Reducing Complexity Assumptions for Statistically-Hiding Commitment

Iftach Haitner* Omer Horvitz†‡ Jonathan Katz†§ Chiu-Yuen Koo†
Ruggero Morselli† Ronen Shaltiel¶

September 29, 2006

Abstract

We revisit the following question: *what are the minimal assumptions needed to construct statistically-hiding commitment schemes*? Naor et al. show how to construct such schemes based on any one-way permutation. We improve upon this by showing a construction based on any *approximable preimage-size* one-way function. These are one-way functions for which it is possible to efficiently approximate the number of pre-images of a given output. A special case is the class of *regular* one-way functions where all points in the image of the function have the same number of pre-images.

We also prove two additional results related to statistically-hiding commitment. First, we prove a (folklore) *parallel composition theorem* showing, roughly speaking, that the statistical hiding property of any such commitment scheme is amplified exponentially when multiple independent executions of the scheme are carried out. Second, we show a *compiler* which transforms any commitment scheme which is statistically hiding against an honest-but-curious receiver into one which is statistically hiding even against a malicious receiver.

1 Introduction

A central focus of modern cryptography has been to investigate the weakest possible assumptions under which various cryptographic primitives exist. This direction of research has been quite fruitful, and minimal assumptions are known for a wide variety of primitives: e.g., it has been sown that one-way functions imply (and are implied by) pseudorandom generators, pseudorandom functions, symmetric-key encryption/message authentication, statistically-binding commitment, and digital signatures [22, 10, 11, 21, 25, 27, 30]. In other cases, black-box separation results exist which indicate the difficulty of constructing "strong" cryptographic protocols (say, key-exchange) from "weak" building blocks (say, one-way permutations; see [23]).

For some cryptographic primitives, however, an exact characterization of the minimal assumption under which the primitive exists is not known. *Statistically-hiding commitment* provides one such example. Informally, a commitment scheme defines a two-phase interactive protocol between

*Department of Computer Science, Weizmann Institute of Science. iftach.haitner@weizmann.ac.il. Research supported by US-Israel Binational Science Foundation grant 2002246.

†Department of Computer Science, University of Maryland. {horvitz,jkatz,cykoo,ruggero}@cs.umd.edu.

‡Research supported by U.S. Army Research Office award DAAD19-01-1-0494.

§Research supported by NSF CAREER award #0447075.

¶Department of Computer Science, University of Haifa. ronen@cs.haifa.ac.il.

a sender $\mathcal{S}$ and a receiver $\mathcal{R}$; after the *commitment phase*, $\mathcal{S}$ is uniquely bound to (at most) one value which is not yet revealed to $\mathcal{R}$, and in the *decommitment phase* $\mathcal{R}$ finally learns this value. The two security properties hinted at in this informal description are known as *binding* (namely, that $\mathcal{S}$ is bound to at most one value after the commitment phase) and *hiding* (namely, that $\mathcal{R}$ does not learn the value to which $\mathcal{S}$ commits before the decommitment phase). In a statistically-hiding commitment scheme the hiding property holds *even against all-powerful receivers* (i.e., hiding holds information-theoretically), while the binding property is required to hold only for computationally-bounded (say, polynomial-time) senders.

Statistically-hiding commitment schemes have been used as a building block in constructions of statistical zero-knowledge arguments [5, 26] or coin-tossing protocols [2, 24]. They are also advantageous when used within protocols in which certain commitments are never revealed; in this case, one can argue that computational binding suffices since it need only be infeasible to violate the binding property *during the period of time the protocol is run*, whereas statistical hiding has the advantage of ensuring that committed values remain hidden *forever* (i.e., regardless of how much time the receiver invests after completion of the protocol). Indeed, this is part of the motivation for statistical zero-knowledge as well. For further discussion, the reader is referred to [28, 29, 26].

Perfectly-hiding[1] commitment schemes were first shown to exist based on specific number-theoretic assumptions [5, 4] or, more generally, based on any collection of claw-free permutations [19] with an efficiently-recognizable index set [13] (see [13] for a weaker variant of statistically-hiding commitment which suffices for some applications and for which an efficiently-recognizable index set is not needed). Naor et al. [26], building on techniques developed earlier by Ostrovsky et al. [28, 29], showed a construction of a perfectly-hiding commitment scheme based on any one-way permutation. Statistically-hiding commitment schemes can also be constructed from collision-resistant hash functions [7, 20]; see [31] for minimal assumptions for the existence of the latter.

1.1 Our Results

We show how to construct a statistically-hiding commitment scheme given any *approximable pre-image-size* one-way function. Informally, this is a one-way function f satisfying the additional property that, given any y in the image of f, the value $|\{x : f(x) = y\}|$ (i.e., the number of pre-images of y) can be efficiently estimated. An interesting special case is an *approximately-regular* one-way function for which every point in the image of f has roughly the same number of pre-images. (We still require that it be feasible to approximate the number of pre-images.) A variety of conjectured one-way functions are in fact regular; we refer the reader to [14] for examples.

Our result may be viewed as an example of the paradigm in which a sequence of works constructs a given primitive from ever-weaker assumptions; e.g., in the cases of pseudorandom generators and universal one-way hash functions/signature schemes (see [8, Chap. 2] and [9, Chap. 6]), constructions were first based on specific, number-theoretic assumptions [3, 19], and then the minimal assumptions were gradually reduced to trapdoor permutations [1], one-way permutations [34, 3, 15, 27], regular one-way functions [14, 32], and (finally) one-way functions [21, 30]. We hope our work will similarly serve as a step toward resolving the question of the minimal assumptions required for statistically-hiding commitment.

We also provide two additional results of independent interest that may be useful for future constructions of statistically-hiding commitment schemes. Before describing these results, we review

[1] Very informally, in a statistically-hiding commitment scheme the receiver learns only a negligible amount of information about the sender's committed value, whereas in a perfectly-hiding commitment scheme the receiver learns *nothing*. Note that any perfectly-hiding scheme is also statistically-hiding.

the standard definition of statistical hiding. Say a commitment scheme $(\mathcal{S}, \mathcal{R})$ is *ρ-hiding against $\mathcal{R}^*$* if the distribution over the view of the (malicious) receiver $\mathcal{R}^*$ when the sender $\mathcal{S}$ commits to '0' is within statistical difference ρ from the distribution over the view of $\mathcal{R}^*$ when $\mathcal{S}$ commits to '1'. The standard definition of statistical hiding requires that for all (even all-powerful) $\mathcal{R}^*$, the commitment scheme should be ε-hiding against $\mathcal{R}^*$ for some negligible function ε. One way of relaxing this is to require only that the scheme be $(1 - \frac{1}{\mathsf{poly}})$-hiding (for all $\mathcal{R}^*$). An alternate relaxation is to require only that the scheme be ε-hiding against the honest receiver $\mathcal{R}$ (this corresponds to the classical cryptographic notion of an *honest-but-curious* adversarial entity). (In all cases, we assume that binding holds with all but negligible probability for any polynomial-time senders.)

We show that a scheme satisfying either of the relaxations above suffices to construct a scheme secure in the standard sense, with minimal increase in the round complexity. Specifically:

1. We prove a *parallel repetition theorem* for statistically-hiding commitment. Given commitment scheme $(\mathcal{S}, \mathcal{R})$, consider the scheme $(\mathcal{S}^q, \mathcal{R}^q)$ in which commitment to a bit b is done as follows: $\mathcal{S}^q$ chooses random bits $b_1, \ldots, b_q$ subject to the constraint $\bigoplus_i b_i = b$, and then runs q parallel executions of $\mathcal{S}$ using input bit b_i in the i^{th} execution. We show that if the initial scheme $(\mathcal{S}, \mathcal{R})$ is ρ-hiding, then the derived scheme $(\mathcal{S}^q, \mathcal{R}^q)$ is ρ^q-hiding. A corollary is that the existence of a $(1 - \frac{1}{\mathsf{poly}})$-hiding scheme implies the existence of a statistically-hiding commitment scheme *using the same number of rounds*.

   We remark that the result is trivial to prove for the case of an honest-but-curious receiver, but (as in the case of analyzing the effect of parallel repetition on soundness of interactive proof systems for a malicious prover [18, Appendix C]) is more difficult to prove for the case of a malicious receiver who may correlate its actions in the different parallel executions.

2. We show a general *compiler* that converts any commitment scheme that is statistically-hiding for an honest-but-curious receiver into one that is statistically-hiding for an arbitrarily-malicious receiver. Our compiler requires only the existence of one-way functions, which are implied anyway by the commitment scheme we start with. The compiler increases the round complexity of the initial scheme by $\omega(1)$ rounds (this could be improved to $O(1)$ rounds given a constant-round zero-knowledge proof based on one-way functions).

Taken together, the above results imply that, in order to construct a statistically-hiding commitment scheme, it suffices to construct a computationally-binding scheme which is $(1 - \frac{1}{\mathsf{poly}})$-hiding against an honest-but-curious receiver.

1.2 Overview of Our Techniques

Our construction is based on the protocol of Naor et al. [26], which is shown there to be perfectly hiding (as well as computationally binding) when based on any one-way permutation. It is natural to ask what happens when this protocol is applied using some other function $f : \{0,1\}^n \to \{0,1\}^\ell$. We first observe that the main proof of [26] shows that the protocol is computationally binding as long as f cannot be efficiently inverted with respect to the uniform distribution U_ℓ (formally, we mean that no efficient algorithm can find an x such that $f(x) = y$, for uniformly-chosen y, with non-negligible probability). We call a function with this property *one-way over its range*. Note that a function with this property is not necessarily one-way; for example, the constant function $f(x) = 0^\ell$ is trivially one-way over its range since the probability that a uniformly-selected y lies in the image of f (that is, the probability that $y = 0^\ell$) is negligible.

As our first main technical result, we show that the protocol of Naor et al. is "weakly hiding" when based on a function f for which the distribution $f(U_n)$ is *balanced* in a sense we will soon

describe. (By "weakly hiding" we mean $(1 - \frac{1}{\mathsf{poly}})$-hiding as defined in the previous section. As discussed there, any such protocol can be used to construct a full-fledged statistically-hiding commitment scheme.) Loosely speaking, a distribution over $\{0,1\}^\ell$ is balanced if it assigns to "most" elements $y \in \{0,1\}^\ell$ a probability that is "close" to $2^{-\ell}$. (In the formal definition we allow some elements to have probability outside this range as long as both the number of such elements and their total weight are small.)

The above shows that statistically-hiding commitment is implied by the existence of a function f that is both balanced and one-way over its range.[2] Section 4 of the paper is devoted to constructing such functions based on any approximately-regular one-way function. Inspired by [21, 30], we use poly-wise independent hashing to achieve this goal. Restricting our attention here to functions f that are *regular*, we define $f'(h,x) = (h, h(f(x)))$ where h is selected from a family of $\Theta(k)$-wise independent hash functions and k is the security parameter. Intuitively, if the output length of h is "small enough" (relative to the regularity parameter[3] of f) then f' will be sufficiently balanced, while if the output length of h is "large enough" then f' will be one-way over its range. We show that it is possible to set the output length "in the middle" and obtain both properties simultaneously. Note that our construction requires that the regularity parameter of f is known; we do not know how to extend the construction to the case where the regularity parameter is unknown.

The same construction (with only a slightly more complex analysis) works also if f is only approximately regular. It is also fairly easy to show how to convert any approximable pre-image-size one-way function into a one-way function that is approximately regular.

1.3 Outline of the Paper

We begin by reviewing some preliminaries and establishing some notation in Section 2. In that section, we also note that any approximable pre-image size one-way function can be converted into an approximately-regular one-way function. In Section 3, we formally define the notion of "balanced" functions described informally earlier, and show that any balanced function that is one-way over its range can be used to construct a statistically-hiding commitment scheme. Our task is thus reduced to constructing such a function starting from any approximately-regular one-way function, and we tackle this in Section 4. This completes the proof of our main result.

In Section 5 we prove a parallel composition theorem for statistically-hiding commitment, and in Section 6 we show a compiler converting any commitment scheme statistically-hiding for an honest-but-curious receiver into one that is statistically-hiding for a malicious receiver.

2 Preliminaries

Throughout this paper, we let k denote the security parameter. If X is a distribution over a finite set $\mathcal{X}$, the *support* of X (denoted $\mathsf{supp}(X)$) consists of those elements having non-zero probability under X. The *min-entropy* of X is defined as:

$$H_\infty(X) \stackrel{\text{def}}{=} \min_{x \in \mathsf{supp}(X)} \log\left(\frac{1}{\Pr_X[x]}\right).$$

[2] We remark that known constructions of "almost-everywhere one-to-one" one-way functions [12], "almost one-to-one" one-way functions [8, Sect. 3.5], and the constructions of [21] do not suffice for our purposes.

[3] That is, the number of pre-images of each value in the image of f.

If X_1 and X_2 are two distributions over a finite set $\mathcal{X}$, their statistical difference (written $\mathsf{SD}(X_1, X_2)$) is defined as:

$$\mathsf{SD}(X_1, X_2) \stackrel{\text{def}}{=} \frac{1}{2} \sum_{x \in \mathcal{X}} |\Pr_{X_1}[x] - \Pr_{X_2}[x]| .$$

Two distribution ensembles $\mathcal{X}_1 = \{X_1(k)\}_{k \in \mathbb{N}}$ and $\mathcal{X}_2 = \{X_2(k)\}_{k \in \mathbb{N}}$ have statistical difference ρ (for ρ a function of k) if $\mathsf{SD}(X_1(k), X_2(k)) \leq \rho(k)$ for all k large enough. If ρ is negligible, we say the ensembles are *statistically indistinguishable*.

We let U_n denote the uniform distribution over $\{0,1\}^n$. Given a function $f : \{0,1\}^n \to \{0,1\}^\ell$, we let $f(U_n)$ denote the distribution over $\{0,1\}^\ell$ induced by choosing x uniformly and outputting $f(x)$. For a function $f : \{0,1\}^n \to \{0,1\}^\ell$, we define $\mathsf{image}(f) \stackrel{\text{def}}{=} \{f(x) \mid x \in \{0,1\}^n\}$. We use 'log' to denote logarithms base 2, and 'ln' to denote the natural logarithm.

2.1 Commitment Schemes

An interactive bit commitment scheme is defined via a triple of PPT algorithms $(\mathcal{S}, \mathcal{R}, \mathcal{V})$. Looking ahead, $\mathcal{S}$ and $\mathcal{R}$ will interact during what is called a *commitment phase*, while $\mathcal{V}$ will be used during the (non-interactive) *decommitment phase*. Formally:

- $\mathcal{S}$ (the *sender*) is an interactive Turing machine (ITM) which receives as initial input the security parameter 1^k and a bit b. Following its interaction, it outputs some information decom (the *decommitment*).
- $\mathcal{R}$ (the *receiver*) is an ITM which receives the security parameter 1^k as initial input. Following its interaction, it outputs some state information com.
- $\mathcal{V}$ (acting as a receiver, in the decommitment phase) is a deterministic algorithm which receives as input state information com and a decommitment decom; it outputs either a bit b or the distinguished value $\perp$.

Denote by $(\mathsf{decom} \mid \mathsf{com}) \leftarrow \langle \mathcal{S}(1^k, b), \mathcal{R}(1^k) \rangle$ the experiment in which $\mathcal{S}$ and $\mathcal{R}$ interact (using the given inputs and uniformly random coins), and then $\mathcal{S}$ outputs decom while $\mathcal{R}$ outputs com. It is required that for all k, all b, and every pair $(\mathsf{decom} \mid \mathsf{com})$ that may be output by $\langle \mathcal{S}(1^k, b), \mathcal{R}(1^k) \rangle$, it is the case that $\mathcal{V}(\mathsf{com}, \mathsf{decom}) = b$.

The security of a commitment scheme can be defined in two complementary ways, protecting against either an all-powerful sender or an all-powerful receiver. Since we are interested in the case of statistically-hiding commitment (i.e., the latter case), we only provide the definition for this case.

Definition 2.1. Commitment scheme $(\mathcal{S}, \mathcal{R}, \mathcal{V})$ is *ρ-hiding* (for ρ a function of k) if the following holds: Given a deterministic ITM $\mathcal{R}^*$, let $\mathsf{view}_{\langle \mathcal{S}(b), \mathcal{R}^* \rangle}(k)$ denote the distribution on the view of $\mathcal{R}^*$ when interacting with $\mathcal{S}(1^k, b)$ (this view simply consists of the sequence of messages it receives from $\mathcal{S}$), where this distribution is taken over the random coins of $\mathcal{S}$. Then we require that for any (even all-powerful) $\mathcal{R}^*$ the ensembles $\{\mathsf{view}_{\langle \mathcal{S}(0), \mathcal{R}^* \rangle}(k)\}$ and $\{\mathsf{view}_{\langle \mathcal{S}(1), \mathcal{R}^* \rangle}(k)\}$ have statistical difference at most $\rho(k)$.

A commitment scheme is *statistically hiding* if it is ρ-hiding for negligible ρ. A 0-hiding scheme is called *perfectly hiding*.

Assuming $\mathcal{R}^*$ to be deterministic is without loss of generality since $\mathcal{R}^*$ may be all-powerful.

Definition 2.2. Commitment scheme $(\mathcal{S}, \mathcal{R}, \mathcal{V})$ is *computationally-binding* if the following is negligible for all PPT $\mathcal{S}^*$:

$$\Pr\left[((\mathsf{decom}, \mathsf{decom}') \mid \mathsf{com}) \leftarrow \langle \mathcal{S}^*(1^k), \mathcal{R}(1^k)\rangle : \begin{array}{c} \mathcal{V}(\mathsf{com}, \mathsf{decom}), \mathcal{V}(\mathsf{com}, \mathsf{decom}') \in \{0,1\} \bigwedge \\ \mathcal{V}(\mathsf{com}, \mathsf{decom}) \neq \mathcal{V}(\mathsf{com}, \mathsf{decom}') \end{array}\right],$$

where the probability is taken over the random coins of both $\mathcal{S}^*$ and $\mathcal{R}$.

Given the above, we now define a statistically-secure commitment scheme:

Definition 2.3. Commitment scheme $(\mathcal{S}, \mathcal{R}, \mathcal{V})$ is *ρ-secure* (resp., *statistically secure*, *perfectly secure*) if it is computationally binding and ρ-hiding (resp., statistically hiding, perfectly hiding).

2.2 One-Way Function Families and Variants

All function families $\mathcal{F} = \{f_k : \{0,1\}^{n(k)} \to \{0,1\}^{\ell(k)}\}_{k\in\mathbb{N}}$ in this paper will have $n, \ell = \mathsf{poly}(k)$, and n, f_k computable in time polynomial in k. We say $\mathcal{F}$ is *one-way* if, for all PPT algorithms A, the following is negligible (in k):

$$\Pr[x \leftarrow \{0,1\}^{n(k)}; y = f_k(x); x' \leftarrow A(1^k, y) : f_k(x') = y].$$

We also consider some additional properties of function families:

- $\mathcal{F}$ is $r(k)$**-regular** if for every k and every $x \in \{0,1\}^{n(k)}$ we have

  $$\left|\{x' \in \{0,1\}^{n(k)} : f_k(x') = f_k(x)\}\right| = 2^{r(k)}.$$

  $\mathcal{F}$ is $r(k)$**-known regular** if, in addition, $r(k)$ is computable in time polynomial in k.

- $\mathcal{F}$ is $(r(k), p(k))$**-approximately-regular** if for every k and every $x \in \{0,1\}^{n(k)}$ we have

  $$\frac{1}{p(k)} \cdot 2^{r(k)} \leq \left|\{x' \in \{0,1\}^{n(k)} : f_k(x') = f_k(x)\}\right| \leq p(k) \cdot 2^{r(k)},$$

  and $r(k), p(k)$ are computable in time polynomial in k. We will be interested in the case where $p(k)$ is upper-bounded by a polynomial in k.

  Note that if f is (r,p)-approximately-regular, then the min-entropy of $D = f(U_n)$ satisfies

  $$n - r - \log p \leq H_\infty(D) \leq n - r + \log p.$$

- $\mathcal{F}$ is **approximable pre-image-size** if the function $\tilde{D}_{f_k}(y) \stackrel{\text{def}}{=} \lceil \log |f_k^{-1}(y)| \rceil$ is computable in time polynomial in k for all y in the image of f_k.

For simplicity, we sometimes drop the explicit dependence on k when clear and write, e.g., $f : \{0,1\}^n \to \{0,1\}^\ell$ rather than $f_k : \{0,1\}^{n(k)} \to \{0,1\}^{\ell(k)}$.

It is relatively straightforward to construct an $(n,2)$-approximately-regular one-way function $f' : \{0,1\}^{2n} \to \{0,1\}^{n+\ell}$ from an approximable-preimage-size one-way function $f : \{0,1\}^n \to \{0,1\}^\ell$ by defining:

$$f'(x\|z) = y \parallel z[1 \ldots \tilde{D}_f(y)] \parallel 0^{n-\tilde{D}_f(y)},$$

where $y = f(x)$, "$\|$" denotes concatenation, and $z[1 \ldots r]$ denotes the first r bits of z. One-wayness of f' is easy to show, and the number of pre-images of an element $y \| \bar{z} \| 0^{n-\tilde{D}_f(y)}$ satisfies

$$\begin{aligned} \left|(f')^{-1}\left(y \parallel \bar{z} \parallel 0^{n-\tilde{D}_f(y)}\right)\right| &= \left|f^{-1}(y)\right| \cdot 2^{n-\tilde{D}_f(y)} \\ &\in \left[\frac{2^{\tilde{D}_f(y)}}{2} \cdot 2^{n-\tilde{D}_f(y)},\ 2^{\tilde{D}_f(y)} \cdot 2^{n-\tilde{D}_f(y)}\right] \\ &= \left[\frac{2^n}{2},\ 2^n\right]. \end{aligned}$$

To prove our main result, it will be thus sufficient for us to construct a statistically-secure commitment scheme starting from any approximately-regular one-way function.

2.3 Universal Hashing and an Extended Chernoff Bound

Let $\mathcal{H} = \{H_k\}_{k \in \mathbb{N}}$ be a collection of function families, where each H_k is a family of functions mapping strings of length $\ell(k)$ to strings of length $v(k)$. We assume that the size of each H_k is a power of 2, and that we can identify each binary string of some appropriate length $s(k)$ with a unique function $h \in H_k$. (In particular, choosing random $h \in H_k$ is identified with choosing random a random string of length $s(k)$.) We say $\mathcal{H}$ is an $n(k)$-*universal hash family* (following [6]) if for each k, any distinct $x_1, \ldots, x_{n(k)} \in \{0,1\}^{\ell(k)}$, and any $y_1, \ldots, y_{n(k)} \in \{0,1\}^{v(k)}$ we have:

$$\Pr_{h \leftarrow H_k}[h(x_1) = y_1 \wedge \cdots \wedge h(x_{n(k)}) = y_{n(k)}] = 2^{-v(k) \cdot n(k)}.$$

Put another way, for any distinct $x_1, \ldots, x_{n(k)}$, the random variables $h(x_1), \ldots, h(x_{n(k)})$ (where $h \leftarrow H_k$) are n-wise independent. Constructions of $n(k)$-universal hash families with $s(k) = O(n(k) \cdot \max(\ell(k), v(k)))$ are known.

The following Chernoff-like bound will be useful in our analysis:

Lemma 2.4. (Extended Chernoff Bound [33, Theorem 5]) *Let X be the sum of (any number of) n-wise independent random variables, each taking values in the interval $[0,1]$, such that $E[X] = \mu$. Then for any $\varepsilon \leq 1$ for which $n \geq \lfloor \varepsilon^2 \mu e^{-1/3} \rfloor$ we have $\Pr[|X - \mu| \geq \varepsilon\mu] \leq e^{-\lfloor \varepsilon^2 \mu/3 \rfloor}$.*

2.4 Interactive Hashing

Interactive hashing was introduced by Ostrovsky, et al. [28, 29], and used by Naor, et al. [26] to construct a statistically-secure (actually, perfectly-secure) commitment scheme based on any one-way permutation family. We review interactive hashing, as well as the resulting commitment scheme, below. In what follows, we let $x \cdot y$ denote $\sum_{i=1}^{\ell} x_i y_i \bmod 2$ for $x, y \in \{0,1\}^{\ell}$.

Construction 2.5 (Interactive hashing). *The protocol is defined by algorithms $\mathcal{S}$ and $\mathcal{R}$, where $\mathcal{S}$ begins with an ℓ-bit value y (with ℓ known to $\mathcal{R}$), and proceeds as follows:*

1. *The parties interact in $\ell - 1$ stages. In stage i (for $i = 1, \ldots, \ell - 1$), $\mathcal{R}$ chooses $r_i \in \{0,1\}^{\ell - i}$ uniformly at random and sends the "query" $q_i = 0^{i-1}1r_i$ to $\mathcal{S}$ (in case $\mathcal{R}$ aborts, $\mathcal{S}$ simply takes q_i to be some default value); in response, $\mathcal{S}$ sends $c_i = q_i \cdot y$.*

2. *At the conclusion of the above, there are exactly two strings $y_0, y_1 \in \{0,1\}^{\ell}$ satisfying the system of equations $\{q_i \cdot X = c_i\}_{1 \leq i \leq \ell - 1}$; let y_0 denote the lexicographically smaller of the two. Both parties compute (y_0, y_1), and $\mathcal{S}$ sets v such that $y = y_v$.*

The output of the protocol to be (y_0, y_1, v) for $\mathcal{S}$ and (y_0, y_1) for $\mathcal{R}$. We denote by $IH(y)$ an execution of the interactive hashing protocol, where $\mathcal{S}$ begins with input y.

The above was used in [26] to construct a perfectly-secure commitment scheme based on one-way permutations via the following approach:

Construction 2.6. *Let $\mathcal{F} = \{f_k : \{0,1\}^{n(k)} \to \{0,1\}^{\ell(k)}\}$ be a function family. Commitment scheme $(\mathcal{S}, \mathcal{R}, \mathcal{V})$ is defined as follows: $\mathcal{S}(1^k, b)$ chooses $x \in \{0,1\}^{n(k)}$ uniformly at random, computes $y = f_k(x)$, and then executes $IH(y)$ with $\mathcal{R}$; this protocol results in output (y_0, y_1, v) for $\mathcal{S}$ and (y_0, y_1) for $\mathcal{R}$. The commitment phase concludes by having $\mathcal{S}$ send $\hat{v} = v \oplus b$ to $\mathcal{R}$. Finally, $\mathcal{S}$ outputs $\mathsf{decom} = x$ while $\mathcal{R}$ outputs state $\mathsf{com} = (y_0, y_1, \hat{v})$.*

In the decommitment phase, $\mathcal{V}((y_0, y_1, \hat{v}), x)$ proceeds as follows: if $f_k(x) = y_0$, output $\hat{v}$; if $f_k(x) = y_1$, output $\hat{v} \oplus 1$; otherwise, output $\perp$.

It is relatively easy to observe that the above protocol is perfectly hiding if $\ell = n$ and $\mathcal{F}$ is a permutation family (regardless of whether $\mathcal{F}$ is one-way). The main result of [26] was to prove that the above is *computationally binding* when $\mathcal{F}$ is a *one-way* permutation family. In fact, careful examination of their proof shows the above commitment scheme is computationally binding under a *weaker* condition on $\mathcal{F}$; it suffices for $\mathcal{F}$ to be *one-way over its range*, defined as follows:

Definition 2.7. Let $\mathcal{F} = \{f_k : \{0,1\}^{n(k)} \to \{0,1\}^{\ell(k)}\}_{k \in \mathbb{N}}$ be a function family. We say $\mathcal{F}$ is *one-way over its range* if, for all PPT A, the following is negligible (in k):

$$\Pr[y \leftarrow \{0,1\}^{\ell(k)}; x \leftarrow A(1^k, y) : f_k(x) = y].$$

We stress that, in contrast to the definition in the case of a (standard) one-way function, here y is chosen uniformly in the range of f_k rather than according to $f(U_n)$.

Theorem 2.8 (Implicit in [26]). *If $\mathcal{F}$ is one-way over its range, then Construction 2.6 is computationally binding.*

3 Statistical Hiding from Balanced Functions

In this section we define a notion of "balance" and show that if a function family $\mathcal{F}$ is sufficiently balanced then Construction 2.6 yields a protocol that is "somewhat hiding." Roughly speaking, a distribution D on $\{0,1\}^\ell$ is balanced if D is "close" to uniform "most" of the time. A function $f : \{0,1\}^n \to \{0,1\}^\ell$ is then defined to be balanced if the distribution $f(U_n)$ is balanced. Formally:

Definition 3.1. Distribution D on $\{0,1\}^\ell$ is (α, δ)-*balanced* if there is a set $\mathsf{Bad} \subset \{0,1\}^\ell$ such that:

1. $|\mathsf{Bad}| \leq \alpha \cdot 2^\ell$.
2. $\Pr_{y \leftarrow D}[y \in \mathsf{Bad}] \leq \alpha$.
3. For every $y_0 \notin \mathsf{Bad}$, $\left|\Pr_{y \leftarrow D}[y = y_0] - \frac{1}{2^\ell}\right| \leq \frac{\delta}{2^\ell}$ (we will always have $\delta < 1$).

Function $f : \{0,1\}^n \to \{0,1\}^\ell$ is (α, δ)-balanced if the distribution $f(U_n)$ is (α, δ)-balanced. Function family $\mathcal{F} = \{f_k : \{0,1\}^{n(k)} \to \{0,1\}^{\ell(k)}\}$ is (α, δ)-balanced if, for all k large enough, f_k is $(\alpha(k), \delta(k))$-balanced.

Our main result of this section is the following:

Theorem 3.2. *If $\mathcal{F} = \{f_k : \{0,1\}^{n(k)} \to \{0,1\}^{\ell(k)}\}$ is an (α,δ)-balanced function family, then Construction 2.6 is ρ-hiding for $\rho = 2\alpha + \delta + \alpha\delta$.*

Proof. Fix k large enough so that f_k is $(\alpha(k),\delta(k))$-balanced; from now on we simply write f, α, δ, ρ without explicitly indicating their dependence on k. For a given execution of the scheme, let τ denote the initial transcript resulting from the interactive hashing sub-protocol; thus, the view of $\mathcal{R}^*$ consists of τ and the bit $\hat{v}$ sent in the final round. Given a particular (deterministic) $\mathcal{R}^*$, we write $\mathsf{Exp}(b)$ to denote the experiment in which $\mathcal{S}$ chooses a uniform random tape and then executes the protocol with $\mathcal{R}^*$ using this random tape and the bit b, resulting in view $(\tau, \hat{v})$ for $\mathcal{R}^*$. Note that the distribution on τ is identical in $\mathsf{Exp}(0)$ and $\mathsf{Exp}(1)$, since the first phase of the commitment scheme is independent of b.

Below, we define a "good" set of initial transcripts Good, and show that:

Claim 3.3. $\Pr_{\mathsf{Exp}(0)}[\tau \in \mathsf{Good}] = \Pr_{\mathsf{Exp}(1)}[\tau \in \mathsf{Good}] \geq 1 - \alpha(2+\delta)$. *Since this probability is independent of the bit b being committed to, we simply write $\Pr_{\mathsf{Exp}}[\tau \in \mathsf{Good}]$ for this probability.*

Claim 3.4. *The following holds for all $\tau^* \in \mathsf{Good}$ and $\hat{v}^* \in \{0,1\}$:*

$$\left| \Pr_{\mathsf{Exp}(0)}[\hat{v} = \hat{v}^* \mid \tau = \tau^*] - \Pr_{\mathsf{Exp}(1)}[\hat{v} = \hat{v}^* \mid \tau = \tau^*] \right| \leq \delta.$$

These claims suffice to prove the theorem, since the statistical difference between the view of $\mathcal{R}^*$ when the sender commits to 0 (i.e., $b = 0$) and the view of $\mathcal{R}^*$ when the sender commits to 1 (i.e., $b = 1$) may be bounded as:

$$\begin{aligned}
&\frac{1}{2} \sum_{\tau^*,\hat{v}^*} \left| \Pr_{\mathsf{Exp}(0)}[(\tau,\hat{v}) = (\tau^*,\hat{v}^*)] \; - \Pr_{\mathsf{Exp}(1)}[(\tau,\hat{v}) = (\tau^*,\hat{v}^*)] \right| \\
&= \quad \frac{1}{2} \sum_{\tau^*,\hat{v}^*} \left| \Pr_{\mathsf{Exp}(0)}[\tau = \tau^*] \Pr_{\mathsf{Exp}(0)}[\hat{v} = \hat{v}^* | \tau = \tau^*] \; - \Pr_{\mathsf{Exp}(1)}[\tau = \tau^*] \Pr_{\mathsf{Exp}(1)}[\hat{v} = \hat{v}^* | \tau = \tau^*] \right| \\
&\leq \quad \Pr_{\mathsf{Exp}}[\tau \notin \mathsf{Good}] + \frac{1}{2} \sum_{\tau^* \in \mathsf{Good},\hat{v}^*} \Pr_{\mathsf{Exp}}[\tau = \tau^*] \cdot \left| \Pr_{\mathsf{Exp}(0)}[\hat{v} = \hat{v}^* | \tau = \tau^*] \; - \Pr_{\mathsf{Exp}(1)}[\hat{v} = \hat{v}^* | \tau = \tau^*] \right| \\
&\leq \quad \alpha(2+\delta) + \frac{1}{2} \sum_{\tau^* \in \mathsf{Good};\hat{v}^*} \Pr[\tau = \tau^*] \cdot \delta \quad \leq \quad \alpha(2+\delta) + \delta.
\end{aligned}$$

We now prove the two stated claims. Let $\mathsf{Bad} \subset \{0,1\}^\ell$ be a subset whose existence is guaranteed by Definition 3.1 (using the fact that f is balanced). Recall that the initial transcript τ defines two strings $y_0^\tau, y_1^\tau \in \{0,1\}^\ell$ (cf. Construction 2.5). We say $\tau \in \mathsf{Good}$ if and only if $y_0^\tau, y_1^\tau \notin \mathsf{Bad}$.

We first bound the probability that $y_v = y$ is in Bad (recall that y_v is the value that the sender starts with; cf. Construction 2.5). Since f is (α,δ)-balanced and y is distributed according to $f(U_n)$ (cf. Construction 2.6), it follows immediately that $y_v \in \mathsf{Bad}$ with probability at most α.

Next, we bound the probability that $y_v \notin \mathsf{Bad}$ but $y_{\bar{v}} \in \mathsf{Bad}$. Since $\mathcal{R}^*$ is deterministic, we have that $y_{\bar{v}}$ is uniquely determined by y_v. Let ϕ be the function mapping the sender's chosen value y_v to the second value $y_{\bar{v}}$ resulting from the interactive hashing protocol. Let $\mathsf{MapToBad} \stackrel{\text{def}}{=} \phi^{-1}(\mathsf{Bad})$. Observe that if $\phi(y) = y'$ then $\phi(y') = y$; this is because, for either of these choices, the sender responds with the exact same answer to each of the receiver's queries during the interactive hashing

sub-protocol. It follows that ϕ is a permutation and $|\mathsf{MapToBad}| = |\mathsf{Bad}|$. We have:

$$\begin{aligned}\Pr\left[y_v \notin \mathsf{Bad} \bigwedge y_{\bar{v}} \in \mathsf{Bad}\right] &= \Pr\left[y_v \in \mathsf{MapToBad} \setminus \mathsf{Bad}\right] \\ &= \sum_{y^* \in \mathsf{MapToBad}\setminus\mathsf{Bad}} \Pr\left[y_v = y^*\right] \\ &\leq \sum_{y^* \in \mathsf{MapToBad}\setminus\mathsf{Bad}} \frac{1+\delta}{2^\ell},\end{aligned}$$

using condition 3 of Definition 3.1 and the fact that $y^* \notin \mathsf{Bad}$. Continuing:

$$\begin{aligned}\sum_{y^* \in \mathsf{MapToBad}\setminus\mathsf{Bad}} \frac{1+\delta}{2^\ell} &= |\mathsf{MapToBad} \setminus \mathsf{Bad}| \cdot \frac{1+\delta}{2^\ell} \\ &\leq |\mathsf{MapToBad}| \cdot \frac{1+\delta}{2^\ell} \\ &= |\mathsf{Bad}| \cdot \frac{1+\delta}{2^\ell} \;\leq\; \alpha \cdot (1+\delta), \qquad (1)\end{aligned}$$

using condition 1 of Definition 3.1. It follows that $\tau \notin \mathsf{Good}$ with probability at most $\alpha \cdot (2+\delta)$, completing the proof of the first claim.

We proceed to prove the second claim. Let $P(\tilde{y}) \stackrel{\text{def}}{=} \Pr_{x\in\{0,1\}^n}[f(x) = \tilde{y}]$. For any τ^* and any $\hat{v}^* \in \{0,1\}$ we have

$$\begin{aligned}\Pr_{\mathsf{Exp}(b)}[\hat{v} = \hat{v}^* \mid \tau = \tau^*] &= \Pr_{\mathsf{Exp}(b)}[v = \hat{v}^* \oplus b \mid \tau = \tau^*] \\ &= \Pr_{\mathsf{Exp}(b)}[y = y^{\tau^*}_{\hat{v}^*\oplus b} \mid \tau = \tau^*] \\ &= \frac{P(y^{\tau^*}_{\hat{v}^*\oplus b})}{P(y^{\tau^*}_0) + P(y^{\tau^*}_1)}.\end{aligned}$$

If $\tau^* \in \mathsf{Good}$, then $y^{\tau^*}_0, y^{\tau^*}_1 \notin \mathsf{Bad}$ and so $P(y^{\tau^*}_0), P(y^{\tau^*}_1)$ lie in the range $[(1-\delta)2^{-\ell}, (1+\delta)2^{-\ell}]$. It follows that when $\tau^* \in \mathsf{Good}$ the following holds for any $\hat{v}^* \in \{0,1\}$:

$$\left|\Pr_{\mathsf{Exp}(0)}[\hat{v} = \hat{v}^* \mid \tau = \tau^*] - \Pr_{\mathsf{Exp}(1)}[\hat{v} = \hat{v}^* \mid \tau = \tau^*]\right| = \frac{\left|P(y^{\tau^*}_0) - P(y^{\tau^*}_1)\right|}{P(y^{\tau^*}_0) + P(y^{\tau^*}_1)} \leq \delta,$$

which proves the claim. This completes the proof of Theorem 3.2. □

Combining the above and Theorem 2.8 we obtain:

Corollary 3.5. *If $\mathcal{F}$ is (α,δ)-balanced function family and one-way over its range, then Construction 2.6 is a $(2\alpha+\delta+\alpha\delta)$-secure commitment scheme.*

We see that if $(2\alpha+\delta+\alpha\delta) = 1 - \frac{1}{\mathsf{poly}(k)}$, then we obtain a "weakly hiding" commitment scheme. This statistical difference can be amplified to give a statistically-hiding scheme (i.e., an ε-hiding scheme for negligible ε) using polynomially-many sequential repetitions (an appropriate sequential composition theorem is easy to prove). In Section 5 we prove a parallel composition theorem which also enables amplification of the statistical hiding property using polynomially-many repetitions (and without increasing the round complexity).

In the following section, we show how to construct an $\mathcal{F}$ with the required properties starting from any $\mathcal{F}$ which is one-way and approximately regular. Applying the observation at the end of Section 2.2, we thus obtain a construction of a statistically-secure commitment scheme from any approximable pre-image-size one-way function.

4 Starting from Approximately-Regular One-Way Functions

As discussed at the very end of the previous section, we show here that given an $(r(k), \mathsf{poly}(k))$-approximately-regular one-way function family $\mathcal{F}$ and an arbitrary constant $\delta \in (0,1)$, it is possible to construct a $(2^{-k}, \delta)$-balanced function family $\mathcal{F}'$ that is also one-way over its range. The construction is as follows:

Construction 4.1. *Let $\mathcal{F} = \{f_k : \{0,1\}^{n(k)} \to \{0,1\}^{\ell(k)}\}_{k\in\mathbb{N}}$ be a family of functions, and let $\mathcal{H} = \{H_k\}$ be a $3k$-universal hash family where each $h \in H_k$ maps strings of length $\ell(k)$ to strings of length $t(k)$, and each such h can be described using $s(k)$ bits. Define*

$$\mathcal{F}' = \left\{ f'_k : \{0,1\}^{s(k)+n(k)} \to \{0,1\}^{s(k)+t(k)} \right\}_{k\in\mathbb{N}}$$

via $f'_k(h,x) = (h, h(f_k(x)))$.

The main result of this section is the following.

Theorem 4.2. *Let $\mathcal{F}$ be an $(r(k), p(k))$-approximately-regular one-way function family, and let $\delta \in (0,1)$ be an arbitrary constant. Set $c = 6\ln 2/\delta^2$, and say there exists a constant γ such that*

$$n(k) - r(k) - \gamma \cdot \log k \leq t(k) \leq n(k) - r(k) - \log p(k) - \log(ck)$$

for all sufficiently-large k. Then $\mathcal{F}'$ as in Construction 4.1 is $(2^{-k}, \delta)$-balanced and one-way over its range.

Using Corollary 3.5 (and the comments following it) in combination with the above, we obtain our main theorem:

Theorem 4.3 (Main Theorem). *If there exists an approximately-regular one-way function family then there exists a statistically-secure commitment scheme.*

We prove Theorem 4.2 in two parts. In Section 4.1 we show that $\mathcal{F}'$ is $(2^{-k}, \delta)$-balanced, and in Section 4.2 we prove that it is one-way over its range.

4.1 Showing that $\mathcal{F}'$ is Balanced

We begin by showing that $\mathcal{F}'$ is $(2^{-k}, \delta)$-balanced. In this proof, we only use the upper bound on $t(k)$, and will prove a more general statement that only relies on the min-entropy of $f(U_n)$.

Lemma 4.4. *Let $\delta \in (0,1)$, $c = 6\ln 2/\delta^2$, and $k \geq 2$. Let H be a $3k$-universal hash family where each $h \in H_k$ maps strings of length ℓ to strings of length t. Then for any distribution Z on $\{0,1\}^\ell$ with $H_\infty(Z) \geq t + \log(ck)$, the distribution $D \stackrel{\text{def}}{=} \{(h, h(z))\}_{h\leftarrow H, z\leftarrow Z}$ is $(2^{-k}, \delta)$-balanced.*

Note that it follows that $\mathcal{F}'$ (as in Theorem 4.2) is $(2^{-k}, \delta)$-balanced, as the distribution $f_k(U_{n(k)})$ has min-entropy at least $n(k) - r(k) - \log p(k) \geq t(k) + \log(ck)$.

Proof. For any $z \in \{0,1\}^\ell$ and $y \in \{0,1\}^t$, define the random variable $X_{z,y}$ (over choice of $h \in H$) as follows:

$$X_{z,y} \stackrel{\text{def}}{=} \begin{cases} 2^{t+\log(ck)} \cdot \Pr_Z[z] & \text{if } h(z) = y \\ 0 & \text{otherwise} \end{cases} .$$

Note that $X_{z,y} \in [0,1]$ since Z has min-entropy at least $t + \log(ck)$. Let $X_y \stackrel{\text{def}}{=} \sum_{z \in \{0,1\}^\ell} X_{z,y}$. For any z, y we have $E[X_{z,y}] = \Pr_{h \leftarrow H}[h(z) = y] \cdot 2^{t+\log(ck)} \cdot \Pr_Z[z] = 2^{-t} \cdot 2^{t+\log(ck)} \cdot \Pr_Z[z] = ck \cdot \Pr_Z[z]$. It follows that

$$\mu \stackrel{\text{def}}{=} E[X_y] = \sum_z E[X_{z,y}] = ck.$$

Furthermore, since H is a $3k$-universal hash family, the random variables $\{X_{z,y}\}_{z \in \{0,1\}^\ell}$ are $3k$-wise independent for arbitrary y. By Lemma 2.4 we thus have, for any y,

$$\Pr_{h \leftarrow H}\Big[\big|X_y - ck\big| \geq \delta ck\Big] \leq e^{-\lfloor \delta^2 ck/3 \rfloor} < 2^{-k}. \tag{2}$$

(Note that $\lfloor \mu \delta^2 e^{-1/3} \rfloor \leq ck\delta^2 e^{-1/3} = 6ke^{-1/3} \ln 2 < 3k$.)

Define $\phi(h,y) \stackrel{\text{def}}{=} 2^{t+\log(ck)} \cdot \sum_{z:h(z)=y} \Pr_Z[z]$ and $\mathsf{Bad} = \{(h,y) : |\phi(h,y) - ck| \geq \delta ck\}$. We show that, setting $\alpha = 2^{-k}$, the set Bad satisfies the three requirements of Definition 3.1. (Note that we are now considering a distribution over $H \times \{0,1\}^t$.) Using the fact that $\phi(h,y) = 2^{t+\log(ck)} \Pr_{z \leftarrow Z}[h(z) = y]$, observe that

$$\begin{aligned} |\mathsf{Bad}| &= \sum_y |H| \cdot \Pr_{h \leftarrow H}[(h,y) \in \mathsf{Bad}] \\ &= \sum_y |H| \cdot \Pr_{h \leftarrow H}\left[\left|2^{t+\log(ck)} \cdot \Pr_{z \leftarrow Z}[h(z) = y] - ck\right| \geq \delta ck\right] \\ &\leq \left(2^t \cdot |H|\right) \cdot 2^{-k}, \end{aligned}$$

using Eq. (2) and the fact that, once h is chosen, $X_y = 2^{t+\log(ck)} \cdot \Pr_{z \leftarrow Z}[h(z) = y]$. This proves property 1 of Definition 3.1.

We move on to property 2. For each $\xi, z \in \{0,1\}^\ell$, define the random variable

$$R_{z,\xi} = \begin{cases} 2^{t+\log(ck)} \cdot \Pr_Z[z] & \text{if } h(z) = h(\xi) \\ 0 & \text{otherwise} \end{cases}.$$

Again, $R_{z,\xi} \in [0,1]$. Let $R_\xi \stackrel{\text{def}}{=} \sum_{z \in \{0,1\}^\ell} R_{z,\xi}$. For an arbitrary $z \in \{0,1\}^\ell \setminus \{\xi\}$ we have $E[R_{z,\xi}] = 2^{-t} \cdot 2^{t+\log(ck)} \cdot \Pr_Z[z] = ck \Pr_Z[z]$; also $R_{\xi,\xi} = 2^{t+\log(ck)} \Pr_Z[\xi]$ with probability 1. It follows that

$$\mu' \stackrel{\text{def}}{=} E[R_\xi] = \sum_z E[R_{z,\xi}] = ck + (2^{t+\log(ck)} - ck)\Pr_Z[\xi]$$

for any ξ. Note that $ck \leq \mu' < ck + 1$. Furthermore, since H is a $3k$-universal hash family, the random variables $\{R_{z,\xi}\}_{z \in \{0,1\}^\ell}$ are $(3k-1)$-wise independent for any ξ. Thus, by Lemma 2.4 we have

$$\Pr_{h \leftarrow H}\left[|R_\xi - \mu'| \geq \frac{3}{4}\delta\mu'\right] \leq e^{-\lfloor 3\delta^2 \mu'/16 \rfloor} \leq 2^{-k}. \tag{3}$$

(Note that $\mu' \left(\frac{3\delta}{4}\right)^2 e^{-1/3} \leq (ck+1)\frac{9}{16}\delta^2 e^{-1/3} \leq 3k - 1$.)

We then derive:

$$\begin{aligned}\Pr_{(h,y)\leftarrow D}[(h,y)\in \mathsf{Bad}] &= \sum_{\xi\in\{0,1\}^\ell} \Pr_Z[\xi]\cdot \Pr_{h\leftarrow H}\left[\left|\phi(h,h(\xi))-ck\right|\geq \delta ck\right]\\ &= \sum_{\xi\in\{0,1\}^\ell} \Pr_Z[\xi]\cdot \Pr_{h\leftarrow H}\left[\left|R_\xi-ck\right|\geq \delta ck\right]\\ &\leq \sum_{\xi\in\{0,1\}^\ell} \Pr_Z[\xi]\cdot \Pr_{h\leftarrow H}\left[\left|R_\xi-E[R_\xi]\right|\geq \tfrac{3}{4}\delta E[R_\xi]\right]\leq 2^{-k},\end{aligned}$$

where the first inequality uses the stated bounds on μ' and the fact that, once h is chosen, $R_\xi = 2^{t+\log(ck)}\cdot \Pr_{z\leftarrow Z}[h(z)=h(\xi)]$, while the second inequality uses Eq. (3). This gives property 2.

Property 3 holds, since for any (h_0,y_0) we have

$$\Pr_{(h,y)\leftarrow D}[(h,y)=(h_0,y_0)] = \Pr_{h\leftarrow H}[h=h_0]\cdot \sum_{z:h_0(z)=y_0}\Pr_Z[z] = \frac{\phi(h_0,y_0)}{|H|\cdot 2^{t+\log(ck)}}.$$

If $(h_0,y_0)\notin \mathsf{Bad}$, this probability is in the range

$$(1\pm\delta)\frac{ck}{|H|\cdot 2^{t+\log(ck)}} = (1\pm\delta)\frac{1}{|H|\cdot 2^t}$$

as needed. □

4.2 Showing that $\mathcal{F}'$ is One-Way Over Its Range

We now show that $\mathcal{F}'$ is one-way over its range (assuming $\mathcal{F}$ is one-way in the standard sense). Here, we use the lower bound on t as well as the fact that $\mathcal{F}'$ is balanced.

Lemma 4.5. *Let $\mathcal{F}=\{f_k:\{0,1\}^{n(k)}\to\{0,1\}^{\ell(k)}\}$ be an $(r(k),p(k))$-approximately-regular one-way function family, and take $\mathcal{F}'$ as in Construction 4.1. If there exist constants $\gamma>0$ and $\delta\in(0,1)$ such that (1) $\mathcal{F}'$ is $(2^{-k},\delta)$-balanced and (2) $t(k)\geq n(k)-r(k)-\gamma\log k$ for all sufficiently-large k, then $\mathcal{F}'$ is one-way over its range.*

Proof. Recall that Construction 4.1 defines $f'_k(h,x)=(h,h(f_k(x))$. Toward establishing that $\mathcal{F}'$ is one way over its range, we first prove that $\mathcal{F}'$ is one way (according to the standard definition). Given a PPT A', its advantage in inverting $\mathcal{F}'$ is given by

$$\mathsf{Adv}_{A',\mathcal{F}'}(k) \stackrel{\text{def}}{=} \Pr[h\leftarrow H_k; x\leftarrow\{0,1\}^{n(k)}; y=h(f_k(x)); x'\leftarrow A'(1^k,h,y) : h(f_k(x'))=y]. \quad (4)$$

(The above implicitly assumes that $A'(1^k,h,y)$ would never output (h',x') with $h'\neq h$; this is without loss of generality since A' can always be modified accordingly without decreasing its advantage.) To avoid visual clutter, we write f and H in place of f_k, H_k from now on.

Construct a PPT adversary A (attempting to invert $\mathcal{F}$) as follows:

$\underline{A(1^k,z)}$
Choose $h\in H_k$ at random, and set $y=h(z)$
Run $A'(1^k,h,y)$ and obtain output x'
Output x'

Note that the distribution over the inputs of A' in the above experiment is identical to the distribution over the inputs of A' in Equation (4). We have:

$$\begin{aligned}
\mathsf{Adv}_{A,\mathcal{F}}(k) &\stackrel{\text{def}}{=} \Pr[x \leftarrow \{0,1\}^{n(k)}; z = f(x); x' \leftarrow A(1^k, z) : f(x') = z] \\
&= \sum_{\hat{x}',\hat{h},\hat{y}} \Pr\left[h \leftarrow H; x \leftarrow \{0,1\}^{n(k)} : \begin{array}{c} h = \hat{h} \bigwedge \hat{h}(f(x)) = \hat{y} \bigwedge \\ A'(1^k, \hat{h}, \hat{y}) = \hat{x}' \bigwedge f(\hat{x}') = f(x) \end{array}\right] \\
&= \sum_{\hat{x}',\hat{h},\hat{y}} \Pr\left[h \leftarrow H; x \leftarrow \{0,1\}^{n(k)} : \begin{array}{c} h = \hat{h} \bigwedge \hat{h}(f(x)) = \hat{y} \bigwedge \\ A'(1^k, \hat{h}, \hat{y}) = \hat{x}' \bigwedge \hat{h}(f(x)) = \hat{h}(f(\hat{x}')) \end{array}\right] \cdot \\
&\quad \Pr\left[x \leftarrow \{0,1\}^{n(k)} : f(x) = f(\hat{x}') \mid \hat{h}(f(x)) = \hat{y} = \hat{h}(f(\hat{x}'))\right]. \qquad (5)
\end{aligned}$$

It is easy to see that

$$\begin{aligned}
\Pr\left[x \leftarrow \{0,1\}^{n(k)} : f(x) = f(\hat{x}') \mid \hat{h}(f(x)) = \hat{y} = \hat{h}(f(\hat{x}'))\right] &= \frac{|\{x : f(x) = f(\hat{x}')\}|}{\left|\{x : \hat{h}(f(x)) = \hat{y}\}\right|} \\
&\geq \frac{2^{r(k)}/p(k)}{\left|\{x : \hat{h}(f(x)) = \hat{y}\}\right|}. \qquad (6)
\end{aligned}$$

By assumption, $\mathcal{F}'$ is $(2^{-k}, \delta)$-balanced; let Bad be as required by Definition 3.1. Restricting summation to $(\hat{h}, \hat{y}) \notin \mathsf{Bad}$ in Eq. (5) and using Eq. (6) we obtain

$$\begin{aligned}
&\mathsf{Adv}_{A,\mathcal{F}}(k) \\
&\geq \sum_{\hat{x}'} \sum_{(\hat{h},\hat{y}) \notin \mathsf{Bad}} |H|^{-1} \cdot \Pr\left[x \leftarrow \{0,1\}^{n(k)} : \begin{array}{c} \hat{h}(f(x)) = \hat{y} = \hat{h}(f(\hat{x}')) \bigwedge \\ A'(1^k, \hat{h}, \hat{y}) = \hat{x}' \end{array}\right] \cdot \frac{2^{r(k)}/p(k)}{\left|\{x : \hat{h}(f(x)) = \hat{y}\}\right|}.
\end{aligned}$$

For $(\hat{h}, \hat{y}) \notin \mathsf{Bad}$, however, we have:

$$\begin{aligned}
\frac{1}{|H|} \cdot \frac{\left|\{x : \hat{h}(f(x)) = \hat{y}\}\right|}{2^{n(k)}} &= \Pr[h \leftarrow H; x \leftarrow \{0,1\}^{n(k)} : h = \hat{h} \bigwedge h(f(x)) = \hat{y}] \\
&\leq (1+\delta)\frac{1}{|H|\, 2^{t(k)}},
\end{aligned}$$

so that $\left|\{x : \hat{h}(f(x)) = \hat{y}\}\right| \leq (1+\delta)2^{n(k)-t(k)}$. Therefore,

$$\begin{aligned}
&\mathsf{Adv}_{A,\mathcal{F}}(k) \\
&\geq \sum_{\hat{x}'} \sum_{(\hat{h},\hat{y}) \notin \mathsf{Bad}} |H|^{-1} \cdot \Pr\left[x \leftarrow \{0,1\}^{n(k)} : \begin{array}{c} \hat{h}(f(x)) = \hat{y} = \hat{h}(f(\hat{x}')) \bigwedge \\ A'(1^k, \hat{h}, \hat{y}) = \hat{x}' \end{array}\right] \cdot \frac{2^{r(k)}/p(k)}{(1+\delta)2^{n(k)-t(k)}} \\
&= \frac{2^{r(k)}}{p(k)(1+\delta)2^{n(k)-t(k)}} \cdot \Pr\left[\begin{array}{c} h \leftarrow H; x \leftarrow \{0,1\}^{n(k)}; \\ y = h(f_k(x)); x' \leftarrow A'(1^k, h, y) \end{array} : h(f_k(x')) = y \bigwedge (h, y) \notin \mathsf{Bad}\right] \\
&\geq \frac{2^{r(k)}}{p(k)(1+\delta)2^{n(k)-t(k)}} \cdot \left(\mathsf{Adv}_{A',\mathcal{F}'}(k) - \Pr\left[h \leftarrow H; x \leftarrow \{0,1\}^{n(k)} : (h, h(f(x)) \in \mathsf{Bad}\right]\right) \\
&\geq \frac{2^{r(k)}}{p(k)(1+\delta)2^{n(k)-t(k)}} \cdot \left(\mathsf{Adv}_{A',\mathcal{F}'}(k) - 2^{-k}\right),
\end{aligned}$$

using property 2 in Definition 3.1 for the final inequality.

Now, since $t(k) \geq n(k) - r(k) - \gamma \log k$,

$$\frac{2^{r(k)}}{(1+\delta)p(k)2^{n(k)-t(k)}} \geq \frac{1}{(1+\delta)\cdot p(k)\cdot k^{\gamma}} = 1/\mathsf{poly}(k).$$

Since $\mathsf{Adv}_{A,\mathcal{F}}(k)$ is negligible by assumption, it must be the case that $\mathsf{Adv}_{A',\mathcal{F}'}(k)$ is negligible as well and thus $\mathcal{F}'$ is one way.

We now show that $\mathcal{F}'$ is one-way over its range. Consider any PPT algorithm A'' inverting $\mathcal{F}'$ "over its range." The advantage of A'' (in this sense) is given by:

$$\begin{aligned}\mathsf{Adv}^{*}_{A'',\mathcal{F}'}(k) &\stackrel{\text{def}}{=} \Pr[h \leftarrow H; y \leftarrow \{0,1\}^{t(k)}; x' \leftarrow A''(1^k, h, y) : h(f(x')) = y] \\ &= \frac{1}{|H|\cdot 2^{t(k)}} \cdot \sum_{h\in H_k} \sum_{y\in\{0,1\}^{t(k)}} \Pr[A'' \text{ inverts } (h,y)],\end{aligned}$$

where "A'' inverts (h,y)" has the obvious meaning.

Consider now the advantage of A'' in inverting $\mathcal{F}'$ in the standard sense:

$$\begin{aligned}\mathsf{Adv}_{A'',\mathcal{F}'}(k) &\stackrel{\text{def}}{=} \Pr[h \leftarrow H; x \leftarrow \{0,1\}^{n(k)} : A'' \text{ inverts } (h, h(f(x)))] \\ &= \frac{1}{|H|\cdot 2^{n(k)}} \sum_{h\in H} \sum_{x\in\{0,1\}^{n(k)}} \Pr[A'' \text{ inverts } (h, h(f(x)))] \\ &= \frac{1}{|H|\cdot 2^{n(k)}} \sum_{h\in H} \sum_{z\in\mathsf{image}(f)} \left|f^{-1}(z)\right| \cdot \Pr[A'' \text{ inverts } (h, h(z))] \\ &\geq \frac{1}{|H|\cdot 2^{n(k)}} \sum_{h\in H} \sum_{z\in\mathsf{image}(f)} \frac{2^{r(k)}}{p(k)} \cdot \Pr[A'' \text{ inverts } (h, h(z))] \\ &= \frac{2^{r(k)}}{p(k)\cdot|H|\cdot 2^{n(k)}} \sum_{h\in H} \sum_{y\in\mathsf{image}(h(f))} \sum_{z\in h^{-1}(y)\cap\mathsf{image}(f)} \Pr[A'' \text{ inverts } (h, h(z))] \\ &\geq \frac{2^{r(k)}}{p(k)\cdot|H|\cdot 2^{n(k)}} \sum_{h\in H} \sum_{y\in\mathsf{image}(h(f))} \Pr[A'' \text{ inverts } (h, y)] \\ &= \frac{2^{r(k)}}{p(k)\cdot|H|\cdot 2^{n(k)}} \sum_{h\in H} \sum_{y\in\{0,1\}^{t(k)}} \Pr[A'' \text{ inverts } (h, y)] \\ &= \frac{2^{r(k)}\cdot 2^{t(k)}}{p(k)\cdot 2^{n(k)}} \cdot \mathsf{Adv}^{*}_{A'',\mathcal{F}'}(k) \;\geq\; \frac{\mathsf{Adv}^{*}_{A'',\mathcal{F}'}(k)}{p(k)\cdot k^{\gamma}}.\end{aligned}$$

Since $\mathsf{Adv}_{A'',\mathcal{F}'}$ is negligible (by one-wayness of $\mathcal{F}'$), $\mathsf{Adv}^{*}_{A'',\mathcal{F}'}$ is negligible as well. This completes the proof that $\mathcal{F}'$ is one-way over its range, and thus completes the proof of the lemma. □

This completes the proof of our main result (Theorem 4.3. As discussed in the introduction, in the following two sections we explore additional questions related to statistically-hiding commitment.

5 Parallel Repetition of Commitments

In this section, we prove a parallel repetition theorem for the case of statistically-hiding commitments. Though such a result appears obvious, and might be considered folklore, a proof is non-trivial and we are unaware of any such proof in the literature.

We first define formally the notion of parallel repetition we consider:

Construction 5.1 (Parallel Repetition). *Let $(\mathcal{S}, \mathcal{R}, \mathcal{V})$ be a commitment scheme and $q = \mathsf{poly}(k)$. Construct commitment scheme $(\mathcal{S}_q, \mathcal{R}_q, \mathcal{V}_q)$ as follows.*

- *On input a bit b, $\mathcal{S}_q$ chooses q bits $b_1, \ldots, b_q$ uniformly at random subject to $b = \bigoplus_{i=1}^{q} b_i$. It then runs (in parallel) q instances of $\mathcal{S}$, where the i^{th} instance commits to b_i. The output of $\mathcal{S}_q$ is $\mathsf{decom} = (\mathsf{decom}_1, \ldots, \mathsf{decom}_q)$, where decom_i is the output of the i^{th} instance of $\mathcal{S}$.*
- *$\mathcal{R}_q$ runs (in parallel) q instances of $\mathcal{R}$. The output of $\mathcal{R}_q$ is $\mathsf{com} = (\mathsf{com}_1, \ldots, \mathsf{com}_q)$, where com_i is the output of the i^{th} instance of $\mathcal{R}$.*
- *$\mathcal{V}_q$, on input $\mathsf{com} = (\mathsf{com}_1, \ldots, \mathsf{com}_q)$ and $\mathsf{decom} = (\mathsf{decom}_1, \ldots, \mathsf{decom}_q)$, computes $b_i = \mathcal{V}(\mathsf{com}_i, \mathsf{decom}_i)$ for all i. If $b_i = \perp$ for any i, $\mathcal{V}_q$ outputs $\perp$; otherwise, it outputs $\bigoplus_{i=1}^{q} b_i$.*

We are now ready to state the result.

Theorem 5.2. *Let $(\mathcal{S}, \mathcal{R}, \mathcal{V})$ and $(\mathcal{S}_q, \mathcal{R}_q, \mathcal{V}_q)$ be as in Construction 5.1. If $(\mathcal{S}, \mathcal{R}, \mathcal{V})$ is a ρ-secure commitment scheme, then $(\mathcal{S}_q, \mathcal{R}_q, \mathcal{V}_q)$ is a ρ^q-secure commitment scheme.*

A straightforward hybrid argument (omitted here) shows that if if $(\mathcal{S}, \mathcal{R}, \mathcal{V})$ is computationally binding then so is $(\mathcal{S}_q, \mathcal{R}_q, \mathcal{V}_q)$. The interesting part of the theorem is that if $(\mathcal{S}, \mathcal{R}, \mathcal{V})$ is ρ-hiding then $(\mathcal{S}_q, \mathcal{R}_q, \mathcal{V}_q)$ is ρ^q-hiding. Although seemingly obvious, it is not easy to prove: the difficulty is that the views of the receiver in the different instances of the basic scheme $(\mathcal{S}, \mathcal{R}, \mathcal{V})$ are not necessarily independent, since a malicious receiver can correlate its messages in each of these executions.[4]

In our proof, we rely on the ideas used to prove an analogous parallel repetition theorem for the *soundness error* in interactive proof systems [18, Appendix C]. As in that case, Theorem 5.2 follows immediately from the following lemma.

Lemma 5.3. *Let $\mathsf{Com}_1 = (\mathcal{S}_1, \mathcal{R}_1, \mathcal{V}_1)$ and $\mathsf{Com}_2 = (\mathcal{S}_2, \mathcal{R}_2, \mathcal{V}_2)$ be two commitment schemes, and construct $\mathsf{Com} = (\mathcal{S}, \mathcal{R}, \mathcal{V})$ by parallel composition of these schemes (in the obvious way, as in Construction 5.1). If Com_1 is ρ_1-hiding and Com_2 is ρ_2-hiding, then Com is $(\rho_1\rho_2)$-hiding.*

Proof. Assume each of the component commitment schemes has an r-round commitment phase (this can always be ensured by sending "dummy messages" as needed), and assume without loss of generality that $\mathcal{S}_1$ (resp., $\mathcal{S}_2$) sends the first and last message. Following the proof of [18, Lemma C.1], we employ the notion of a *game tree* T, defined for any commitment scheme as follows:

- The root of T is a node at level 0 denoted ε. This corresponds to the beginning of an execution of the scheme.
- Each node v at an *even* level ℓ corresponds to a point when the *honest* sender makes a move. This node has children at level $\ell + 1$ corresponding to all possible (legal) messages of the sender (i.e., for all possible random coins and either possible input bit).

[4] This difficulty goes away if either (1) we use *sequential* repetition; or (2) the receiver is *honest-but-curious*. In either of these cases, a parallel composition theorem is easy to prove.

- Each node v at an *odd* level $\ell < r$ corresponds to a point when the malicious receiver makes a move. This node has children at level $\ell + 1$ corresponding to all possible messages of the receiver.

We identify a node v with the partial view of the receiver up to that point. As a special case, if v is a leaf then v corresponds to a possible view of the receiver when the commitment scheme is run to completion. For a node v, we let $\mathsf{ch}(v)$ denote the set of children of v.

Let T_1, T_2, T denote the game trees for $\mathsf{Com}_1, \mathsf{Com}_2, \mathsf{Com}$, respectively. Let $\mathcal{C}_1$ denote the size of the space of random coins for $\mathcal{S}_1$ (i.e., if $\mathcal{S}_1$ uses s_1 random coins, then $\mathcal{C}_1 = 2^{s_1}$), and define $\mathcal{C}_2, \mathcal{C}$ analogously. For a partial transcript v_1 of an execution of Com_1, let $C_1(v_1; b)$ denote the size of the set of random coins for $\mathcal{S}_1$ consistent with input bit b and v_1; $C_2(v_2; b)$ and $C(v; b)$ are defined analogously.

We now define a *value* val for each node in a game tree. We focus on game tree T corresponding to Com, but the value $\mathsf{val}_1, \mathsf{val}_2$ of a node in T_1, T_2 is defined analogously. The value of a node in T is defined inductively:

- If v is a leaf, then $\mathsf{val}(v) \stackrel{\text{def}}{=} \frac{|C(v;0) - C(v;1)|}{2\mathcal{C}}$.
- If v is an even node, then $\mathsf{val}(v) \stackrel{\text{def}}{=} \sum_{w \in \mathsf{ch}(v)} \mathsf{val}(w)$.
- If v is an odd node, then $\mathsf{val}(v) \stackrel{\text{def}}{=} \max_{w \in \mathsf{ch}(v)} \mathsf{val}(w)$.

If v is a partial transcript and $\widehat{\mathsf{view}}$ is a full transcript, we say that $\widehat{\mathsf{view}}$ *is consistent with* v if v is a prefix of $\widehat{\mathsf{view}}$. It is not hard to see that for any (unbounded) receiver $\mathcal{R}^*$ and every node v:

$$\mathsf{val}(v) \geq \frac{1}{2} \cdot \sum_{\widehat{\mathsf{view}} \text{ consistent with } v} \left| \Pr[\mathsf{view}_{\langle \mathcal{S}(b), \mathcal{R}^* \rangle} = \widehat{\mathsf{view}}] - \Pr[\mathsf{view}_{\langle \mathcal{S}(1), \mathcal{R}^* \rangle} = \widehat{\mathsf{view}}] \right|.$$

Furthermore, there exists an unbounded receiver $\mathcal{R}^*$ for which equality holds. By the assumption of the lemma, then, we have $\mathsf{val}_1(\varepsilon) \leq \rho_1$ and $\mathsf{val}_2(\varepsilon) \leq \rho_2$. Moreover, we can prove the theorem by showing that $\mathsf{val}(\varepsilon) \leq \rho_1 \rho_2$.

Note that every node (i.e., partial view) v in T corresponds in the natural way to a tuple $(v_1, v_2) \in T_1 \times T_2$ (in particular, ε corresponds to $(\varepsilon, \varepsilon)$). From now on, we write $\mathsf{val}(v_1, v_2)$ in place of $\mathsf{val}(v)$. The desired bound on $\mathsf{val}(\varepsilon, \varepsilon)$ is immediate from the following, more general claim

Claim 5.4. *For all v_1, v_2:*

$$\mathsf{val}(v_1, v_2) = \mathsf{val}_1(v_1) \cdot \mathsf{val}_2(v_2).$$

Proof. We prove this by induction on the level of the tree, starting from the bottom. The base case occurs when $v = (v_1, v_2)$ is a leaf in T. By construction of Com, we have $\mathcal{C} = 2 \cdot \mathcal{C}_1 \cdot \mathcal{C}_2$ (the sender $\mathcal{S}$ uses an additional random bit to determine the inputs to $\mathcal{S}_1, \mathcal{S}_2$) and furthermore

$$C(v; 0) = C_1(v_1; 0) \cdot C_2(v_2; 0) + C_1(v_1; 1) \cdot C_2(v_2; 1)$$

and

$$C(v; 1) = C_1(v_1; 0) \cdot C_2(v_2; 1) + C_1(v_1; 1) \cdot C_2(v_2; 0).$$

So:

$$\begin{aligned}\frac{|C(v;0)-C(v;1)|}{2\mathcal{C}} &= \frac{|C_1(v_1;0)-C_1(v_1;1)|\cdot|C_2(v_2;0)-C_2(v_2;1)|}{2\cdot 2\cdot \mathcal{C}_1,\cdot\mathcal{C}_2}\\ &= \frac{|C_1(v_1;0)-C_1(v_1;1)|}{2\mathcal{C}_1}\cdot\frac{|C_2(v_2;0)-C_2(v_2;1)|}{2\mathcal{C}_2}\\ &= \mathsf{val}_1(v_1)\cdot\mathsf{val}_2(v_2).\end{aligned}$$

Now, assume the claim is true for all nodes at level $\ell+1$ and consider a node $v=(v_1,v_2)$ at level ℓ. There are two case. If ℓ is odd, then

$$\begin{aligned}\mathsf{val}(v) &= \max_{w\in\mathsf{ch}(v)}\mathsf{val}(w)\\ &= \max_{w_1\in\mathsf{ch}(v_1),w_2\in\mathsf{ch}(v_2)}\mathsf{val}(w_1,w_2)\\ &= \max_{w_1\in\mathsf{ch}(v_1),w_2\in\mathsf{ch}(v_2)}\mathsf{val}_1(w_1)\cdot\mathsf{val}_2(w_2)\\ &= \left(\max_{w_1\in\mathsf{ch}(v_1)}\mathsf{val}_1(w_1)\right)\cdot\left(\max_{w_2\in\mathsf{ch}(v_2)}\mathsf{val}_2(w_2)\right)\\ &= \mathsf{val}_1(v_1)\cdot\mathsf{val}_2(v_2).\end{aligned}$$

If, on the other hand, ℓ is even:

$$\begin{aligned}\mathsf{val}(v) &= \sum_{w\in\mathsf{ch}(v)}\mathsf{val}(w)\\ &= \sum_{w_1\in\mathsf{ch}(v_1),w_2\in\mathsf{ch}(v_2)}\mathsf{val}_1(w_1)\cdot\mathsf{val}_2(w_2)\\ &= \left(\sum_{w_1\in\mathsf{ch}(v_1)}\mathsf{val}_1(w_1)\right)\cdot\left(\sum_{w_2\in\mathsf{ch}(v_2)}\mathsf{val}_2(w_2)\right)\\ &= \mathsf{val}_1(v_1)\cdot\mathsf{val}_2(v_2).\end{aligned}$$

This completes the proof of the claim, and hence the theorem. □

□

6 Honest-but-Curious vs. Malicious Receivers

In this section, we demonstrate a *compiler* that converts any commitment scheme that is statistically-secure against an *honest-but-curious* (i.e., semi-honest) receiver into a statistically-secure commitment scheme (i.e., where hiding holds even against a *malicious* receiver). Our compiler requires only the existence of one-way functions, which are implied by the existence of a statistically-secure commitment scheme against an honest-but-curious receiver. Thus, we obtain:

Theorem 6.1. *The existence of a statistically-secure commitment scheme against an honest-but-curious receiver implies the existence of a statistically-secure commitment scheme.*

Our compiler increases the round complexity of the original scheme by a constant number of rounds plus the number of rounds needed for a zero-knowledge proof system (with negligible soundness

error) for NP. Currently, the best known constructions of the latter based on one-way functions require $\omega(1)$ rounds [17]. (In particular, the constant-round construction of [13] requires a statistically-hiding commitment scheme, the very primitive we are trying to construct.)

For completeness, we provide a definition of security against a semi-honest receiver; such a definition follows easily from Definitions 2.1 and 2.3.

Definition 6.2. Commitment scheme $(\mathcal{S}, \mathcal{R}, \mathcal{V})$ is *ρ-hiding against an honest-but-curious receiver* if the following holds: Let $\mathsf{view}_{\langle \mathcal{S}(b), \mathcal{R} \rangle}(k)$ be the view of an *honest* receiver $\mathcal{R}$. Then we require that the ensembles $\{\mathsf{view}_{\langle \mathcal{S}(0), \mathcal{R} \rangle}(k)\}$ and $\{\mathsf{view}_{\langle \mathcal{S}(1), \mathcal{R} \rangle}(k)\}$ have statistical difference at most ρ.

A commitment scheme is *statistically hiding against an honest-but-curious receiver* if it is ρ-hiding for an honest-but-curious receiver with ρ negligible. $(\mathcal{S}, \mathcal{R}, \mathcal{V})$ is statistically-secure against an honest-but-curious receiver if it is statistically-hiding against an honest-but-curious receiver and computationally-binding (as in Definition 2.2).

We remark that since we consider the view of the honest receiver $\mathcal{R}$, we must allow $\mathcal{R}$ to be probabilistic. The view of $\mathcal{R}$ consists of its random coins as well as the messages sent by $\mathcal{S}$.

Our compiler uses a coin-tossing protocol and zero-knowledge (ZK) proofs (in a way similar to [16]) to "force" honest behavior on the part of the receiver. However, we do not require "simulatable" coin-tossing (as in [2, 16, 24]) or ZK proofs of correctness following each round (as in [16]); instead, we show that a weaker variant of coin-tossing along with a single ZK proof at the end suffice. (The latter in particular is essential for obtaining a round-efficient compiler.)

Informally, our compiler proceeds as follows: the receiver first uses a statistically-*binding* commitment scheme (e.g., [25]) to commit to a sufficiently-long string r_1, and the sender responds with a string r_2 of the same length. The sender and receiver then execute the original protocol, with the receiver using $r_1 \oplus r_2$ as its random tape and the sender committing to a *random* bit b'. At the conclusion of the original protocol, the receiver uses a ZK proof to show that each of the messages it sent during the course of the protocol is consistent with the messages sent by $\mathcal{S}$ as well as the random tape $r_1 \oplus r_2$ (we stress that r_1 is never revealed to $\mathcal{S}$). Finally (assuming $\mathcal{S}$ accepts the proof), $\mathcal{S}$ concludes the protocol by sending $b' \oplus b$ (where b is the bit that $\mathcal{S}$ wants to commit to).

Before giving a formal description and proof of security for our compiler, some brief remarks are in order. First, note that one-way functions are sufficient for both statistically-binding commitment [25] as well as zero-knowledge proofs for all of NP [17, 25]. Second, since we have the receiver provide a ZK proof of correctness only at the conclusion of the protocol we must take into account the fact that the receiver may cheat during the course of the protocol, learn some information about the bit committed to by $\mathcal{S}$, and then abort. To prevent this, we have $\mathcal{S}$ commit to a *random* bit b' in the initial phase of the protocol (i.e., before the ZK proof); the only portion of the transcript that depends on the input bit of $\mathcal{S}$ is sent *after* the receiver successfully proves correctness of its actions. We now provide a detailed description of our compiler.

Construction 6.3. *Let $(\mathcal{S}, \mathcal{R}, \mathcal{V})$ be a commitment scheme, and construct commitment scheme $(\mathcal{S}^*, \mathcal{R}^*, \mathcal{V}^*)$ as follows: The sender $\mathcal{S}^*$ with input bit b interacts with the receiver $\mathcal{R}^*$ as follows:*

1. *Let $\ell = \ell(k)$ denote the length of the random tape used by $\mathcal{R}$. Then $\mathcal{R}^*$ uses a (possibly interactive) statistically-binding commitment scheme to commit to a random string $r_1 \in \{0,1\}^\ell$. Let* com *denote the resulting commitment (known to both $\mathcal{S}^*$ and $\mathcal{R}^*$) and let* decom *be the decommitment (known to $\mathcal{R}^*$).*

2. *$\mathcal{S}^*$ sends a random string $r_2 \in \{0,1\}^\ell$. This defines a string $r \stackrel{\text{def}}{=} r_1 \oplus r_2$ which is known to $\mathcal{R}^*$ (but not to $\mathcal{S}^*$).*

3. *$\mathcal{S}^*$ chooses a random bit b', and then $\mathcal{S}^*$ and $\mathcal{R}^*$ run protocols $\mathcal{S}(b')$ and $\mathcal{R}$, respectively, where the latter is run using random tape r. At the conclusion of this stage, $\mathcal{S}^*$ obtains decom' as output from $\mathcal{S}$, while $\mathcal{R}^*$ obtains com' as output from $\mathcal{R}$.*

4. *$\mathcal{R}^*$ provides a ZK proof (with negligible soundness error) that it acted correctly throughout the previous stage. Formally, $\mathcal{R}^*$ proves that there exists (decom, r_1) such that com is a commitment to r_1 (with decommitment decom) and all the messages sent by $\mathcal{R}^*$ in the previous stage are consistent with the messages sent by $\mathcal{S}^*$ and the random tape $r = r_1 \oplus r_2$.*

5. *If $\mathcal{S}^*$ rejects the proof given by $\mathcal{R}^*$, it aborts. Otherwise, $\mathcal{S}^*$ sends $\hat{b} = b \oplus b'$ and outputs decom'; the receiver $\mathcal{R}^*$ outputs $(\mathsf{com}', \hat{b})$.*

In the decommitment phase, $\mathcal{V}^*$ proceeds as follows: it runs $\mathcal{V}(\mathsf{com}', \mathsf{decom}')$ to obtain a bit b' (if the output of $\mathcal{V}$ is $\perp$, then $\mathcal{V}^*$ outputs $\perp$ as well), and then outputs $\hat{b} \oplus b'$.

Theorem 6.4. *If $(\mathcal{S}, \mathcal{R}, \mathcal{V})$ is a statistically-secure commitment scheme against an honest-but-curious receiver, then $(\mathcal{S}^*, \mathcal{R}^*, \mathcal{V}^*)$ as generated by the above compiler is a statistically-secure commitment scheme (i.e., even against a malicious receiver).*

In proving the theorem, we consider the hiding and binding properties individually.

Claim 6.5. *If $(\mathcal{S}, \mathcal{R}, \mathcal{V})$ is ρ-hiding against an honest-but-curious receiver, then $(\mathcal{S}^*, \mathcal{R}^*, \mathcal{V}^*)$ as in Construction 6.3 is $(\rho + \varepsilon)$-hiding for a negligible function ε.*

Proof. Let $\mathcal{R}^{**}$ denote an unbounded malicious receiver who interacts with $\mathcal{S}^*$, and assume that $\mathcal{R}^{**}$ is deterministic without loss of generality. We say a transcript of an execution of $\mathcal{S}^*$ with $\mathcal{R}^{**}$ is *non-aborting* if $\mathcal{S}^*$ sends the final bit in the protocol; i.e., $\mathcal{R}^{**}$ gave a successful ZK proof that it acted correctly. (We say it is *aborting* otherwise.) We say a transcript is *good* if (1) the commitment com in the transcript indeed commits $\mathcal{R}^{**}$ to a single value r_1; and (2) $\mathcal{R}^{**}$ indeed acted correctly in its execution of the sub-routine $\mathcal{R}$; that is, each message sent by $\mathcal{R}^{**}$ in this transcript is consistent with $r_1 \oplus r_2$ (note that r_1 is uniquely defined by (1), and r_2 is explicit in the transcript). (We say that it is *bad* otherwise.) Note that the probablity of a transcript that is bad and non-aborting is negligible.

The statistical difference between distributions $\mathsf{view}_{\langle \mathcal{S}^*(0), \mathcal{R}_1^{**} \rangle}(k)$ and $\mathsf{view}_{\langle \mathcal{S}^*(1), \mathcal{R}_1^{**} \rangle}(k)$ is

$$\mathsf{SD}^*(k) \stackrel{\text{def}}{=} \frac{1}{2} \cdot \sum_{\mathsf{view}} |\Pr^*_{b=0}[\mathsf{view}] - \Pr^*_{b=1}[\mathsf{view}]| \,, \tag{7}$$

where $\Pr^*_{b=0}[\cdot]$ denotes the probability taken over coins of the sender when its input bit is 0 (the case of $b = 1$ is defined analogously). When view is aborting, $\Pr^*_{b=0}[\mathsf{view}] = \Pr^*_{b=1}[\mathsf{view}]$ (since only the final message depends on the input bit). On the other hand, as we have alredy noted, the probability of obtaing a bad and non-aborting view is negligible. It follows that

$$\overline{\mathsf{SD}}^* \stackrel{\text{def}}{=} \frac{1}{2} \cdot \sum_{\substack{\mathsf{view} \text{ non-aborting} \\ \text{and good}}} |\Pr^*_{b=0}[\mathsf{view}] - \Pr^*_{b=1}[\mathsf{view}]|$$

is negligibly close to $\mathsf{SD}^*(k)$.

Any non-aborting view can be parsed as an initial portion view' and the bit $\hat{b}$ sent in the final round. Note

$$\begin{aligned}\overline{\mathsf{SD}}^*(k) &= \frac{1}{2} \cdot \sum_{\substack{\mathsf{view}' \text{ non-aborting} \\ \text{and good}}} \sum_{\hat{b}} \left| \Pr^*_{b=0}[\mathsf{view}' \| \hat{b}] - \Pr^*_{b=1}[\mathsf{view}' \| \hat{b}] \right| \\ &= \frac{1}{2} \cdot \sum_{\substack{\mathsf{view}' \text{ non-aborting} \\ \text{and good}}} \left| \Pr^*_{b'=0}[\mathsf{view}'] - \Pr^*_{b'=1}[\mathsf{view}'] \right| . \qquad (8)\end{aligned}$$

(The notions of 'non-aborting' and 'good' depend only on the initial portion, and so are meaningfully defined above.) Above, $\Pr_{b'=0}[\mathsf{view}']$ denotes the probability of view' conditioned on the random bit b' of the sender being equal to 0 (with the case $b' = 1$ defined analogously). Note that $\Pr_{b'=0}[\mathsf{view}']$ is independent of the input bit b.

We show the existence of a randomized (but not polynomial-time) procedure ψ, outputting either a partial transcript or $\perp$, with the following property. Let $D(b)$ be the distribution defined by $\psi\left(\mathsf{view}_{\langle \mathcal{S}(b), \mathcal{R} \rangle}(k)\right)$ (i.e., ψ applied to a randomly-generated view of $\mathcal{R}$ interacting with $\mathcal{S}(b)$), and let $\Pr_{D(b)}[\mathsf{view}']$ be the probability of view' with respect to distribution $D(b)$. Then if view' is non-aborting and good

$$\Pr_{D(b)}[\mathsf{view}'] = \Pr_{b'=b}[\mathsf{view}'], \qquad (9)$$

while for any other view' we have $\Pr_{D(b)}[\mathsf{view}'] = 0$. That is, any view output by ψ is good and non-aborting (and is furthermore output with probability as in Eq. (9)), and ψ outputs $\perp$ otherwise. The statistical difference between $D(0)$ and $D(1)$ is exactly given by Eq. (8); on the other hand, since the statistical difference between $\mathsf{view}_{\langle \mathcal{S}(b), \mathcal{R} \rangle}(k)$ and $\mathsf{view}_{\langle \mathcal{S}(1), \mathcal{R} \rangle}(k)$ is at most ρ (as is guaranteed by the ρ-hiding of $(\mathcal{S}, \mathcal{R}, \mathcal{V})$ against an honest-but-curious receiver), this statistical difference can be at most ρ. We conclude that $\overline{\mathsf{SD}}^*(k) \leq \rho(k)$, and so $\mathsf{SD}^*(k) \leq \rho(k) + \varepsilon(k)$ for negligible ε.

It remains to show ψ. Procedure ψ, on input a tuple $(m_1, \ldots, m_i, r)$ (where $m_1, \ldots, m_i$ denote the messages of the sender $\mathcal{S}$ and r denotes the random coins used by honest-but-curious $\mathcal{R}$), proceeds as follows:

1. Begin interacting with $\mathcal{R}^{**}$, simulating an execution of $\mathcal{S}^*$.
2. When ψ obtains a commitment com from $\mathcal{R}^{**}$, it computes (in exponential time) a unique string r_1 consistent with this commitment. In case no such r_1 exists or multiple such r_1 exist, ψ outputs $\perp$. (Note that in the first case, the ZK proof of $\mathcal{R}^{**}$ will fail with all but negligible probabilit; the second case occurs with negligible probability by statistical binding of the commitment scheme used here.)
3. ψ_0 sends the string $r_2 = r \oplus r_1$ to $\mathcal{R}^{**}$.
4. ψ_0 interacts with $\mathcal{R}^{**}$ by sending messages $m_1, \ldots, m_r$ in response to the messages of $\mathcal{R}^{**}$. If $\mathcal{R}^{**}$ ever sends a message inconsistent with random tape r and these messages, ψ outputs $\perp$.
5. Finally, ψ acts as a verifier in an execution of the ZK proof with $\mathcal{R}^{**}$. If the proof succeeds, then ψ_0 outputs the entire view view' to this point. If the proof fails, ψ outputs $\perp$.

It is immediate that ψ never outputs a view' that is bad or aborting. It is also not hard to see that for any view' that is non-aborting and good, the probability that view' is output by $\psi\left(\mathsf{view}_{\langle \mathcal{S}(b),\mathcal{R}\rangle}(k)\right)$ is exactly $\Pr^*_{b'=b}[\mathsf{view}']$. The claim follows. □

We next consider the binding property.

Claim 6.6. *If $\Pi = (\mathcal{S}, \mathcal{R}, \mathcal{V})$ is computationally-binding, then $\Pi^* = (\mathcal{S}^*, \mathcal{R}^*, \mathcal{V}^*)$ as in Construction 6.3 is computationally-binding as well.*

Proof. Given a PPT sender $\hat{\mathcal{S}}^*$ violating the binding property of Π^* with non-negligible probability, we construct a PPT sender $\hat{\mathcal{S}}$ violating the binding property of Π with non-negligible probability. $\hat{\mathcal{S}}$ is defined as follows:

1. $\hat{\mathcal{S}}$ interacts with an honest receiver $\mathcal{R}$, and runs a copy of $\hat{\mathcal{S}}^*$ internally. $\hat{\mathcal{S}}$ begins by sending a random commitment to the string $r_1 = 0^\ell$ to $\hat{\mathcal{S}}^*$, who responds with a string $r_2 \in \{0,1\}^\ell$.
2. $\hat{\mathcal{S}}$ then relays messages faithfully between $\mathcal{R}$ and $\hat{\mathcal{S}}^*$. At the conclusion of this phase, no more messages are sent to $\mathcal{R}$.
3. Finally, $\hat{\mathcal{S}}$ simulates a ZK proof of correct behavior with $\hat{\mathcal{S}}^*$ acting as the potentially-dishonest verifier. ($\hat{\mathcal{S}}^*$ then sends a final bit, which $\hat{\mathcal{S}}$ ignores.)
4. If $\hat{\mathcal{S}}^*$ outputs valid decommitments to two different bits, then $\hat{\mathcal{S}}$ does so as well.

To complete the proof, we argue that the probability that $\hat{\mathcal{S}}^*$ outputs two valid decommitments in its interaction with $\hat{\mathcal{S}}$, above, is negligibly-close to the probability that $\hat{\mathcal{S}}^*$ outputs two valid decommitments when interacting with an honest receiver $\mathcal{R}^*$. This is relatively straightforward to show via a hybrid argument, and we only sketch the proof. Consider a sequence of experiments, and let $\Pr_i[\mathsf{NoBind}]$ denote the probability that $\hat{\mathcal{S}}^*$ outputs two valid decommitments in experiment i):

Experiment 0. This is the original experiment, where $\hat{\mathcal{S}}^*$ interacts with $\mathcal{R}^*$.

Experiment 1. Here, we have $\mathcal{R}^*$ act exactly as in Experiment 0, except that it simulates the final ZK proof of correct behavior. By the ZK property of the proof system (against computationally-bounded verifiers), $|\Pr_0[\mathsf{NoBind}] - \Pr_1[\mathsf{NoBind}]|$ is negligible.

Experiment 2. Now, we have $\mathcal{R}^*$ act as in the previous experiment, except that its initial commitment is to 0^ℓ rather than to a random $r_1 \in \{0,1\}^\ell$. (However, it uses random tape $r_1 \oplus r_2$ in computing its messages to send.) Computational hiding of the commitment scheme implies that $|\Pr_2[\mathsf{NoBind}] - \Pr_1[\mathsf{NoBind}]|$ is negligible.

Experiment 2 corresponds exactly to an interaction of $\hat{\mathcal{S}}^*$ with $\hat{\mathcal{S}}$. □

Acknowledgments

We are grateful to Yan Zong Ding, Virgil Gligor, Oded Goldreich, Danny Harnik, Omer Reingold, and Alon Rosen for helpful conversations and for reading preliminary versions of this manuscript. We thank the anonymous referees for helpful comments that improved the presentation.

References

[1] M. Bellare and S. Micali. How to sign given any trapdoor permutation. *J. ACM*, 39(1):214–233, 1992.

[2] M. Blum. Coin flipping by phone. In *IEEE COMPCOM*, 1982.

[3] M. Blum and S. Micali. How to generate cryptographically-strong sequences of pseudorandom bits. *SIAM J. Computing*, 13(4):850–864, 1984.

[4] J.F. Boyar, S.A. Kurtz, and M.W. Krentel. Discrete logarithm implementation of perfect zero-knowledge blobs. *Journal of Cryptology*, 2(2):63–76, 1990.

[5] G. Brassard, D. Chaum, and C. Crépeau. Minimum disclosure proofs of knowledge. *J. Computer and System Sciences*, 37(2):156–189, 1988.

[6] J.L. Carter and M.N. Wegman. Universal classes of hash functions. *J. Computer and System Sciences*, 18(2):143–154, 1979.

[7] I. Damgård, T. Pedersen, and B. Pfitzmann. On the existence of statistically-hiding bit commitment and fail-stop signatures. In *Crypto*, 1993.

[8] O. Goldreich. *Foundations of Cryptography, vol. 1: Basic Tools*. Cambridge University Press, 2001.

[9] O. Goldreich. *Foundations of Cryptography, vol. 2: Basic Applications*. Cambridge University Press, 2004.

[10] O. Goldreich, S. Goldwasser, and S. Micali. On the cryptographic applications of random functions. In *Crypto '84*.

[11] O. Goldreich, S. Goldwasser, and S. Micali. How to construct random functions. *J. ACM*, 33(4):792–807, 1986.

[12] O. Goldreich, R. Impagliazzo, L. Levin, R. Venkatesan, and D. Zuckerman. Security preserving amplification of hardness. In *FOCS*, 1990.

[13] O. Goldreich and A. Kahan. How to construct constant-round zero-knowledge proof systems for NP. *Journal of Cryptology*, 9(3):167–190, 1996.

[14] O. Goldreich, H. Krawczyk, and M. Luby. On the existence of pseudorandom generators. *SIAM J. Computing*, 22(6):1163–1175, 1993.

[15] O. Goldreich and L.A. Levin. Hard-core predicates for any one-way function. In *STOC*, 1989.

[16] O. Goldreich, S. Micali, and A. Widgerson. How to play any mental game — a completeness theorem for protocols with honest majority. In *19th STOC*, 1987.

[17] O. Goldreich, S. Micali, and A. Widgerson. Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. *J. ACM*, 38(1):691–729, 1991.

[18] Oded Goldreich. *Modern cryptography, probabilistic proofs and pseudorandomness*. Springer-Verlag, Berlin, 1999.

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

[20] S. Halevi and S. Micali. Practical and provably-secure commitment schemes from collision-free hashing. In *Crypto*, 1996.

[21] J. Håstad, R. Impagliazzo, L.A. Levin, and M. Luby. A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28(4):1364–1396, 1999.

[22] R. Impagliazzo and M. Luby. One-way functions are essential for complexity-based cryptography. In *FOCS*, 1989.

[23] R. Impagliazzo and S. Rudich. Limits on the provable consequences of one-way permutations. In *STOC*, 1989.

[24] Y. Lindell. Parallel coin-tossing and constant-round secure two-party computation. *Journal of Cryptology*, 16(3):143–184, 2003.

[25] M. Naor. Bit commitment using pseudorandomness. *Journal of Cryptology*, 4(2):151–158, 1991.

[26] M. Naor, R. Ostrovsky, R. Venkatesan, and M. Yung. Perfect zero-knowledge arguments for NP using any one-way permutation. *J. Crypto.*, 11(2):87–108, 1998.

[27] M. Naor and M. Yung. Universal one-way hash functions and their cryptographic application. In *STOC*, 1989.

[28] R. Ostrovsky, R. Venkatesan, and M. Yung. Secure commitment against a powerful adversary. In *STACS*, 1992.

[29] R. Ostrovsky, R. Venkatesan, and M. Yung. Fair games against an all-powerful adversary. In *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, volume 13, 1993.

[30] J. Rompel. One-way functions are necessary and sufficient for secure signatures. In *STOC*, 1990.

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

[32] A. De Santis and M. Yung. On the design of provably-secure cryptographic hash functions. In *Eurocrypt*, 1990.

[33] J.P. Schmidt, A. Siegel, and A. Srinivasan. Chernoff-Hoeffding bounds for applications with limited independence. *SIAM J. Discrete Math.*, 8(2):223–250, 1995.

[34] A. C.-C. Yao. Theory and application of trapdoor functions. In *23rd IEEE Symposium on Foundations of Computer Science*, pages 80–91, 1982.