

A Linear Lower Bound on the Communication Complexity of Single-Server Private Information Retrieval

Iftach Haitner*

Jonathan J. Hoch*

Gil Segev*

Abstract

We study the communication complexity of single-server Private Information Retrieval (PIR) protocols that are based on fundamental cryptographic primitives in a black-box manner. In this setting, we establish a tight lower bound on the number of bits communicated by the server in any polynomially-preserving construction that relies on trapdoor permutations. More specifically, our main result states that in such constructions $\Omega(n)$ bits must be communicated by the server, where n is the size of the server's database. This improves the $\Omega\left(\frac{n}{\log n}\right)$ lower bound due to Haitner et al. (FOCS '07). Therefore, in the very natural setting under consideration, the naive solution in which the user downloads the entire database turns out to be optimal up to constant multiplicative factors. Moreover, while single-server PIR protocols with poly-logarithmic communication complexity were shown to exist based on specific number-theoretic assumptions, the lower bound we provide identifies a substantial gap between black-box and non-black-box constructions of single-server PIR.

Technically speaking, this paper consists of two main contributions from which our lower bound is obtained. First, we derive a tight lower bound on the number of bits communicated by the sender during the commit stage of any black-box construction of a statistically-hiding bit-commitment scheme from a family of trapdoor permutations. This lower bound asymptotically matches the upper bound provided by the scheme of Naor, Ostrovsky, Venkatesan and Yung (CRYPTO '92). Second, we significantly improve the efficiency of the reduction of statistically-hiding commitment schemes to non-trivial single-server PIR, due to Beimel, Ishai, Kushilevitz and Malkin (STOC '99). In particular, we present a reduction that essentially preserves the communication complexity of the underlying single-server PIR protocol.

*Department of Computer Science and Applied Mathematics, Weizmann Institute of Science, Rehovot 76100, Israel. Email: {iftach.haitner,yaakov.hoch,gil.segev}@weizmann.ac.il.

1 Introduction

A single-server Private Information Retrieval (PIR) scheme is a protocol between a server and a user. The server holds a database $x \in \{0, 1\}^n$ and the user holds an index $i \in [n]$ to an entry of the database. Very informally, the user wishes to retrieve the i^{th} entry of the database, without revealing the index i to the server. The notion of PIR was introduced by Chor, Goldreich, Kushilevitz and Sudan [4] to model applications that enable users to query public databases without revealing any information on the specific data that the users wish to retrieve. Chor et al. showed that in the information-theoretic setting any single-server PIR protocol has the server communicating at least n bits. Therefore in this setting the naive solution in which the user downloads the entire database is optimal.

Kushilevitz and Ostrovsky [23] were the first to construct a non-trivial single-server PIR protocol relying on computational assumptions. Their result initiated a sequence of papers showing that there exist single-server PIR protocols with poly-logarithmic communication complexity based on *specific* number-theoretic assumptions (see, for example, [2, 3, 10, 25, 23, 35], and a recent survey by Ostrovsky and Skeith [29]). The only non-trivial construction based on *general* computational assumptions is due to Kushilevitz and Ostrovsky [24], and is based on trapdoor permutations. In their construction, however, the server is required to communicate $n - o(n)$ bits to the user.

Motivated by this ever-growing line of work, we study the communication complexity of single-server PIR protocols that are based on fundamental cryptographic primitives. We establish a linear lower bound on the number of bits communicated by the server in such constructions that rely on trapdoor permutations in a black-box manner. Therefore, in the very natural setting under consideration in this paper, the naive solution in which the user downloads the entire database turns out to be optimal up to constant multiplicative factors. In the following paragraphs, we briefly describe the setting in which our lower bound is proved.

Black-box reductions. As previously mentioned, under widely believed number-theoretic assumptions, there are very efficient single-server PIR protocols. Therefore, if any of these specific assumptions holds, the existence of trapdoor permutations implies the existence of very efficient single-server PIR protocols in a trivial sense. Faced with similar difficulties, Impagliazzo and Rudich [19] presented a paradigm for proving impossibility results under a restricted, yet very natural and important, subclass of reductions called *black-box reductions*. Informally, a black-box reduction of a primitive P to a primitive Q is a construction of P out of Q that ignores the internal structure of the implementation of Q and merely uses it as a “subroutine” (i.e., as a black-box). In addition, in the case of fully-black-box reductions [30], the proof of security (showing that an adversary that breaks the implementation of P implies an adversary that breaks the implementation of Q), is black-box as well, that is, the internal structure of the adversary that breaks the implementation of P is ignored.

The strength of cryptographic reductions. Luby [26] provides a classification of the strength of cryptographic reductions into three classes: linearly-preserving reductions, polynomially-preserving reductions and weakly-preserving reductions. In our setting, this classification comes into play when comparing the size of the server’s database and the domain of the trapdoor permutations. Very informally, a reduction of single-server PIR for an n -bit database to a family of trapdoor permutations is linearly-preserving or polynomially-preserving if it uses trapdoor permutations over $\Omega(n)$ bits. Such a reduction is weakly-preserving if it uses trapdoor permutations over $\Omega(n^\epsilon)$ bits for some constant $0 < \epsilon \leq 1$. In linearly-preserving and polynomially-preserving reductions we are guaranteed that breaking the constructed primitive is essentially as hard as breaking the underlying primitive. However, in weakly-preserving reductions, we are only guaranteed that breaking the constructed primitive is as hard as breaking the underlying primitive for polynomially smaller security parameters. We refer the reader to [26] for a more comprehensive and complete discussion.

1.1 Related Work

Non-trivial single-server PIR is one of the fundamental primitives in the foundations of cryptography. For example, it was shown to imply the existence of Oblivious Transfer protocols [5, 28], and any non-interactive non-trivial single-server PIR was shown to imply collision-resistance hash functions [20]. In addition, it was shown to be tightly related to several other aspects of cryptography and complexity theory (see, for example, [6, 17, 21]). As it is not the goal of the current paper to make justice with the importance and applications of single-server PIR, the reader is referred to a recent in-depth survey [29] for a more comprehensive discussion.

In the context of black-box reductions, Impagliazzo and Rudich [19] showed that there are no black-box reductions of key-agreement protocols to one-way permutations, and substantial additional work in this line followed (see, for example, [11, 12, 31, 33]). Kim, Simon and Tetali [22] initiated a new line of impossibility results, by providing a lower bound on the *efficiency* of black-box reductions (rather than on their feasibility). They proved a lower bound on the efficiency, in terms of the number of calls to the underlying primitive, of any black-box reduction of universal one-way hash functions to one-way permutations. This result was later improved, to match the known upper bound, by Gennaro and Trevisan [9], which together with Gennaro et al. [8] provided tight lower bounds on the efficiency of several other black-box reductions. Building upon the technique developed by [9], Horvitz and Katz [18] provided lower bounds on the efficiency of black-box reductions of statistically-hiding and computationally-binding commitment schemes to one-way permutations. In all the above results the measure of efficiency under consideration is the number of calls to the underlying primitives.

Very recently, Haitner et al. [16], improving upon previous works [7, 36], proved that any fully-black-box reduction of a statistically-hiding bit-commitment scheme to trapdoor permutations has $\Omega\left(\frac{n}{\log n}\right)$ communication rounds (where n is the security parameter of the scheme). As a corollary, they showed that any polynomially-preserving fully-black-box reduction of single-server PIR to trapdoor permutations has $\Omega\left(\frac{n}{\log n}\right)$ communication rounds, where n is the size of the server's database. In particular, the server is required to communicate $\Omega\left(\frac{n}{\log n}\right)$ bits to the user. Haitner et al. also establish similar lower bounds on the communication complexity of oblivious transfer that guarantees statistical security for one of the parties and for interactive hashing.

1.2 Our Results

We study the class of black-box constructions of single-server PIR from trapdoor permutations, and establish a tight lower bound on the number of bits communicated by the server in such constructions. Our main result is the following:

Main Theorem (Informal). *In any polynomially-preserving fully-black-box construction of a single-server PIR protocol from a family of trapdoor permutations the server communicates $\Omega(n)$ bits, where n is the size of the server's database.*

Our lower bound holds for constructions which are polynomially-preserving. We note that the construction of Kushilevitz and Ostrovsky [24], which is based on trapdoor permutations in an fully-black-box manner and in which the server communicates $n - o(n)$ bits, is only weakly-preserving (i.e., it is significantly easier to break their protocol than to break the security of the underlying family of trapdoor permutations¹). Thus, the question of whether a tight linear lower bound can be established for weakly-preserving constructions as well remains open.

The main technical contributions. This paper consists of two main contributions from which our lower bound is immediately obtained. First, we derive a tight lower bound on the communication complexity of

¹Thought the security guarantees of the two primitives are still polynomially-related.

black-box constructions of statistically-hiding bit-commitment schemes from trapdoor permutations. Very recently, Haitner et al. [16] proved that any fully-black-box construction of statistically-hiding bit-commitment scheme from a family of trapdoor permutations has $\Omega\left(\frac{n}{\log n}\right)$ communication rounds, where n is the security parameter of the scheme. In particular, this implies a lower bound on the number of bits communicated by the sender. In this paper we manage to improve their lower bound for the communication complexity. Specifically, we prove the following theorem:

Theorem (Informal) 1.1. *In any polynomially-preserving fully-black-box construction of a statistically-hiding bit-commitment scheme from a family of trapdoor permutations the sender communicates $\Omega(n)$ bits during the commit stage, where n is the security parameter of the scheme.*

This lower bound asymptotically matches the upper bound given by the statistically-hiding commitment scheme of Naor, Ostrovsky, Venkatesan and Yung [27].

In addition, we significantly improve the efficiency of the reduction of statistically-hiding commitment schemes to non-trivial single-server PIR, presented by Beimel, Ishai, Kushilevitz and Malkin [1]. Our reduction essentially preserves both the round complexity and the communication complexity of the underlying single-server PIR protocol. As stating this result turns out to involve subtle technical details, here we only state a very informal statement:

Theorem (Informal) 1.2. *There exists a linearly-preserving fully-black-box reduction of statistically-hiding commitment scheme to single-server PIR, which preserves both the round complexity and the communication complexity of the underlying single-server PIR protocol.*

1.3 Paper Organization

In Section 2 we briefly present the notations and formal definitions used in this paper. In Section 3 we prove our tight lower bound on the number of bits communicated by the sender during the commit stage of statistically-hiding commitment schemes. In Section 4 we describe an improved reduction of statistically-hiding commitment scheme to non-trivial single-server PIR. Finally, in Section 5 we establish the lower bound for single-server PIR by combining our main technical contributions.

2 Preliminaries

We denote by Π_n the set of all permutations over $\{0, 1\}^n$. For an integer n , we denote by U_n the uniform distribution over the set $\{0, 1\}^n$. For a finite set X , we denote by $x \leftarrow X$ the experiment of choosing an element of X according to the uniform distribution. Similarly, for a distribution \mathcal{D} over a set X , we denote by $x \leftarrow \mathcal{D}$ the experiment of choosing an element of X according to the distribution \mathcal{D} . For a distribution \mathcal{D} we denote by $\text{supp}(\mathcal{D})$ set of elements having non-zero probability under \mathcal{D} . The min-entropy of \mathcal{D} is defined as:

$$H_\infty(\mathcal{D}) = \min_{x \in \text{supp}(\mathcal{D})} \left(\log \frac{1}{\Pr_{\mathcal{D}}[x]} \right) .$$

The statistical distance between two distributions X and Y over Ω is denoted $\text{SD}(X, Y)$, and defined as

$$\text{SD}(X, Y) = \frac{1}{2} \sum_{\omega \in \Omega} |\Pr_X[\omega] - \Pr_Y[\omega]| .$$

Definition 2.1. A function $E : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is a (k, ϵ) -extractor if for every distribution X over $\{0, 1\}^n$ with $H_\infty(X) \geq k$ the distribution $E(X, U_d)$ is ϵ -close to uniform. E is a *strong* (k, ϵ) -extractor if the function $E'(x, y) = y \circ E(x, y)$ is a (k, ϵ) -extractor (where \circ denotes concatenation).

In our construction of a statistically-hiding commitment scheme from single-server PIR we will be using the following explicit construction of strong extractors, which is obtained as a corollary of [34, Corollary 3.4].

Proposition 2.2. For any $k \in \omega(\log(n))$, there exists an explicit strong $(k, 2^{1-k})$ -extractor $\text{EXT} : \{0, 1\}^n \times \{0, 1\}^{3k} \rightarrow \{0, 1\}^{k/2}$.

The following standard fact (see, for example [32, Fact 2.6]) will be useful for us in analyzing statistically-close distributions.

Fact 2.3. If X and Y are two distributions such that $\text{SD}(X, Y) < \epsilon$, then with probability at least $1 - 2\sqrt{\epsilon}$ over $x \leftarrow X$ it holds that

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

2.1 One-Way Permutations and Trapdoor Permutations

We briefly present the notions of one-way permutations and trapdoor (one-way) permutations which are used in this paper. For a more comprehensive discussion we refer the reader to [13].

Definition 2.4. A collection of permutations $\pi = \{\pi_n\}_{n=1}^{\infty}$, where $\pi_n \in \Pi_n$ for every n , is $s(n)$ -hard if for every probabilistic Turing-machine A that runs in time $s(n)$, and for all sufficiently large n ,

$$\Pr[A(1^n, y) = \pi_n^{-1}(y)] \leq \frac{1}{s(n)} ,$$

where the probability is taken uniformly over all the possible choices of $y \in \{0, 1\}^n$ and over all the possible outcomes of the internal coin tosses of A .

In our setting, whenever such a collection π is given as an oracle, we denote by A^π a circuit or a Turing-machine A with oracle access to π . In addition, when we consider the probability of an event over the choice of π , we mean that for every integer n , a permutation π_n is chosen uniformly at random from Π_n and independently of all other permutations.

A collection of trapdoor permutations is represented as a triplet $\tau = (G, F, F^{-1})$. Informally, G corresponds to a key generation procedure, which is queried on a string td (intended as the ‘‘trapdoor’’) and produces a corresponding public key pk . The procedure F is the actual permutation, which is queried on a public key pk and an input x . Finally, the procedure F^{-1} is the inverse of F : If $G(td) = pk$ and $F(pk, x) = y$, then $F^{-1}(td, y) = x$. In this paper, since we are concerned with providing a lower bound, we do not consider the most general definition of a collection of trapdoor permutations. Instead, we denote by T_n the set of all triplets $\tau_n = (G_n, F_n, F_n^{-1})$ of the following form:

1. $G_n \in \Pi_n$.
2. $F_n : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ is a function such that $F_n(pk, \cdot) \in \Pi_n$ for every $pk \in \{0, 1\}^n$.
3. $F_n^{-1} : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ is a function such that $F_n^{-1}(td, y)$ returns the unique $x \in \{0, 1\}^n$ for which $F_n(G_n(td), x) = y$.

Our lower bound proof is based on analyzing random instances of such collections. A uniformly distributed $\tau_n \in T_n$ can be chosen as follows: G_n is chosen uniformly at random from Π_n , and for each $pk \in \{0, 1\}^n$ a permutation $F_n(pk, \cdot)$ is chosen uniformly and independently at random from Π_n . As above, we do not consider a single collection τ_n : we consider a family $\tau = \{\tau_n\}_{n=1}^{\infty}$ of collection of trapdoor permutations where $\tau_n \in T_n$ for every n . Whenever such a family τ is given as an oracle, we denote by A^τ a circuit or a Turing-machine A with oracle access to τ . In addition, when we consider the probability of an event over the choice of τ , we mean that for every integer n , a collection of trapdoor permutation τ_n is chosen uniformly at random from T_n and independently of all other collections.

Definition 2.5. A family of trapdoor permutations $\tau = \{\tau_n = (G_n, F_n, F_n^{-1})\}_{n=1}^{\infty}$ is $s(n)$ -hard if for every probabilistic Turing-machine A that runs in time $s(n)$, and for all sufficiently large n ,

$$\Pr \left[A^\tau(1^n, G_n(td), y) = F_n^{-1}(td, y) \right] \leq \frac{1}{s(n)} ,$$

where the probability is taken uniformly over all the possible choices of $td \in \{0, 1\}^n$ and $y \in \{0, 1\}^n$, and over all the possible outcomes of the internal coin tosses of A .

Note that Definition 2.5 refers to the difficulty of inverting a random permutation $F(pk, \cdot)$ on a uniformly distributed image y , when given only $pk = G(td)$ and y . Some applications, however, require enhanced hardness conditions. For example, it may be required (cf. [14, Appendix C]) that it is hard to invert $F(pk, \cdot)$ on y even given the random coins used in the generation of y . Note that our formulation captures such hardness condition as well and therefore the impossibility results proved in this paper hold also for enhanced trapdoor permutations.²

2.2 Single-Server Private Information Retrieval

A single-server Private Information Retrieval (PIR) scheme is a protocol between a server and a user. The server holds a database $x \in \{0, 1\}^n$ and the user holds an index $i \in [n]$ to an entry of the database. Very informally, the user wishes to retrieve the i^{th} entry of the database, without revealing the index i to the server. More formally, a single-server PIR scheme is defined via a pair of probabilistic polynomial-time Turing-machines $(\mathcal{S}, \mathcal{U})$ such that:

- \mathcal{S} receives as input a string $x \in \{0, 1\}^n$. Following its interaction it does not have any output.
- \mathcal{U} receives as input an index $i \in [n]$. Following its interaction it outputs a value $b \in \{0, 1, \perp\}$.

Denote by $b \leftarrow \langle \mathcal{S}(x), \mathcal{U}(i) \rangle$ the experiment in which \mathcal{S} and \mathcal{U} interact (using the given inputs and uniformly chosen random coins), and then \mathcal{U} outputs the value b . It is required that there exists a negligible function $\nu(n)$, such that for all sufficiently large n , and for every string $x = x_1 \circ \dots \circ x_n \in \{0, 1\}^n$, it holds that $x_i \leftarrow \langle \mathcal{S}(x), \mathcal{U}(i) \rangle$ with probability at least $1 - \nu(n)$ over the random coins of both \mathcal{S} and \mathcal{U} .

In order to define the security properties of such schemes, we first introduce the following notation. Given a single-server PIR scheme $(\mathcal{S}, \mathcal{U})$ and a Turing-machine \mathcal{S}^* (a malicious server), we denote by $\text{view}_{\langle \mathcal{S}^*, \mathcal{U}(i) \rangle}(n)$ the distribution on the view of \mathcal{S}^* when interacting with $\mathcal{U}(i)$ where $i \in [n]$. This view consists of its random coins and of the sequence of messages it receives from \mathcal{U} , and the distribution is taken over the random coins of both \mathcal{S}^* and \mathcal{U} .

Definition 2.6. A single-server PIR scheme $(\mathcal{S}, \mathcal{U})$ is secure if for every probabilistic polynomial-time Turing-machines \mathcal{S}^* and \mathcal{D} , and for every two sequences of indices $\{i_n\}_{i=1}^{\infty}$ and $\{j_n\}_{i=1}^{\infty}$ where $i_n, j_n \in [n]$ for every n , it holds that

$$\left| \Pr \left[v \leftarrow \text{view}_{\langle \mathcal{S}^*, \mathcal{U}(i_n) \rangle}(n) : \mathcal{D}(v) = 1 \right] - \Pr \left[v \leftarrow \text{view}_{\langle \mathcal{S}^*, \mathcal{U}(j_n) \rangle}(n) : \mathcal{D}(v) = 1 \right] \right| \leq \nu(n) ,$$

for some negligible function $\nu(k)$ and for all sufficiently large n .

²A different enhancement, used by [15], requires the permutations' domain to be polynomially dense in $\{0, 1\}^n$. Clearly, our impossibility result holds for such an enhancement as well.

2.3 Commitment Schemes

A commitment scheme is a two-stage interactive protocol between a sender and a receiver. Informally, after the first stage of the protocol, which is referred to as the *commit stage*, the sender is bound to at most one value, not yet revealed to the receiver. In the second stage, which is referred to as the *reveal stage*, the sender reveals its committed value to the receiver. More formally, a commitment scheme is defined via a triplet of probabilistic polynomial-time Turing-machines $(\mathcal{S}, \mathcal{R}, \mathcal{V})$ such that:

- \mathcal{S} receives as input the security parameter 1^n and a string $x \in \{0, 1\}^k$. Following its interaction, it outputs some information decom (the decommitment).
- \mathcal{R} receives as input the security parameter 1^n . Following its interaction, it outputs a state information com (the commitment).
- \mathcal{V} (acting as the receiver in the reveal stage³) receives as input the security parameter 1^n , a commitment com and a decommitment decom . It outputs either a string $x' \in \{0, 1\}^k$ or \perp .

Denote by $(\text{decom}|\text{com}) \leftarrow \langle \mathcal{S}(1^n, x), \mathcal{R}(1^n) \rangle$ the experiment in which \mathcal{S} and \mathcal{R} interact (using the given inputs and uniformly chosen random coins), and then \mathcal{S} outputs decom while \mathcal{R} outputs com . It is required that for all n , every string $x \in \{0, 1\}^k$, and every pair $(\text{decom}|\text{com})$ that may be output by $\langle \mathcal{S}(1^n, x), \mathcal{R}(1^n) \rangle$, it holds that $\mathcal{V}(\text{com}, \text{decom}) = x$.⁴ In the remainder of the paper, it will often be convenient for us to identify \mathcal{V} with \mathcal{R} , and refer to a commitment scheme as a pair $(\mathcal{S}, \mathcal{R})$.

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. In this paper, we deal with commitment schemes of the latter type, which are referred to as *statistically-hiding* commitment schemes. In order to define the security properties of such schemes, we first introduce the following notation. Given a commitment scheme $(\mathcal{S}, \mathcal{R})$ and a Turing-machine \mathcal{R}^* , we denote by $\text{view}_{\langle \mathcal{S}(x), \mathcal{R}^* \rangle}(n)$ the distribution on the view of \mathcal{R}^* when interacting with $\mathcal{S}(1^n, x)$. This view consists of \mathcal{R}^* 's random coins and of the sequence of messages it receives from \mathcal{S} . The distribution is taken over the random coins of both \mathcal{S} and \mathcal{R}^* . Note that whenever no computational restrictions are assumed on \mathcal{R}^* , without loss of generality we can assume that \mathcal{R}^* is deterministic.

Definition 2.7. A commitment scheme $(\mathcal{S}, \mathcal{R})$ is $\rho(n)$ -*hiding* if for every deterministic Turing-machine \mathcal{R}^* , and for every two sequences of strings $\{x_n\}_{i=1}^\infty$ and $\{x'_n\}_{i=1}^\infty$ where $x_n, x'_n \in \{0, 1\}^{k(n)}$ for every n the ensembles $\{\text{view}_{\langle \mathcal{S}(x_n), \mathcal{R}^* \rangle}(n)\}$ and $\{\text{view}_{\langle \mathcal{S}(x'_n), \mathcal{R}^* \rangle}(n)\}$ have statistical difference at most $\rho(n)$ for all sufficiently large n . Such a scheme is *statistically-hiding* if it is $\rho(n)$ -hiding for some negligible function $\rho(n)$.

Our lower bound for commitment schemes holds in fact under a weaker hiding requirement. We derive our results even for commitment schemes in which the sender is statistically protected only against an honest receiver. Such schemes are referred to as *statistically-hiding honest-receiver* commitment schemes. Formally, it is only required that the statistical difference between ensembles $\{\text{view}_{\langle \mathcal{S}(x_n), \mathcal{R} \rangle}(n)\}$ and $\{\text{view}_{\langle \mathcal{S}(x'_n), \mathcal{R} \rangle}(n)\}$ is some negligible function of n .

Definition 2.8. A commitment scheme $(\mathcal{S}, \mathcal{R}, \mathcal{V})$ is $\mu(n)$ -*binding* if for every probabilistic polynomial-time Turing-machine \mathcal{S}^* it holds that the probability that $((\text{decom}, \text{decom}')|\text{com}) \leftarrow \langle \mathcal{S}^*(1^n), \mathcal{R}(1^n) \rangle$ (where the probability is over the random coins of both \mathcal{S}^* and \mathcal{R}) such that $\mathcal{V}(\text{com}, \text{decom}) \neq \mathcal{V}(\text{com}, \text{decom}')$ and $\mathcal{V}(\text{com}, \text{decom}), \mathcal{V}(\text{com}, \text{decom}') \neq \perp$ is negligible in n for all sufficiently large n . Such a scheme is *computationally-binding* if it is $\mu(n)$ -binding for some negligible function $\mu(n)$, and is *weakly-binding* if it is $(1 - 1/p(n))$ -binding for some polynomial $p(n)$.

³Note that there is no loss of generality in assuming that the reveal stage is non-interactive. This is since any such interactive stage can be replaced with a non-interactive one as follows: The sender sends its internal state to the receiver, who then simulates the sender in the interactive stage.

⁴Although we assume perfect completeness, it is not essential for our results.

2.4 Black-Box Reductions

A reduction of a primitive P to a primitive Q is a construction of P out of Q . Such a construction consists of showing that if there exists an implementation C of Q , then there exists an implementation M_C of P . This is equivalent to showing that for every adversary that breaks M_C , there exists an adversary that breaks C . Such a reduction is *semi-black-box* if it ignores the internal structure of Q 's implementation, and it is *fully-black-box* if the proof of correctness is black-box as well, i.e., the adversary for breaking Q ignores the internal structure of both Q 's implementation and of the (alleged) adversary breaking P . Semi-black-box reductions are less restricted and thus more powerful than fully-black-box reductions. A taxonomy of black-box reductions was provided by Reingold, Trevisan and Vadhan [30], and the reader is referred to their paper for a more complete and formal view of these notions.

We now formally define the class of constructions considered in this paper. Our results in the current paper are concerned with the particular setting of fully-black-box constructions of single-server PIR and of statistically-hiding commitment schemes from trapdoor permutations. We focus here on specific definitions for these particular primitives and we refer the reader to [30] for a more general definition.

When examining efficiency measures of fully-black-box constructions, an essential parameter for such characterizations, as introduced by Haitner et al. [16], is the *security-parameter-expansion* of the construction. Consider, for example, a fully-black-construction of a commitment scheme from a family of trapdoor permutations. One ingredient of such a construction is a machine A that attempts to break the security of the trapdoor permutation family given oracle access to any malicious sender \mathcal{S}^* that breaks the security of the commitment scheme. Then, A receives a security parameter 1^n (and possibly some additional inputs) and invokes \mathcal{S}^* in a black-box manner. The standard definition does not restrict the range of security parameters that A is allowed to invoke \mathcal{S}^* on. For example, A may invoke \mathcal{S}^* on security parameter 1^{n^2} , or even on security parameter $1^{\Theta(s(n))}$, where $s(n)$ is the running time of A . In this paper, we will use the notion $\ell(n)$ -expanding for short, and note that according to Luby's classification [26], any polynomially-preserving reduction is $O(n)$ -expanding in our terminology.

Definition 2.9. A fully-black-box $\ell(n)$ -expanding construction of a single-server PIR scheme from an $s(n)$ -hard family of trapdoor permutations is a triplet of probabilistic oracle Turing-machines $(\mathcal{S}, \mathcal{U}, A)$ for which the following hold:

1. **Correctness:** For every family τ of trapdoor permutations, $(\mathcal{S}^\tau, \mathcal{U}^\tau)$ is a single-server PIR scheme.
2. **Black-box proof of security:** For every family $\tau = \{\tau_n = (G_n, F_n, F_n^{-1})\}_{n=1}^\infty$ of trapdoor permutations and for every probabilistic polynomial-time Turing-machine \mathcal{S}^* , if \mathcal{S}^* with oracle access to τ breaks the security of $(\mathcal{S}^\tau, \mathcal{U}^\tau)$, then

$$\Pr \left[A^{\tau, \mathcal{S}^*}(1^n, G_n(td), y) = F_n^{-1}(td, y) \right] > \frac{1}{s(n)},$$

for infinitely many values of n , where A runs in time $s(n)$ and invokes \mathcal{S}^* on security parameters which are at most $1^{\ell(n)}$. The probability is taken uniformly over all the possible choices of $td \in \{0, 1\}^n$ and $y \in \{0, 1\}^n$, and over all the possible outcomes of the internal coin tosses of A .

Definition 2.10. A fully-black-box $\ell(n)$ -expanding construction of a weakly-binding and statistically-hiding honest-receiver commitment scheme from an $s(n)$ -hard family of trapdoor permutations is a triplet of probabilistic oracle Turing-machines $(\mathcal{S}, \mathcal{R}, A)$ for which the following hold:

1. **Correctness:** For every family τ of trapdoor permutations, $(\mathcal{S}^\tau, \mathcal{R}^\tau)$ is a statistically-hiding honest-receiver commitment scheme.

2. **Black-box proof of binding:** For every family $\tau = \{\tau_n = (G_n, F_n, F_n^{-1})\}_{n=1}^{\infty}$ of trapdoor permutations and for every probabilistic polynomial-time Turing-machine \mathcal{S}^* , if \mathcal{S}^* with oracle access to τ breaks the binding of $(\mathcal{S}^\tau, \mathcal{R}^\tau)$, then

$$\Pr \left[A^{\tau, \mathcal{S}^*}(1^n, G_n(td), y) = F_n^{-1}(td, y) \right] > \frac{1}{s(n)},$$

for infinitely many values of n , where A runs in time $s(n)$ and invokes \mathcal{S}^* on security parameters which are at most $1^{\ell(n)}$. The probability is taken uniformly over all the possible choices of $td \in \{0, 1\}^n$ and $y \in \{0, 1\}^n$, and over all the possible outcomes of the internal coin tosses of A .

We remark that the above correctness requirements are very strict and are not essential for our results. For example, in the setting of commitment schemes, for every τ such that the protocol $(\mathcal{S}^\tau, \mathcal{R}^\tau)$ is a weakly-binding statistically-hiding honest-receiver commitment scheme, we construct a malicious sender \mathcal{S}^* which breaks the binding property of the scheme. Therefore, we could have dealt with weaker correctness requirements as well, but stating such a weaker requirement in a meaningful way turns out to be quite subtle.

3 Communication Lower Bound for Statistically-Hiding Commitment Schemes

In this section we prove a lower bound on the communication complexity of fully-black-box constructions of a statistically-hiding commitment scheme from trapdoor permutations. We establish a lower bound on the number of bits communicated by the sender during the commit stage of any such scheme. Since we are interested in proving an impossibility result for commitment schemes, it will be sufficient for us to deal with bit-commitment schemes, i.e., commitment schemes in which the committed value is only one bit. We prove the following theorem:

Theorem 3.1. *In any fully-black-box $O(n)$ -expanding construction of a weakly-binding statistically-hiding honest-receiver bit-commitment scheme from a family of trapdoor permutations, the sender communicates $\Omega(n)$ bits during the commit stage.*

The proof of Theorem 3.1 follows the approach and technique of Haitner et al. [16] who constructed a ‘‘collision-finding’’ oracle in order to derive a lower bound on the round complexity of statistically-hiding commitment schemes. Given any fully-black-box $O(n)$ -expanding construction $(\mathcal{S}, \mathcal{R}, A)$ of a weakly-binding statistically-hiding honest-receiver bit-commitment scheme from a family of trapdoor permutations τ , we show that relative to their oracle the following holds: there exists a malicious sender \mathcal{S}^* that breaks the binding of the scheme $(\mathcal{S}^\tau, \mathcal{R}^\tau)$, but if the sender communicates $o(n)$ bits during the commit stage of $(\mathcal{S}^\tau, \mathcal{R}^\tau)$, then the machine A (with oracle access to \mathcal{S}^*) fails to break the security of τ .

In the remainder of this section we formally describe the oracle constructed in [16], and show that it can be used to break the binding of any statistically-hiding commitment scheme in which the sender communicates $o(n)$ bits during the commit stage.

3.1 The Oracle

We briefly describe the oracle constructed by Haitner et al. [16], and state one of its properties that will be used to imply our impossibility result. The oracle is of the form $\mathcal{O} = (\tau, \text{Sam}^\tau)$, where τ is a family of trapdoor permutations (i.e., $\tau = \{\tau_n\}_{n=1}^{\infty}$, where $\tau_n \in T_n$ for every n), and Sam^τ is an oracle that, very informally, receives as input a description of a circuit C (which may contain τ -gates) and a string z , and outputs a uniformly distributed preimage of z under the mapping defined by C . As discussed in [16], several essential restrictions are imposed on the querying of Sam that will prevent it from assisting in inverting τ .

Description of Sam. The oracle Sam receives as input a query $Q = (C_{\text{next}}^\tau, C^\tau, z)$, and outputs a pair (w', z') where w' is a uniformly distributed preimage of z under the mapping defined by the circuit C^τ , and $z' = C_{\text{next}}^\tau(w')$. We impose the following restrictions:

1. z was the result of a previous query with C^τ as the next-query circuit (note that this imposes a forest-like structure on the queries).
2. The circuit C_{next}^τ is a *refinement* of the circuit C^τ , where by a refinement we mean that $C_{\text{next}}^\tau(w) = (C^\tau(w), \tilde{C}^\tau(w))$ for some circuit \tilde{C}^τ and for every w . In particular, this implies that C^τ and C_{next}^τ have the same input length. Given a query Q , we denote this input length by $m(Q)$, and when the query Q is clear from the context we will write only m .
3. Each query contains a security parameter 1^n , and Sam answers queries only up to depth $\text{depth}(n)$, for some “depth restriction” function $\text{depth} : \mathbb{N} \rightarrow \mathbb{N}$ which is part of the description of Sam. The security parameter is set such that a query with security parameter 1^n is allowed to contain circuits with queries to permutations on up to n bits. Note that although different queries may have different security parameters, we ask that in the same “query-tree”, all queries will have the same security parameter (hence the depth of the tree is already determined by the root query).

In order to impose these restrictions, Sam is equipped with a family $\text{sign} = \{\text{sign}_k\}_{k=1}^\infty$ of (random) functions $\text{sign}_k : \{0, 1\}^k \rightarrow \{0, 1\}^{2k}$ that will be used as “signatures” for identifying legal queries as follows: in addition to outputting (w', z') , Sam will also output the value $\text{sign}(1^n, C_{\text{next}}^\tau, z', \text{dep} + 1)$, where dep is the depth of the query, 1^n is the security parameter of the query, and by applying the “function” sign we actually mean that we apply the function sign_k for the correct input length. Each query of the form $Q = (1^n, C_{\text{next}}^\tau, C^\tau, z, \text{dep}, \text{sig})$ is answered by Sam if and only if C_{next}^τ is a refinement of C^τ , $\text{dep} \leq \text{depth}(n)$ and $\text{sig} = \text{sign}(1^n, C^\tau, z, \text{dep})$.

Finally, Sam is provided with a family of (random) permutations $\mathcal{F} = \{f_Q\}$, where for every possible query Q a permutation f_Q is chosen uniformly at random from $\Pi_{m(Q)}$. Given a query $Q = (1^n, C_{\text{next}}^\tau, C^\tau, z, \text{dep}, \text{sig})$, the oracle Sam uses the permutation $f_Q \in \mathcal{F}$ in order to sample w' as follows: it outputs $w' = f_Q(t)$ for the lexicographically smallest $t \in \{0, 1\}^m$ such that $C^\tau(f_Q(t)) = z$. Note that whenever the permutation f_Q is chosen from Π_m uniformly at random, and independently of all other permutations in \mathcal{F} , then w' is indeed a uniformly distributed preimage of z . In this paper, whenever we consider the probability of an event over the choice of the family \mathcal{F} , we mean that for each query Q a permutation f_Q is chosen uniformly at random from $\Pi_{m(Q)}$ and independently of all other permutations. A complete and formal description of the oracle is provided in Figure 1.




Figure 1: The oracle Sam.

As mentioned above, the restrictions impose a forest-like structure on any sequence of queries: each query of the form $Q = (1^n, C_{\text{next}}^\tau, \perp, \perp, \perp, \perp)$ serves as a root of a tree. For any other “legal” query $Q =$

$(1^n, C_{\text{next}}^\tau, C^\tau, z, \text{dep}, \text{sig})$, there exists a previous query Q' which resulted in output z and contained C^τ as its next-query circuit. The query Q' is identified as the parent of Q in the query forest and is denoted $Q' = p(Q)$. If there is more than one such Q' , then we choose the first Q' according to some fixed ordering of the queries. When dealing with Turing-machines, we can identify the queries according to their chronological order.⁵

Notation 3.2. We say that a circuit A queries the oracle $\text{Sam}_{\text{depth}}^{\tau, \mathcal{F}, \text{sign}}$ up to depth d , if for every Sam-query $Q = (1^n, C_{\text{next}}^\pi, C^\pi, z, \text{dep}, \text{sig})$ that A makes, it holds that $\text{dep} \leq d$.

Random permutations are hard to invert even with Sam. One of the main properties of the oracle Sam, as proved in [16], is the following: any circuit with oracle access to Sam that tries to invert a random trapdoor permutation, fails with high probability. More specifically, Haitner et al. managed to relate this success probability to the maximal depth of the Sam-queries made by the circuit, and to the size of the circuit. They proved the following theorem:

Theorem 3.3 ([16]). *For every circuit A of size $s(n)$ that queries Sam up to depth $d(n)$ such that $s(n)^{3d(n)+2} < 2^{n/8}$, for every depth restriction function depth and for all sufficiently large n , it holds that*

$$\Pr_{\substack{td \leftarrow \{0,1\}^n, \tau, \mathcal{F} \\ y \leftarrow \{0,1\}^n, \text{sign}}} \left[A^{\tau, \text{Sam}_{\text{depth}}^{\tau, \mathcal{F}, \text{sign}}}(G_n(td), y) = F_n^{-1}(td, y) \right] \leq \frac{2}{s(n)} .$$

3.2 Breaking Low-Communication Statistically-Hiding Commitment Schemes

We show that a random instance of the oracle Sam can be used to break the binding of any weakly-binding statistically-hiding honest-receiver bit-commitment scheme. For every bit-commitment scheme $(\mathcal{S}, \mathcal{R})$ which is weakly-binding, statistically-hiding against an honest-receiver, and has oracle access to a family τ of trapdoor permutations, we construct a malicious sender \mathcal{S}^* which has oracle access to $\text{Sam}_{\text{depth}}^{\tau, \mathcal{F}, \text{sign}}$, and breaks the binding of $(\mathcal{S}^\tau, \mathcal{R}^\tau)$ with high probability over the choices of τ, \mathcal{F} and sign . The main idea in our proof is that if the sender (with security parameter 1^n) communicates $c(n)$ bits during the commit stage, then \mathcal{S}^* needs to query $\text{Sam}_{\text{depth}}^{\tau, \mathcal{F}, \text{sign}}$ only up to depth $\lceil \frac{c(n)}{\log n} \rceil + 1$. Formally, the following theorem is proved:

Theorem 3.4. *For any statistically-hiding honest-receiver bit-commitment scheme $(\mathcal{S}, \mathcal{R}, \mathcal{V})$ with oracle access to a family of trapdoor permutations in which the sender communicates at most $c(n)$ bits during the commit stage, and for any polynomial $p(n)$, there exists a polynomial-time malicious sender \mathcal{S}^* such that*

$$\Pr_{\tau, \mathcal{F}, \text{sign}, r_{\mathcal{R}}} \left[\begin{array}{l} ((\text{decom}, \text{decom}') | \text{com}) \leftarrow \left\langle \mathcal{S}^* \text{Sam}_{\text{depth}}^{\tau, \mathcal{F}, \text{sign}}(1^n), \mathcal{R}^\tau(1^n, r_{\mathcal{R}}) \right\rangle : \\ \mathcal{V}^\tau(\text{com}, \text{decom}) = 0, \mathcal{V}^\tau(\text{com}, \text{decom}') = 1 \end{array} \right] > 1 - \frac{1}{p(n)} ,$$

for all sufficiently large n , where $\text{depth}(n) = \lceil \frac{c(n)}{\log n} \rceil + 1$.

In what follows we introduce the notation used in this section. We proceed with a brief presentation of the main ideas underlying the proof of Theorem 3.4. Then, we formally describe the malicious sender \mathcal{S}^* and analyze its success probability in order to prove Theorem 3.4.

Notations. Let $(\mathcal{S}, \mathcal{R})$ be a bit-commitment scheme with oracle access to a family of trapdoor permutations. We denote by $b \in \{0, 1\}$ and $r_{\mathcal{S}}, r_{\mathcal{R}} \in \{0, 1\}^*$ the input bit of the sender and the random coins of the sender and the receiver, respectively. We denote by $c(n)$ the maximal number of bits communicated from the sender to the receiver in the commit stage with security parameter 1^n . In addition we denote by $d(n)$ the number of

⁵However, when dealing with circuits we will have to identify the queries according to a some topological order which is consistent with their forest structure.

communication rounds in the scheme with security parameter 1^n , and without loss of generality we assume that the receiver makes the first move. Each communication round consists of a message sent from the receiver to the sender followed by a message sent from the sender to the receiver. We denote by q_i and a_i the messages sent by the receiver and the sender in the i -th round, respectively, and denote by a_{d+1} the message sent by the sender in the reveal stage. Finally, we let $\bar{a}_i = (a_1, \dots, a_i)$ and $\bar{q}_i = (q_1, \dots, q_i)$. A generic d -round bit-commitment scheme is described in Figure 2.




Figure 2: A d -round bit-commitment scheme.

Although the sender is a probabilistic polynomial-time *Turing-machine*, in order to interact with the oracle Sam we need to identify the sender with a sequence of polynomial-size *circuits* S_1, \dots, S_{d+1} as follows. In the first round, S sends a_1 by computing $a_1 = S_1(b, r_S, q_1)$. Similarly, in the following rounds, S sends a_i by computing $a_i = S_i(b, r_S, \bar{q}_i)$.

Finally, in order to simplify the notation regarding the input and output of the oracle Sam, in this section we ignore parts of the input and output of Sam: we ignore the security parameter and the “signatures” (since our malicious sender S^* will only ask legal queries), and consider queries of a simplified form $Q = (C_{\text{next}}^\tau, C^\tau, z)$, and answers that consist only of w' (i.e., an answer consists only of a uniformly distributed preimage of z under the mapping defined by C^τ).

A brief overview. Informally, recall that the oracle Sam described in Section 3.1 acts as follows: Sam is given as input a query $Q = (C_{\text{next}}, C, z)$, and outputs a pair (w', z') where w' is a uniformly distributed preimage of z under the mapping defined by the circuit C , and $z' = C_{\text{next}}(w')$. In addition, we imposed the restriction that there was a previous query (C, \cdot, \cdot) that was answered by (w, z) (note that this imposes a forest-like structure on the queries), and we only allow querying Sam up to depth $d(n) = O\left(\frac{n}{\log n}\right)$.

Given a statistically-hiding bit-commitment scheme in which the sender communicates $c(n)$ bits during the commit stage, we assume without loss of generality that the commit stage of the scheme has $c(n)$ communication rounds, where in each round the sender communicates one bit to the receiver. The malicious sender S^* operates as follows: it chooses a random input w (consisting of random coins and a random committed bit), and during the first $\log n$ rounds it simulates the honest sender. In these $\log n$ rounds, it receives $\log n$ messages $q_1, \dots, q_{\log n}$ from the receiver. Then, S^* constructs the circuit $C_{q_1, \dots, q_{\log n}}$ that receives as input a sender’s input w and outputs the $\log n$ sender’s messages corresponding to the receiver’s messages $q_1, \dots, q_{\log n}$. This circuit is used to query Sam for a random input w_1 . It may be the case, however, that w_1 is not consistent with the actual messages $a_1, \dots, a_{\log n}$ that S^* sent in the first $\log n$ rounds. In this case, S^* rewinds Sam for a polynomial number of times, and since the total length of the sender’s messages in these $\log n$ rounds is only $\log n$ bits, then with sufficiently high probability S^* will obtain a consistent w_1 . Now, in the next $\log n$ rounds the malicious sender S^* simulates the honest sender with input w_1 , and in the end of these $\log n$ rounds it

will query (and rewind) Sam again for another consistent input $w_{\log n}$, and so on. Finally, after completing the commit stage, \mathcal{S}^* queries Sam to obtain two random inputs $w_{c(n)}$ and $w'_{c(n)}$ which are consistent with the transcript of the commit stage. Since the commitment scheme is statistically-hiding, then with probability roughly half they can be used to break the binding of the protocol. A crucial point in this description, is that \mathcal{S}^* queries Sam only up to depth $\frac{c(n)}{\log n}$. Therefore, if $c(n) = o(n)$, then such access to Sam cannot be used to invert a random trapdoor permutations, according to Theorem 3.3.

A formal Description of \mathcal{S}^* . Given a bit-commitment scheme $(\mathcal{S}, \mathcal{R})$ in which the sender communicates $c(n)$ bits during the commit stage, we assume without loss of generality (and for simplicity of the presentation) that the scheme has $c(n)$ communication rounds (i.e., $d(n) = c(n)$) where in each round during the commit stage the sender communicates one bit to the receiver (i.e., each of $a_1, \dots, a_{d(n)}$ is one bit). Furthermore, in order to simplify the description of \mathcal{S}^* , we assume that $\log n$ is an integral value (where 1^n is the security parameter given as input to \mathcal{S}^*) and that $c(n) = k \log n + 1$ for some integer $k = k(n)$. We stress that these assumptions are not at all essential, but avoiding them will result in a more complicated description.

On input 1^n , the malicious sender \mathcal{S}^* with oracle access to $\text{Sam}_{\text{depth}}^{\tau, \mathcal{F}, \text{sign}}$ interacts with the honest receiver \mathcal{R} as follows.

1. The commit stage:

- (a) In the first round the malicious sender \mathcal{S}^* receives \mathcal{R} 's message q_1 , and computes the description of the circuit $C_1 = S_1(\cdot, \cdot, q_1)$ obtained from the circuit S_1 by fixing q_1 as its third input. Then, \mathcal{S}^* queries $\text{Sam}_{\text{depth}}^{\tau, \mathcal{F}, \text{sign}}$ with (C_1, \perp, \perp) , receives an answer $w_1 = (b_1, r_1)$ and sends $a_1 = S_1(b_1, r_1, q_1)$ to \mathcal{R} .
- (b) In every round $i \in \{2, \dots, \log n\}$ the malicious sender \mathcal{S}^* simulates the honest sender \mathcal{S} with input w_1 . That is, \mathcal{S}^* receives \mathcal{R} 's message q_i and replies with $a_i = S_i(b_1, r_1, \bar{q}_i)$.
- (c) In round $\log n + 1$ the malicious sender \mathcal{S}^* receives \mathcal{R} 's message $q_{\log n+1}$, and computes the description of the circuit $C_{\log n+1} = S_{\log n+1}(\cdot, \cdot, \bar{q}_{\log n+1})$ obtained from the circuit $S_{\log n+1}$ by fixing $\bar{q}_{\log n+1}$ as its third input. Then, \mathcal{S}^* queries $\text{Sam}_{\text{depth}}^{\tau, \mathcal{F}, \text{sign}}$ with $(C_{\log n+1}, C_1, w_1)$ for $t = 2n^5 c(n) p(n)$ times and receives t answers. If one of these answers is consistent with the transcript of the protocol so far, then denote the first such answer by $w_{\log n+1} = (b_{\log n+1}, r_{\log n+1})$, and in this case \mathcal{S}^* sends $a_{\log n+1} = S_{\log n+1}(b_{\log n+1}, r_{\log n+1}, \bar{q}_{\log n+1})$ to \mathcal{R} . Otherwise, \mathcal{S}^* aborts the execution of the protocol.
- (d) In the remainder of the commit stage \mathcal{S}^* acts as follows:
 - i. For every integer k and in every round $i \in \{(k-1)\log n + 2, \dots, k \log n\}$, the malicious sender \mathcal{S}^* simulates the honest sender \mathcal{S} with input w_k .
 - ii. For every integer k and in every round $k \log n + 1$ the malicious sender \mathcal{S}^* receives \mathcal{R} 's message $q_{k \log n+1}$, and computes the description of the circuit $C_{k \log n+1} = S_{k \log n+1}(\cdot, \cdot, \bar{q}_{k \log n+1})$ obtained from the circuit $S_{k \log n+1}$ by fixing $\bar{q}_{k \log n+1}$ as its third input. Then, \mathcal{S}^* queries $\text{Sam}_{\text{depth}}^{\tau, \mathcal{F}, \text{sign}}$ with $(C_{k \log n+1}, C_{(k-1)\log n+1}, w_{(k-1)\log n+1})$ for $t = 2n^5 c(n) p(n)$ times and receives t answers. If one of these answers is consistent with the transcript of the protocol so far, then denote the first such answer by $w_{k \log n+1} = (b_{k \log n+1}, r_{k \log n+1})$, and in this case \mathcal{S}^* sends $a_{k \log n+1} = S_{k \log n+1}(b_{k \log n+1}, r_{k \log n+1}, \bar{q}_{k \log n+1})$ to \mathcal{R} . Otherwise, \mathcal{S}^* aborts the execution of the protocol.

2. The reveal stage:

- (a) \mathcal{S}^* queries $\text{Sam}_{\text{depth}}^{\tau, \mathcal{F}, \text{sign}}$ with $(\perp, C_{d(n)}, w_{d(n)})$ for n times, and receives n pairs $\left\{ \left(b_{d(n)+1}^{(j)}, r_{d(n)+1}^{(j)} \right) \right\}_{j=1}^n$. If there exist $j_0, j_1 \in [n]$ such that $b_{d(n)+1}^{(j_0)} = 0$ and $b_{d(n)+1}^{(j_1)} = 1$, then \mathcal{S}^* outputs $\text{decom} = S_{d(n)+1} \left(b_{d(n)+1}^{(j_0)}, r_{d(n)+1}^{(j_0)}, \bar{q}_{d(n)} \right)$ and $\text{decom}' = S_{d(n)+1} \left(b_{d(n)+1}^{(j_1)}, r_{d(n)+1}^{(j_1)}, \bar{q}_{d(n)} \right)$. Otherwise, \mathcal{S}^* aborts the execution of the protocol.

Two minor technical details were omitted from the description of \mathcal{S}^* . First, according to the description of Sam (Section 3.1), whenever Sam is queried multiple times with the same input, it returns the exact same answer. Thus, whenever \mathcal{S}^* queries Sam more than once with the same input, \mathcal{S}^* has to make sure that the queries are all different (for example, by artificially embedding the query number to one of the circuits in the query). Second, in order for \mathcal{S}^* 's queries to be legal, it should hold that the circuit $C_{k \log n + 1}$ is a refinement of the circuit $C_{(k-1) \log n + 1}$ for every integer k (as discussed in Section 3.1). This can be done very easily by embedding the description of each $C_{(k-1) \log n + 1}$ inside each $C_{k \log n + 1}$ (i.e., the output of C_i is the sequence of bits \bar{a}_i instead of only the bit a_i).

We proceed by arguing that the malicious sender \mathcal{S}^* successfully completes the commit stage with high probability. Then, given that \mathcal{S}^* has successfully completed the commit stage, we prove that the transcript of the commit stage is distributed identically to the transcript of the commit stage in an honest execution of the protocol. This enables us to use the fact that the commitment scheme is statistically-hiding, and therefore a random transcript can be revealed both as a commitment to $b = 0$ and as a commitment to $b = 1$, with almost equal probabilities.

Lemma 3.5. *The malicious sender \mathcal{S}^* successfully completes the commit stage with probability at least $1 - 1/(n^3 p(n))$ over the choices of $\tau, \mathcal{F}, \text{sign}$ and $r_{\mathcal{R}}$.*

Proof. The malicious sender \mathcal{S}^* may abort the commit stage only in rounds of the form $k \log n + 1$. For every integer $1 \leq k \leq \frac{c(n)-1}{\log n}$ we denote by \mathcal{E}_k the event in which \mathcal{S}^* aborts in round $k \log n + 1$ of the commit stage. Then, the probability that \mathcal{S}^* fails to complete the commit stage is

$$\Pr \left[\bigcup_{k=1}^{\frac{c(n)-1}{\log n}} \mathcal{E}_k \right] \leq \sum_{k=1}^{\frac{c(n)-1}{\log n}} \Pr [\mathcal{E}_k] ,$$

where the probability is taken over the choices of $\tau, \mathcal{F}, \text{sign}$ and $r_{\mathcal{R}}$. We show that for every $1 \leq k \leq \frac{c(n)-1}{\log n}$ it holds that $\Pr [\mathcal{E}_k] \leq 1/(n^3 c(n) p(n))$, which yields the correctness of the lemma. For simplicity, we first consider the case $k = 1$, and then show that the exact same argument generalizes for general k in a straightforward manner.

At the beginning of the protocol, after receiving q_1 from the receiver, \mathcal{S}^* queries Sam with $Q_1 = (C_1, \perp, \perp)$ and receives an answer $w_1 = (b_1, r_1)$. The description of Sam implies that w_1 is uniformly distributed among all possible inputs of the sender. \mathcal{S}^* then uses w_1 to simulate the honest sender during the first $\log n$ rounds by sending the bit $a_i = S_i(b_1, r_1, \bar{q}_i)$ in each of these rounds. In round $\log n + 1$, the malicious sender \mathcal{S}^* queries Sam with $(C_{\log n + 1}, C_1, w_1)$ for $t = 2n^5 c(n) p(n)$ times and receives t answers. We claim that since each a_i is a bit and we consider here only $\log n$ of them, then at least one of these answers will be consistent with the transcript of the protocol so far with high probability. Moreover, we show that this holds for any random coins of the receiver, and therefore from this point on we fix the random coins of the receiver. Note that by the description of Sam and the circuit C_1 , these t answers are chosen independently and uniformly at random from all possible inputs of the sender. Since the random coins of the receiver are fixed, the values $a_1, \dots, a_{\log n}$ can be viewed as a deterministic function of the input w_1 . Let us denote this function by $h : \{0, 1\}^{q(n)} \rightarrow \{0, 1\}^{\log n}$, where $q(n)$ is the bit-length of the sender's input. Then, it remains to analyze success probability of \mathcal{S}^* in the following experiment:

- $t + 1$ values $w_1, w_{\log n+1}^{(1)}, \dots, w_{\log n+1}^{(t)} \in \{0, 1\}^{q(n)}$ are chosen independently and uniformly at random.
- \mathcal{S}^* is successful if $h(w_1) = h(w_{\log n+1}^{(i)})$ for some $i \in [t]$.

In order to analyze this experiment, we consider a set of “bad” inputs for h . This set consists of all inputs w for which the set $h^{-1}(h(w))$ is very small relative to $\{0, 1\}^{q(n)}$ (less than some polynomial fraction). In case that w_1 is not in this bad set, then \mathcal{S}^* has a very high success probability, and the probability that w_1 is in this set is rather low. More formally, let

$$\text{BAD} = \left\{ w \in \{0, 1\}^{q(n)} : \frac{|h^{-1}(h(w))|}{2^{q(n)}} \leq \frac{1}{2^{n^4 c(n) p(n)}} \right\},$$

then since the range of h contains at most n elements, we have that

$$\Pr[w_1 \in \text{BAD}] \leq n \cdot \frac{1}{2^{n^4 c(n) p(n)}} = \frac{1}{2^{n^3 c(n) p(n)}}.$$

Therefore, the probability that \mathcal{S}^* aborts in round $\log n + 1$ can be upper bounded as follows

$$\begin{aligned} \Pr[\mathcal{E}_1] &\leq \Pr[w_1 \in \text{BAD}] + \Pr[\mathcal{E}_1 \mid w_1 \notin \text{BAD}] \\ &\leq \frac{1}{2^{n^3 c(n) p(n)}} + \left(1 - \frac{1}{2^{n^4 c(n) p(n)}}\right)^t \\ &= \frac{1}{2^{n^3 c(n) p(n)}} + \left(1 - \frac{1}{2^{n^4 c(n) p(n)}}\right)^{2^{n^5 c(n) p(n)}} \\ &\leq \frac{1}{2^{n^3 c(n) p(n)}} + \exp(-n) \\ &\leq \frac{1}{n^3 c(n) p(n)}. \end{aligned}$$

More generally, in every round of form $k \log n + 1$ for $k > 1$, the malicious sender \mathcal{S}^* holds some input $w_{(k-1) \log n+1}$ which is uniformly distributed among all inputs of the sender. This $w_{(k-1) \log n+1}$ was used by \mathcal{S}^* to simulate the honest sender in rounds $(k-1) \log n + 1, \dots, k \log n$. Then, \mathcal{S}^* uses **Sam** to sample independently and uniformly at random t elements from the set of all inputs which are consistent with the transcript of the protocol in the first $(k-1) \log n$ rounds. Therefore, it is only required that one of these inputs will be consistent with $w_{(k-1) \log n+1}$ on the answers it provided in rounds $(k-1) \log n + 1, \dots, k \log n$ and the same argument as before goes through, with the only difference that in this case the function h is defined only over the set of inputs which are consistent with the first $(k-1) \log n$ rounds (and not over the whole set $\{0, 1\}^{q(n)}$). ■

In the following lemma we show that given that \mathcal{S}^* has successfully completed the commit stage, the transcript of the commit stage is distributed identically to the transcript of the commit stage in an honest execution of the protocol. Formally, we define two the following two distributions:

- $\mathcal{D}_n^* = \text{view}_{\langle \mathcal{S}^*, \mathcal{R} \rangle}(n)$ is the distribution of the view of \mathcal{R} in the commit stage when interacting with the malicious sender $\mathcal{S}^*(1^n)$. This view consists of \mathcal{R} 's random coins and of the sequence of messages it receives from \mathcal{S}^* . The distribution is taken over \mathcal{R} 's random coins and over the uniform choice of τ, \mathcal{F} and sign .
- $\mathcal{D}_n = \text{view}_{\langle \mathcal{S}, \mathcal{R} \rangle}(n)$ is the distribution of the view of \mathcal{R} in the commit stage when interacting with the honest sender $\mathcal{S}(1^n, b, r_S)$. This view consists of \mathcal{R} 's random coins and of the sequence of messages it receives from \mathcal{S} . The distribution is taken over the random coins of \mathcal{R} and \mathcal{S} , and over the uniform choice of $b \in \{0, 1\}$ and τ .

Lemma 3.6. *Given that \mathcal{S}^* successfully completed the commit stage, the distributions \mathcal{D}_n and \mathcal{D}_n^* are identical.*

Proof. We show that the distributions \mathcal{D}_n and \mathcal{D}_n^* assign equal probabilities to every triplet $(r_{\mathcal{R}}, \bar{q}_d, \bar{a}_d)$ given that \mathcal{S}^* did not abort during the commit stage. More specifically, we prove by induction on $1 \leq i \leq d$ that $\Pr_{\mathcal{D}_n}[r_{\mathcal{R}}, \bar{q}_d, \bar{a}_d] = \Pr_{\mathcal{D}_n^*}[r_{\mathcal{R}}, \bar{q}_d, \bar{a}_d]$.

For $i = 1$, clearly we have that $\Pr_{\mathcal{D}_n}[r_{\mathcal{R}}, q_1] = \Pr_{\mathcal{D}_n^*}[r_{\mathcal{R}}, q_1]$ since $r_{\mathcal{R}}$ is distributed exactly the same in the two cases, and q_1 is a deterministic function of $r_{\mathcal{R}}$. Therefore we only have to show that $\Pr_{\mathcal{D}_n}[a_1|r_{\mathcal{R}}, q_1] = \Pr_{\mathcal{D}_n^*}[a_1|r_{\mathcal{R}}, q_1]$. In the first round, the malicious sender \mathcal{S}^* queries $\text{Sam}_{\text{depth}}^{\tau, \mathcal{F}, \text{sign}}$ with $Q = (C_1, \perp, \perp)$, and receives $w_1 = (b_1, r_1)$. Note that by the description of $\text{Sam}_{\text{depth}}^{\tau, \mathcal{F}, \text{sign}}$ and of \mathcal{F} , there is a random permutation f_Q which corresponds to Q , and $\text{Sam}_{\text{depth}}^{\tau, \mathcal{F}, \text{sign}}$ outputs $(b_1, r_1) = f_Q(0^m)$, which is a uniformly distributed value. That is, \mathcal{S}^* sends $a_1 = S_1(b_1, r_1, q_1)$ for a uniformly distributed pair (b_1, r_1) exactly as the honest sender \mathcal{S} should do.

Assume now that the claim holds for $i - 1$, i.e., $\Pr_{\mathcal{D}_n}[r_{\mathcal{R}}, \bar{q}_{i-1}, \bar{a}_{i-1}] = \Pr_{\mathcal{D}_n^*}[r_{\mathcal{R}}, \bar{q}_{i-1}, \bar{a}_{i-1}]$. Again, we have that $\Pr_{\mathcal{D}_n}[q_i|r_{\mathcal{R}}, \bar{q}_{i-1}, \bar{a}_{i-1}] = \Pr_{\mathcal{D}_n^*}[q_i|r_{\mathcal{R}}, \bar{q}_{i-1}, \bar{a}_{i-1}]$, since in both cases q_i is a deterministic function of $r_{\mathcal{R}}$, \bar{q}_{i-1} and \bar{a}_{i-1} . It remains to show that $\Pr_{\mathcal{D}_n}[a_i|r_{\mathcal{R}}, \bar{q}_i, \bar{a}_{i-1}] = \Pr_{\mathcal{D}_n^*}[a_i|r_{\mathcal{R}}, \bar{q}_i, \bar{a}_{i-1}]$. At this point we have to distinguish between two possible cases. The first case is that in the current round \mathcal{S}^* computes a_i by simulating the honest sender using an input w which has already been sampled in an earlier round. Therefore the distribution of the resulting a_i is exactly as if the honest sender \mathcal{S} had input w to begin with, and the lemma follows inductively. The second case is that in the current round \mathcal{S}^* queries $\text{Sam}_{\text{depth}}^{\tau, \mathcal{F}, \text{sign}}$ multiple times with some query Q and obtains some w which is consistent with the transcript of the protocol up to this point. Note that by the description of $\text{Sam}_{\text{depth}}^{\tau, \mathcal{F}, \text{sign}}$ and of \mathcal{F} , the permutation f_Q which corresponds to Q was chosen uniformly at random from Π_m and independently of all the other permutations in \mathcal{F} . Therefore, w is uniformly distributed among all inputs which are consistent with the protocol's transcript until this point, and therefore the distribution of the resulting a_i is exactly as if the honest sender \mathcal{S} had input w to begin with. Thus, $\Pr_{\mathcal{D}_n}[a_i|r_{\mathcal{R}}, \bar{q}_i, \bar{a}_{i-1}] = \Pr_{\mathcal{D}_n^*}[a_i|r_{\mathcal{R}}, \bar{q}_i, \bar{a}_{i-1}]$, which yields the correctness of the lemma. \blacksquare

We conclude the proof of Theorem 3.4 by combining Lemmata 3.5 and 3.6, and by exploiting the statistical-hiding property of the commitment scheme.

Proof of Theorem 3.4. Assuming that the malicious sender \mathcal{S}^* has successfully completed the commit stage, then in the reveal stage \mathcal{S}^* uses $\text{Sam}_{\text{depth}}^{\tau, \mathcal{F}, \text{sign}}$ in order to sample uniformly and independently at random n input pairs $\{(b_{d+1}^{(j)}, r_{d+1}^{(j)})\}_{j=1}^n$ from the set of all input pairs which are consistent with the transcript of the commit stage. We prove that with overwhelming probability these inputs enable \mathcal{S}^* to reveal both to $b = 0$ and to $b = 1$.

Denote by $\mathcal{D}_n^0 = \text{view}_{\langle \mathcal{S}(0), \mathcal{R} \rangle}(n)$ the distribution of the honest receiver's view in the commit stage when interacting with the honest sender $\mathcal{S}(1^n, 0, r_{\mathcal{S}})$. This view consists of its random coins and of the sequence of messages it receives from \mathcal{S} , and the distribution is taken over the random coins of \mathcal{R} and \mathcal{S} and over the choice of τ . Similarly, let $\mathcal{D}_n^1 = \text{view}_{\langle \mathcal{S}(1), \mathcal{R} \rangle}(n)$. Then, the assumption that the commitment scheme is statistically-hiding against an honest receiver, implies that the statistical difference between the distributions \mathcal{D}_n^0 and \mathcal{D}_n^1 is some negligible function $\rho(n)$.

We define a set of “good” transcripts. This set consists of all transcripts of the commit stage which enable \mathcal{S}^* to reveal both to $b = 0$ and to $b = 1$ with overwhelming probability. We show that with overwhelming probability the transcript is in this set. Formally, we define

$$\text{GOOD} = \left\{ \text{trans} : \left(1 - \sqrt{\rho(n)}\right) \cdot \Pr_{\mathcal{D}_n^0}[\text{trans}] < \Pr_{\mathcal{D}_n^1}[\text{trans}] < \left(1 + \sqrt{\rho(n)}\right) \cdot \Pr_{\mathcal{D}_n^0}[\text{trans}] \right\} .$$

Note that for every transcript trans of the commit stage and for every $j \in [n]$, it holds that

$$\frac{\Pr_{\tau, \mathcal{F}, r_{\mathcal{R}}} \left[b_{d+1}^{(j)} = 0 \mid \text{trans} \right]}{\Pr_{\tau, \mathcal{F}, r_{\mathcal{R}}} \left[b_{d+1}^{(j)} = 1 \mid \text{trans} \right]} = \frac{\Pr_{\tau, \mathcal{F}, r_{\mathcal{R}}} \left[b_{d+1}^{(j)} = 0 \wedge \text{trans} \right]}{\Pr_{\tau, \mathcal{F}, r_{\mathcal{R}}} \left[b_{d+1}^{(j)} = 1 \wedge \text{trans} \right]} = \frac{\Pr_{\mathcal{D}_n^0} [\text{trans}]}{\Pr_{\mathcal{D}_n^1} [\text{trans}]},$$

where the second equality follows from Lemma 3.6. The definition of the set GOOD implies that if $\text{trans} \in \text{GOOD}$, then for all sufficiently large n it holds that

$$\min \left\{ \Pr_{\tau, \mathcal{F}, r_{\mathcal{R}}} \left[b_{d+1}^{(j)} = 0 \mid \text{trans} \right], \Pr_{\tau, \mathcal{F}, r_{\mathcal{R}}} \left[b_{d+1}^{(j)} = 1 \mid \text{trans} \right] \right\} > 1/3.$$

Therefore,

$$\Pr_{\tau, \mathcal{F}, r_{\mathcal{R}}} \left[\mathcal{S}^* \text{ fails in the reveal stage} \mid \text{trans} \in \text{GOOD} \right] < 2 \cdot \left(\frac{2}{3} \right)^n,$$

since a failure occurs only in the case that all n input pairs sampled in the reveal stage have $b_{d+1}^{(j)} = 0$, or that they all have $b_{d+1}^{(j)} = 1$. It remains to show that the transcript is in GOOD with overwhelming probability. Lemma 3.6 and the fact that the statistical distance between the distributions \mathcal{D}_n^0 and \mathcal{D}_n^1 is at most $\rho(n)$ imply that

$$\begin{aligned} \Pr_{\tau, \mathcal{F}, r_{\mathcal{R}}} [\text{trans} \in \text{GOOD}] &= \Pr_{\mathcal{D}_n} [\text{trans} \in \text{GOOD}] \\ &= \frac{1}{2} \cdot (\Pr_{\mathcal{D}_n^0} [\text{trans} \in \text{GOOD}] + \Pr_{\mathcal{D}_n^1} [\text{trans} \in \text{GOOD}]) \\ &\geq \frac{1}{2} \cdot (2 \cdot \Pr_{\mathcal{D}_n^0} [\text{trans} \in \text{GOOD}] - \rho(n)) \\ &> 1 - 2\sqrt{\rho(n)} - \frac{\rho(n)}{2}, \end{aligned}$$

where the last inequality follows from Fact 2.3. Therefore,

$$\begin{aligned} \Pr [\mathcal{S}^* \text{ fails in the reveal stage}] &\leq \Pr [\text{trans} \notin \text{GOOD}] + \Pr \left[\mathcal{S}^* \text{ fails in the reveal stage} \mid \text{trans} \in \text{GOOD} \right] \\ &\leq 2\sqrt{\rho(n)} + \frac{\rho(n)}{2} + 2 \cdot \left(\frac{2}{3} \right)^n. \end{aligned}$$

Finally, Lemma 3.5 states that \mathcal{S}^* successfully completes the commit stage with probability at least $1 - 1/(n^3 p(n))$, and therefore

$$\begin{aligned} \Pr_{\tau, \mathcal{F}, \text{sign}, r_{\mathcal{R}}} &\left[\begin{array}{l} ((\text{decom}, \text{decom}') | \text{com}) \leftarrow \left\langle \mathcal{S}^* \text{Sam}_{\text{depth}}^{\tau, \mathcal{F}, \text{sign}}(1^n), \mathcal{R}^{\tau}(1^n, r_{\mathcal{R}}) \right\rangle : \\ \mathcal{V}^{\tau}(\text{com}, \text{decom}) = 0, \mathcal{V}^{\tau}(\text{com}, \text{decom}') = 1 \end{array} \right] \\ &> 1 - \left(\frac{1}{n^3 p(n)} + 2\sqrt{\rho(n)} + \frac{\rho(n)}{2} + 2 \cdot \left(\frac{2}{3} \right)^n \right) \\ &> 1 - \frac{1}{p(n)}, \end{aligned}$$

for all sufficiently large n . ■

3.3 Proof of Theorem 3.1

In this short section we combine Theorems 3.3 and 3.4 and derive the proof of Theorem 3.1. Let $(\mathcal{S}, \mathcal{R}, \mathcal{V}, A)$ be a fully-black-box $O(n)$ -expanding construction of a weakly-binding statistically-hiding honest-receiver bit-commitment scheme from an $s(n)$ -hard family of trapdoor permutations, in which the sender communicates at most $c(n)$ bits during the commit stage. Denote by $p(n)$ the polynomial for which the scheme is $(1 - 1/p(n))$ -binding. From this point on, we fix the depth restriction function $\text{depth} : \mathbb{N} \rightarrow \mathbb{N}$ of the oracle Sam to be the function $\text{depth}(n) = \lceil \frac{c(n)}{\log n} \rceil + 1$. Theorem 3.4 states that there exists a polynomial-time malicious sender \mathcal{S}^* such that

$$\Pr_{\tau, \mathcal{F}, \text{sign}, r_{\mathcal{R}}} \left[\left((\text{decom}, \text{decom}') | \text{com} \right) \leftarrow \left\langle \mathcal{S}^* \text{Sam}_{\text{depth}}^{\tau, \mathcal{F}, \text{sign}}(1^n), \mathcal{R}^{\tau}(1^n, r_{\mathcal{R}}) \right\rangle : \mathcal{V}^{\tau}(\text{com}, \text{decom}) = 0, \mathcal{V}^{\tau}(\text{com}, \text{decom}') = 1 \right] > 1 - \frac{1}{p(n)},$$

for all sufficiently large n . Thus, the fully-black-box construction guarantees that

$$\Pr_{\substack{td \leftarrow \{0,1\}^n, \tau, \mathcal{F} \\ y \leftarrow \{0,1\}^n, \text{sign}}} \left[A^{\tau, \mathcal{S}^*, \text{Sam}_{\text{depth}}^{\tau, \mathcal{F}, \text{sign}}}(G_n(td), y) = F_n^{-1}(td, y) \right] > \frac{1}{s(n)},$$

for infinitely many values of n , where A runs in time $s(n)$, and the probability is taken also over all the possible outcomes of the internal coin tosses of A . By converting the Turing-machine A to a circuit family, and by incorporating the description of \mathcal{S}^* into this family, we obtain that there exists a circuit A^* of size at most, say, $s^*(n) = (s(n))^2$ such that

$$\Pr_{\substack{td \leftarrow \{0,1\}^n, \tau, \mathcal{F} \\ y \leftarrow \{0,1\}^n, \text{sign}}} \left[A^*{}^{\tau, \text{Sam}_{\text{depth}}^{\tau, \mathcal{F}, \text{sign}}}(G_n(td), y) = F_n^{-1}(td, y) \right] > \frac{1}{s(n)} > \frac{2}{s^*(n)},$$

for infinitely many values of n . The assumption that the construction is $O(n)$ -expanding (i.e., that A when given security parameter 1^n invokes \mathcal{S}^* on security parameters which are at most $1^{O(n)}$), guarantees that A uses \mathcal{S}^* in a way such that Sam is queried up to depth at most $\text{depth}(n) = O\left(\frac{c(n)}{\log n}\right)$. This means that also the circuit A^* queries Sam up to depth at most $\text{depth}(n)$. We conclude the proof by observing that if $s^*(n)^{3\text{depth}(n)+2} < 2^{n/8}$, then the existence of the circuit A^* contradicts Theorem 3.3, and therefore $s^*(n)^{3\text{depth}(n)+2} \geq 2^{n/8}$, i.e., $c(n) = \Omega\left(\frac{n \log n}{\log s(n)}\right) = \Omega(n)$. \blacksquare

4 Refining the Relation Between Single-Server PIR and Commitment Schemes

The relation between single-server PIR and commitment schemes was first explored by Beimel, Ishai, Kushilevitz and Malkin [1], who showed that any single-server PIR protocol in which the server communicates at most $n/2$ bits to the user (where n is the size of the server's database), can be used to construct a weakly-binding statistically-hiding bit-commitment scheme. In particular, this served as the first indication that the existence of low-communication PIR protocols implies the existence of one-way functions. In this section, we refine the relation between these two fundamental primitives by significantly improving their reduction. More specifically, our improvements are the following:

1. The construction of Beimel et al. preserves the round complexity of the underlying single-server PIR, but it does not preserve its communication complexity. In their construction the sender is always required to send $\Omega(n)$ bits during the commit stage of the commitment scheme. We show that it is possible to preserve both the round complexity and the communication complexity. In our construction the number of bits sent by the sender during the commit stage of the commitment scheme is essentially the number of bits sent by the server in the PIR protocol.

2. The construction of Beimel et al. requires an execution of the single-server PIR protocol for every committed bit (that is, they constructed a bit-commitment scheme). We show that it is possible to commit to a super-logarithmic number of bits while executing the underlying single-server PIR protocol only once.
3. The construction of Beimel et al. was presented for single-server PIR protocols in which the server communicates at most $n/2$ bits. Our construction can deal with single-server PIR protocols in which the server communicates up to $n - \omega(\log n)$ bits.

In what follows we state our main theorem in the current section, and then turn to formally describe the construction and prove Theorem 4.1.

Theorem 4.1. *Let $d(n) \in \omega(\log n)$, $k(n) \geq 2d(n)$, and let \mathcal{P} be a single-server PIR protocol in which the server communicates $n - k(n)$ bits, where n is the size of the server's database. Then, there exists a weakly-binding statistically-hiding commitment scheme $COM^{\mathcal{P}}$ for $d(n)/6$ bits, in which the sender communicates less than $n - k(n) + 2d(n)$ bits during the commit stage. Moreover, the construction is fully-black-box and linearly-preserving.*

The construction. Fix $d(n)$, $k(n)$ and \mathcal{P} as in Theorem 4.1. Figure 2 describes our construction of the commitment scheme $COM^{\mathcal{P}} = (\mathcal{S}, \mathcal{R})$. In the construction we use a strong $\left(\frac{d(n)}{3}, 2^{1-\frac{d(n)}{3}}\right)$ -extractor $EXT : \{0, 1\}^n \times \{0, 1\}^{d(n)} \rightarrow \{0, 1\}^{d(n)/6}$ whose existence is guaranteed by Proposition 2.2.

Protocol $COM^{\mathcal{P}} = (\mathcal{S}, \mathcal{R})$

Joint input: security parameter 1^n .
Sender's input: $s \in \{0, 1\}^{d(n)/6}$.

Commit stage:

1. \mathcal{S} chooses a uniformly distributed $x \in \{0, 1\}^n$.
2. \mathcal{R} chooses a uniformly distributed index $i \in [n]$.
3. \mathcal{S} and \mathcal{R} execute the single-server PIR protocol \mathcal{P} for database of length n , where \mathcal{S} acts as the server with input x and \mathcal{R} acts as the user with input i . As a result, \mathcal{R} obtains a bit $x_i \in \{0, 1\}$.
4. \mathcal{S} chooses a uniformly distributed seed $t \in \{0, 1\}^{d(n)}$, computes $y = EXT(x, t) \oplus s$, and sends (t, y) to \mathcal{R} .

Reveal stage:

1. \mathcal{S} sends (s, x) to \mathcal{R} .
2. If the i^{th} bit of x equals x_i and $y = EXT(x, t) \oplus s$, then \mathcal{R} outputs s . Otherwise, \mathcal{R} outputs \perp .

Figure 2: A construction of a commitment scheme from any low-communication single-server PIR protocol.

The correctness of $COM^{\mathcal{P}}$ follows directly from the correctness of \mathcal{P} . In addition, notice that the total number of bits communicated by the sender in the commit stage is the total number of bits that the server communicates in \mathcal{P} plus the seed length and the output length of the extractor EXT . Thus, the sender communicates less than $n - k(n) + 2d(n)$ bits during the commit stage. In Lemma 4.2 we prove that $COM^{\mathcal{P}}$ is

statistically-hiding, and in Lemma 4.4 we prove that $COM^{\mathcal{P}}$ is weakly-binding. We note that the proof of hiding does not rely on any computational properties of the underlying PIR protocol \mathcal{P} , but only on the assumed bound on the number of bits communicated by the server in \mathcal{P} .

Lemma 4.2. *The scheme $COM^{\mathcal{P}}$ is statistically hiding.*

Proof. We have to show that for any computationally unbounded receiver \mathcal{R}^* and for any two strings s_0 and s_1 , the statistical distance between the distributions $\{\text{view}_{\langle \mathcal{S}(s_0), \mathcal{R}^* \rangle}(n)\}$ and $\{\text{view}_{\langle \mathcal{S}(s_1), \mathcal{R}^* \rangle}(n)\}$ (see Definition 2.7) is negligible in n . The transcript of the commit stage consists of the transcript $\text{trans}_{\mathcal{P}}$ of the execution of \mathcal{P} and of the pair $(t, \text{EXT}(x, t) \oplus s)$, where s is the committed string. Note that since $\text{trans}_{\mathcal{P}}$ is independent of the committed string, it is sufficient to prove that the statistical distance between the distribution of $(t, \text{EXT}(x, t))$ given $\text{trans}_{\mathcal{P}}$ and the uniform distribution is negligible in n .

We argue that due to the bound on the number of bits communicated by the server in \mathcal{P} , then even after executing \mathcal{P} , the database x still has sufficient min-entropy in order to guarantee that $(t, \text{EXT}(x, t))$ is sufficiently close to uniform. More specifically, let \mathcal{R}^* be an all-powerful receiver (recall that without loss of generality such an \mathcal{R}^* is deterministic), and denote by X the random variable corresponding to the value x in the scheme $COM^{\mathcal{P}}$. The following claim states the with high probability X has high min-entropy from \mathcal{R}^* 's point of view.

Claim 4.3. *It holds that*

$$\Pr_{\text{trans}_{\mathcal{P}} \leftarrow COM^{\mathcal{P}}} \left[H_{\infty}(X \mid \text{trans}_{\mathcal{P}}) < \frac{k(n)}{6} \right] < 2^{-\frac{k(n)}{4}},$$

where $\text{trans}_{\mathcal{P}}$ is the transcript of the embedded execution of \mathcal{P} in $COM^{\mathcal{P}}$.

Proof. For any value of r , the random coins used by \mathcal{S} in the execution of \mathcal{P} , let $f_r : \{0, 1\}^n \mapsto \{0, 1\}^{n-k(n)}$ be the function that maps x to the value of $\text{trans}_{\mathcal{P}}$ generated by the interaction of $(\mathcal{S}(x, r), \mathcal{R}^*)$, and let $\text{Col}(x, r) \stackrel{\text{def}}{=} \{x' \in \{0, 1\}^n : f_r(x') = f_r(x)\}$. Since f_r has at most $2^{n-k(n)}$ possible outputs, it follows that

$$\Pr_{x,r} \left[|\text{Col}(x, r)| < 2^{\frac{k(n)}{2}+1} \right] < \frac{2^{n-k(n)} \cdot 2^{\frac{k(n)}{2}+1}}{2^n} = 2^{1-\frac{k(n)}{2}}. \quad (4.1)$$

Let

$$\text{BAD} = \left\{ \text{trans}_{\mathcal{P}} : \Pr_{x,r} \left[|\text{Col}(x, r)| < 2^{\frac{k(n)}{2}+1} \mid \text{trans}_{\mathcal{P}} \right] > 2^{\frac{k(n)}{4}} \cdot 2^{1-\frac{k(n)}{2}} \right\},$$

then a standard averaging argument yields

$$\Pr_{\text{trans}_{\mathcal{P}} \leftarrow COM^{\mathcal{P}}} [\text{trans}_{\mathcal{P}} \in \text{BAD}] \leq 2^{-\frac{k(n)}{4}}.$$

Denote by U_r the random variable corresponding to r in the execution of $COM^{\mathcal{P}}$. Then, the following holds every value of x and $\text{trans}_{\mathcal{P}}$:

$$\begin{aligned} & \Pr[X = x \mid \text{trans}_{\mathcal{P}}] \\ &= \Pr \left[X = x \wedge |\text{Col}(X, U_r)| < 2^{\frac{k(n)}{2}+1} \mid \text{trans}_{\mathcal{P}} \right] + \Pr \left[X = x \wedge |\text{Col}(X, U_r)| \geq 2^{\frac{k(n)}{2}+1} \mid \text{trans}_{\mathcal{P}} \right] \\ &\leq \Pr \left[|\text{Col}(X, U_r)| < 2^{\frac{k(n)}{2}+1} \mid \text{trans}_{\mathcal{P}} \right] + 2^{-\left(\frac{k(n)}{2}+1\right)}. \end{aligned} \quad (4.2)$$

Note that if $H_{\infty}(X \mid \text{trans}_{\mathcal{P}}) < k(n)/6$ for some $\text{trans}_{\mathcal{P}}$, then there exists an x for which

$$\Pr[X = x \mid \text{trans}_{\mathcal{P}}] \geq 2^{-\frac{k(n)}{6}},$$

and therefore Equation 4.2 implies that

$$\Pr \left[|\text{Col}(X, U_r)| < 2^{\frac{k(n)}{2}+1} \mid \text{trans}_{\mathcal{P}} \right] > 2^{-\frac{k(n)}{6}} - 2^{-\left(\frac{k(n)}{2}+1\right)} > 2^{1-\frac{k(n)}{4}} .$$

Thus,

$$\begin{aligned} & \Pr_{\text{trans}_{\mathcal{P}} \leftarrow \text{COM}^{\mathcal{P}}} \left[H_{\infty}(X \mid \text{trans}_{\mathcal{P}}) < \frac{k(n)}{6} \right] \\ & \leq \Pr_{\text{trans}_{\mathcal{P}} \leftarrow \text{COM}^{\mathcal{P}}} \left[\Pr \left[|\text{Col}(X, U_r)| < 2^{\frac{k(n)}{2}+1} \mid \text{trans}_{\mathcal{P}} \right] > 2^{1-\frac{k(n)}{4}} \right] \\ & \leq \Pr_{\text{trans}_{\mathcal{P}} \leftarrow \text{COM}^{\mathcal{P}}} [\text{trans}_{\mathcal{P}} \in \text{BAD}] \\ & \leq 2^{-\frac{k(n)}{4}} . \end{aligned}$$

■

Now, since $d(n) \in \omega(\log n)$ and $k(n)/6 \geq d(n)/3$, Claim 4.3 implies that with probability $1 - \text{neg}(n)$, the extractor EXT guarantees that the statistical distance between the pair $(t, \text{EXT}(x, t))$ (given $\text{trans}_{\mathcal{P}}$) and the uniform distribution is at most $2^{1-d(n)/3}$ (which is again negligible in n). Therefore the scheme $\text{COM}^{\mathcal{P}}$ is statistically-hiding. More specifically, for every string $s \in \{0, 1\}^{d(n)/6}$ it holds that

$$\begin{aligned} & \text{SD}(\{\text{trans}_{\mathcal{P}}, t, \text{EXT}(X, t) \oplus s\}, \{\text{trans}_{\mathcal{P}}, U_{7d(n)/6}\}) \\ & \leq \Pr \left[H_{\infty}(X \mid \text{trans}_{\mathcal{P}}) < \frac{k(n)}{6} \right] \\ & \quad + \text{SD} \left(\{\text{trans}_{\mathcal{P}}, t, \text{EXT}(X, t) \oplus s\}, \{\text{trans}_{\mathcal{P}}, U_{7d(n)/6}\} \mid H_{\infty}(X \mid \text{trans}_{\mathcal{P}}) \geq \frac{k(n)}{6} \right) \\ & \leq 2^{-\frac{k(n)}{4}} + 2^{1-\frac{d(n)}{3}} . \end{aligned}$$

Therefore, for any two strings $s_0, s_1 \in \{0, 1\}^{d(n)/6}$ we have

$$\begin{aligned} & \text{SD}(\{\text{view}_{\langle S(s_0), \mathcal{R}^* \rangle}(n)\}, \{\text{view}_{\langle S(s_1), \mathcal{R}^* \rangle}(n)\}) = \text{SD}(\{\text{trans}_{\mathcal{P}}, t, \text{EXT}(X, t) \oplus s_0\}, \{\text{trans}_{\mathcal{P}}, t, \text{EXT}(X, t) \oplus s_1\}) \\ & \leq 2 \cdot \left(2^{-\frac{k(n)}{4}} + 2^{1-\frac{d(n)}{3}} \right) , \end{aligned}$$

which is negligible in n as required. ■

Lemma 4.4. *The scheme $\text{COM}^{\mathcal{P}}$ is weakly binding.*

Proof. We show that the scheme $\text{COM}^{\mathcal{P}}$ is $(1 - 1/n^2)$ -binding. Given any malicious sender Snd^* that violates the binding of the commitment scheme $\text{COM}^{\mathcal{P}}$ with probability at least $1 - 1/n^2$, we construct a malicious server Srv^* that breaks the security of the single-server PIR protocol \mathcal{P} .

Let Snd^* be a polynomial-time malicious sender that violates the binding of $\text{COM}^{\mathcal{P}}$ with probability at least $1 - 1/n^2$. As an intermediate step, we first construct a malicious server that has a non-negligible advantage in predicting a uniformly chosen index held by the user in \mathcal{P} . More specifically, we construct a malicious server Srv^* and a predictor \mathcal{D}' such that

$$\Pr \left[v \leftarrow \text{view}_{\langle \text{Srv}^*, \mathcal{U}(i) \rangle}(n) : \mathcal{D}'(v) = i \right] \geq \frac{1}{n} + \frac{1}{n^2} ,$$

where the probability is taken over the uniform choice of $i \in [n]$ and over the coin tosses of Srv^* , \mathcal{D}' and \mathcal{U} . Recall that $\text{view}_{\langle \text{Srv}^*, \mathcal{U}(i) \rangle}(n)$ denotes the distribution on the view of Srv^* when interacting with $\mathcal{U}(i)$ where $i \in [n]$. This view consists of its random coins and of the sequence of messages it receives from \mathcal{U} .

The malicious server Srv^* follows the malicious sender Snd^* in the embedded execution of \mathcal{P} in $\text{COM}^{\mathcal{P}}$. Following the interaction, Srv^* proceeds the execution of Snd^* to obtain a pair (t, y) and two decommitments (x_1, s_1) and (x_2, s_2) . If $x_1 = x_2$, then Srv^* fails. Otherwise, denote by $j \in [n]$ the minimal index such that $x_1[j] \neq x_2[j]$. Now, the predictor \mathcal{D}' outputs a uniformly distributed value i' from the set $[n] \setminus \{j\}$.

In order to analyze the success probability in predicting i , note that if (x_1, s_1) and (x_2, s_2) are valid decommitments and $s_1 \neq s_2$ (i.e., \mathcal{S}^* broke the binding of $\text{COM}^{\mathcal{P}}$), then it must hold that $x_1 \neq x_2$. In this case, let $j \in [n]$ be the minimal index such that $x_1[j] \neq x_2[j]$, then it must be the case that $i \neq j$, as otherwise \mathcal{R} will not accept the two decommitments. Therefore, when the predictor \mathcal{D}' outputs a uniformly distributed $i' \in [n] \setminus \{j\}$ it will output i with probability $1/(n-1)$. Thus,

$$\begin{aligned} \Pr \left[v \leftarrow \text{view}_{\langle \text{Srv}^*, \mathcal{U}(i) \rangle}(n) : \mathcal{D}'(v) = i \right] &\geq \left(1 - \frac{1}{n^2} \right) \cdot \frac{1}{n-1} \\ &= \frac{n+1}{n^2} \\ &= \frac{1}{n} + \frac{1}{n^2} . \end{aligned}$$

In the remainder of the proof, we apply a rather standard argument in order to be fully consistent with Definition 2.6 of the security of single-server PIR. That is, we need to show that there exists a pair of indices $i, j \in [n]$, a malicious server Srv^* and a distinguisher \mathcal{D} such that

$$\left| \Pr \left[v \leftarrow \text{view}_{\langle \text{Srv}^*, \mathcal{U}(i) \rangle}(n) : \mathcal{D}(v) = 1 \right] - \Pr \left[v \leftarrow \text{view}_{\langle \text{Srv}^*, \mathcal{U}(j) \rangle}(n) : \mathcal{D}(v) = 1 \right] \right| \geq \frac{1}{p(n)} ,$$

for some polynomial $p(n)$. We prove that this holds for independently and uniformly chosen $i, j \in [n]$ (and therefore there exist i and j for which this holds) where Srv^* is the malicious server described above, and $\mathcal{D} = \mathcal{D}_{i,j}$ is a distinguisher that uses \mathcal{D}' as follows:

- If \mathcal{D}' outputs i , then \mathcal{D} outputs 1.
- If \mathcal{D}' outputs j , then \mathcal{D} outputs 0.
- Otherwise, \mathcal{D} outputs a uniformly distributed $b \in \{0, 1\}$.

Then,

$$\begin{aligned} &\Pr \left[v \leftarrow \text{view}_{\langle \text{Srv}^*, \mathcal{U}(i) \rangle}(n) : \mathcal{D}(v) = 1 \right] \\ &= \Pr \left[v \leftarrow \text{view}_{\langle \text{Srv}^*, \mathcal{U}(i) \rangle}(n) : \mathcal{D}'(v) = i \right] + \frac{1}{2} \cdot \Pr \left[v \leftarrow \text{view}_{\langle \text{Srv}^*, \mathcal{U}(i) \rangle}(n) : \mathcal{D}'(v) \notin \{i, j\} \right] \\ &\geq \frac{1}{n} + \frac{1}{n^2} + \frac{1}{2} \cdot \Pr \left[v \leftarrow \text{view}_{\langle \text{Srv}^*, \mathcal{U}(i) \rangle}(n) : \mathcal{D}'(v) \notin \{i, j\} \right] , \end{aligned}$$

and

$$\begin{aligned} &\Pr \left[v \leftarrow \text{view}_{\langle \text{Srv}^*, \mathcal{U}(j) \rangle}(n) : \mathcal{D}(v) = 1 \right] \\ &= \Pr \left[v \leftarrow \text{view}_{\langle \text{Srv}^*, \mathcal{U}(j) \rangle}(n) : \mathcal{D}'(v) = i \right] + \frac{1}{2} \cdot \Pr \left[v \leftarrow \text{view}_{\langle \text{Srv}^*, \mathcal{U}(j) \rangle}(n) : \mathcal{D}'(v) \notin \{i, j\} \right] \\ &= \frac{1}{n} + \frac{1}{2} \cdot \Pr \left[v \leftarrow \text{view}_{\langle \text{Srv}^*, \mathcal{U}(j) \rangle}(n) : \mathcal{D}'(v) \notin \{i, j\} \right] , \end{aligned}$$

where the last equality holds since both i and j are independently chosen. Finally, note that

$$\Pr \left[v \leftarrow \text{view}_{\langle \text{Srv}^*, \mathcal{U}(i) \rangle}(n) : \mathcal{D}'(v) \notin \{i, j\} \right] = \Pr \left[v \leftarrow \text{view}_{\langle \text{Srv}^*, \mathcal{U}(j) \rangle}(n) : \mathcal{D}'(v) \notin \{i, j\} \right] ,$$

and therefore

$$\left| \Pr \left[v \leftarrow \text{view}_{\langle \text{Srv}^*, \mathcal{U}^{(i)} \rangle}(n) : \mathcal{D}(v) = 1 \right] - \Pr \left[v \leftarrow \text{view}_{\langle \text{Srv}^*, \mathcal{U}^{(j)} \rangle}(n) : \mathcal{D}(v) = 1 \right] \right| \geq \frac{1}{n^2} .$$

■

5 Communication Lower Bound for Single-Server PIR

In this section we combine the results from sections 3 and 4 derive an immediate proof of our main result, formally stated as follows:

Theorem 5.1. *In any fully-black-box $O(n)$ -expanding construction of a single-server PIR protocol from a family of trapdoor permutations, the server communicates $\Omega(n)$ bits to the user, where n is the size of the server’s database.*

Proof. Assume towards a contradiction that there exists a fully-black-box $O(n)$ -expanding construction of a single-server PIR protocol from a family of trapdoor permutations in which the server communicates $o(n)$ bits, where n is the size of the server’s database. By applying Theorem 4.1 with parameters $k(n) = n - o(n)$ and $d(n) = \log^2 n$ (actually any $d(n) = \omega(\log n)$ suffices) we obtain a fully-black-box $O(n)$ -expanding weakly-binding statistically-hiding bit-commitment scheme from a family of trapdoor permutations, in which the sender communicates $o(n)$ bits during the commit stage, where n is the security parameter of the scheme. However, the existence of such a scheme contradicts Theorem 3.1. ■

On extending the lower bound to weakly-preserving constructions. Theorem 5.1 does not rule out weakly-preserving (fully-black-box) constructions of single-server PIR from trapdoor permutations in which the sender communicates $o(n)$ bits to the user. We note that although weakly-preserving reductions guarantee much weaker security than polynomially-preserving reductions, investigating lower bounds for such reductions is still a very interesting research topic. Even more so as the sole construction to date of a single-server PIR protocol from trapdoor permutations uses such a reduction. A possible step towards tightening our bound is to first provide an improved lower bound on the communication complexity of statistically-hiding commitment schemes that allow the sender to commit to more than a single bit. Whereas in Section 4 we proved that any low-communication single-server PIR implies a statistically-hiding commitment scheme that allows the sender to commit to a relatively long string, our lower bound on the communication complexity of statistically-hiding commitment schemes in Section 3 serves as a bottleneck: it does not take into consideration the number of committed bits (the lower bound is only in terms of the security parameter). It is quite possible that a much tighter lower bound can be proved for string-commitment schemes. Such a lower bound may extend the result of the current paper to the setting of weakly-preserving reductions, and prove the optimality of the single-server PIR protocol of Kushilevitz and Ostrovsky [24].

Acknowledgements

We thank Yuval Ishai and Omer Reingold for their useful suggestions.

References

- [1] A. Beimel, Y. Ishai, E. Kushilevitz, and T. Malkin. One-way functions are essential for single-server private information retrieval. In *Proceedings of the 31st Annual ACM Symposium on Theory of Computing*, pages 89–98, 1999.

- [2] C. Cachin, S. Micali, and M. Stadler. Computationally private information retrieval with polylogarithmic communication. In *Advances in Cryptology - EUROCRYPT '99*, pages 402–414, 1999.
- [3] Y. Chang. Single database private information retrieval with logarithmic communication. In *Proceedings of the 9th Australasian Conference on Information Security and Privacy*, pages 50–61, 2004.
- [4] B. Chor, O. Goldreich, E. Kushilevitz, and M. Sudan. Private information retrieval. In *Proceedings of 36th Annual IEEE Symposium on Foundations of Computer Science*, pages 41–50, 1995.
- [5] G. D. Crescenzo, T. Malkin, and R. Ostrovsky. Single database private information retrieval implies oblivious transfer. In *Advances in Cryptology - EUROCRYPT '00*, pages 122–138, 2000.
- [6] S. Dziembowski and U. M. Maurer. On generating the initial key in the bounded-storage model. In *Advances in Cryptology - EUROCRYPT '04*, pages 126–137, 2004.
- [7] M. Fischlin. On the impossibility of constructing non-interactive statistically-secret protocols from any trapdoor one-way function. In *Topics in Cryptology - The Cryptographers' Track at the RSA Conference*, pages 79–95, 2002.
- [8] R. Gennaro, Y. Gertner, and J. Katz. Lower bounds on the efficiency of encryption and digital signature schemes. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing*, pages 417–425, 2003.
- [9] R. Gennaro and L. Trevisan. Lower bounds on the efficiency of generic cryptographic constructions. In *Proceedings of the 41st Annual IEEE Symposium on Foundations of Computer Science*, pages 305–313, 2000.
- [10] C. Gentry and Z. Ramzan. Single-database private information retrieval with constant communication rate. In *Proceedings of the 32nd International Colloquium on Automata, Languages and Programming*, pages 803–815, 2005.
- [11] Y. Gertner, S. Kannan, T. Malkin, O. Reingold, and M. Viswanathan. The relationship between public key encryption and oblivious transfer. In *Proceedings of the 41st Annual IEEE Symposium on Foundations of Computer Science*, pages 325–335, 2000.
- [12] Y. Gertner, T. Malkin, and O. Reingold. On the impossibility of basing trapdoor functions on trapdoor predicates. In *Proceedings of the 42nd Annual IEEE Symposium on Foundations of Computer Science*, pages 126–135, 2001.
- [13] O. Goldreich. *Foundations of Cryptography – Volume 1: Basic Tools*. Cambridge University Press, 2001.
- [14] O. Goldreich. *Foundations of Cryptography – Volume 2: Basic Applications*. Cambridge University Press, 2004.
- [15] I. Haitner. Implementing oblivious transfer using collection of dense trapdoor permutations. In *Proceedings of the 1st Theory of Cryptography Conference*, pages 394–409, 2004.
- [16] I. Haitner, J. J. Hoch, O. Reingold, and G. Segev. Finding collisions in interactive protocols – A tight lower bound on the round complexity of statistically-hiding commitments. To appear in *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science*, 2007.

- [17] D. Harnik and M. Naor. On the compressibility of NP instances and cryptographic applications. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science*, pages 719–728, 2006.
- [18] O. Horvitz and J. Katz. Bounds on the efficiency of “black-box” commitment schemes. In *Proceedings of the 32nd International Colloquium on Automata, Languages and Programming*, pages 128–139, 2005.
- [19] R. Impagliazzo and S. Rudich. Limits on the provable consequences of one-way permutations. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing*, pages 44–61, 1989.
- [20] Y. Ishai, E. Kushilevitz, and R. Ostrovsky. Sufficient conditions for collision-resistant hashing. In *Proceedings of the 2nd Theory of Cryptography Conference*, pages 445–456, 2005.
- [21] Y. T. Kalai and R. Raz. Succinct non-interactive zero-knowledge proofs with preprocessing for LOGSNP. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science*, pages 355–366, 2006.
- [22] J. H. Kim, D. R. Simon, and P. Tetali. Limits on the efficiency of one-way permutation-based hash functions. In *Proceedings of the 40th Annual IEEE Symposium on Foundations of Computer Science*, pages 535–542, 1999.
- [23] E. Kushilevitz and R. Ostrovsky. Replication is NOT needed: SINGLE database, computationally-private information retrieval. In *Proceedings of the 38th Annual IEEE Symposium on Foundations of Computer Science*, pages 364–373, 1997.
- [24] E. Kushilevitz and R. Ostrovsky. One-way trapdoor permutations are sufficient for non-trivial single-server private information retrieval. In *Advances in Cryptology - EUROCRYPT '00*, pages 104–121, 2000.
- [25] H. Lipmaa. An oblivious transfer protocol with log-squared communication. In *Proceedings of the 8th International Conference on Information Security*, pages 314–328, 2005.
- [26] M. Luby. *Pseudorandomness and Cryptographic Applications*. Princeton University Press, 1996.
- [27] M. Naor, R. Ostrovsky, R. Venkatesan, and M. Yung. Perfect zero-knowledge arguments for NP using any one-way permutation. *Journal of Cryptology*, 11(2):87–108, 1998.
- [28] M. Naor and B. Pinkas. Oblivious transfer and polynomial evaluation. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing*, pages 245–254, 1999.
- [29] R. Ostrovsky and W. E. Skeith. A survey of single database PIR: Techniques and applications. Cryptology ePrint Archive, Report 2007/059, 2007.
- [30] O. Reingold, L. Trevisan, and S. P. Vadhan. Notions of reducibility between cryptographic primitives. In *Proceedings of the 1st Theory of Cryptography Conference*, pages 1–20, 2004.
- [31] S. Rudich. *Limits on the provable consequences of one-way functions*. PhD thesis, EECS Department, University of California, Berkeley, 1988.
- [32] A. Sahai and S. P. Vadhan. A complete problem for statistical zero knowledge. *Journal of the ACM*, 50(2):196–249, 2003.
- [33] D. R. Simon. Finding collisions on a one-way street: Can secure hash functions be based on general assumptions? In *Advances in Cryptology - EUROCRYPT '98*, pages 334–345, 1998.

- [34] A. Srinivasan and D. Zuckerman. Computing with very weak random sources. *SIAM Journal on Computing*, 28(4):1433–1459, 1999.
- [35] J. P. Stern. A new efficient all-or-nothing disclosure of secrets protocol. In *Advances in Cryptology - ASIACRYPT '98*, pages 357–371, 1998.
- [36] H. Wee. One-way permutations, interactive hashing and statistically hiding commitments. In *Proceedings of the 4th Theory of Cryptography Conference*, pages 419–433, 2007.