On the Security Loss in Cryptographic Reductions*

Chi-Jen Lu

Institute of Information Science, Academia Sinica, Taipei, Taiwan.
cjlu@iis.sinica.edu.tw

Abstract. Almost all the important cryptographic protocols we have today base their security on unproven assumptions, which all imply $\mathsf{NP} \neq \mathsf{P}$, and thus having unconditional proofs of their security seems far beyond our reach. One research effort then is to identify more basic primitives and prove the security of these protocols by reductions to the security of these primitives. However, in doing so, one often observes some security loss in the form that the security of the protocols is measured against weaker adversaries, e.g., adversaries with a smaller running time. Is such a security loss avoidable? We study two of the most basic cryptographic reductions: hardness amplification of one-way functions and constructing pseudorandom generators from one-way functions. We show that when they are done in a certain black-box way, such a security loss is in fact unavoidable.

1 Introduction

Although we have many protocols today for all kinds of interesting and important cryptographic tasks, almost all of these protocols have their security based on some assumptions. These assumptions all imply $\mathsf{P} \neq \mathsf{NP}$, so having unconditional proofs of their security seems far beyond our reach. One line of research then is to identify the weakest possible assumptions or primitives from which one can build more advanced cryptographic protocols. One such primitive is one-way function (OWF), a function which is easy to compute but hard to invert, with respect to polynomial time computation. It is now known that from a OWF, one can construct other cryptographic primitives such as pseudo-random generator (PRG), pseudo-random function, private-key encryption, bit commitment, zero-knowledge proof, and digital signature. In fact, all these primitives are known to be equivalent in the sense that they can all be built from each other [21, 6, 9, 7, 11, 18, 17, 10]. According to [5], these primitives may be categorized as in the world of private cryptography. There

* This work supported was in part by the National Science Council under the Grant NSC97-2221-E-001-012-MY3.

are other primitives, including public-key encryption, oblivious transfer, private information retrieval, and key agreement, which may be categorized as in the world of public cryptography. Primitives in the world of public cryptography seem to require a stronger assumption, and it has been shown that trapdoor one-way permutations can be used to build all of them. The relationships among primitives in public cryptography are more complicated, but most of them have been settled [13, 12, 5, 2, 14, 4, 1].

From a theoretical perspective, we seem to have obtained a good understanding of the relationships among these primitives. However, from a practical point of view, there are still issues to be resolved. The first is that even when we can construct one primitive from another, the construction may not be as efficient as we desire. For example, although we can use any one-way function to construct all the primitives in the world of private cryptography, the constructions often do not appear efficient enough to have a practical impact. Can we improve the efficiency of such constructions? Some negatives results have been obtained recently for the tasks of amplifying hardness of OWF [15, 16], constructing PRG from OWF [3, 20, 16], and constructing encryption or signature scheme [3] from (trapdoor) one-way permutation.

The second issue is that when constructing a primitive from another, one often suffers some kind of security loss. For example, although one can construct a PRG from a OWF, the proofs currently available can only guarantee the security of the PRG for weaker adversaries having a smaller running time (or circuit size) than that for OWF. Therefore, if we want to have a PRG with a certain security level, we need to start from a OWF with a much higher security level, which would require a substantial cost to implement and make it less attractive in practice. Similar problems also occur in other constructions, and people have tried to improve these constructions, but with limited success so far. Again, one may wonder whether or not such improvements are indeed impossible. Not much seems to be known, and our goal is to show that such security losses are basically unavoidable. We would like to start from two of the most basic primitives: OWF and PRG, and study the task of hardness amplification for OWF and the task of constructing PRG from OWF.

We say that a function f is ε-hard (to invert) for time t (or size t), if any algorithm running in time t (or circuit of size t) must fail to invert $f(x)$ for at least ε fraction of x. The task of hardness amplification is to transform a function f which is ε-hard for time t into a function $\bar{f}$ which is $(1-\delta)$-hard for time $\bar{t}$, for some small δ and ε. According to [21, 8], this is

possible with $\bar{t} = t/\gamma_1$, for some $\gamma_1 = ((1/\delta)/\log(1/\varepsilon))^{O(1)}$ (it seems that a more careful analysis can give $\gamma_1 = O((1/\delta)/\log(1/\varepsilon))$). That is, the hardness of the new function $\bar{f}$ is now measured against algorithms with a running time (or circuit size) smaller by a γ_1 factor than that for the initial function f. Therefore, when we want to transform a weakly OWF (with hardness $n^{-O(1)}$) into a strongly OWF (with hardness $1-n^{-\omega(1)}$), we lose a polynomial factor in the running time (or circuit size) of adversaries. We say that a function $g : \{0,1\}^n \to \{0,1\}^m$ is ε-random for time t (or size t) if for any algorithm C running in time t (or circuit of size t), the probabilities of $C(u) = 1$ and $C(g(x)) = 1$, over random u and random x respectively, differ by at most ε. According to [10], one can construct such a function g with $m > n$ (a PRG), which is ε-random for time t/γ_2 (or size t/γ_2), from any function f which is $(1 - n^{-\Omega(1)})$-hard for time t (or size t), for some $\gamma_2 = (n/\varepsilon)^{O(1)}$. From [7], one can have $\gamma_2 = n^{O(1)}/\varepsilon^2$, for the simpler case when $m = n + 1$ and f is a permutation.

We would like to show the impossibility of having a hardness amplification of OWF or a construction of PRG from OWF which can avoid such a loss of security. However, it is not clear how to establish the impossibility of transforming one primitive P to another primitive Q, especially given the possibility that the primitive Q may indeed exist. Therefore, one can only expect to have such impossibility results for a certain restricted types of transformations. Here, we consider transformations which are done in some black-box way.

Black-Box Reductions. The standard notion of black-box transformation from a primitive P to a primitive Q consists of two oracle algorithms $T^{(\cdot)}$ and $R^{(\cdot)}$ satisfying the following two conditions: (1) correctness: for any N that implements P, T^N implements Q, and (2) security: for any N that implements P and for any A that breaks T^N (as an implementation of Q), $R^{A,N}$ breaks N (as an implementation of P).

Although this may look restricted, almost all the known transformations between primitives in cryptography (including those we discussed before) are done in such a black-box way. In this paper, we consider a more general model, in which we drop the first condition and keep only the second one, namely, only the security proof is required to be done in a black-box way. We call this the weakly-black-box model, and note that impossibility results on such a more general model become stronger. We consider two transformations in this model: hardness amplification for OWF and constructing PRG from OWF. In the case of weakly-black-box hardness amplification, there exists an oracle algorithm R (an adversary)

such that for any M which breaks the hardness condition of the new function $\bar{f}$, R using M as an oracle can break the hardness condition of the initial function f. In the case of weakly-black-box PRG construction, there exists an oracle algorithm R (an adversary) such that for any D which breaks the randomness condition of the resulting generator g, R using D as an oracle can break the hardness condition of the initial function f. Here we consider the more general case in which R can be non-uniform by allowing it to have an advice string (or seeing R as a collection of circuits with oracle gates), and again this makes our impossibility results stronger.

Our Results. We first consider the task of weakly-black-box hardness amplification for OWF, which transforms ε-hard functions into $(1-\delta)$-hard functions. Our first two results show that any algorithm R realizing such a hardness amplification must make at least $q_1 = \Omega((1/\delta)/\log(1/\varepsilon))$ queries to the oracle, unless it can use a long advice string. More precisely, our first result shows that for any R which is allowed to make adaptive oracle queries, it must make at least q_1 oracle queries or use some linear-size advice. This implies that when doing hardness amplification in this way and considering adversaries as uniform (or slightly non-uniform) algorithms, one can only guarantee the hardness of the new function $\bar{f}$ against adversaries with a computation time smaller by a q_1 factor, so a security loss of this factor is in fact unavoidable. Our second result shows that for any R which can only make non-adaptive queries, it must again makes at least q_1 oracle queries or now use an advice of exponential length. This implies that when doing hardness amplification in this way and considering adversaries as non-uniform circuits of some small exponential size, one can only guarantee the hardness of the new function $\bar{f}$ against adversaries with a circuit size smaller by a q_1 factor, so a security loss of this factor is again unavoidable.

We next consider the task of weakly-black-box construction of PRG from OWF, which transforms $(1-\delta)$-hard functions into ε-random functions. Our third and forth results show that any algorithm R realizing such a construction must make at least $q_2 = \Omega(n/\varepsilon^2)$ queries, unless it can use a long advice. More precisely, our third result shows that for any R which is allowed to make adaptive oracle queries, it must make at least q_2 oracle queries or use some linear-size advice. Again, this implies that when constructing PRG in this way and considering adversaries as uniform (or slightly non-uniform) algorithms, a security loss of a q_2 factor is in fact unavoidable. Finally, our forth result shows that for any R which

can only make non-adaptive queries, it must again make at least q_2 oracle queries or now use an advice of exponential length. Again, this implies that when constructing PRG in this way and considering adversaries as non-uniform circuits of some small exponential size, a security loss of a q_2 factor is also unavoidable.

We remark that in a different setting, Shaltiel and Viola [19] recently showed that for the task of amplifying the hardness of computing Boolean functions (instead of inverting one-way functions), a security loss in terms of circuit size is also unavoidable when this is done in a black-box way. However, they only considered the case that the oracle algorithm R makes non-adaptive queries, and there seem to be further complications when dealing with inverting one-way functions. On the other hand, our proof (for hardness amplification, with R making non-adaptive queries) is different in spirit from theirs, and our proof can be modified to give an alternative (and perhaps more intuitive) proof to their result.

2 Preliminaries

For $n \in \mathbb{N}$, let $[n]$ denote the set $\{1, 2, \ldots, n\}$ and let $\mathcal{U}_n$ denote the uniform distribution over $\{0,1\}^n$. We will consider the computational model of non-uniform oracle algorithms. For such an algorithm R, let $R^{f;\alpha}$ denote the algorithm R using f as an oracle and α as an advice.

For a many-to-one function $f : \{0,1\}^n \to \{0,1\}^m$ and an algorithm $M : \{0,1\}^m \to \{0,1\}^n$, we say that M inverts $f(x)$, denoted as $M(f(x)) \equiv x$, if $M(f(x)) \in f^{-1}(f(x))$, and we say that M can α-invert f, if

$$\Pr_{x \in \mathcal{U}_n} [M(f(x)) \equiv x] \geq \alpha.$$

For a function $g : \{0,1\}^n \to \{0,1\}^m$ and an algorithm $D : \{0,1\}^m \to \{0,1\}$, we say that D can ε-distinguish g if

$$\left| \Pr_{x \in \mathcal{U}_n} [D(g(x)) = 1] - \Pr_{u \in \mathcal{U}_m} [D(u) = 1] \right| \geq \varepsilon.$$

One can allow M and D to be probabilistic, in which case the probabilities above are taken also over their randomness.

Next, let us introduce two types of black-box transformations which we will study in this paper. Note that in a usual black-box model, both the construction and the security proof are required to be done in a black-box way. Here we consider weaker models which only require the security proof to be done in a black-box way.

Definition 1. *A weakly-black-box hardness amplification from (ε, n, m)-hardness to $(\bar{\varepsilon}, \bar{n}, \bar{m})$-hardness consists of a non-uniform oracle algorithm R satisfying the following condition. For any function $f : \{0,1\}^n \to \{0,1\}^m$, there exists a function $\bar{f} : \{0,1\}^{\bar{n}} \to \{0,1\}^{\bar{m}}$ such that*

- *given any $M : \{0,1\}^{\bar{m}} \to \{0,1\}^{\bar{n}}$ which can $(1-\bar{\varepsilon})$-invert $\bar{f}$, there exists an advice α such that $R^{f,M;\alpha}$ can $(1-\varepsilon)$-invert f.*

Definition 2. *A weakly-black-box transformation from (ε, n, m)-hardness to $(\bar{\varepsilon}, \bar{n}, \bar{m})$-randomness consists of a non-uniform oracle algorithm R satisfying the following condition. For any function $f : \{0,1\}^n \to \{0,1\}^m$, there exists a function $g : \{0,1\}^{\bar{n}} \to \{0,1\}^{\bar{m}}$, with $\bar{m} \geq \bar{n} + 1$, such that*

- *given any $D : \{0,1\}^{\bar{m}} \to \{0,1\}$ which can $\bar{\varepsilon}$-distinguish g, there exists an advice α such that $R^{f,D;\alpha}$ can $(1-\varepsilon)$-invert f.*

In the two definitions above, the oracle algorithm R in general is allowed to make *adaptive* queries, which can depend on the answers from previous queries. We will also consider the case requiring that R only makes *non-adaptive* queries, which do not depend on the answers from previous queries but can depend on the input and the advice.

We will need the following (known) fact that a randomly chosen function is likely to be hard to invert. The proof can be modified from those in, e.g., [3, 22], which we omit here.

Lemma 1. *Let c any constant such that $0 < c < 1$, and let C be any non-uniform oracle algorithm which uses an advice of length at most 2^{cn} and makes at most 2^{cn} queries to the oracle. Then there is a constant $c_1 > 0$ such that if we sample a random function $f : \{0,1\}^n \to \{0,1\}^m$,*

$$\Pr_f \left[\exists \alpha : C^{f;\alpha} \text{ can } 2^{-c_1 n}\text{-invert } f \right] \leq 2^{-2^{\Omega(n)}}.$$

Finally, we will rely on the following lemma, which gives a large deviation bound for a sequence of random variables with a sparse dependency relationship. This may have some interest of its own.

Lemma 2. *Suppose $Z_1, \ldots, Z_k$ is a sequence of binary random variables such that for each $i \in [k]$, $\mathrm{E}[Z_i] = \mu_i$ and Z_i is mutually independent of all but $d-1$ other random variables. Then for any even $t \in \mathbb{N}$ and for any $A \in \mathbb{N}$, $\Pr[|\sum_{i\in[k]} Z_i - \sum_{i\in[k]} \mu_i| \geq A] \leq 2(4tdk/A^2)^{t/2}$.*

Due to the space limitation, we omit the proof here. The idea is that $\Pr[|\sum_{i\in[k]} Z_i - \sum_{i\in[k]} \mu_i| \geq A] \leq \mathrm{E}[(\sum_{i\in[k]}(Z_i - \mu_i))^t]/A^t$, and the numerator equals $\sum_{i_1,\ldots,i_t\in[k]} \mathrm{E}[\prod_{i\in\{i_1,\ldots,i_t\}}(Z_i - \mu_i)]$ which have most of the terms equal to zero.

3 Hardness Amplification

In this section, we show that any algorithm R realizing a weakly-black-box hardness amplification must make many queries, unless it can use a long advice. Our first result, Theorem 1 below, shows such a query lower bound for any R which is allowed to use an advice of linear length and to make adaptive queries. We will give the proof in Subsection 3.1.

Theorem 1. *Suppose an algorithm R uses an advice of length $\bar{\ell}$ and realizes a weakly-black-box hardness amplification from (ε, n, m)-hardness to $((1-\delta), \bar{n}, \bar{m})$-hardness, with $2^{-cn} \leq \varepsilon, \delta \leq c$ and $\bar{\ell} \leq cn$ for a small enough constant $c > 0$. Then R must make at least $\Omega((1/\delta)\log(1/\varepsilon))$ oracle queries.*

Our second result, Theorem 2 below, shows a query lower bound for any R which is even allowed an advice of exponential length but can only make non-adaptive queries. We will give the proof in Subsection 3.2.

Theorem 2. *Suppose an algorithm R uses an advice of length ℓ and realizes a weakly-black-box hardness amplification from (ε, n, m)-hardness to $((1-\delta), \bar{n}, \bar{m})$-hardness, with $2^{-cn} \leq \varepsilon, \delta \leq c$ and $\ell \leq 2^{cn}$ for a small enough constant $c > 0$. Then R must make at least $\Omega((1/\delta)\log(1/\varepsilon))$ oracle queries, if it only makes* non-adaptive *queries.*

3.1 Proof of Theorem 1

Consider any non-uniform oracle algorithm R which realizes such a hardness amplification. Assume that R makes at most $q \leq (c_0/\delta)\log(1/\varepsilon)$ oracle queries, for a small enough constant c_0, and we will show that this leads to a contradiction. The basic idea is that when R makes only a small number of queries, it is easy to get confused between some useful oracle $M : \{0,1\}^{\bar{m}} \to \{0,1\}^{\bar{n}}$ (which is correlated with f) and a useless one $\bar{0} : \{0,1\}^{\bar{m}} \to \{0,1\}^{\bar{n}}$ (which is independent of f). Here, we take $\bar{0}$ to be the all-zero function, where $\bar{0}(\bar{y}) = 0^{\bar{n}}$ for any $\bar{y} \in \{0,1\}^{\bar{m}}$. We will first describe a natural approach which will encounter two obstacles, and we will then show how to modify the approach to overcome the obstacles.

First, we would like to pick a function $f : \{0,1\}^n \to \{0,1\}^m$, such that f is hard to invert by $R^{f,\bar{0};\alpha}$ for any advice α, and the corresponding harder function $\bar{f}$ does not map many inputs into a small subset. Its existence is guaranteed by the following lemma.

Lemma 3. *There exists a function $f : \{0,1\}^n \to \{0,1\}^m$ satisfying the following two conditions:*

1. *for any advice* $\alpha \in \{0,1\}^{\bar{\ell}}$, $\Pr_{x \in \mathcal{U}_n}\left[R^{f,\bar{0};\alpha}(f(x)) \equiv x\right] \leq 2^{-\Omega(n)}$, *and*
2. *for any set* $S \subseteq \{0,1\}^{\bar{m}}$ *with* $|S| \leq 2^{(3/4)n}$, $\Pr_{\bar{x} \in \mathcal{U}_{\bar{n}}}\left[\bar{f}(\bar{x}) \in S\right] \leq \delta$.

Proof. First, note that giving R the oracle $\bar{0}$ does not help as any query to it can be answered by R itself without actually querying $\bar{0}$. Next, observe that for any f such that the corresponding function $\bar{f}$ does not satisfy the second condition, witnessed by the set S, the function C, defined as

$$C(\bar{y}) = \begin{cases} \text{any element from } \bar{f}^{-1}(\bar{y}) & \text{if } \bar{y} \in S, \\ 0^{\bar{n}} & \text{otherwise,} \end{cases}$$

can δ-invert $\bar{f}$ which implies that $R^{f,C;\alpha}$ can $(1-\varepsilon)$-invert f for some advice $\alpha \in \{0,1\}^{\bar{\ell}}$. Note that such a function C can be described by $|S|(\bar{m}+\bar{n})$ bits, so it can be replaced by an additional advice of that length. Thus, we can obtain from R another algorithm $\bar{R}$ which uses an advice of length at most $\bar{\ell}+|S|(\bar{m}+\bar{n}) \leq 2^{(4/5)n}$ such that if a function f fails on one of the conditions, then we have $\Pr_x[\bar{R}^{f;\bar{\alpha}}(f(x)) \equiv x] > 2^{-\Omega(n)}$ for some advice $\bar{\alpha}$. By Lemma 1, the fraction of such f's is less than one, which implies the existence of an f satisfying both conditions. □

Fix one such function f guaranteed by the lemma, and let $\bar{f} : \{0,1\}^{\bar{n}} \to \{0,1\}^{\bar{m}}$ be the corresponding harder function. To find an inverter for $\bar{f}$, we start from the function $\bar{f}^{-1}$, which clearly 1-inverts $\bar{f}$, and since it suffices to δ-invert $\bar{f}$, we can afford to destroy most of its outputs. More precisely, let $M : \{0,1\}^{\bar{m}} \to \{0,1\}^{\bar{n}}$ be the *probabilistic* function (or equivalently, a distribution over deterministic functions) such that independently for any $\bar{y} \in \{0,1\}^{\bar{m}}$,

$$M(\bar{y}) = \begin{cases} \text{any element from } \bar{f}^{-1}(\bar{y}) & \text{with probability } 3\delta, \\ 0^{\bar{n}} & \text{with probability } 1-3\delta, \end{cases}$$

where we let $\bar{f}^{-1}(\bar{y}) = \{0^{\bar{n}}\}$ when $\bar{y} \notin \textsc{Image}(\bar{f})$. Then by a Markov inequality, we can have the following lemma showing that with a good probability, M inverts $\bar{f}$ well and thus can help R for inverting f.

Lemma 4. $\Pr_M\left[M \text{ } 2\delta\text{-}inverts \text{ } \bar{f}\right] \geq \delta$.

On the other hand, we would like to show that M is unlikely to help R for inverting f. More precisely, we would like to show that for any advice $\alpha \in \{0,1\}^{\bar{\ell}}$, the probability (over M) that $R^{f,M;\alpha}$ $(1-\varepsilon)$-inverts f is very small. The idea is that when R only makes a small number of queries, it has some chance of confusing the (useful) oracle M with the (useless) oracle $\bar{0}$.

Let us fix an advice $\alpha \in \{0,1\}^{\bar{\ell}}$ now. Consider the binary random variables V_x^α, for $x \in \{0,1\}^n$, defined as

- $V_x^\alpha = 1$ if and only if $R^{f,M;\alpha}(f(x)) \not\equiv x$.

Then we would like to give an upper bound on the probability

$$\Pr_M\left[\Pr_x\left[R^{f,M;\alpha}(f(x)) \not\equiv x\right] \le \varepsilon\right] = \Pr_M\left[\sum_{x\in\{0,1\}^n} V_x^\alpha \le \varepsilon 2^n\right].$$

However, a Markov inequality can only give an upper bound about $1-2\varepsilon$ (one can show that $\mathrm{E}[V_x^\alpha] \ge 3\varepsilon$ for most x), which is too large, while our goal is to have an upper bound of $2^{-\Omega(n)}$. For this we would like to apply Lemma 2.

However, there seem to be some obstacles preventing us from applying Lemma 2. First, many of these random variables may all depend on each other because the corresponding computations of R may all query M on some common entry. Second, even for a fixed x, the queries made by $R^{f,M;\alpha}(f(x))$ can vary for different M, as we allow R to make adaptive queries which can depend on the answers of previous queries.

To deal with the second obstacle, we consider another set of random variables Z_x^α, for $x \in \{0,1\}^n$, defined as

- $Z_x^\alpha = 1$ if and only if $M(\bar{y}) = 0^{\bar{n}}$ for any query $\bar{y}$ made by $R^{f,\bar{0};\alpha}(f(x))$.

Note that $R^{f,M;\alpha}(f(x)) = R^{f,\bar{0};\alpha}(f(x))$ if $Z_x^\alpha = 1$. Thus, for any x in the set $\textsc{Bad}^\alpha$, defined as

$$\textsc{Bad}^\alpha = \left\{x \in \{0,1\}^n : R^{f,\bar{0};\alpha}(f(x)) \not\equiv x\right\},$$

we have $Z_x^\alpha \le V_x^\alpha$, because $V_x^\alpha = 1$ if $Z_x^\alpha = 1$. Furthermore, by Lemma 3, $|\textsc{Bad}^\alpha| \ge 2^n(1-2^{-\Omega(n)}) \ge 2^n/2$.

Even when working with the variables Z_x^α's, we still face the first obstacle discussed above. To deal with this, we fix the values of M at those frequently queried entries. Call $\bar{y} \in \{0,1\}^{\bar{m}}$ *heavy* for an advice α if

$$\Pr_x\left[R^{f,\bar{0};\alpha}(f(x)) \text{ queries } \bar{0} \text{ at } \bar{y}\right] \ge w,$$

where we choose $w = 2^{-(2/3)n}$. Let $\hat{M}$ be the restriction of M defined as

$$\hat{M}(\bar{y}) = \begin{cases} 0^{\bar{n}} & \text{if } \bar{y} \text{ is heavy for some } \alpha, \\ M(\bar{y}) & \text{otherwise.} \end{cases}$$

Note that for each α, the number of heavy entries for α is at most q/w, because this number times w is at most the average number of queries made by $R^{f,\bar{0};\alpha}(f(x))$ over $x \in \{0,1\}^n$, which is at most q. Thus, the total number of all such heavy entries (over all α's) is at most $(q/w)2^{\bar{\ell}} \leq 2^{(3/4)n}$, since $2^{-cn} \leq \delta, \varepsilon$ and $\bar{\ell} \leq c\bar{n}$ for a small enough constant c. Let $\hat{V}_x^\alpha$ and $\hat{Z}_x^\alpha$ denote the random variables corresponding to V_x^α and Z_x^α, respectively, over the distribution of $\hat{M}$. Observe that now for each α, every variable $\hat{Z}_x^\alpha$ is mutually independent of all but $qw2^n$ other such variables, so it becomes possible to apply Lemma 2.

From now on, we will work with the distribution $\hat{M}$ and the random variables $\hat{V}_x^\alpha$ and $\hat{Z}_x^\alpha$, for $x \in \{0,1\}^n$. Let us see how this affects the arguments before when considering M, V_x^α, and Z_x^α. First, just as Lemma 4, we can show that such $\hat{M}$ also has a good chance of inverting $\bar{f}$ well.

Lemma 5. $\Pr_{\hat{M}}[\hat{M}$ *δ-inverts* $\bar{f}] \geq \delta$.

Proof. Let S be the set of heavy entries over all α's, which has $|S| \leq 2^{(3/4)n}$, and we know that any $\bar{y}$ such that $\hat{M}(\bar{y}) \neq M(\bar{y})$ is contained in S. Since f satisfies the second condition of Lemma 3, we have

$$\Pr_{\bar{x}}\left[\hat{M}(\bar{f}(\bar{x})) \neq M(\bar{f}(\bar{x}))\right] \leq \Pr_{\bar{x}}\left[\bar{f}(\bar{x}) \in S\right] \leq \delta.$$

Thus, if M can 2δ-invert $\bar{f}$, $\hat{M}$ can δ-invert $\bar{f}$. Then from Lemma 4, we have the lemma. □

From this lemma and the guarantee of R, we have

$$\Pr_{\hat{M}}\left[\exists \alpha : R^{f,\hat{M};\alpha} \text{ } (1-\varepsilon)\text{-inverts } f\right] \geq \delta$$

which implies the existence of an advice $\alpha \in \{0,1\}^{\bar{\ell}}$ such that

$$\Pr_{\hat{M}}\left[R^{f,\hat{M};\alpha} \text{ } (1-\varepsilon)\text{-inverts } f\right] \geq \delta 2^{-\bar{\ell}}. \tag{1}$$

Let's fix one such α. On the other hand, $\Pr_{\hat{M}}[R^{f,\hat{M};\alpha} \text{ } (1-\varepsilon)\text{-inverts } f]$ is

$$\Pr_{\hat{M}}\left[\sum_{x\in\{0,1\}^n} \hat{V}_x^\alpha \leq \varepsilon 2^n\right] \leq \Pr_{\hat{M}}\left[\sum_{x\in \textsc{Bad}^\alpha} \hat{V}_x^\alpha \leq \varepsilon 2^n\right],$$

and since we still have $\hat{Z}_x^\alpha \leq \hat{V}_x^\alpha$ for any $x \in \textsc{Bad}^\alpha$, the above is at most

$$\Pr_{\hat{M}}\left[\sum_{x\in \textsc{Bad}^\alpha} \hat{Z}_x^\alpha \leq \varepsilon 2^n\right] \leq \Pr_{\hat{M}}\left[\sum_{x\in \textsc{Bad}^\alpha} \hat{Z}_x^\alpha \leq 2\varepsilon \left|\textsc{Bad}^\alpha\right|\right]$$

as $|\textsc{Bad}^\alpha| \geq 2^n/2$. Then we bound the last probability by the following.

Lemma 6. $\Pr_{\hat{M}}[\sum_{x\in \textsc{Bad}^\alpha} \hat{Z}_x^\alpha \le 2\varepsilon|\textsc{Bad}^\alpha|] < \delta 2^{-\bar{\ell}}$.

Proof. Note that for any $x \in \{0,1\}^n$,

$$\Pr_{\hat{M}}\left[\hat{Z}_x^\alpha = 1\right] \ge (1-3\delta)^q \ge 3\varepsilon,$$

as we assume $\varepsilon, \delta \le c$ and $q \le (c_0/\delta)\log(1/\varepsilon)$, for small enough constants c, c_0. By Lemma 2 with $k = |\textsc{Bad}^\alpha| \ge 2^n/2$, $A = \varepsilon k$, $d = 1+qw2^n \le 2^{n/2}$, and $t = \varepsilon^2 2^{n/2}/16$, we have

$$\Pr_{\hat{M}}\left[\sum_{x\in \textsc{Bad}^\alpha} \hat{Z}_x^\alpha \le 2\varepsilon\,|\textsc{Bad}^\alpha|\right] \le 2\left(\frac{\varepsilon^2 2^n}{4\varepsilon^2 k}\right)^{t/2} \le 2\left(\frac{1}{2}\right)^{t/2} < \delta 2^{-\bar{\ell}},$$

since $2^{-cn} \le \delta, \varepsilon$ and $\bar{\ell} \le cn$ for a small enough constant c. (Note that the bound still holds even if we allow $\bar{\ell} \le 2^{cn}$.) □

This leads to a contradiction to the bound in (1). Therefore, the assumption we made at the beginning cannot hold, and R must make $\Omega((1/\delta)\log(1/\varepsilon))$ oracle queries, which proves Theorem 1.

3.2 Proof of Theorem 2

Note that what is different from Theorem 1 is that now we require that R makes only *non-adaptive* queries and as a result we allow R a much longer advice. Again, assume that R makes at most $q \le (c_0/\delta)\log(1/\varepsilon)$ oracle queries, for a small enough constant $c_0 > 0$, and we will show that this leads to a contradiction.

Observe that in the proof of Theorem 1, why we can only have $\bar{\ell} \le O(\bar{n})$ is that the restriction $\hat{M}$ fixes all the heavy entries over all α's at once, but there are $(q/w)2^{\bar{\ell}}$ such entries which can not exceed $2^{\bar{n}}$. Here, instead, we will consider different restrictions for different α's separately. Call $\bar{y} \in \{0,1\}^{\bar{m}}$ heavy for an advice α if

$$\Pr_x\left[R^{f,M;\alpha}(f(x)) \text{ queries } M \text{ at } \bar{y}\right] \ge w,$$

where $w = 2^{-(2/3)n}$, and note that this definition actually is independent of the choice of f and M as we assume that R makes only non-adaptive queries. As in the proof of Theorem 1, one can show that the number of heavy entries for any advice α is at most q/w. We will consider restricting an oracle function M (or $\bar{0}$) by fixing its values on those heavy entries for

some α according to some function $\rho : \{0,1\}^{\bar{m}} \to \{0,1\}^{\bar{n}}$. For an advice α and a function $\rho : \{0,1\}^{\bar{m}} \to \{0,1\}^{\bar{n}}$, let M_ρ^α denote such a restriction of M, defined as $M_\rho^\alpha(\bar{y}) = \rho(\bar{y})$ if $\bar{y}$ is heavy for α and $M_\rho^\alpha(\bar{y}) = M(\bar{y})$ otherwise. Similarly, let $\bar{0}_\rho^\alpha$ denote such a restriction of $\bar{0}$. As in Lemma 3, we have the following lemma. Due to the space limitation, we omit its proof (which follows from Lemma 1, as $\bar{0}_\rho^\alpha$ has a short description and can be replaced by a short additional advice).

Lemma 7. *There exists a function $f : \{0,1\}^n \to \{0,1\}^m$ such that for any advice $\alpha \in \{0,1\}^\ell$ and any function $\rho : \{0,1\}^{\bar{m}} \to \{0,1\}^{\bar{n}}$,*

$$\Pr_x \left[R^{f,\bar{0}_\rho^\alpha;\alpha}(f(x)) \equiv x \right] \leq 2^{-\Omega(n)}.$$

Let us pick one such function f guaranteed by the lemma, let $\bar{f}$ be the corresponding harder function, and let M be the probabilistic function defined according to this $\bar{f}$. As in Lemma 4, we have $\Pr_M \left[M \text{ } \delta\text{-inverts } \bar{f} \right] \geq \delta$, and as before, this implies the existence of an advice $\alpha \in \{0,1\}^\ell$ such that $\Pr_M \left[R^{f,M;\alpha} \text{ } (1-\varepsilon)\text{-inverts } f \right] \geq \delta 2^{-\ell}$. Let us fix one such advice α. By an average argument, there must exist a restriction M_ρ^α of M which fixes the values of those heavy entries for α, such that

$$\Pr_{M_\rho^\alpha} \left[R^{f,M_\rho^\alpha;\alpha} \text{ } (1-\varepsilon)\text{-inverts } f \right] \geq \delta 2^{-\ell}. \tag{2}$$

On the other hand, we will show a contradiction to this bound. Consider the set

$$\textsc{Bad}^\alpha = \left\{ x \in \{0,1\}^n : R^{f,\bar{0}_\rho^\alpha;\alpha}(f(x)) \not\equiv x \right\},$$

which by Lemma 7 has $|\textsc{Bad}^\alpha| \geq 2^n/2$. Let $\hat{V}_x^\alpha$ denote the binary random variable, over the distribution of M_ρ^α, defined as

- $\hat{V}_x^\alpha = 1$ if and only if $R^{f,M_\rho^\alpha;\alpha}(f(x)) \not\equiv x$.

Note that each variable $\hat{V}_x^\alpha$ is mutually independent of all but $qw2^n$ other variables. Then one can again show that for any $x \in \textsc{Bad}^\alpha$, $\Pr_{M_\rho^\alpha}[\hat{V}_x^\alpha = 1] \geq (1-3\delta)^q \geq 3\varepsilon$, and $\Pr_{M_\rho^\alpha} \left[R^{f,M_\rho^\alpha;\alpha} \text{ } (1-\varepsilon)\text{-inverts } f \right]$ is

$$\Pr_{M_\rho^\alpha} \left[\sum_{x \in \{0,1\}^n} \hat{V}_x^\alpha \leq \varepsilon 2^n \right] \leq \Pr_{M_\rho^\alpha} \left[\sum_{x \in \textsc{Bad}^\alpha} \hat{V}_x^\alpha \leq 2\varepsilon \left| \textsc{Bad}^\alpha \right| \right] < \delta 2^{-\ell},$$

by Lemma 2. This contradicts the bound in (2). Thus, R must make at least $\Omega((1/\delta)\log(1/\varepsilon))$ queries, which proves Theorem 2.

4 Randomness from Hardness

In this section, we show that any algorithm R realizing a weakly-black-box transformation from hardness to randomness must make many queries, unless it can use a long advice string. Our first result, Theorem 3 below, shows such a query lower bound for any R which is allowed to use an advice of linear length and to make adaptive queries. We will give the proof in Subsection 4.1.

Theorem 3. *Suppose an algorithm R uses an advice of length $\bar{\ell}$ and realizes a weakly-black-box transformation from $((1-\delta), n, m)$-hardness to $(\varepsilon, \bar{n}, \bar{m})$-randomness, with $2^{-cn} \leq \varepsilon, \delta \leq c$ and $\bar{\ell} \leq cn$ for a small enough constant $c > 0$. Then R must make at least $\Omega(n/\varepsilon^2)$ oracle queries.*

Our second result, Theorem 4 below, shows a query lower bound for any R which is even allowed an advice of exponential length but can only make non-adaptive queries. We will give the proof in Subsection 4.2.

Theorem 4. *Suppose an algorithm R uses an advice of length ℓ and realizes a weakly-black-box transformation from $((1-\delta), n, m)$-hardness to $(\varepsilon, \bar{n}, \bar{m})$-randomness, with $2^{-cn} \leq \varepsilon, \delta \leq c$ and $\ell \leq 2^{c\bar{n}}$ for a small enough constant $c > 0$. Then R must make at least $\Omega(n/\varepsilon^2)$ oracle queries, if it only makes* non-adaptive *queries.*

4.1 Proof of Theorem 3

Consider any R which realizes such a weakly-black-box transformation. Assume that R makes at most $q \leq c_0 n/\varepsilon^2$ oracle queries, for a small enough constant $c_0 > 0$, and we will show that this leads to a contradiction. The basic idea is that when R makes only a small number of queries, it is easy to get confused between some useful oracle $D : \{0,1\}^{\bar{m}} \to \{0,1\}$ (which is correlated with f) and a useless one $B : \{0,1\}^{\bar{m}} \to \{0,1\}$ (which is independent of f). Here, we take B to be a random function.

First, we would like to pick a function f which is hard to invert by $R^{f,B;\alpha}$ for any α. The existence of such a function is guaranteed by the following lemma, which follows from Lemma 1 by observing that the oracle B, which is independent of f, can be simulated by R itself without needing to query it.

Lemma 8. *There exists a function $f : \{0,1\}^n \to \{0,1\}^m$ such that for any advice $\alpha \in \{0,1\}^{\bar{\ell}}$, $\Pr_{x,B}\left[R^{f,B;\alpha}(f(x)) \equiv x\right] \leq 2^{-\Omega(n)}$.*

Fix one such function f guaranteed by the lemma, and let $g = g_f : \{0,1\}^{\bar{n}} \to \{0,1\}^{\bar{m}}$ be the resulting generator. To find a distinguisher for g, let us start from the characteristic function of $\textsc{Image}(g)$, denoted as T (i.e., $T(u) = 1$ if and only if $u \in \textsc{Image}(g)$), which clearly can $(1 - 2^{-(\bar{m}-\bar{n})})$-distinguish g, and since we only need to ε-distinguish g, we can afford to add some noise to T. More precisely, let $N : \{0,1\}^{\bar{m}} \to \{0,1\}$ be a noise function such that independently for any $u \in \{0,1\}^{\bar{m}}$,

$$N(u) = \begin{cases} 1 & \text{with probability } \frac{1-4\varepsilon}{2}, \\ 0 & \text{with probability } \frac{1+4\varepsilon}{2}, \end{cases}$$

and consider the *probabilistic* distinguisher (or equivalently a distribution over deterministic distinguishers) $D = T \oplus N$ such that for any $u \in \{0,1\}^{\bar{m}}$, $D(u) = T(u) \oplus N(u)$. The following shows that D has a good chance of distinguishing g well.

Lemma 9. $\Pr_D[D$ *ε-distinguishes* $g] \geq \varepsilon$.

Proof. From the definition, $\Pr_{x,D}[D(g(x)) = 1] - \Pr_{u,D}[D(u) = 1]$ is

$$\frac{1+4\varepsilon}{2} - \left(2^{-(\bar{m}-\bar{n})} \cdot \frac{1+4\varepsilon}{2} + \left(1 - 2^{-(\bar{m}-\bar{n})}\right) \cdot \frac{1-4\varepsilon}{2}\right) \geq 2\varepsilon,$$

since $\bar{m} \geq \bar{n} + 1$ and hence $2^{-(\bar{m}-\bar{n})} \leq \frac{1}{2}$. Then by a Markov inequality, we have the lemma. □

As before, from this lemma and the guarantee of R, one can show the existence of an advice $\alpha \in \{0,1\}^{\bar{\ell}}$ such that

$$\Pr_D\left[R^{f,D;\alpha} \text{ } \delta\text{-inverts } f\right] \geq \varepsilon 2^{-\bar{\ell}}. \tag{3}$$

Let us fix one such advice α.

On the other hand, we will show that this bound cannot hold as D is unlikely to help R for inverting f. The idea is that when R only makes a small number of queries, it behaves similarly according to the two different oracles D and B. More precisely, we will show that for most input x,

$$\Pr_D\left[R^{f,D;\alpha}(f(x)) \equiv x\right] \text{ is small } \textit{if } \Pr_B\left[R^{f,B;\alpha}(f(x)) \equiv x\right] \text{ is small.}$$

Here we choose not to show that the two probabilities have a small difference, as it would then only give a much smaller query lower bound.

Let c_1 be a small enough constant, and consider the set

$$\textsc{Bad}^{\alpha} = \left\{ x \in \{0,1\}^n : \Pr_B \left[R^{f,B;\alpha}(f(x)) \equiv x \right] \leq 2^{-c_1 n} \right\}. \tag{4}$$

From Lemma 8 and a Markov inequality, we have $\Pr_x [x \notin \textsc{Bad}^{\alpha}] \leq 2^{-\Omega(n)}$. Furthermore, we have the following.

Lemma 10. *For any $x \in \textsc{Bad}^{\alpha}$, $\Pr_D[R^{f,D;\alpha}(f(x)) \equiv x] \leq 2^{-\Omega(n)}$.*

Proof. Fix any $x \in \textsc{Bad}^{\alpha}$. First, let us consider the computations of $R^{f,B;\alpha}(f(x))$ over all possible instances of the oracle B, which can be seen as a tree in a natural way as follows. Each internal node corresponds to a query $u \in \{0,1\}^{\bar{m}}$ to the oracle B, which has two edges coming out for the two possible answers of $B(u)$, and each leaf contains the output of $R^{f,B;\alpha}(f(x))$ following the corresponding path of computation. We can assume without loss of generality that R always makes exactly q queries to the oracle B (by making dummy queries if necessary), and it never makes the same query twice on any path of computation (by remembering all previous queries and their answers). The tree has exactly 2^q leaves, and the bound of (4) implies that at most $L = 2^{-c_1 n} \cdot 2^q$ leaves have $R^{f,B;\alpha}(f(x)) \equiv x$.

Now let us see what happens to the probability bound in (4) when we change the oracle from B to D. While over a random B, each leaf is reached with the same probability 2^{-q}, this no longer holds if now we measure the probability over the distribution of D. Note that each edge corresponds to the bit $D(u) = T(u) \oplus N(u)$, for some u, and since T is fixed (as we have fixed f and thus g), we can label that edge by the bit $N(u)$. Then a leaf on a path with i labels of 1 is now reached with probability $\left(\frac{1-4\varepsilon}{2}\right)^i \left(\frac{1+4\varepsilon}{2}\right)^{q-i}$, which increases as i decreases. Observe that the sequences of labels on the 2^q paths to the leaves are all different. Thus, the probability $\Pr_D \left[R^{f,D;\alpha}(f(x)) \equiv x \right]$ can be bounded from above by the sum of the L largest probabilities among the leaves, which is at most

$$\sum_{i=0}^{t} \binom{q}{i} \left(\frac{1-4\varepsilon}{2}\right)^i \left(\frac{1+4\varepsilon}{2}\right)^{q-i}, \tag{5}$$

for any t such that $\sum_{i=0}^{t} \binom{q}{i} \geq L$. We can take $t = (\frac{1-5\varepsilon}{2})q$, which gives

$$\sum_{i=0}^{t} \binom{q}{i} \geq 2^{-O(\varepsilon^2 q)} \cdot 2^q \geq 2^{-c_1 n} \cdot 2^q \geq L,$$

since $q \leq c_0 n/\varepsilon^2$ and we can take c_0 to be much smaller than c_1. Then the bound in (5) is at most $2^{-\Omega(\varepsilon^2 q)} \leq 2^{-\Omega(n)}$ (using, for example, a Chernoff bound), which proves the lemma. □

Now let V_x^α, for $x \in \{0,1\}^n$, denote the binary random variable, over D, such that

– $V_x^\alpha = 1$ if and only if $R^{f,D;\alpha}(f(x)) \equiv x$.

We know from Lemma 10 that for any $x \in \textsc{Bad}^\alpha$, $\Pr_D[V_x^\alpha = 1] \leq 2^{-\Omega(n)}$. Then we have

$$\mathop{\mathrm{E}}_D\left[\sum_{x\in\{0,1\}^n} V_x^\alpha\right] \leq \sum_{x\in\textsc{Bad}^\alpha} \mathop{\mathrm{E}}_D[V_x^\alpha] + \sum_{x\notin\textsc{Bad}^\alpha} 1 \leq 2^{-\Omega(n)}\cdot 2^n < \varepsilon 2^{-\bar{\ell}}\delta\cdot 2^n,$$

since we assume $2^{-cn} \leq \varepsilon, \delta$ and $\bar{\ell} \leq cn$ for a small enough constant c. By a Markov inequality, we have

$$\Pr_D\left[R^{f,D;\alpha} \text{ } \delta\text{-inverts } f\right] = \Pr_D\left[\sum_x V_x^\alpha \geq \delta 2^n\right] < \varepsilon 2^{-\bar{\ell}}, \tag{6}$$

which contradicts the bound in (3). This implies that R must make at least $\Omega(n/\varepsilon^2)$ oracle queries, which proves Theorem 3.

4.2 Proof of Theorem 4

Note that what is different from Theorem 3 is that now we require that R makes only *non-adaptive* queries and as a result we allow R a much longer advice. Again, assume that R makes at most $q \leq c_0 n/\varepsilon^2$ oracle queries, for a small enough constant $c_0 > 0$, and we will show that this leads to a contradiction. The proof here will follow closely that for Theorem 3, except that at the end we will manage to get a smaller bound than that in inequality (6), by applying Lemma 2 instead of a Markov inequality. The idea is similar to that in the proof of Theorem 2.

Let $w = 2^{-(2/3)n}$, call an entry $\bar{y} \in \{0,1\}^{\bar{m}}$ heavy for an advice α if $\Pr_x[R^{f,D;\alpha}(f(x))$ queries D at $\bar{y}] \geq w$, and again one can show that the number of heavy entries for any α is at most q/w. Note that this definition is independent of the choice of D and f as we assume that R only makes non-adaptive queries. As in the proof of Theorem 2, we will consider restricting the functions D and B by fixing their values on those heavy entries for some α according to some function $\rho : \{0,1\}^{\bar{m}} \to \{0,1\}^{\bar{n}}$, and let D_ρ^α and B_ρ^α denote the corresponding restrictions, respectively.

Then similar to Lemma 7 and Lemma 8, one can show the existence of a function $f : \{0,1\}^n \to \{0,1\}^m$ such that for any advice $\alpha \in \{0,1\}^\ell$ and any function $\rho : \{0,1\}^{\bar{m}} \to \{0,1\}$,

$$\Pr_{x, B_\rho^\alpha}\left[R^{f, B_\rho^\alpha;\alpha}(f(x)) \equiv x\right] \leq 2^{-\Omega(n)}. \tag{7}$$

Let f be such a function guaranteed above, let $g = g_f$ be the resulting generator, and let D be the probabilistic distinguisher defined according to this g. As in Lemma 9, we have $\Pr_D[D \text{ } \varepsilon\text{-distinguishes } g] \geq \varepsilon$. Then as in the proof of Theorem 2, one can show that this implies the existence of an advice α and a restriction D_ρ^α of D which fixes the values of those heavy entries for α, such that

$$\Pr_{D_\rho^\alpha}\left[R^{f, D_\rho^\alpha;\alpha} \text{ } \delta\text{-inverts } f\right] \geq \varepsilon 2^{-\ell}. \tag{8}$$

Let us fix one such α and ρ.

On the other hand, we will show that the bound above can not hold. Let c_1 be a small enough constant, and consider the set:

$$\textsc{Bad}^\alpha = \left\{x \in \{0,1\}^n : \Pr_{B_\rho^\alpha}\left[R^{f, B_\rho^\alpha;\alpha}(f(x)) \equiv x\right] \leq 2^{-c_1 n}\right\}.$$

By the bound in (7) and a Markov inequality, we have $\Pr_x[x \in \textsc{Bad}^\alpha] \geq 1/2$. Let $\hat{V}_x^\alpha$ denote the binary random variable, over the distribution of D_ρ^α, defined as

- $\hat{V}_x^\alpha = 1$ if and only if $R^{f, D_\rho^\alpha;\alpha}(f(x)) \equiv x$.

Note that each variable $\hat{V}_x^\alpha$ is mutually independent of all but $qw2^n$ other variables. Then using an argument similar to that in the proof of Theorem 3, one can show that for any $x \in \textsc{Bad}^\alpha$, $\Pr_{D_\rho^\alpha}[\hat{V}_x^\alpha = 1] \leq 2^{-\Omega(n)} \leq \delta/3$, and by Lemma 2, $\Pr_{D_\rho^\alpha}[R^{f, D_\rho^\alpha;\alpha} \text{ } \delta\text{-inverts } f]$ is

$$\Pr_{D_\rho^\alpha}\left[\sum_{x \in \{0,1\}^n} \hat{V}_x^\alpha \geq \delta 2^n\right] \leq \Pr_{D_\rho^\alpha}\left[\sum_{x \in \textsc{Bad}^\alpha} \hat{V}_x^\alpha \geq (\delta/2)\left|\textsc{Bad}^\alpha\right|\right] < \varepsilon 2^{-\ell},$$

which contradicts the bound in (8). As a result, R must make at least $\Omega(n/\varepsilon^2)$ queries, which proves Theorem 4.

References

1. Y.-C. Chang, C.-Y. Hsiao, and C.-J. Lu. The impossibility of basing one-way permutations on central cryptographic primitives. *J. Cryptology*, 19(1), pp. 97–114, 2006.
2. G. Di Crescenzo, T. Malkin, and R. Ostrovsky. Single database private information retrieval implies oblivious transfer. In *Proc. EUROCRYPT'00*, pp. 122–138.
3. R. Gennaro, Y. Gertner, J. Katz, and L. Trevisan. Bounds on the efficiency of generic cryptographic constructions. *SIAM J. Comput.*, 35(1), pp. 217–246, 2005.
4. Y. Gertner, T. Malkin, and O. Reingold. On the impossibility of basing trapdoor functions on trapdoor predicates. In *Proc. IEEE FOCS'01*, pp. 126–135.
5. Y. Gertner, S. Kannan, T. Malkin, O. Reingold, and M. Viswanathan. The relationship between public key encryption and oblivious transfer. In *Proc. IEEE FOCS'00*, pp. 325–335.
6. O. Goldreich, S. Goldwasser, and S. Micali. How to construct random functions. *J. ACM*, 33(4), pp. 792–807, 1986.
7. O. Goldreich and L. Levin. A hard-core predicate for all one-way functions. In *Proc. ACM STOC'89*, pp. 25–32.
8. O. Goldreich, R. Impagliazzo, L. Levin, Ramarathnam Venkatesan, and David Zuckerman. Security preserving amplification of hardness. In *Proc. IEEE FOCS'90*, pp. 318–326.
9. O. Goldreich, S. Micali, and A. Wigderson. Proofs that yield nothing but their validity and a methodology of cryptographic protocol design. In *Proc. IEEE FOCS'86*, pp. 174–187.
10. J. Håstad, R. Impagliazzo, L. A. Levin, and M. Luby. A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28(4), pp. 1364–1396, 1999.
11. R. Impagliazzo and M. Luby. One-way functions are essential for complexity based cryptography. In *Proc. IEEE FOCS'89*, pp. 230–235.
12. R. Impagliazzo and S. Rudich. Limits on the provable consequences of one-way permutations. In *Proc. ACM STOC'89*, pp. 44–61.
13. J. Kilian. Founding cryptography on oblivious transfer. In *Proc. ACM STOC'98*, pp. 20–31.
14. E. Kushilevitz and R. Ostrovsky. One-way trapdoor permutations are sufficient for non-trivial single-server private information retrieval. In *Proc. EUROCRYPT'00*, pp. 104–121.
15. H. Lin, L. Trevisan, and H. Wee. On hardness amplification of one-way functions. In *Proc. TCC'05*, pp. 34–49.
16. C.-J. Lu. On the complexity of parallel hardness amplification for one-way functions. In *Proc. TCC'06*, pp. 462–481.
17. M. Naor. Bit commitment using pseudorandomness. *J. Cryptology*, 4(2), pp. 151–158, 1991.
18. J. Rompel. One-way functions are necessary and sufficient for secure signatures. In *Proc. ACM STOC'90*, pp. 387–394.
19. R. Shaltiel and E. Viola. Hardness amplification proofs require majority. In *Proc. ACM STOC'08*, pp. 589–598.
20. E. Viola. On constructing parallel pseudorandom generators from one-way functions. In *Proc. IEEE CCC'05*, pp. 183–197.
21. A. C.-C. Yao. Theory and applications of trapdoor functions. In *Proc. IEEE FOCS'82*, pp. 80–91.
22. M. Zimand. Exposure-resilient extractors and the derandomization of probabilistic sublinear time. *Computational Complexity*, 17(2), pp. 220–253, 2008.