An Equivalence between Zero Knowledge and Commitments

Shien Jin Ong* and Salil Vadhan**

Harvard University, School of Engineering & Applied Sciences
33 Oxford Street, Cambridge, MA 02138, USA
{shienjin,salil}@eecs.harvard.edu

Abstract. We show that a language in NP has a zero-knowledge protocol if and only if the language has an "instance-dependent" commitment scheme. An instance-dependent commitment schemes for a given language is a commitment scheme that can depend on an instance of the language, and where the hiding and binding properties are required to hold only on the YES and NO instances of the language, respectively. The novel direction is the *only if* direction. Thus, we confirm the widely held belief that commitments are not only sufficient for zero knowledge protocols, but *necessary* as well. Previous results of this type either held only for restricted types of protocols or languages, or used nonstandard relaxations of (instance-dependent) commitment schemes.

1 Introduction

From the early days in the study of zero knowledge, it has seemed that *commitment schemes* are the heart of *zero-knowledge protocols*. Indeed, the first construction of zero-knowledge proofs for all of NP, due to Goldreich, Micali, and Wigderson [GMW], shows that commitment schemes *suffice* for zero knowledge. Moreover, there have been a number of partial converses to this result, showing how to obtain certain kinds of commitments from certain kinds of zero-knowledge protocols for certain kinds of languages. In this paper, we present a complete equivalence between zero knowledge protocols and *instance-dependent* commitment schemes [BMO,IOS], in which the protocol depends on a given instance of a language (or promise problem). Specifically, we show that for every language $L \in \mathrm{NP}$, *L has a zero-knowledge protocol if and only if L has an instance-dependent commitment scheme.* Thus, we confirm the intuition that commitments are not only sufficient for zero knowledge, but necessary as well.

1.1 Review of Zero Knowledge and Commitments

In zero-knowledge protocols [GMR], a *prover* tries to convince a *verifier* that an assertion is true, namely that some string x is a YES instance of a (promise)

* Supported by NSF grant CNS-0430336.

** Supported by NSF grant CNS-0430336 and US-Israel BSF grant 2002246. Work done in part when visiting UC Berkeley, supported by a Guggenheim Fellowship and the Miller Institute for Basic Research in Science.

problem Π,[1] without leaking any additional knowledge. Zero-knowledge protocols have two security requirements. Informally, *soundness* says that a cheating prover should not be able to convince the verifier of a false statement, and *zero knowledge* says that a cheating verifier should not be able to learn anything from the interaction other than the fact that the assertion being proven is true. Both security requirements come in two flavors — *statistical*, whereby we require security to hold even against computationally unbounded cheating strategies (except with negligible probability[2]), and *computational*, whereby we only require security against polynomial-time strategies (except with negligible probability). Protocols with statistical soundness are typically called interactive *proof systems* (which constitute the original model proposed by [GMR]), and those with computational soundness are typically called *argument systems* (which were introduced by [BCC]). Considering all combinations of computational and statistical versions of soundness and zero knowledge rise to four main flavors of zero knowledge protocols, and thus four complexity classes consisting of the problems Π having zero-knowledge protocols of a particular flavor. We denote these complexity classes SZKP, CZKP, SZKA, and CZKA, with the prefix of S or C denoting statistical or computational zero knowledge, and the suffix of P or A denoting proof systems (statistical soundness) or argument systems (computational soundness).

A commitment scheme is the cryptographic analogue of a locked box. It is a two-stage interactive protocol between a pair of probabilistic polynomial-time parties, the *sender* and the *receiver*. In the first stage, the sender "commits" to a string m, corresponding to locking an object in the box. In the second stage, the sender "reveals" m to the receiver, corresponding to opening the box. Like zero-knowledge protocols, commitment schemes have two security properties. Informally, *hiding* says that a cheating receiver should not be able to learn anything about m during the commit stage, and *binding* says that a cheating sender should not be able to reveal two different messages after the commit stage. Again, each of these properties can be statistical (holding against computationally unbounded cheating strategies, except with negligible probability) or computational (holding against polynomial-time cheating strategies, except with negligible probability). Thus we again get four flavors of commitment schemes, but it is easily seen to be impossible to simultaneously achieve statistical security for both hiding and binding. However, it is known that if one-way functions exist, then we can achieve statistical security for either one of the security properties [HILL,Nao,NOV,HR]. Conversely, commitment schemes, even with both properties computational, imply one-way functions [IL].

[1] A promise problem Π is a pair (Π_Y, Π_N) of disjoint sets of strings, corresponding to the YES instances and NO instances. Given a string x that is "promised" to be in $\Pi_Y \cup \Pi_N$, the task is to decide whether $x \in \Pi_Y$ or $x \in \Pi_N$.

[2] An even stronger notion that statistical security is *perfect security*, where the clause "except with negligible probability" is removed. We will not consider perfect security in this paper.

1.2 Previous Work

The classic construction of Goldreich, Micali, and Wigderson [GMW] shows how to construct zero-knowledge protocols for NP given any commitment scheme. Moreover, the security properties of the commitment scheme translate to the security properties of the zero-knowledge protocol: a statistically (resp., computationally) hiding commitment scheme yields statistical (resp., computational) zero knowledge, and a statistically (resp., computationally) binding commitment scheme yields a proof (resp. argument) system.[3] Thus, if one-way functions exist, CZKP, SZKA, and CZKA are very powerful in that they contain NP (and even the classes MA or IP [IY,BGG$^+$], depending on whether or not we require the honest prover to be efficient).

Several papers, beginning with Damgård [Dam1], gave results of a converse nature, culminating in two theorems of Ostrovsky and Wigderson [OW]. The first theorem shows that a zero-knowledge protocol (of any type[4]) for a *hard-on-average* problem implies the existence of one-way functions. The second theorem shows that a zero-knowledge protocol for any problem that cannot be solved in probabilistic polynomial time (BPP) implies a "weak form" of one-way functions. (For problems in BPP, we do not expect to obtain any implication, since every problem in BPP has a trivial zero-knowledge proof in which the prover sends nothing and the verifier decides on its own.) These results suggest that the nontriviality of zero knowledge is equivalent to the existence of one-way functions, which in turn is equivalent to the existence of commitment schemes [HILL,Nao,IL]. However, they are only partial converses to [GMW], and do not provide an exact characterization of the power of zero knowledge. This is because for problems that are neither hard on average nor in BPP, a zero-knowledge protocol only implies the "weak form" of one-way functions in the second result, which seems too weak to construct commitment schemes and thus zero-knowledge protocols. Finally, note that first direction (one-way functions imply that zero knowledge is powerful) seems to say nothing about SZKP: to get an SZKP protocol out of [GMW], one would need a commitment scheme that is both statistically hiding and statistically binding, which is impossible.

The above difficulties no longer seem inherent, however, if one turns away from one-way functions and standard commitments to *instance-dependent* commitments [BMO,IOS]. These are commitment protocols where the sender and receiver both receive an instance x of some promise problem Π as an auxiliary input. We only require the commitment scheme to be hiding when x is a YES instance and binding when x is a NO instance. For example, Graph Isomorphism has a simple instance-dependent commitment scheme: when the auxiliary input is $x = (G_0, G_1)$, the sender commits to a bit $b \in \{0, 1\}$ by sending a random isomorphic copy H of G_b, and reveals b by sending the isomorphism between H

[3] In [GMW], only computational zero-knowledge proof systems were considered. The original construction of statistical zero-knowledge arguments for NP [BCC] used stronger cryptographic primitives than commitment schemes.

[4] The results of [OW] are stated for CZKP, but are easily seen to hold even for the most general class CZKA.

and G_b. This protocol is perfectly hiding when $G_0 \cong G_1$, and perfectly binding when $G_0 \ncong G_1$. It is possible to achieve both perfect hiding and perfect binding because we do not require the properties to hold at the same time.

As shown by Itoh, Ohta, and Shizuya [IOS], this relaxation of commitment schemes remains useful for constructing zero-knowledge protocols, because in many constructions, the hiding property is used for zero knowledge (which is required only when x is a YES instance) and the binding property is used for soundness (which is required ony when x is a NO instance). For example, using [GMW], we see that if a promise problem $\Pi \in$ NP has an instance-dependent commitment scheme, then Π has a zero-knowledge protocol, where the hiding property (statistical or computational) translates to the zero-knowledge property and the binding property translates to the soundness property.

In the last few years, there has been substantial progress on proving the converse: if a problem has a zero-knowledge protocol, then it has an instance-dependent commitment scheme. This progress started with SZKP, where both security properties are statistical.

- It was conjectured in [MV] that every problem in SZKP has an instance-dependent commitment scheme. As a first step, they constructed an instance-dependent commitment scheme for a *restricted version* of Statistical Difference, one of the complete problems for SZKP [SV].
- In [Vad], it was shown that SZKP consists exactly of the problems with instance-dependent commitment schemes in which the sender is computationally unbounded rather than polynomial time. The unbounded sender renders the result useless for the study of zero knowledge with efficient honest provers, which was the motivation of [MV]. But the result was useful for the study of CZKP; see below.
- In [NV], the sender was made efficient, at the price of working with a new variant of commitments, called *1-out-of-2-binding commitments* (denoted as $\binom{2}{1}$-binding). These $\binom{2}{1}$-binding commitments were shown to be sufficient for constructing zero-knowledge proofs for NP, but are otherwise cumbersome and of unclear value as cryptographic primitives on their own.
- In [Vad,OV], the classes involving computational security, namely CZKP, SZKA, and CZKA, were characterized in terms of SZKP and "instance-dependent one-way functions." Thus, combining the above types of instance-dependent commitments for SZKP with constructions of commitments from one-way functions [HILL,Nao,NOV,HR], the classes CZKP, SZKA, and CZKA could be characterized in terms of instance-dependent commitments, but inheriting the deficiencies of [Vad,OV] (namely, an unbounded sender or $\binom{2}{1}$-binding).[5] These instance-dependent commitments played a crucial role in the characterization of the classes CZKP, SZKA, and CZKA, and in proving various unconditional results about these classes (such as equivalence

[5] The proceedings version of [OV] actually quotes the main result (Theorem 1) of this present paper. However, this was done only to simplify the presentation there, and the main results of [OV] were actually obtained prior to Theorem 1.

of honest-verifier and cheating-verifier zero knowledge and closure under union).

– Instance-dependent commitments for a restricted class of zero-knowledge proofs, namely *3-round public-coin* zero-knowledge proofs, were implicit in the works of Damgård [Dam1,Dam2]. Indeed, Kapron, Malka, and Srinivasan [KMS] used Damgård's techniques to show that 3-round public-coin zero-knowledge proofs where the verifier just sends a single random bit — called *V-bit protocols* — are exactly characterized by *noninteractive* instance-dependent commitments.[6]

1.3 Our Results

In this paper, we show that zero knowledge proofs are equivalent to standard instance-dependent commitments, where the sender is efficient and there is no non-standard $\binom{2}{1}$-binding property. The main technical contribution is the construction for SZKP:

Theorem 1. *For every promise problem* Π, $\Pi \in \mathrm{SZKP}$ *if and only if* Π *has an instance-dependent commitment scheme that is statistically hiding on the YES instances and statistically binding on the NO instances. Moreover, every* $\Pi \in \mathrm{SZKP}$ *has an instance-dependent commitment scheme that is public coin and is constant round.*

As mentioned previously, a construction of instance-dependent commitments for SZKP implies ones for the other classes (by their characterizations in terms SZKP and instance-dependent one-way functions [Vad,OV] together with the constructions of commitments from one-way functions [HILL,Nao,NOV,HR]).

Corollary 1. *The following hold for every problem* $\Pi \in \mathrm{NP}$*:*[7]

1. $\Pi \in \mathrm{CZKP}$ *if and only if* Π *has an instance-dependent commitment scheme that is computationally hiding on the YES instances and statistically binding on the NO instances. Moreover, this instance-dependent commitment scheme is public coin and is constant round.*
2. $\Pi \in \mathrm{SZKA}$ *if and only if* Π *has an instance-dependent commitment scheme that is statistically hiding on the YES instances and computationally binding on the NO instances. Moreover, this instance-dependent commitment scheme is public coin.*

[6] *Noninteractive commitments* are commitments where the sender commits to a message in the commit stage by sending a *single* message to the receiver; hence, the receiver does not send any message, both in the commit and reveal stages.

[7] We state the result for problems in NP for simplicity. The direction stating that zero-knowledge implies instance-dependent commitments (which is our main contribution) actually holds without any constraint on Π other than being in the stated zero-knowledge class. The other direction actually generalizes to problems in MA when the honest prover is required to be efficient and IP when the honest prover is allowed to be computationally unbounded.

3. $\Pi \in$ CZKA *if and only if* Π *has an instance-dependent commitment scheme that is computationally hiding on the YES instances and computationally binding on the NO instances. Moreover, this instance-dependent commitment scheme is public coin.*

Note that for the case of *proof* systems (i.e., statistical binding), our instance-dependent commitment schemes are constant round. (For arguments, the polynomial round complexity is inherited from the statistically hiding commitments based on one-way functions [NOV,HR].) This enables us to resolve some open questions regarding the round complexity of zero-knowledge proofs. For example:

Corollary 2. *Every problem in* SZKP *(resp.,* CZKP $\cap$ NP[8]*) has a constant-round, public-coin statistical (resp., computational) zero-knowledge proof system with soundness error* $1/\operatorname{poly}(n)$ *and a black-box simulator.*[9]

It was known how to achieve constant rounds for CZKP under the assumption that one-way functions exist, but it was not known for SZKP under any assumption. Previously, it was only known that SZKP had constant-round *honest-verifier* statistical zero-knowledge proofs, and these were private coin [Oka].

Since SZKP is closed under complement [Oka], we can also obtain instance-dependent commitments in which the security properties are reversed (i.e., statistically binding on YES instances and statistically hiding on NO instances). Such commitments are useful for implementing commitments from the verifier. Using such commitments in the protocol of [GK1] (or, more easily, [Ros]), we obtain:

Corollary 3. *Every problem in* SZKP *has a constant-round (private-coin) zero-knowledge proof system with negligible soundness error.*

Following [MOSV], a potential application of our instance-dependent commitments is to show that every problem in SZKP has a *concurrent* statistical zero-knowledge proof system with $\omega(\log n)$ rounds. However, the analysis of [MOSV] is given only for noninteractive commitments, so it would need to be extended to handle our interactive commitments.

1.4 Overview of Our Techniques

Our proof of Theorem 1 uses techniques from Nguyen and Vadhan [NV] and Haitner and Reingold [HR]. Recall that [NV] constructed instance-dependent $\binom{2}{1}$-binding commitments for SZKP. In the standard, non-instance-dependent setting, [HR] showed how to convert $\binom{2}{1}$-binding commitments into standard

[8] This result actually generalizes to CZKP $\cap$ AM. No further restriction is needed for SZKP because SZKP $\subseteq$ AM [AH].

[9] Using [GMW] would yield a poor soundness error of $1-1/\operatorname{poly}(n)$. To obtain soundness error $1/\operatorname{poly}(n)$, we use an $O(\log n)$-fold parallel repetition of [Blu]. Negligible soundness error cannot be achieved with public coins and black-box simulation for problems outside BPP [GK2].

commitments using *universal one-way hash functions* [NY], whose existence is equivalent to that of one-way functions [Rom,KK]. Thus, we obtain our result by constructing an instance-dependent analogue of universal one-way hash functions for every problem in SZKP, and then applying the Haitner & Reingold transformation.

It is not immediately clear, however, how to define instance-dependent universal one-way hash functions in a way that allows for statistical security (as we need for Theorem 1). The standard definition of a universal one-way hash family is as a family $\mathcal{H}$ of *length-decreasing* functions $h\colon \{0,1\}^n \to \{0,1\}^m$, such that for every fixed $y \in \{0,1\}^n$, if we are given a random $h \xleftarrow{\text{R}} \mathcal{H}$, it is infeasible to find an $y' \neq y$ such that $h(y') = h(y)$. Note that the latter property is necessarily computational. Since the hash functions are length-decreasing, $h(y)$ will have many preimages y' with high probability over a random y and h, and thus an unbounded adversary could find a collision easily. We observe, however, that the Haitner & Reingold transformation [HR] does *not* really require a length-decreasing function. They only use the fact that $h(y)$ typically has many preimages, and they only use this to establish the hiding property of the resulting commitment scheme. For the binding property, they use infeasibility of finding collisions; for statistical security, this amounts to the functions being nearly injective. With these observations, our notion of an instance-dependent universal one-way hash family $\mathcal{H}_x$ is as a family of functions (typically not length decreasing) that also depend on an instance x of some promise problem Π. When x is a YES instance, a random hash function from the family has large preimages with high probability, and when x is a NO instance, a random hash function is nearly injective with high probability. We show that every problem in SZKP has an instance-dependent universal one-way hash family of this type, and thus are able to apply the Haitner & Reingold transformation to the $\binom{2}{1}$-binding commitments of [NV] to obtain our result.

1.5 Organization

In Section 2, we provide definitions to terminologies used in this paper. We prove our main result, Theorem 1, in Section 3. The proof of Corollary 1 can be found in [OV], and the proofs of the other corollaries will appear in the full version of this paper.

2 Preliminaries

If X is a random variable taking values in a finite set $\mathcal{U}$, then we write $x \xleftarrow{\text{R}} X$ to indicate that x is selected according to X. If S is a subset of $\mathcal{U}$, then $x \xleftarrow{\text{R}} S$ means that x is selected according to the uniform distribution on S. We adopt the convention that when the same random variable occurs several times in an expression, they refer to a single sample. For example, $\Pr[f(X) = X]$ is defined to be the probability that when $x \xleftarrow{\text{R}} X$, we have $f(x) = x$. We write U_n to denote the random variable distributed uniformly over $\{0,1\}^n$.

A function $\varepsilon : \mathbb{N} \to [0,1]$ is called *negligible* if $\varepsilon(n) = n^{-\omega(1)}$. We let $\mathrm{neg}(n)$ denote an arbitrary negligible function (i.e., when we say that $f(n) < \mathrm{neg}(n)$ we mean that *there exists* a negligible function $\varepsilon(n)$ such that for every n, $f(n) < \varepsilon(n)$). Likewise, $\mathrm{poly}(n)$ denotes an arbitrary function $f(n) = n^{O(1)}$.

PPT refers to probabilistic algorithms (i.e., Turing machines) that run in *strict* polynomial time. A *nonuniform* PPT algorithm is a pair $(A, \bar{z})$, where $\bar{z} = z_1, z_2, \ldots$ is an infinite sequence of strings where $|z_n| = \mathrm{poly}(n)$, and A is a PPT algorithm that receives pairs of inputs of the form $(x, z_{|x|})$. (The string z_n is the called the *advice string* for A for inputs of length n.) Nonuniform PPT algorithms are equivalent to (nonuniform) families of polynomial-sized Boolean circuits.

Promise problems. Roughly speaking, a *promise problem* [ESY] is a decision problem where some inputs are excluded. Formally, a promise problem is specified by two disjoint sets of strings $\Pi = (\Pi_{\mathrm{Y}}, \Pi_{\mathrm{N}})$, where we call Π_{Y} the set of *YES instances* and Π_{N} the set of *NO instances*. Such a promise problem is associated with the following computational problem: given an input that is "promised" to lie in $\Pi_{\mathrm{Y}} \cup \Pi_{\mathrm{N}}$, decide whether it is in Π_{Y} or in Π_{N}. Note that languages are a special case of promise problems (namely, a language L over alphabet Σ corresponds to the promise problem $(L, \Sigma^* \setminus L)$). Thus working with promise problems makes our results more general. Moreover, even to prove our results just for languages, it turns out to be extremely useful to work with promise problems along the way. All of the complexity classes in this paper are taken to be classes of promise problems. We refer the reader to the recent survey of Goldreich [Gol] for more on the utility and subtleties of promise problems.

2.1 Instance-Dependent Cryptographic Primitives

Instance-dependent functions. It will be very useful for us to work with cryptographic primitives that may depend on an instance x of a problem $\Pi = (\Pi_{\mathrm{Y}}, \Pi_{\mathrm{N}})$, and where the security condition will hold only if x is in some particular set $I \subseteq \{0,1\}^*$. We begin our discussion of instance-dependent primitives with the following definition.

Definition 1. *An* instance-dependent function *is a family* $\mathcal{F} = \{f_x \colon \{0,1\}^{n(|x|)} \to \{0,1\}^{m(|x|)}\}_{x \in \{0,1\}^*}$*, where* $n(\cdot)$ *and* $m(\cdot)$ *are polynomials. We call* $\mathcal{F}$ polynomial-time computable *if there is a deterministic polynomial-time algorithm* F *such that for every* $x \in \{0,1\}^*$ *and* $y \in \{0,1\}^{n(|x|)}$*, we have* $F(x,y) = f_x(y)$*.*

To simplify notation, we often write $f_x \colon \{0,1\}^{n(|x|)} \to \{0,1\}^{m(|x|)}$ to mean the family $\{f_x \colon \{0,1\}^{n(|x|)} \to \{0,1\}^{m(|x|)}\}_{x \in \{0,1\}^*}$.

Indistinguishability of instance-dependent ensembles. The notions of statistical and computational indistinguishability have instance-dependent analogues. But first, we define an instance-dependent analogue of probability ensembles.

Definition 2. *An* instance-dependent probability ensemble *is a collection of random variables* $\{A_x\}_{x\in\{0,1\}^*}$*, where* A_x *takes values in* $\{0,1\}^{p(|x|)}$ *for some polynomial* p*. We call such an ensemble* samplable *if there is a probabilistic polynomial-time algorithm* M *such that for every* x*, the output* $M(x)$ *is distributed according to* A_x*.*

Definition 3. *Two instance-dependent probability ensembles* $\{A_x\}_{x\in\{0,1\}^*}$ *and* $\{B_x\}_{x\in\{0,1\}^*}$ *are* computationally indistinguishable *on* $I \subseteq \{0,1\}^*$ *if for every nonuniform PPT* D*, there exists a negligible function* ε *such that for all* $x \in I$*,*

$$|\Pr[D(x, A_x) = 1] - \Pr[D(x, B_x) = 1]| \leq \varepsilon(|x|) \ .$$

Similarly, we say that $\{A_x\}_{x\in\{0,1\}^*}$ *and* $\{B_x\}_{x\in\{0,1\}^*}$ *are* statistically indistinguishable *on* $I \subseteq \{0,1\}^*$ *if the above is required for all functions* D*, instead of only nonuniform PPT ones. Equivalently,* $\{A_x\}_{x\in\{0,1\}^*}$ *and* $\{B_x\}_{x\in\{0,1\}^*}$ *are statistically indistinguishable on* I *iff* A_x *and* B_x *are have statistical distance at most* $\varepsilon(|x|)$ *for some negligible function* ε *and all* $x \in I$*. We write* $\approx_c$ *and* $\approx_s$ *to denote computational and statistical indistinguishability, respectively.*

Instance-dependent commitments. We give a definition of instance-dependent commitment schemes that extends the standard (that is, non-instance dependent) definition of commitment schemes in a natural way. Note that in our definition below, the reveal stage is *noninteractive* (that is, consisting of a single message from the sender to the receiver). This because in the reveal stage, without loss of generality, we can have the sender provide the receiver the random coin tosses it used in the commit stage, and have the receiver verify consistency.

Definition 4. *An* instance-dependent commitment scheme *is a family of protocols* $\{\mathsf{Com}_x\}_{x\in\{0,1\}^*}$ *with the following properties:*

1. *Scheme* Com_x *proceeds in two stages: a* commit stage *and a* reveal stage. *In both stages, the* sender *and* receiver *receive instance* x *as common input, and hence we denote the sender and receiver as* S_x *and* R_x*, respectively, and write* $\mathsf{Com}_x = (S_x, R_x)$*.*
2. *At the beginning of the commit stage, sender* S_x *receives a private input* $b \in \{0,1\}$*, which denotes the bit that* S *is supposed to commit to. At the end of the commit stage, both sender* S_x *and receiver* R_x *output a* commitment c*.*
3. *In the reveal stage, sender* S_x *sends a pair* (b, d)*, where* d *is the* decommitment *string for bit* b*. Receiver* R_x *accepts or rejects based on* x*,* b*,* d*, and* c*.*
4. *The sender* S_x *and receiver* R_x *algorithms are computable in polynomial time (in* $|x|$*), given* x *as auxiliary input.*
5. *For every* $x \in \{0,1\}^*$*,* R_x *will always accept (with probability 1) if both sender* S_x *and receiver* R_x *follow their prescribed strategy.*

Instance-dependent commitment scheme $\{\mathsf{Com}_x = (S_x, R_x)\}_{x\in\{0,1\}^*}$ *is* public coin *if for every* $x \in \{0,1\}^*$*, all messages sent by* R_x *are independent random coins.*

To simplify notation, we write Com_x or (S_x, R_x) to denote instance-dependent commitment scheme $\{\mathsf{Com}_x = (S_x, R_x)\}_{x\in\{0,1\}^*}$.

The hiding and binding properties of standard commitments extend in a natural way to their instance-dependent analogues.

Definition 5. *Instance-dependent commitment scheme* $\mathsf{Com}_x = (S_x, R_x)$ *is* statistically [resp., computationally] hiding on $I \subseteq \{0,1\}^*$ *if for every [resp., nonuniform PPT]* R^**, the ensembles* $\{\mathrm{view}_{R^*}(S_x(0), R^*)\}_{x\in I}$ *and* $\{\mathrm{view}_{R^*}(S_x(1), R^*)\}_{x\in I}$ *are statistically [resp., computationally] indistinguishable, where random variable* $\mathrm{view}_{R^*}(S_x(b), R^*)$ *denotes the view of* R^* *in the commit stage interacting with* $S_x(b)$*. For a problem* $\Pi = (\Pi_{\mathrm{Y}}, \Pi_{\mathrm{N}})$*, an instance-dependent commitment scheme* Com_x *for* Π *is* statistically [resp., computationally] hiding on the YES instances *if* Com_x *is statistically [resp., computationally] hiding on* Π_{Y}.

Definition 6. *Instance-dependent commitment scheme* $\mathsf{Com}_x = (S_x, R_x)$ *is* statistically [resp., computationally] binding on $I \subseteq \{0,1\}^*$ *if for every [resp., nonuniform PPT]* S^**, there exists a negligible function* ε *such that for all* $x \in I$*, the malicious sender* S^* *succeeds in the following game with probability at most* $\varepsilon(|x|)$.

> S^* *interacts with* R_x *in the commit stage obtaining commitment* c*. Then* S^* *outputs pairs* $(0, d_0)$ *and* $(1, d_1)$*, and* succeeds *if in the reveal stage,* $R_x(0, d_0, c) = R_x(1, d_1, c) = \mathtt{accept}$.

For a problem $\Pi = (\Pi_{\mathrm{Y}}, \Pi_{\mathrm{N}})$*, an instance-dependent commitment scheme* Com_x *for* Π *is* statistically [resp., computationally] binding on the NO instances *if* Com_x *is statistically [resp., computationally] binding on* Π_{N}.

1-out-of-2-binding commitments. A *1-out-of-2-binding commitment scheme*—denoted as $\binom{2}{1}$-binding—is a commitment schemes with two *sequential* and *related* phases such that in each phase, the sender commits to and reveals a value. (They are related in the sense that the protocol for the second phase takes the transcript of the first phase as a common input to both the sender and receiver, and the sender may maintain private state between the two phases.) The hiding property of such commitments is strong: we require that at the end of each commit stage, the receiver has not learned anything about the value to which the sender is committing. The binding property, however, is relatively weak. It only says that it is infeasible for a cheating sender to break the commitment in *both* phases. That is, with high probability over the first commit stage, there is at most one value to which the sender can open that will result in the second phase being non-binding. A formal definition can be found in [NV, Sect. 2] (cf., [Ong, Sect. 3.4.1]).

3 Instance-Dependent Commitments for SZKP

Our goal in this section is to prove Theorem 1. We begin by recalling the result of Nguyen and Vadhan [NV], which is the starting point for our work. Their construction started off from the SZKP-complete problem Entropy Difference [GV], $\mathrm{ED} = (\mathrm{ED_Y}, \mathrm{ED_N})$, defined as:

$$\mathrm{ED_Y} = \{(X,Y) : \mathrm{H}(X) \geq \mathrm{H}(Y) + 1\};$$
$$\mathrm{ED_N} = \{(X,Y) : \mathrm{H}(X) \leq \mathrm{H}(Y) - 1\},$$

where X and Y are random variables specified by circuits that same from them (by evaluating the circuit on a uniformly random input), and $\mathrm{H}(\cdot)$ denotes the *(Shannon) entropy*, i.e., $\mathrm{H}(Z) = \mathrm{E}_{z \xleftarrow{\mathrm{R}} Z}[\log(1/\Pr[Z = z])]$. We assume, without loss of generality, that the size of the circuits X and Y are upper bounded by the square of their respective input lengths. (This can be guaranteed by padding dummy input variables to circuits.)

The [NV] construction of instance-dependent schemes for ED does not provide a commitment scheme with a standard binding property, but rather with the weaker $\binom{2}{1}$ *binding* property (cf., Sect. 2.1). These commitments, even though with a weaker binding property, suffice for getting efficient-prover statistical zero-knowledge proofs for all of SZKP $\cap$ NP [NV].

Our construction of instance-dependent commitments for all of SZKP will follow the same approach as [NV], except at the place where they get stuck with $\binom{2}{1}$-binding commitments, we convert them into commitments with the standard binding property using the ideas of Haitner and Reingold [HR]. Specifically, we use an instance-dependent variant of the Haitner & Reingold transformation to convert $\binom{2}{1}$-binding commitments into commitments with the standard binding property.

The commitments of [NV] were not constructed directly from ED, but instead utilized a Cook reduction from ED to a *restricted version* of the Entropy Approximation [GSV] problem, denoted as $\mathrm{EA'} = (\mathrm{EA'_Y}, \mathrm{EA'_N})$, and defined below:[10]

$$\mathrm{EA'_Y} = \{(X,t) : \mathrm{H}(X) \geq t + 1, \text{ and } |X| \leq n^2\};$$
$$\mathrm{EA'_N} = \{(X,t) : t - 1/n^{14} \leq \mathrm{H}(X) \leq t, \text{ and } |X| \leq n^2\}.$$

Here n denotes the number of input gates to the circuit encoding X, and $|X|$ is the size of that circuit. The condition $|X| \leq n^2$ simply allows us to use n as the security parameter, even though the security properties of instance-dependent commitment schemes are defined in terms of the size of the instance (X,t).

The problem EA' is considered a restricted version of Entropy Approximation because (unrestricted version of) the Entropy Approximation problem $\mathrm{EA} = (\mathrm{EA_Y}, \mathrm{EA_N})$ does not lower-bound the entropy in the case of the NO

[10] The definition of EA' in [NV] has an additional 'security parameter' k, which is eventually set to $\max\{n^{14}, |X|\}$. For convenience, we have restricted to the case that $|X| \leq n^2$; this is without loss of generality for ED and is preserved in the reduction from ED to EA' in Proposition 1 below. Under this restriction, we can simply set $k = n^14$, resulting in our definition of EA'.

instances. EA is defined as follows:

$$\mathrm{EA_Y} = \{(X,t) : \mathrm{H}(X) \geq t+1\};$$
$$\mathrm{EA_N} = \{(X,t) : \mathrm{H}(X) \leq t\}.$$

For instances of EA, we will assume, without loss of generality, that the size of the circuit X is upper bounded by the square of its input length (similar to what we assumed for instances of ED).

The Cook reduction from ED to EA' is established by the following proposition.

Proposition 1. *(Cook Reduction from* ED *to* EA'*; from [NV, Lem. 4.9], which builds on [GSV].) Let* (X,Y) *be an instance of the* ENTROPY DIFFERENCE *problem* $\mathrm{ED} = (\mathrm{ED_Y}, \mathrm{ED_N})$*, where the circuits encoding the random variables* X *and* Y *both have input length* n *and are of size at most* n^2 *(wlog). The Cook reduction from* ED *to* EA' *is as follows:*

$$(X,Y) \in \mathrm{ED_Y} \Rightarrow \bigvee_{i=0}^{n\cdot k} \left((Y, i/k) \in \overline{\mathrm{EA'}}_{\mathrm{Y}} \wedge \bigwedge_{j=0}^{i} (X, j/k) \in \mathrm{EA'_Y} \right) \quad ;$$
$$(X,Y) \in \mathrm{ED_N} \Rightarrow \bigwedge_{i=0}^{n\cdot k} \left((Y, i/k) \in \overline{\mathrm{EA'}}_{\mathrm{N}} \vee \bigvee_{j=0}^{i} (X, j/k) \in \mathrm{EA'_N} \right) \quad ,$$

where $k = n^{14}$.

Note that the reduction from ED to EA' in the above proposition does not alter the circuits; hence, the size of the circuits in both problems remain upper bounded by the square of their respective input lengths, which is what we require.

Using Proposition 1, [NV] noted that it suffices to construct instance-dependent commitments for both EA' and its complement $\overline{\mathrm{EA'}}$ in order to obtain instance-dependent commitments for ED and hence all of SZKP. We capture that observation in the following lemma.

Lemma 1. *Suppose that both the special case of the* ENTROPY APPROXIMATION *problem* EA' *and its complement* $\overline{\mathrm{EA'}}$ *have instance-dependent commitments. That is,*

- *there exist instance-dependent commitments that are statistically hiding on instances in* $\mathrm{EA'_Y}$ *and statistically binding on instances in* $\mathrm{EA'_N}$*, and*
- *there exist instance-dependent commitments that are statistically hiding on instances in* $\mathrm{EA'_N}$ *and statistically binding on instances in* $\mathrm{EA'_Y}$*.*

Then the ENTROPY DIFFERENCE *problem* ED *(and hence, every problem in* SZKP*) has an instance-dependent commitment scheme that is* statistically hiding *on the YES instances and* statistically binding *on the NO instances.*

Indeed, [NV] constructed instance-dependent schemes for both EA' and $\overline{\text{EA'}}$. Their scheme for EA' is a standard instance-dependent commitment scheme, but for $\overline{\text{EA'}}$, they only managed to only get a weaker $\binom{2}{1}$-binding commitment scheme (cf., Sect. 2.1).

Lemma 2. *(From [NV, Thm. 4.4].) The problem* EA' *has an instance-dependent commitment scheme that is statistically hiding on YES instances (namely, instances in* EA'_Y*) and statistically binding on NO instances (namely, instances in* EA'_N*). Moreover, this scheme is public coin and constant round.*

Lemma 3. *(From [NV, Thm. 4.5].) The problem* $\overline{\text{EA'}}$*has an instance-dependent 2-phase commitment scheme that is statistically hiding on the YES instances (namely, instances in* EA'_N*) and statistically* $\binom{2}{1}$ *binding on NO instances (namely, instances in* EA'_Y*). Moreover, this scheme is public coin and constant round.*

3.1 The Haitner & Reingold Transformation

To obtain instance-dependent commitments (with the standard binding property) for $\overline{\text{EA'}}$, we use an instance-dependent variant of the Haitner & Reingold transformation [HR], which we informally describe now. (A detailed description can be found in [HNO$^+$, Sect. 7].)

Overview of [HR]. The $\binom{2}{1}$ binding property of 2-phase commitment schemes states that it is infeasible for an adversarial sender S^* to break both phases of the commitment, but nonetheless it might be possible for S^* to break one of the two phases of its choice. With this in mind, suppose that after the first commitment phase, receiver R flips a coin $phase \leftarrow \{1, 2\}$. If $phase = 1$, the first commitment phase is used to do the commitment. On the other hand, if $phase = 2$, the second commitment phase is used to do the commitment (this is done by S^* revealing its first-phase commitment, and then proceeding to the second phase with R). Intuitively, this would make the scheme binding (with probability 1/2) if S^* chooses which of the two phases it wants to break in advance. The problem, however, is that S^* could choose the phase that it wants to break after seeing the value of *phase*.

A way to overcome this problem is to force the adversary S^* to decide which of the two phases it wants to break before seeing the value of *phase*. Haitner and Reingold [HR] achieved this by having S^* send back a value $y = f(\sigma)$ before the value of *phase* is announced by the receiver R, where σ is the message committed to by S^* in the first phase, and f is a random hash function from a universal one-way hash family. A *universal one-way hash family* [NY] is a family of length-decreasing hash functions such that it is hard to find collisions with any particular value of x specified in advance. In other words, for a value of σ announced before a random hash function f is selected from that family, any efficient algorithm will not be able to find another σ' such that $f(\sigma') = f(\sigma)$. This property of a universal one-way hash family is termed *target collision resistance* by Bellare and Rogaway [BR].

We first argue the hiding property of this new scheme. Before y is sent, the value of σ, the message committed in the first phase, is hidden. If hash function f is *compressing enough*, then the value of $y = f(\sigma)$ leaks at most a few bits of *information* about σ, so the *entropy* of σ given y is still large. This means that we can apply a pairwise-independent hash on σ to get an almost uniform value (by the Leftover Hash Lemma [HILL]). Thus, this new scheme is hiding when $phase = 1$. When $phase = 2$, the sender reveals σ and proceeds on to the second phase, which is used for the commitment. In this case, the hiding property of this new scheme follows from the hiding property of the second commitment phase.

Next, we argue the binding property of this new scheme by making the following observation: the $\binom{2}{1}$ binding property says that after the first commitment phase, there exists at most one value of σ^* that allows an adversarial sender S^* to cheat in the second phase. In other words, if S^* reveals to a value other than σ^*, the second phase will be binding.

When it is the sender's turn to send y, after receiving a random hash function f from receiver R, sender S^* could decide to either send $y = f(\sigma^*)$ or send $y \neq f(\sigma^*)$. If it decides to send $y = f(\sigma^*)$, and if R selects $phase = 1$ following that, then S^* is bound to a single value, since to decommit to two different values it will have to reveal a $\sigma' \neq \sigma^*$ with $f(\sigma') = y = f(\sigma^*)$, and this is infeasible by the target collision resistance property of f. (The value of σ^* is determined by the first-phase commitment, which is completed before a random f is selected.) Instead if it decides to send $y \neq f(\sigma^*)$, and if R selects $phase = 2$ following that, then S^* will have to reveal to a value other than σ^* for its first-phase commitment. In this case, the commitments are done in the second phase, and by the $\binom{2}{1}$ binding property, this phase is guaranteed to be binding. Since the value of *phase* is independent of y, both cases happen with probability 1/2, which would make our scheme binding with probability close to 1/2.

3.2 Instance-Dependent UOWHFs

Our approach to construction standard instance-dependent commitments for $\overline{\text{EA'}}$ and hence all of SZKP is to carry out an instance-dependent analogue of the Haitner & Reingold transformation. To do this, we want to construct an instance-dependent analogue of universal one-way hash functions $\overline{\text{EA'}}$ and apply it to the $\binom{2}{1}$-binding commitments of Lemma 3. Since we want instance-dependent commitments with statistical security (and are not making any complexity assumptions), we need to formulate the properties of universal one-way hash functions in a way that allows for statistical security. The properties used in the Haitner & Reingold transformation are that the functions should be compressing (used for hiding) and target collision-resistant (used for binding). Thus the first attempt would be to require that our instance-dependent universal one-way hash functions are compressing on YES instances and statistically target collision-resistant on NO instances. However, it seems unlikely that this is possible. Indeed, it would imply that $\overline{\text{EA'}}$ is in BPP: statistical target collision resistance implies that the functions are not compressing, so we could distinguish

YES and NO instances simply be checking whether the functions are compressing or not. To get around this difficulty, we observe that the hiding analysis sketched above only requires that σ retains a lot of entropy given $y = f(\sigma)$. This property can hold for non-compressing functions; it simply says that f has large preimage sizes.

This leads to the following definition.

Definition 7. *Problem* $\Pi = (\Pi_{\mathrm{Y}}, \Pi_{\mathrm{N}})$ *has an* instance-dependent universal one-way hash family *if there exists a polynomial-time computable family* $\mathcal{F} = \bigcup_x \mathcal{F}_x = \{f \colon \{0,1\}^{n(|x|)} \to \{0,1\}^{m(|x|)}\}$*, where* $n(\cdot)$ *and* $m(\cdot)$ *are polynomials, such that the following two conditions hold.*

- *The family* $\mathcal{F}_{\mathrm{Y}} = \bigcup_{x \in \Pi_{\mathrm{Y}}} \mathcal{F}_x$ *has the* large preimages property*: there exists a function* $\alpha(\cdot) = \omega(1)$ *and a negligible function* ε*, such that the following holds for all* $x \in \Pi_{\mathrm{Y}}$ *and every function* $f \in \mathcal{F}_x$*:*

$$\Pr_{y \leftarrow \{0,1\}^{n(|x|)}} \left[\left| f^{-1}(f(y)) \right| \geq |x|^{\alpha(|x|)} \right] \geq 1 - \varepsilon(|x|) \ .$$

- *The family* $\mathcal{F}_{\mathrm{N}} = \bigcup_{x \in \Pi_{\mathrm{N}}} \mathcal{F}_x$ *has* statistical target collision resistance*: there exists a negligible function* ε *such that for every* A*, the following holds for all* $x \in \Pi_{\mathrm{Y}}$ *and every* $y \in \{0,1\}^{n(|x|)}$*:*

$$\Pr_{f \leftarrow \mathcal{F}_x} \left[|f^{-1}(f(y))| = 1 \right] \geq 1 - \varepsilon(|x|) \ .$$

Following the discussion above, this definition allows $m(|x|) > n(|x|)$, and only insist that the family has the large preimages property on the YES instances. In fact, our construction of an instance-dependent universal one-way hash family for $\overline{\text{EA'}}$ will be such that $m(|x|)$ is much larger than $n(|x|)$.

With the above definition, we have the following instance-dependent analogue of the Haitner & Reingold transformation, obtained as corollary of Theorem 7.20 in [HNO+]:

Proposition 2. *(Corollary of [HNO+, Thm. 7.20].) Let* $\Pi = (\Pi_{\mathrm{Y}}, \Pi_{\mathrm{N}})$ *be a promise problem, and suppose that the following two conditions hold:*

- *there exists an instance-dependent universal one-way hash family* $\mathcal{F} = \bigcup_x \mathcal{F}_x$ *for* Π*, and*
- *there is an instance-dependent 2-phase commitment scheme* $(\mathsf{S}_x, \mathsf{R}_x)$ *for* Π *that is statistically hiding on the YES instances, and statistically* $\binom{2}{1}$ *binding on NO instances.*

Then, there is an instance-dependent commitment scheme (S_x, R_x) *for* Π *that is statistically hiding on the YES instances, and statistically binding on NO instances. Moreover,* (S_x, R_x) *is public coin if* $(\mathsf{S}_x, \mathsf{R}_x)$ *is.*

Based on the above proposition, it suffices to construct an instance-dependent universal one-way hash family for $\overline{\text{EA'}}$ in order to get instance-dependent commitments for $\overline{\text{EA'}}$.

3.3 Instance-dependent UOWHF for $\overline{\mathrm{EA}}$.

Although we just need an instance-dependent universal one-way hash family for $\overline{\mathrm{EA}}$', we will construct one for the slightly more general problem of $\overline{\mathrm{EA}}$.

Working directly with $\overline{\mathrm{EA}}$ on an instance (X, t) is difficult since we do not know any structure of the random variable X other than its entropy bound. To get more structure, we *flatten* X by taking multiple independent samples of it and outputting all of them. Let X' denote this new random variable. Doing this makes the probability masses of X' concentrated around $2^{-\mathrm{H}(X')}$, and this is why we call it flattening the random variable. (This is also known as the Asymptotic Equipartition Property in the information theory literature; see [CT].) Following [GV], we give a quantitative definition of *flatness* as follows:

Definition 8. *Random variable X is Θ-*flat *if for every $r \geq 1$,*

$$\Pr_{x \leftarrow X}\left[2^{-r\cdot\Theta} < \frac{\Pr[X = x]}{2^{\mathrm{H}(X)}} < 2^{r\cdot\Theta}\right] > 1 - 2^{-r^2} \quad .$$

Consider the flattened version of the ENTROPY APPROXIMATION problem, denoted as $\textsc{FlatEA} = (\textsc{FlatEA}_{\mathrm{Y}}, \textsc{FlatEA}_{\mathrm{N}})$, and defined as follows:

$$\textsc{FlatEA}_{\mathrm{Y}} = \{(X, t) : \mathrm{H}(X) \geq t + n^{14/15},\ X \text{ is } n^{8/15}\text{-flat, and } |X| \leq n^2\}$$
$$\textsc{FlatEA}_{\mathrm{N}} = \{(X, t) : \mathrm{H}(X) \leq t,\ X \text{ is } n^{8/15}\text{-flat, and } |X| \leq n^2\}$$

Here n denotes the number of input gates to the circuit encoding X, and $|X|$ is the size of that circuit. Recall that the condition $|X| \leq n^2$ simply allows us to use n as the security parameter, even though the security properties of instance-dependent commitment schemes are defined in terms of the size of the instance (X, t). Note that the entropy gap between the two cases is close to being linear in n, whereas the deviation from flatness is close to being $\sqrt{n}$.

It is clear that FLATEA is polynomial time reducible to EA, and the reverse reduction from EA to FLATEA is achieved by taking many (e.g. n^{28}) independent copies of X (cf., the Flattening Lemma of [GV, Lem. 3.5]). Hence, constructing an instance-dependent universal one-way hash family for $\overline{\mathrm{EA}}$ is equivalent to constructing one for $\overline{\textsc{FlatEA}}$, and we do this next.

Theorem 2. *The complement of the flattened version of the* ENTROPY APPROXIMATION *problem, namely $\overline{\textsc{FlatEA}}$ has an instance-dependent universal one-way hash family.*

In the remaining of this section, we abuse notation by using $X : \{0,1\}^n \to \{0,1\}^m$ to denote the circuit that samples random variable X.

Proof Idea of Theorem 2

For the problem FLATEA, we will need to construct an instance-dependent (family of) functions that have statistical target collision resistance on the YES instances and large preimages property on the NO instances. These are reversed properties because we want to prove that the complement $\overline{\text{FLATEA}}$ has an instance-dependent universal one-way hash family.

For the YES instances of FLATEA, X has entropy at least $t + \gamma$, where $\gamma = n^{14/15}$. Since X is a *nearly flat* random variable, most of its preimages are small, i.e., their sizes are $\lesssim 2^{n-t-\gamma}$. So with high probability over a random $y \leftarrow \{0,1\}^n$, the preimage size of $X(y)$ is $\lesssim 2^{n-t-\gamma}$. By applying a pairwise-independent hash $h\colon \{0,1\}^n \to \{0,1\}^\beta$ to y, for $\beta \gtrsim n-t-\gamma$, it would make the function $g_h(y) = (X(y), h(y))$ *almost injective*, in that for almost every element in the range has a unique preimage. (An injective function is, by definition, collision resistant.)

The adversary, however, need not choose y uniformly at random; in particular, it could choose an element y such that $X^{-1}(X(y))$ is large, making $f(y) = (X(y), h(y))$ no longer injective. To prevent the adversary from gaining, we add a *shift* $s \in \{0,1\}^n$ to the circuit X. Specifically, let the new function be $f_{s,h}(y) = (X(y \oplus s), h(y))$. Since y is now randomly shifted by s, the preimage size of $X(y \oplus s)$ is small with high probability over a random $s \leftarrow \{0,1\}^n$. Thus, we can conclude that $f_{s,h}(y)$ is *almost injective* even for an adversarially chosen y. This will give us the desired target collision resistance property for $\beta \gtrsim n - t - \gamma$.

For the NO instances of FLATEA, X has entropy at most t. Since X is a *nearly-flat* random variable, most of its preimages are large, i.e., their sizes are $\gtrsim 2^{n-t}$. Restricting to a hash $h\colon \{0,1\}^n \to \{0,1\}^\beta$ will shrink the size of the preimages by a factor of approximately $2^{-\beta}$. So if $\beta \lesssim n - t$, the size of the preimages will still be large enough to satisfy the large preimages property.

The fact that the entropy gap $\gamma = n^{14/15}$ between the YES and NO instances is much greater than the deviation $\Theta = n^{8/15}$ from flatness is what allows us to find an appropriate value of β between $n - t - \gamma$ and $n - t$ that satisfies both cases. A complete proof will be given in the full version of this paper.

Acknowledgments

We are grateful to Iftach Haitner and Omer Reingold for a stimulating collaboration during the time that the results of this paper were obtained.

References

[AH] William Aiello and Johan Håstad. Statistical zero-knowledge languages can be recognized in two rounds. *J. Comput. Syst. Sci.*, 42(3):327–345, 1991.

[BCC] Gilles Brassard, David Chaum, and Claude Crépeau. Minimum disclosure proofs of knowledge. *J. Comput. Syst. Sci.*, 37(2):156–189, 1988.

[BGG+] Michael Ben-Or, Oded Goldreich, Shafi Goldwasser, Johan Håstad, Joe Kilian, Silvio Micali, and Phillip Rogaway. Everything provable is provable in zero-knowledge. In *Proc. CRYPTO '88*, pages 37–56, 1988.

[Blu] Manuel Blum. How to prove a theorem so no one else can claim it. In *Proc. International Congress of Mathematicians*, pages 1444–1451, 1987.

[BMO] Mihir Bellare, Silvio Micali, and Rafail Ostrovsky. Perfect zero-knowledge in constant rounds. In *Proc. 22nd STOC*, pages 482–493, 1990.

[BR] Mihir Bellare and Phillip Rogaway. Collision-resistant hashing: towards making UOWHFs practical. In *Proc. CRYPTO '97*, pages 470–484, 1997.

[CT] Thomas M. Cover and Joy A. Thomas. *Elements of information theory*. Wiley-Interscience, New York, NY, USA, second edition, 2006.

[Dam1] Ivan Damgård. On the existence of bit commitment schemes and zero-knowledge proofs. In *Proc. CRYPTO '89*, pages 17–27, 1989.

[Dam2] Ivan B. Damgård. Interactive hashing can simplify zero-knowledge protocol design without computational assumptions. In *Proc. CRYPTO '93*, pages 100–109, 1993.

[ESY] Shimon Even, Alan L. Selman, and Yacov Yacobi. The complexity of promise problems with applications to public-key cryptography. *Inform. Control*, 61(2):159–173, 1984.

[GK1] Oded Goldreich and Ariel Kahan. How to construct constant-round zero-knowledge proof systems for NP. *J. Cryptol.*, 9(3):167–190, 1996.

[GK2] Oded Goldreich and Hugo Krawczyk. On the composition of zero-knowledge proof systems. *SIAM Journal on Computing*, 25(1):169–192, 1996.

[GMR] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, 1989.

[GMW] Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. *J. ACM*, 38(1):691–729, 1991.

[Gol] Oded Goldreich. On promise problems (a survey in memory of Shimon Even [1935-2004]). Technical Report TR05–018, Electronic Colloquium on Computational Complexity, February 2005.

[GSV] Oded Goldreich, Amit Sahai, and Salil Vadhan. Can statistical zero-knowledge be made non-interactive?, or On the relationship of SZK and NISZK. In *Proc. CRYPTO '99*, pages 467–484, 1999.

[GV] Oded Goldreich and Salil P. Vadhan. Comparing entropies in statistical zero knowledge with applications to the structure of SZK. In *Proc. 14th Computational Complexity*, pages 54–73, 1999.

[HILL] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28(4):1364–1396, 1999.

[HNO+] Iftach Haitner, Minh-Huyen Nguyen, Shien Jin Ong, Omer Reingold, and Salil Vadhan. Statistically hiding commitments and statistical zero-knowledge arguments from any one-way function. Preliminary versions appeared as [NOV] and [HR]. Available at `http://eecs.harvard.edu/~salil/papers/SHcommit-abs.html`. In submission, 2007.

[HR] Iftach Haitner and Omer Reingold. Statistically-hiding commitment from any one-way function. In *Proc. 39th STOC*, pages 1–10, 2007.

[IL] Russell Impagliazzo and Michael Luby. One-way functions are essential for complexity based cryptography. In *Proc. 30th FOCS*, pages 230–235, 1989.

[IOS] Toshiya Itoh, Yuji Ohta, and Hiroki Shizuya. A language-dependent cryptographic primitive. *J. Cryptol.*, 10(1):37–49, 1997.

[IY] Russell Impagliazzo and Moti Yung. Direct minimum-knowledge computations (extended abstract). In *Proc. CRYPTO '87*, volume 293, pages 40–51, 1987.

[KK] Jonathan Katz and Chiu-Yuen Koo. On constructing universal one-way hash functions from arbitrary one-way functions. Technical Report 2005/328, Cryptology ePrint Archive, 2005.

[KMS] Bruce Kapron, Lior Malka, and Venkatesh Srinivasan. A characterization of non-interactive instance-dependent commitment-schemes (NIC). In *Proc. ICALP 2007*, pages 328–339, 2007.

[MOSV] Daniele Micciancio, Shien Jin Ong, Amit Sahai, and Salil Vadhan. Concurrent zero knowledge without complexity assumptions. In *Proc. TCC 2004*, pages 1–20, 2006.

[MV] Daniele Micciancio and Salil Vadhan. Statistical zero-knowledge proofs with efficient provers: lattice problems and more. In *Proc. CRYPTO 2003*, pages 282–298, 2003.

[Nao] Moni Naor. Bit commitment using pseudorandomness. *J. Cryptol.*, 4(2):151–158, 1991.

[NOV] Minh-Huyen Nguyen, Shien Jin Ong, and Salil Vadhan. Statistical zero-knowledge arguments for NP from any one-way function. In *Proc. 47th FOCS*, pages 3–14, 2006.

[NV] Minh-Huyen Nguyen and Salil Vadhan. Zero knowledge with efficient provers. In *Proc. 38th STOC*, pages 287–295, 2006.

[NY] Moni Naor and Moti Yung. Universal one-way hash functions and their cryptographic applications. In *Proc. 21st STOC*, pages 33–43, 1989.

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

[Ong] Shien Jin Ong. *Unconditional Relationships within Zero Knowledge*. PhD thesis, Harvard University, Cambridge, MA, USA, May 2007.

[OV] Shien Jin Ong and Salil Vadhan. Zero knowledge and soundness are symmetric. In *Proc. EUROCRYPT 2007*, pages 187–209, 2007. Earlier version appeared as TR06-139 in the Electronic Colloquium on Computational Complexity.

[OW] Rafail Ostrovsky and Avi Wigderson. One-way functions are essential for non-trivial zero-knowledge. In *Proceedings of the 2nd Israel Symposium on Theory of Computing Systems*, pages 3–17. IEEE Computer Society, 1993.

[Rom] John Rompel. One-way functions are necessary and sufficient for secure signatures. In *Proc. 22nd STOC*, pages 387–394, 1990.

[Ros] Alon Rosen. A note on constant-round zero-knowledge proofs for NP. In *Proc. TCC 2004*, pages 191–202, 2004.

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

[Vad] Salil P. Vadhan. An unconditional study of computational zero knowledge. *SIAM J. Comput.*, 36(4):1160–1214, 2006.