

One-Way Permutations, Interactive Hashing and Statistically Hiding Commitments

Hoeteck Wee*

Computer Science Division,
University of California,
Berkeley
hoeteck@cs.berkeley.edu

Abstract. We present a lower bound on the round complexity of a natural class of black-box constructions of statistically hiding commitments from one-way permutations. This implies a $\Omega(\frac{n}{\log n})$ lower bound on the round complexity of a computational form of interactive hashing, which has been used to construct statistically hiding commitments (and related primitives) from various classes of one-way functions, starting with the work of Naor, Ostrovsky, Venkatesan and Yung (J. Cryptology, 1998). Our lower bound matches the round complexity of the protocol studied by Naor et al.

Keywords: Interactive hashing, statistically hiding commitments, black-box lower bounds.

1 Introduction

A *zero-knowledge proof* is a protocol wherein one party, the prover, convinces another party, the verifier, of the validity of an assertion while revealing no additional knowledge. Introduced by Goldwasser, Micali and Rackoff in the 1980s [GMR89], zero-knowledge proofs have played a central role in the design and study of cryptographic protocols. In these applications, it is important to construct constant-round zero-knowledge protocols for NP under minimal assumptions. In many cases, a computational zero-knowledge argument system suffices, and we know how to construct such protocols for NP under the (essentially) minimal assumption of one-way functions [BJY97, OW93]. On the other hand, there are cases wherein we need stronger guarantees, namely a computational zero-knowledge proof system, or a statistical zero-knowledge argument system.¹ Surprisingly, the main bottleneck to reducing the assumptions for known constructions of both constant-round computational

* Work done while visiting Tsinghua University, Beijing; IBM T.J. Watson Research Center; and IPAM, Los Angeles.

¹ It is unlikely that every language in NP has a statistical zero-knowledge proof system [F89, AH91, BHZ87].

zero-knowledge proof systems and statistical zero-knowledge argument systems [BCY91, GK96a] is statistically hiding commitments.²

We know how to construct constant-round statistically-hiding commitments from collision-resistant hash functions [DPP98, NY89] and from claw-free permutations [GK96a]. In 1992, Naor, Ostrovsky, Venkatesan and Yung [NOVY98] showed that one-way permutations are sufficient for statistically hiding commitments wherein the round complexity is linear in the security parameter. This was very recently extended to one-way functions by Haitner and Reingold [HR06b]. Both works use the powerful tool of interactive hashing [OVY93], a 2-party protocol for choosing a small set of strings, with binding and hiding requirements similar to those in commitment schemes. An intriguing open problem (posed in [NOVY98] and reiterated in [DHRS04, KS06, HR06a]) is whether some variant of interactive hashing could yield a constant-round statistically hiding commitment from one-way permutations. In fact, even a $n^{o(1)}$ -round commitment would be interesting. The restriction to interactive hashing may seem limiting, but it is the only technique that we presently know of. Moreover, Ding, et al. [DHRS04] exhibited a constant-round interactive hashing protocol satisfying a weaker binding guarantee, which indicates that interactive hashing may not be the bottleneck.

1.1 Our Contributions and Techniques

We study a natural class of black-box constructions of statistically hiding commitments from one-way permutations that include several generalizations of the NOVY construction, and show that any such construction yields a commitment scheme with at least $\Omega(n/\log n)$ rounds. This matches the round complexity of a variant of the main NOVY construction ([KS06, HR06a]). Specifically, our lower bound holds for constructions wherein the sender (in the commitment scheme) evaluates the one-way permutation only at the start of the commit phase, and does so on independent random inputs. The sender then uses the output values, her private input to the commitment scheme, and possibly additional randomness in the rest of the commit phase and does not use the inputs to the one-way permutation until the reveal phase.

We derive as a corollary, a $\Omega(n/\log n)$ lower bound on a computational form of interactive hashing presented in [NOV06, HR06a], based on an abstraction of the way interactive hashing is used in the NOVY construction and the subsequent works of Haitner et al. [HHK⁺05, NOV06, HR06b]. The same abstraction also applies to the use of interactive hashing in the transformation of honest-verifier zero-knowledge arguments into cheating-verifier zero-knowledge arguments [D93, OY93]. The lower bound tells us that we need to avoid the standard notion of interactive hashing if we want round-efficient versions of these applications.

² It is not surprising that we need statistically hiding commitments for statistical zero-knowledge arguments; what is surprising is that the only known approach for constructing constant-round zero-knowledge proof systems [GK96a] requires statistically hiding commitments to guarantee soundness, because the verifier begins by committing to her challenges.

Our lower bound for statistically hiding commitments only holds for fully black-box reductions [RTV04], namely, we require not only that the construction treats the one-way permutation as a black-box, but also that the reduction in the proof of security uses black-box access to a cheating sender that breaks the binding property to invert the permutation with noticeable probability. At a high level, our lower bound follows the paradigm of Gennaro and Trevisan [GT00] for proving lower bounds on efficiency of black-box cryptographic constructions, which is in turn based on the Impagliazzo-Rudich framework [IR89] for separating cryptographic primitives. The proof techniques and ideas are otherwise largely inspired by lower bounds for black-box zero-knowledge from the work of Goldreich and Krawczyk [GK96b].

Roughly speaking, a fully black-box reduction guarantees an efficient procedure that by interacting and rewinding the cheating sender, produces transcripts of the commitment scheme with a certain outcome. Using the repeated sampling technique from [IR89], we can ensure that the probability that a partial transcript has the outcome is exponentially small in the length of the sender's last message. This means that the sender sends $O(\log n)$ bits in each round of protocol. On the other hand, the sender must send a total of $\Omega(n)$ bits in the protocol (so that there is a different transcript for every possible challenge for the one-way permutation), which means the protocol must have $\Omega(n/\log n)$ rounds. This simplified and slightly inaccurate sketch overlooks several technical difficulties.

1.2 Perspective

Notions and limitations of interactive hashing. The last few years has witnessed a lot of work on the use of interactive hashing protocols in cryptography with two main notions of security: computationally binding, and binding for static sets [NOV06]. The latter is used in building and studying oblivious transfer protocols in the bounded storage model and over noisy channels [CCM98, DHRS04, CS06], in constructing variants of statistically binding commitments [NV06], and in transforming honest-verifier zero-knowledge proofs into cheating-verifier zero-knowledge proofs [D93, DGOW95, GSV98]. It was noted in [NOV06, CCM98] that the computational binding implies binding for static sets; our lower bound implies that the converse is not true. Specifically, the constant-round protocol of [DHRS04] does not satisfy the computational formulation (which answers an open problem in [DHRS04] in the negative).

Efficiency of cryptographic reductions. Previous work establishing lower bounds for efficiency of black-box cryptographic reductions has focused on the query complexity and randomness complexity of these reductions [KST99, GGKT05, LTW05, HK05] whereas our work focuses on round complexity. Upon closer inspection, our work is also qualitatively very different (apart from studying a different computational resource) as the works of [GGKT05, LTW05, HK05] rule out weakly black-box reductions (unless $P = NP$), wherein the proof of security may exploit the code of the adversary (in a non-black-box manner). As mentioned earlier, our main result only rules out fully

black-box reductions and uses fairly different techniques. We stress that all known reductions between cryptographic primitives - with the exception of the non-black-box techniques used in zero-knowledge and multi-party protocols, e.g. [B01], but including the non-black-box constructions in [AIK04] - do not exploit the code of the adversary in the proof of security. As such, ruling out fully black-box reductions is almost as meaningful as ruling weakly black-box reductions.

Information-theoretic analogues. Many black-box cryptographic constructions apart from interactive hashing-based commitments have an information-theoretic analogue which is easier to achieve, in that it does not have some kind of “simulateable” requirement, namely, an efficient procedure for simulating random transcripts with a certain outcome. This was articulated in [DGW95], using random selection as a case study. Such connections been exploited in both directions, the most remarkable being the construction of extractors from Nisan-Wigderson pseudo-random generators [T01]. In [LTW05], the connection between hardness amplification and combinatorial hitters was used to derive lower bounds on query and randomness complexity of the former. While the resulting lower bounds on query complexity are tight, those for randomness complexity are far from the best-known constructions. The information-theoretic analogue for computational interactive hashing would be interactive hashing with binding for static sets, for which we cannot expect to prove a super-constant lower bound (again, due to the constant-round protocol in [DHRS04]). Indeed, we exploit the “simulateable” requirement for our main result.

Trade-offs between interaction and assumptions. The NOVY construction demonstrated the feasibility of trading off higher interaction costs in order to build a cryptographic primitive under weaker cryptographic assumptions (specifically, perfectly hiding commitments with a linear number of rounds assuming one-way permutations, versus a constant number of rounds assuming claw-free permutations). Rudich’s work [R91] shows that this trade-off is necessary in relation to secret key agreement and trapdoor functions. Our main result shows that the trade-off is also necessary for the NOVY construction. While the trade-off is an additive constant in Rudich’s work, our lower bound yields a gap between constant and almost-linear number of rounds.

Implications for protocol design. One could view this work quite broadly as providing a simple informal criterion for reasoning about the round complexity of classes of fully-black-box constructions (of protocols with a “simulatable” requirement) and formal techniques towards establishing a lower bound. The former is especially useful for protocol design in identifying and ruling out inefficient constructions. We stress here that our lower bounds do not apply to the black-box constructions of commitments from various classes of one-way functions in the works of Haitner et al. [HHK⁺05, NOV06, HR06b], in two different ways. One is the use of one-way functions in [HHK⁺05] to implement coin-tossing and zero-knowledge proofs to transform commitments that are hiding against

honest receivers into commitments that are hiding against arbitrary receivers. We note that our lower bound holds assuming merely hiding against honest receivers. The second is that the inputs to the one-way functions are used again in the commit phase. This is only needed to handle the lack of structure in general one-way functions. In particular, all the constructions are much simpler and requires fewer rounds when optimized for one-way permutations - they “collapse” to the NOVY construction. In short, the ways in which these constructions bypass our lower bounds do not provide much insight into how we may bypass the lower bounds for one-way permutations.

1.3 Additional Related Work

Fischlin [F02] showed that there is no black-box construction of 2-message statistically hiding from one-way permutations (or even trapdoor permutations). The result follows quite readily from Simon’s oracle separating collision-resistant hash functions and one-way permutations [S98]. On the other hand, Harnik and Naor [HN06] gave a non-black-box construction of a 2-message statistically hiding commitment from one-way functions under a non-standard assumption on compressibility of NP instances. From what we understand, there is no strong evidence either supporting or refuting the assumption.

2 Definitions and Preliminaries

We use PPT to denote both probabilistic polynomial-time Turing machines and probabilistic polynomial-time interactive Turing machines. The *round complexity* of a 2-party protocol is number of pairs of messages exchanged by both parties (in both directions). Unless otherwise stated, we use 1^n as the security parameter.

2.1 One-Way Permutations

Definition 1. A function $f : \{0,1\}^* \rightarrow \{0,1\}^*$ is a $s(n)$ -secure one-way function if f is computable in polynomial time and for every nonuniform PPT A ,

$$\Pr_{x \in \{0,1\}^n} [A(1^n, f(x)) \in f^{-1}(f(x))] < 1/s(n)$$

A function f is a one-way permutation if for every n , f restricted to $\{0,1\}^n$ is a permutation, and for all polynomials $s(n)$ and all sufficiently large n , f is $s(n)$ -secure.

A random permutation π is exponentially one-way even if the adversary is given access to a π^{-1} oracle, as long as it cannot query π^{-1} on the challenge. Here, $\pi_{\neq y}^{-1}$ is an oracle that on input y' , returns $\pi^{-1}(y')$ if $y' \neq y$, and \perp otherwise.

Lemma 1 (implicit in [GT00]). Fix $s(n) = 2^{n/5}$. For all sufficiently large n , there exists a permutation π on $\{0,1\}^n$ such that for all circuits A of size $s(n)$,

$$\Pr_{y \in \{0,1\}^n} [A^{\pi, \pi_{\neq y}^{-1}}(y) = \pi^{-1}(y)] < \frac{1}{s(n)}$$

Moreover, the statement relativizes.

2.2 Statistically Hiding Commitments

We present the definition for bit commitment. To commit to multiple bits, we may simply run a bit commitment scheme in parallel.

Definition 2. A (bit) commitment scheme $(\mathcal{S}, \mathcal{R})$ is an efficient two-party protocol consisting of two stages. Throughout, both parties receive the security parameter 1^n as input.

COMMIT. The sender \mathcal{S} has a private input $b \in \{0, 1\}$, which she wishes to commit to the receiver \mathcal{R} , and a sequence of coin tosses σ . At the end of this stage, both parties receive as common output a commitment z .

REVEAL. Both parties receive as input a commitment z . \mathcal{S} also receives the private input b and coin tosses σ for z . This stage is non-interactive: \mathcal{S} sends a single message to \mathcal{R} , and \mathcal{R} either outputs a bit and accepts or rejects.

Definition 3. A commitment scheme $(\mathcal{S}, \mathcal{R})$ is perfectly hiding if

COMPLETENESS. If both parties are honest, then for any input bit $b \in \{0, 1\}$ that \mathcal{S} gets, \mathcal{R} outputs b and accepts at the end of the decommit stage.

STATISTICALLY HIDING. For every unbounded deterministic strategy \mathcal{R}^* , the distributions of the view of \mathcal{R}^* in the commit stage while interacting with an honest \mathcal{S} are identical for $b = 0$ and $b = 1$. If the distributions are statistically indistinguishable, we obtain a statistically hiding commitment.

COMPUTATIONALLY BINDING. For every nonuniform PPT \mathcal{S}^* , \mathcal{S}^* succeeds in the following game (breaks the commitment) with negligible probability:

- \mathcal{S}^* interacts with an honest \mathcal{R} and outputs a commitment z .
- \mathcal{S}^* outputs two messages τ_0, τ_1 such that for both $b = 0$ and $b = 1$, \mathcal{R} on input (z, τ_b) accepts and outputs b .

3 Constructing Commitments from One-Way Permutations

In this section, we provide formal definitions of the various classes of constructions of commitments from one-way permutations we consider in this paper.

3.1 Fully Black-Box Constructions

Definition 4. A fully black-box construction of a statistically hiding commitment scheme from one-way permutations is a triplet of polynomial time computable oracle procedures $(\mathcal{S}, \mathcal{R}, M)$ for which there exists a polynomial T and a constant c satisfying the following properties:

EFFICIENCY. *The running times of $\mathcal{S}, \mathcal{R}, M$ are bounded by T .*

FUNCTIONALITY. *For every family of permutations π , $(\mathcal{S}^\pi, \mathcal{R}^\pi)$ is a statistically hiding commitment scheme.*

SECURITY. *For every $\varepsilon = 1/\text{poly}(n)$, for all sufficiently large n , every permutation $\pi : \{0, 1\}^n \rightarrow \{0, 1\}^n$ and every adversary \mathcal{S}^* , if \mathcal{S}^* breaks $(\mathcal{S}^\pi, \mathcal{R}^\pi)$ with probability ε , then*

$$\Pr_{y \in \{0, 1\}^n} [M^{\mathcal{S}^*, \pi}(y) = \pi^{-1}(y)] \geq \left(\frac{\varepsilon}{T}\right)^c$$

3.2 Interactive Hashing

Interactive hashing is a 2-party protocol between a sender and a receiver, similar to a commitment scheme. The sender begins with a private input $y \in \{0, 1\}^q$ and goal is for both parties to select a set of 2^k strings in $\{0, 1\}^q$ (specified by a circuit $C : \{0, 1\}^k \rightarrow \{0, 1\}^q$) containing y . The hiding property stipulates that the receiver does not learn which of the 2^k strings equals y , and the binding property stipulates that the sender can “control” at most one of the 2^k strings. The computational formulation (introduced explicitly in [NOV06] along with selecting many instead of merely 2 outputs) guarantees an efficient reduction from breaking the binding property to solving some computational problem on random instances.

Definition 5 ([NOV06]). *A computational interactive hashing scheme (with multiple outputs) is an efficient protocol $(\mathcal{S}_{\text{IH}}, \mathcal{R}_{\text{IH}})$ where both parties receive common inputs $(1^q, 1^k)$, \mathcal{S}_{IH} receives a private input $y \in \{0, 1\}^q$, with the common output being a circuit $C : \{0, 1\}^k \rightarrow \{0, 1\}^q$ and the private output of \mathcal{S}_{IH} being a string $z \in \{0, 1\}^k$. The protocol satisfies the following properties:*

CORRECTNESS. *For all \mathcal{R}^* and all $y \in \{0, 1\}^q$, let C, z be the common and private output of \mathcal{S}_{IH} in the protocol $(\mathcal{S}_{\text{IH}}, \mathcal{R}^*)(1^q, 1^k)$. Then, $C(z) = y$.*

PERFECTLY HIDING. *For all \mathcal{R}^* , (V, Z) is distributed identically to (V, U_k) , where $V = \text{view}_{\mathcal{R}^*}(\mathcal{S}_{\text{IH}}(U_q), \mathcal{R}^*)$.*

COMPUTATIONALLY BINDING. *There exists an oracle PPT A such that for every \mathcal{S}^* and any relation W , letting $C, ((x_0, z_0), (x_1, z_1))$ be the common and private output of \mathcal{S}_{IH} in the protocol $(\mathcal{S}_{\text{IH}}, \mathcal{R}^*)(1^q, 1^k)$, if it holds that*

$$\Pr[(x_0, C(z_0)) \in W \wedge (x_1, C(z_1)) \in W \wedge z_0 \neq z_1] > \varepsilon,$$

where the above probability is over the coin tosses of \mathcal{R}_{IH} and \mathcal{S}^ , then we have that*

$$\Pr_{y \in \{0, 1\}^q} [(A^{\mathcal{S}^*}(y, 1^q, 1^k, \varepsilon), y) \in W] > 2^{-k} \cdot (\varepsilon/q)^{O(1)}.$$

Nguyen et al. [NOV06] presented a protocol satisfying the above definition with $q - k$ rounds, obtained by ending the NOVY protocol $k - 1$ rounds earlier. The protocol is very simple: the receiver chooses $q - k$ linearly independent vectors v_1, \dots, v_k over $\{0, 1\}^q$. In round i , the receiver sends v_i and the sender responds with bit-wise dot product $v_i \cdot y$. We may reduce the round complexity by a factor of $O(\log q)$ by having the receiver send a pairwise independent hash function $h_i : \{0, 1\}^q \rightarrow \{0, 1\}^{O(\log q)}$ in round i and the sender responding with $h_i(y)$ [HR06a]. Note that the sender is deterministic, and the protocol is public-coin. Our lower bound shows that using a randomized sender or a private-coin protocol or q -wise independent hash functions will not further improve the round complexity (beyond constant factors).

Returning to the above definition, note that it refers to general relations W that may not be polynomial-time computable, and it does not give A oracle access to the relation W , which strengthens the security guarantee of the [NOV06] protocol. Our lower bound holds even if A has oracle access to the relation W , which is a weaker guarantee and thus a stronger lower bound. We also note that we may use the techniques in [LTW05] to show that this weaker guarantee also implies binding for static sets, thereby strengthening an observation made in [NOV06].

Naor et al. [NOV98] showed that any computational interactive hashing scheme $(\mathcal{S}_{\text{IH}}, \mathcal{R}_{\text{IH}})$ yields a fully black-box construction of a perfectly hiding commitment scheme $(\mathcal{S}, \mathcal{R})$ from any one-way permutation π with essentially the same round complexity.³ The construction is as follows:

COMMIT. To commit to a bit b , \mathcal{S} chooses a random $\sigma \in \{0, 1\}^n$, where n is the security parameter. Then, \mathcal{S} and \mathcal{R} run as a sub-protocol $(\mathcal{S}_{\text{IH}}(\pi(\sigma), \mathcal{R}_{\text{IH}})(1^n, 1^1)$, playing the roles $\mathcal{S}_{\text{IH}}, \mathcal{R}_{\text{IH}}$ respectively. Let C, z be the common and private outputs of \mathcal{S} in the sub-protocol. \mathcal{S} then sends $b' = b \oplus z$.

DECOMMIT. \mathcal{S} sends (b, σ) . \mathcal{R} accepts and outputs b if $C(b \oplus b') = \pi(\sigma)$, and rejects otherwise.

We stress that in the construction, \mathcal{S} queries π exactly once, to compute $\pi(\sigma)$, and does not need σ again except for decommitment.

As noted in the introduction, Damgård [D93] showed how any computational interactive hashing scheme can be used to transform constant-round honest-verifier public-coin zero-knowledge arguments into cheating-verifier public-coin zero-knowledge arguments unconditionally. The transformation may also be made more efficient by exploiting interactive hashing with multiple outputs so that a single application of interactive hashing yields a cheating-verifier zero-knowledge argument with soundness to $1/\text{poly}(n)$ (instead of $1/2$).

³ More precisely, Naor et al. showed how to construct a perfectly hiding commitment scheme from any one-way permutation using the interactive hashing protocol in [OVY93]. Implicit in the proof of correctness and security is a proof that the [OVY93] protocol satisfies Definition 5 for $k = 1$.

3.3 π -Oblivious Constructions

We describe the syntactic constraints on the class of fully black-box constructions for which we prove a lower bound. We consider constructions wherein the sender evaluates the one-way permutation only at the start of the commit phase, and does so on independent random inputs. The sender then uses the values (and not the inputs to the permutation), its input bit and possibly additional randomness in the rest of the commit phase. To decommit, the sender sends its input bit and its random tape, including the inputs to the permutation. We allow the receiver to query the permutation at any point in the protocol.

More formally,

Definition 6. *A fully black-box construction $(\mathcal{S}, \mathcal{R}, M)$ of a statistically hiding commitments from one-way permutations is π -oblivious if there exists some interactive PPT \mathcal{S}_{ob} such that for any permutation π on $\{0, 1\}^n$, to commit to a bit b with coin tosses σ , \mathcal{S} parses $\sigma = (\mathbf{z}, \tilde{\sigma})$, where $\mathbf{z} = (z_1, \dots, z_t) \in (\{0, 1\}^n)^t$, and proceeds according to $\mathcal{S}_{\text{ob}}(b, \sigma')$, where $\sigma' = (\mathbf{z}', \tilde{\sigma})$ and $\mathbf{z}' = \pi(\mathbf{z}) = (\pi(z_1), \dots, \pi(z_t))$. In particular, \mathcal{S}_{ob} never queries π . To decommit, \mathcal{S} sends a single message (b, σ) .*

Clearly, the NOVY construction is a π -oblivious; there, $t = 1$ and $\mathcal{S}_{\text{ob}} = \mathcal{S}_{\text{IH}}$ gets input $\pi(z_1)$, and $\tilde{\sigma}$ is the empty string since \mathcal{S}_{IH} is deterministic. Other candidates of π -oblivious constructions include variants of the NOVY construction wherein we run n^2 copies of some variant of interactive hashing in parallel either on the same $t = 1$ input $\pi(z_1)$ or on $t = n^2$ independent inputs $\pi(z_1), \dots, \pi(z_t)$, or a single copy of interactive hashing on the tn -bit string $\pi(z_1), \dots, \pi(z_t)$.

On the other hand, the construction of statistically hiding commitments from one-way functions in [HR06b] is not π -oblivious. This is because the sender will query π at some point z_1 and send both $h_1(\pi(z_1))$ and $h_2(z_1)$ during the commit phase, for some hash functions h_1, h_2 .

4 Main Result: Lower Bound for Commitments

Now, we state and prove our main result:

Theorem 1. *Any π -oblivious fully black-box construction of a statistically hiding commitment scheme from one-way permutations yields a commitment scheme with $\Omega(\frac{n}{\log n})$ rounds. This holds even if the hiding property for commitment scheme only holds for the honest receiver. More generally, if we assume that permutation is s -secure one-way, then we have an $\Omega(\frac{n}{\log s})$ lower bound.*

Our lower bound is tight:

Theorem 2 ([NOVY98, KS06, HR06a]). *There is a π -oblivious fully black-box construction of a perfectly hiding commitment scheme from s -secure one-way permutations with $O(\frac{n}{\log s})$ rounds.*

4.1 Proof Intuition

First, we point out at a high level how we exploit the fact that the construction is fully black-box. We use as the one-way permutation the one guaranteed by Lemma 1, which remains one-way even under a “chosen challenge” attack. This means that in order for the reduction M to successfully invert a challenge y , it must get a cheating sender \mathcal{S}^* to invert π on y itself. However, M is only given black-box access to \mathcal{S}^* , so it is limited to sending \mathcal{S}^* different inputs and possibly rewinding \mathcal{S}^* .

For concreteness, consider the NOVY construction of commitment schemes from one-way permutation using computational interactive hashing as a sub-protocol. When trying to invert a challenge y , the reduction M tries to get the sender to generate a commitment that is consistent with her input to interactive hashing protocol being y (otherwise, the decommitments will not help to invert y). At each round of commit phase, the honest \mathcal{S}_{IH} reveals some information about her input $\pi(\sigma)$. At the end of the commit phase, she should have revealed $n - 1$ bits of information about her input (since we’re using interactive hashing to choose 2 strings). We claim, at each round, she can only reveal $O(\log n)$ bits of information about her input, which yields a $\Omega(n/\log n)$ lower bound on the number of rounds. Suppose there is some round where \mathcal{S}_{IH} reveals $\omega(\log n)$ bits of information. This means that there are $n^{\omega(1)}$ inputs to the interactive hashing protocol that are consistent with the partial transcript. Consider a cheating sender that at each round samples a random input y' that is consistent with the partial transcript and responds as though her input to the interactive hashing protocol is y' , then the probability that the reduction observes a transcript that is consistent with y is negligible. It is important that \mathcal{S}_{IH} does not query π , so that we may sample consistent partial transcripts using a PSPACE oracle. If \mathcal{S}_{IH} is deterministic, it is straight-forward to quantify “information” about the sender’s input and turn this outline into a proof.

For general π -oblivious constructions, we construct the cheating sender in essentially the same way: at each round (for both the commit and reveal phases), the sender samples a random (b, σ') that is consistent with the partial transcript and responds as though her input to \mathcal{S}_{ob} is (b, σ') (where $\sigma' = (\mathbf{z}', \tilde{\sigma})$). The main technical difficulty in the analysis is in quantifying “information” about the sender’s input. Indeed, how much information a message reveals about \mathbf{z} depends on both b and $\tilde{\sigma}$. Also, for a fixed partial transcript, the set (and number) of \mathbf{z}' ’s that are consistent with the given transcript may vary with different choices of $b, \tilde{\sigma}$.

4.2 Proof of Theorem 1

We may assume that the commitment scheme $(\mathcal{S}, \mathcal{R})$ runs in r rounds, with \mathcal{R} going first. Let T, c be the polynomial and constant guaranteed by the fully black-box reduction. We will show that $r \gtrsim \frac{n - \log t}{8c \log T} = \Omega\left(\frac{n}{\log n}\right)$. Suppose otherwise, and take π to be the permutation guaranteed by Lemma 1.

Conventions regarding M . Recall that the reduction M has oracle access to a sender \mathcal{S}^* with which it inverts the permutation π . It can query \mathcal{S}^* on sequences of messages of the form $\mathbf{q}_i = (q_1, \dots, q_i)$ corresponding to the first i messages from \mathcal{R} in the commit phase, or a message of the form $(\mathbf{q}_r, \text{decommit})$, requesting for a decommit to a previous commitment. M runs for at most T steps, and therefore makes at most T queries to \mathcal{S}^* . In addition, we may adopt WLOG the following simplifying assumptions on M by modifying M appropriately (as is the case with lower bounds for black-box zero-knowledge [GK96b]):

1. It never asks the same query twice.
2. If M queries the oracle with \mathbf{q}_i , it has queried the oracle with all proper prefixes of \mathbf{q}_i (namely all sequences of the form (q_1, \dots, q_j) for $j \leq i$.)

Notations. We introduce some notations:

- $\mathcal{S}_{\text{ob}}(b, \sigma', \mathbf{q}_i)$ denotes the \mathcal{S}_{ob} 's response with input b, σ' and the first i messages from \mathcal{R} being \mathbf{q}_i .
- Given a partial transcript $(\mathbf{q}_i, \mathbf{a}_i) = (q_1, \dots, q_i, a_1, \dots, a_i)$ and $y \in \{0, 1\}^n$, $\text{Con}(\mathbf{q}_i, \mathbf{a}_i)$ is the set of inputs (b, σ') to \mathcal{S}_{ob} that would yield the transcript $(\mathbf{q}_i, \mathbf{a}_i)$; formally,

$$\text{Con}(\mathbf{q}_i, \mathbf{a}_i) = \{(b, \sigma') \mid \mathcal{S}_{\text{ob}}(b, \sigma', q_1, \dots, q_i) = a_j, \forall j = 1, 2, \dots, i\}$$

and

$$\text{Con}_y(\mathbf{q}_i, \mathbf{a}_i) = \{(b, \mathbf{z}', \tilde{\sigma}) \in \text{Con}(\mathbf{q}_i, \mathbf{a}_i) \mid \exists j : z'_j = y\}$$

In particular, $|\text{Con}_y(\epsilon)|/|\text{Con}(\epsilon)| = 1 - (1 - 2^{-n})^t \leq t2^{-n}$, where ϵ is the empty string (transcript).

Sender strategy \mathcal{S}^* . Consider the following sender strategy \mathcal{S}^* :

- Upon receiving a query of the form (\mathbf{q}_{i-1}, q_i) , look up previous replies \mathbf{a}_{i-1} . (For $i = 1$, $(\mathbf{q}_{i-1}, \mathbf{a}_{i-1}) = \epsilon$.) Sample uniformly at random⁴ (b, σ') from the set $\text{Con}(\mathbf{q}_{i-1}, \mathbf{a}_{i-1})$, and respond with $a_i = \mathcal{S}_{\text{ob}}(b, \sigma', \mathbf{q}_i)$.
- Upon receiving a query of the form $(\mathbf{q}_r, \text{decommit})$, look up previous replies \mathbf{a}_r . Sample uniformly and independently at random $(b_0, \mathbf{z}_0, \tilde{\sigma}_0), (b_1, \mathbf{z}_1, \tilde{\sigma}_1)$ from the set $\text{Con}(\mathbf{q}_r, \mathbf{a}_r)$, and send $(b_0, \pi^{-1}(\mathbf{z}_0), \tilde{\sigma}_0), (b_1, \pi^{-1}(\mathbf{z}_1), \tilde{\sigma}_1)$.

Note that in an interaction with an honest receiver \mathcal{R} , \mathcal{S}^* breaks the commitment with probability $1/2 - \text{neg}(n) > 1/4$. This is because the hiding property of the commitment scheme guarantees that a random decommitment is almost equally likely to be a 0 and a 1. Hence,

$$\Pr_{y \in \{0, 1\}^n} \left[M^{\mathcal{S}^*, \pi}(y) = \pi^{-1}(y) \right] > \left(\frac{1}{4T} \right)^c$$

⁴ \mathcal{S}^* can be made stateless by using a rT -wise independent family of hash functions, namely apply a hash function to the queries and use the output as randomness for uniform sampling [GK96b].

Analysis. Note that a PSPACE oracle suffices for simulating \mathcal{S}^* in the commit phase, whereas a PSPACE oracle and a π^{-1} oracle suffice in the reveal phase. Fix an input y to M . We want to show that with high probability, we may efficiently simulate the computation $M^{\mathcal{S}^*, \pi}(y)$ given oracle access to PSPACE, $\pi, \pi_{\neq y}^{-1}$.

We say that a partial transcript $(\mathbf{q}_i, \mathbf{a}_i)$ is *heavy* if

$$\frac{|\text{Con}_y(\mathbf{q}_i, \mathbf{a}_i)|}{|\text{Con}(\mathbf{q}_i, \mathbf{a}_i)|} > \gamma^{r+1-i}, \quad \text{where } \gamma = \left(\frac{t}{2^n}\right)^{\frac{1}{r+1}};$$

otherwise, we say that $(\mathbf{q}_i, \mathbf{a}_i)$ is *light*. In particular, ϵ is light, since $\frac{|\text{Con}_y(\epsilon)|}{|\text{Con}(\epsilon)|} \leq \gamma^{r+1}$. Informally, the quantity $\frac{|\text{Con}_y(\cdot)|}{|\text{Con}(\cdot)|}$ applied to a transcript $(\mathbf{q}_i, \mathbf{a}_i)$ is the density of “favorable” outcomes for the reduction M , wherein an outcome is favorable if in the decommitment, \mathcal{S}^* inverts π on y . We want to show that with high probability, every transcript generated by \mathcal{S}^* (in its interaction with M) is light, that is, the density of favorable outcomes is low.

Consider the queries M makes to \mathcal{S}^* :

- A commit phase query of the form $\mathbf{q}_i = (\mathbf{q}_{i-1}, q_i)$. Let \mathbf{a}_{i-1} be \mathcal{S}^* 's answers to the prefixes. Observe that

$$\begin{aligned} \frac{|\text{Con}_y(\mathbf{q}_{i-1}, \mathbf{a}_{i-1})|}{|\text{Con}(\mathbf{q}_{i-1}, \mathbf{a}_{i-1})|} &= \sum_{a_i} \frac{|\text{Con}(\mathbf{q}_i, \mathbf{a}_{i-1}, a_i)|}{|\text{Con}(\mathbf{q}_{i-1}, \mathbf{a}_{i-1})|} \cdot \frac{|\text{Con}_y(\mathbf{q}_i, \mathbf{a}_{i-1}, a_i)|}{|\text{Con}(\mathbf{q}_i, \mathbf{a}_{i-1}, a_i)|} \\ &= \sum_{a_i} \Pr[\mathcal{S}^*(\mathbf{q}_i) = a_i] \cdot \frac{|\text{Con}_y(\mathbf{q}_i, \mathbf{a}_{i-1}, a_i)|}{|\text{Con}(\mathbf{q}_i, \mathbf{a}_{i-1}, a_i)|} \\ &> \Pr[\mathcal{S}^*(\mathbf{q}_i) \rightarrow a_i; (\mathbf{q}_i, \mathbf{a}_{i-1}, a_i) \text{ is heavy}] \cdot \gamma^{r+1-i} \end{aligned}$$

This implies

$$\Pr[\mathcal{S}^*(\mathbf{q}_i) \rightarrow a_i; (\mathbf{q}_i, \mathbf{a}_{i-1}, a_i) \text{ is heavy} \mid (\mathbf{q}_{i-1}, \mathbf{a}_{i-1}) \text{ is light}] < \gamma$$

- A reveal phase query of the form $(\mathbf{q}_r, \text{decommit})$. Let \mathbf{a}_r be \mathcal{S}^* 's answers to \mathbf{q}_r . If $(\mathbf{q}_r, \mathbf{a}_r)$ is light, that is, $\frac{|\text{Con}_y(\mathbf{q}_r, \mathbf{a}_r)|}{|\text{Con}(\mathbf{q}_r, \mathbf{a}_r)|} \leq \gamma$, then with probability $1 - 2\gamma$, we can generate two independent random decommitments without inverting π on y .

Applying a union bound over that rT commit phase queries that M makes to \mathcal{S}^* , we have: with probability at least $1 - rT\gamma$, in every reveal phase query $(\mathbf{q}_r, \text{decommit})$ that M makes to \mathcal{S}^* , the transcript $(\mathbf{q}_r, \mathbf{a}_r)$ is light. Taking another union bound, we deduce that with probability $1 - (r+2)T\gamma$, we may efficiently simulate $M^{\mathcal{S}^*, \pi}$ on input y with oracle access to PSPACE, $\pi, \pi_{\neq y}^{-1}$. Hence, there is an oracle PPT \tilde{M} running in time $\text{poly}(T, n)$ such that

$$\Pr_{y \in \{0,1\}^n} \left[\tilde{M}^{\text{PSPACE}, \pi, \pi_{\neq y}^{-1}}(y) = \pi^{-1}(y) \right] > \left(\frac{1}{4T}\right)^c - (r+2)T\gamma > \frac{1}{2} \left(\frac{1}{4T}\right)^c$$

a contradiction to π being one-way. \square

4.3 Lower Bounds for Interactive Hashing

Using the connection between commitment schemes and computational interactive hashing described in Section 3.2, we derive a tight lower bound for the latter [NOV06, HR06a]:

Theorem 3. *Any computational interactive hashing scheme on common input $(1^n, 1^k)$ has $\Omega(\frac{n}{\log n})$ rounds, for $k = o(1)$.*

We believe that our techniques and analysis extend readily to yield lower bounds on efficiency of the security reduction for computational interactive hashing (an open problem posed in [HR06a]) and the round complexity of random selection [DGW95, DGOW95, GSV98]. We will explore these extensions in the full version of this paper.

Acknowledgements

I would like to thank Rosario Gennaro and Luca Trevisan for very helpful technical discussions, as well as Oded Goldreich and Chiu-Yuen Koo for pointers and discussions on previous related work. I would also like to thank Salil Vadhan and the anonymous referees for helpful suggestions and constructive feedback on this work. Finally, I am grateful towards Andy Yao, Luca, the IBM Cryptography Group, and the organizers of the IPAM Workshop on Securing Cyberspace for hosting my delightful stays in Beijing, San Francisco, New York and Los Angeles, during which this work was conceived, continued, and completed.

References

- [AH91] W. Aiello and J. Håstad. Statistical zero-knowledge languages can be recognized in two rounds. *JCSS*, 42(3):327–345, 1991.
- [AIK04] B. Applebaum, Y. Ishai, and E. Kushilevitz. Cryptography in NC^0 . In *Proc. 45th FOCS*, 2004.
- [B01] B. Barak. How to go beyond the black-box simulation barrier. In *Proc. 42nd FOCS*, 2001.
- [BCY91] G. Brassard, C. Crépeau, and M. Yung. Constant-round perfect zero-knowledge computationally convincing protocols. *Theoretical Computer Science*, 84(1):23–52, 1991.
- [BHZ87] R. B. Boppana, J. Håstad, and S. Zachos. Does co-NP have short interactive proofs? *IPL*, 25(2):127–132, 1987.
- [BJY97] M. Bellare, M. Jakobsson, and M. Yung. Round-optimal zero-knowledge arguments based on any one-way function. In *Proc. Eurocrypt '97*, 1997.
- [CCM98] C. Cachin, C. Crépeau, and J. Marcil. Oblivious transfer with a memory-bounded receiver. In *Proc. 39th FOCS*, 1998.
- [CS06] C. Crépeau and G. Savvides. Optimal reductions between oblivious transfers using interactive hashing. In *Proc. Eurocrypt '06*, 2006.
- [D93] I. Damgård. Interactive hashing can simplify zero-knowledge protocol design without computational assumptions. In *Proc. Crypto '93*, 1993.

- [DGOW95] I. Damgård, O. Goldreich, T. Okamoto, and A. Wigderson. Honest verifier vs dishonest verifier in public-coin zero-knowledge proofs. In *Proc. Crypto '95*, 1995.
- [DGW95] I. Damgård, O. Goldreich, and A. Wigderson. Information theory versus complexity theory: Another test case. manuscript, 1995.
- [DHRS04] Y. Z. Ding, D. Harnik, A. Rosen, and R. Shaltiel. Constant-round oblivious transfer in the bounded storage model. In *Proc. 1st TCC*, 2004.
- [DPP98] I. Damgård, T. P. Pedersen, and B. Pfitzmann. Statistical secrecy and multibit commitments. *IEEE Transactions on Information Theory*, 4(3):1143–1151, 1998.
- [F89] L. Fortnow. The complexity of perfect zero-knowledge. *Advances in Computing Research*, 5:429–442, 1989.
- [F02] M. Fischlin. On the impossibility of constructing non-interactive statistically-secret protocols from any trapdoor one-way function. In *Proc. CT-RSA '02*, 2002.
- [GGKT05] R. Gennaro, Y. Gertner, J. Katz, and L. Trevisan. Bounds on the efficiency of generic cryptographic constructions. *SIAM Journal on Computing*, 35(1):217–246, 2005.
- [GK96a] O. Goldreich and A. Kahan. How to construct constant-round zero-knowledge proof systems for NP. *J. Cryptology*, 9(3):167–190, 1996.
- [GK96b] O. Goldreich and H. Krawczyk. On the composition of zero-knowledge proof systems. *SIAM Journal on Computing*, 25(1):169–192, 1996.
- [GMR89] S. Goldwasser, S. Micali, and C. Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, 1989.
- [GSV98] O. Goldreich, A. Sahai, and S. Vadhan. Honest-verifier statistical zero-knowledge equals general statistical zero-knowledge. In *Proc. 30th STOC*, 1998.
- [GT00] R. Gennaro and L. Trevisan. Lower bounds on efficiency of generic cryptographic constructions. In *Proc. 41st FOCS*, 2000.
- [HHK⁺05] I. Haitner, O. Horvitz, J. Katz, C.-Y. Koo, R. Morselli, and R. Shaltiel. Reducing complexity assumptions for statistically-hiding commitment. In *Proc. Eurocrypt '05*, 2005.
- [HK05] O. Horvitz and J. Katz. Bounds on the efficiency of "black-box" commitment schemes. In *Proc. 32nd ICALP*, 2005.
- [HN06] D. Harnik and M. Naor. On the compressibility of NP instances and cryptographic applications. In *Proc. 47th FOCS*, 2006.
- [HR06a] I. Haitner and O. Reingold. A new interactive hashing theorem. ECCC TR06-096, 2006.
- [HR06b] I. Haitner and O. Reingold. Statistically-hiding commitment from any one-way function. Cryptology ePrint Archive, Report 2006/436, 2006.
- [IR89] R. Impagliazzo and S. Rudich. Limits on the provable consequences of one-way permutations. In *Proc. 21st STOC*, 1989.
- [KS06] T. Koshiha and Y. Seri. Round-efficient one-way permutation based perfectly concealing bit commitment scheme. ECCC TR06-093, 2006.
- [KST99] J. H. Kim, D. R. Simon, and P. Tetali. Limits on the efficiency of one-way permutation-based hash functions. In *Proc. 40th FOCS*, 1999.
- [LTW05] H. Lin, L. Trevisan, and H. Wee. On hardness amplification of one-way functions. In *Proc. 2nd TCC*, 2005.
- [NOV06] M.-H. Nguyen, S. J. Ong, and S. Vadhan. Statistical zero-knowledge arguments for NP from any one-way function. In *Proc. 47th FOCS*, 2006.

- [NOVY98] M. Naor, R. Ostrovsky, R. Venkatesan, and M. Yung. Perfect zero-knowledge arguments for NP using any one-way permutation. *J. Cryptology*, 11(2):87–108, 1998.
- [NV06] M.-H. Nguyen and S. Vadhan. Zero knowledge with efficient provers. In *Proc. 38th STOC*, 2006.
- [NY89] M. Naor and M. Yung. Universal one-way hash functions and their cryptographic applications. In *Proc. 20th STOC*, 1989.
- [OVY93] R. Ostrovsky, R. Venkatesan, and M. Yung. Fair games against an all-powerful adversary. In *AMS DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 1993.
- [OW93] R. Ostrovsky and A. Wigderson. One-way functions are essential for non-trivial zero-knowledge. In *ISTCS*, 1993.
- [R91] S. Rudich. The use of interaction in public cryptosystems. In *Proc. Crypto '91*, 1991.
- [RTV04] O. Reingold, L. Trevisan, and S. Vadhan. Notions of reducibility between cryptographic primitives. In *Proc. 1st TCC*, 2004.
- [S98] D. R. Simon. Finding collisions on a one-way street: Can secure hash functions be based on general assumptions? In *Proc. Eurocrypt '98*, 1998.
- [T01] L. Trevisan. Extractors and pseudorandom generators. *JACM*, 48(4):860–879, 2001.