. 1994. On monotone formula closure of SZK. In *Proceedings of the 35th Annual Symposium on Foundations of Computer Science* (Santa Fe, N.M., Nov. 20–22). IEEE Computer Society Press, Los Alamitos, Calif., pp. 454–465.
- DE SANTIS, A., DI CRESCENZO, G., PERSIANO, G., AND YUNG, M. 1998. Image Density is complete for non-interactive-SZK. In *Automata, Languages and Programming, 25th International Colloquium* (Aalborg, Denmark, July 13–17). Lecture Notes in Computer Science, Springer-Verlag, New York, pp. 784–795. (See also preliminary draft of full version, May 1999.)
- DI CRESCENZO, G., OKAMOTO, T., AND YUNG, M. 1997. Keeping the SZK-verifier honest unconditionally. In *Advances in Cryptology—CRYPTO '97*, Burton S. Kaliski Jr., Ed. Lecture Notes in Computer Science, vol. 1294. Springer-Verlag, New York, 31–45.
- DI CRESCENZO, G., SAKURAI, K., AND YUNG, M. 2000. On zero-knowledge proofs: “From membership to decision”. In *Proceedings of the 32nd Annual ACM Symposium on Theory of Computing* (Portland, Ore., May). ACM, New York, pp. 255–264.
- EVEN, S., SELMAN, A. L., AND YACOBI, Y. 1984. The complexity of promise problems with applications to public-key cryptography. *Inf. Cont.* 61, 2 (May), 159–173.
- FORTNOW, L. 1989. The complexity of perfect zero-knowledge. In *Advances in Computing Research*, vol. 5, Silvio Micali, Ed. JAC Press, Inc., pp. 327–343.
- GOLDREICH, O. 1990. A note on computational indistinguishability. *Inf. Proc. Lett.* 34, 6 (May), 277–281.
- GOLDREICH, O. 2001. *Foundations of Cryptography: Basic Tools*. Cambridge University Press.
- GOLDREICH, O., AND GOLDWASSER, S. 2000. On the limits of nonapproximability of lattice problems. *J. Comp. Syst. Sci.* 60, 3, 540–563.
- GOLDREICH, O., AND KRAWCZYK, H. 1996. On the composition of zero-knowledge proof systems. *SIAM J. Comput.* 25, 1 (Feb.), 169–192.
- GOLDREICH, O., AND KUSHILEVITZ, E. 1993. A perfect zero-knowledge proof system for a problem equivalent to the discrete logarithm. *J. Cryptology* 6, 97–116.
- GOLDREICH, O., MICALI, S., AND WIGDERSON, A. 1991. Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. *J. ACM* 38, 1, 691–729.
- GOLDREICH, O., NISAN, N., AND WIGDERSON, A. 1995. On Yao’s XOR lemma. Tech. Rep. TR95–050. Electronic Colloquium on Computational Complexity. Mar. <http://www.eccc.uni-trier.de/eccc>.
- GOLDREICH, O., AND OREN, Y. 1994. Definitions and properties of zero-knowledge proof systems. *J. Cryptology* 7, 1 (Winter), 1–32.
- GOLDREICH, O., OSTROVSKY, R., AND PETRANK, E. 1998. Computational complexity and knowledge complexity. *SIAM J. Comput.* 27, 4 (Aug.), 1116–1141.
- GOLDREICH, O., AND PETRANK, E. 1999. Quantifying knowledge complexity. *Comput. Complex.* 8, 1, 50–98.
- GOLDREICH, O., SAHAI, A., AND VADHAN, S. 1998. Honest verifier statistical zero-knowledge equals general statistical zero-knowledge. In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing* (Dallas, Tex., May 23–26). ACM, New York, pp. 399–408.

- GOLDREICH, O., SAHAI, A., AND VADHAN, S. 1999. Can statistical zero-knowledge be made non-interactive?, or On the relationship of SZK and NISZK. In *Advances in Cryptology—CRYPTO '99* (Aug. 15–19). Lecture Notes in Computer Science, Springer-Verlag, New York.
- GOLDREICH, O., AND VADHAN, S. 1999. 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, Ga., May). IEEE Computer Society Press, Los Alamitos, Calif., pp. 54–73.
- GOLDREICH, O., VADHAN, S., AND WIGDERSON, A. 2001. On interactive proofs with a laconic prover. In *Automata, Languages and Programming, 28th International Colloquium*. (Crete, Greece, July 7–11). Lecture Notes in Computer Science, vol. 2076. Springer-Verlag, New York, pp. 334–345.
- GOLDWASSER, S., AND MICALI, S. 1984. Probabilistic encryption. *J. Comput. Syst. Sci.* 28, 2 (Apr.), 270–299.
- GOLDWASSER, S., MICALI, S., AND RACKOFF, C. 1989. The knowledge complexity of interactive proof systems. *SIAM J. Comput.* 18, 1 (Feb.), 186–208.
- GOLDWASSER, S., AND SIPSER, M. 1989. Private coins versus public coins in interactive proof systems. In *Advances in Computing Research*, vol. 5. Silvio Micali, Ed. JAC Press, Inc., pp. 73–90.
- GUTFREUND, D., AND BEN-OR, M. 2000. Increasing the power of the dealer in non-interactive zero-knowledge proof systems. In *Advances in Cryptology—ASIACRYPT '00* (Kyoto, Japan). Springer-Verlag, Berlin, Germany. To appear.
- HÅSTAD, J., IMPAGLIAZZO, R., LEVIN, L. A., AND LUBY, M. 1999. A pseudorandom generator from any one-way function. *SIAM J. Comput.* 28, 4, 1364–1396 (electronic).
- HOFRI, M. 1995. *Analysis of Algorithms: Computational Methods & Mathematical Tools*. Oxford University Press.
- KANNAN, S. 1989. *Program Checkers for Algebraic Problems*. Ph.D. dissertation. Univ. of California, Berkeley, Berkeley, Calif.
- KARP, R. M. 1972. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, J. W. Thatcher and R. E. Miller, Eds. Plenum Press, New York, pp. 85–103.
- LADNER, R. E., LYNCH, N. A., AND SELMAN, A. L. 1975. A comparison of polynomial time reducibilities. *Theoret. Comput. Sci.* 1, 2 (Dec.), 103–123.
- LEVIN, L. A. 1973. Universal'nyĕ perebornyĕ zadachi (Universal search problems : in Russian). *Problemy Peredachi Informatsii* 9, 3, 265–266.
- LUND, C., FORTNOW, L., KARLOFF, H., AND NISAN, N. 1992. Algebraic methods for interactive proof systems. *J. ACM* 39, 4 (Oct.), 859–868.
- OKAMOTO, T. 2000. On relationships between statistical zero-knowledge proofs. *J. Comput. Syst. Sci.* 60, 1 (Feb.), 47–108.
- OSTROVSKY, R. 1991. One-way functions, hard on average problems, and statistical zero-knowledge proofs. In *Proceedings of the 6th Annual Structure in Complexity Theory Conference* (Chicago, Ill. June 30–July 3). IEEE Computer Society Press, Los Alamitos, Calif., pp. 133–138.
- OSTROVSKY, R., VENKATESAN, R., AND YUNG, M. 1993. Interactive hashing simplifies zero-knowledge protocol design. In *Proceedings of Eurocrypt '93*. Lecture Notes in Computer Science. Springer-Verlag, New York.
- OSTROVSKY, R., AND WIGDERSON, A. 1993. One-way functions are essential for non-trivial zero-knowledge. In *Proceedings of the 2nd Israel Symposium on Theory of Computing and Systems*.
- PAPADIMITRIOU, C. H. 1994. *Computational Complexity*. Addison-Wesley, Reading, Mass.
- PETRAND, E., AND TARDOS, G. 1996. On the knowledge complexity of \mathcal{NP} . In *Proceedings of the 37th Annual Symposium on Foundations of Computer Science* (Burlington, Vt. Oct. 14–16). IEEE Computer Society Press, Los Alamitos, Calif., pp. 494–503.
- SAHAI, A. 2000. *Frontiers in Zero Knowledge*. Ph.D. dissertation. Massachusetts Institute of Technology, Cambridge, Mass.
- SAHAI, A., AND VADHAN, S. 1999. Manipulating statistical difference. In *Randomization Methods in Algorithm Design (DIMACS Workshop, December 1997)*, Panos Pardalos, Sanguthevar Rajasekaran, and José Rolim, Eds. DIMACS Series in Discrete Mathematics and Theoretical Computer Science, vol. 43. American Mathematical Society, Providence, R.I., pp. 251–270.
- SAHAI, A., AND VADHAN, S. P. 1997. A complete promise problem for statistical zero-knowledge. In *Proceedings of the 38th Annual Symposium on Foundations of Computer Science* (Miami Beach, Fla. Oct. 20–22). IEEE Computer Society Press, Los Alamitos, Calif., pp. 448–457.
- SHAMIR, A. 1992. $IP = PSPACE$. *J. ACM* 39, 4 (Oct.), 869–877.
- VADHAN, S. P. 1999. *A Study of Statistical Zero-Knowledge Proofs*. Ph.D. dissertation. Massachusetts Institute of Technology, Cambridge, Mass.

- VADHAN, S. P. 2000. On transformations of interactive proofs that preserve the prover's complexity. In *Proceedings of the 32nd Annual ACM Symposium on Theory of Computing* (Portland, Ore., May). ACM, New York, pp. 200–207.
- YAO, A. C. 1982. Theory and applications of trapdoor functions (extended abstract). In *Proceedings of the 23rd Annual Symposium on Foundations of Computer Science* (Chicago, Ill., Nov. 3–5). IEEE Computer Society Press, Los Alamitos, Calif., pp. 80–91.

RECEIVED OCTOBER 2000; REVISED OCTOBER 2002; ACCEPTED OCTOBER 2002