On the Power of Nonuniformity in Proofs of Security

Kai-Min Chung
Cornell University
Ithaca, NY, USA
chung@cs.cornell.edu

Mohammad Mahmoody*
Cornell University
Ithaca, NY, USA
mohammad@cs.cornell.edu

Huijia Lin†
MIT and Boston University
Cambridge, MA, USA
huijia@csail.mit.edu

Rafael Pass‡
Cornell University
Ithaca, NY, USA
rafael@cs.cornell.edu

ABSTRACT

Nonuniform proofs of security are common in cryptography, but traditional black-box separations consider only uniform security reductions. In this paper, we initiate a formal study of the power and limits of nonuniform black-box proofs of security. We first show that a known protocol (based on the existence of one-way permutations) that uses a nonuniform proof of security, and it *cannot* be proven secure through a uniform security reduction. Therefore, nonuniform proofs of security are indeed provably more powerful than uniform ones.

We complement this result by showing that many known black-box separations in the uniform regime actually do extend to the nonuniform regime. We prove our results by providing general techniques for extending certain types of black-box separations to handle nonuniformity.

General Terms

Theory, Security

Keywords

Black-Box Separations, Proofs of Security, Nonuniformity

*Supported in part by NSF Award CNS-1217821.

†Supported by DARPA grant FA8750-11-2-0225 and NSF grant CCF-1018064.

‡Supported in part by a Alfred P. Sloan Fellowship, Microsoft New Faculty Fellowship, NSF Award CNS-1217821, NSF CAREER Award CCF-0746990, AFOSR YIP Award FA9550-10-1-0093, and DARPA and AFRL under contract FA8750-11-2- 0211. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the US Government

ITCS'13, January 9–12, 2012, Berkeley, California, USA.

1. INTRODUCTION

Most of cryptography relies on unproved hardness assumptions, such as the existence of one-way functions and one-way permutations, the hardness of factoring, etc. Understanding the minimal assumptions needed for proving the security of cryptographic tasks is thus of fundamental importance. Understanding *barriers* to such proofs of security has been an active line of research in the last decades since the seminal work of Impagliazzo and Rudich [34]. Their work demonstrates barriers to providing a "black-box construction" of key-agreement from one-way permutations—a black-box construction of a primitive $\mathcal{Q}$ (e.g., key-agreement) uses a primitive $\mathcal{P}$ (e.g., one-way permutations) as an oracle, while the specific details of the implementation of $\mathcal{P}$ is ignored by the construction; additionally, the proof of security is also black-box in the sense that the security reduction only uses the presumed attacker to $\mathcal{Q}$ as a black-box in order to violate the security of $\mathcal{P}$. Subsequently, many other black-box separations between cryptographic primitives have been established (e.g., [1, 10, 23, 23, 24, 35, 50, 52]); this paradigm has also been used to demonstrate lower bounds on the *efficiency* of black-box constructions (e.g., [5,6,13,20,29,36,38]). Very recently, several works [22,44,45] demonstrated barriers to proofs of security that apply even when the construction is non-black-box (that is, the implementation of $\mathcal{Q}$ may use the code of $\mathcal{P}$ instead of just treating it as an oracle) as long as just the proof of security is black-box.

The result of [34] and its follow-ups, however, suffer from the following restriction: They rule out only constructions with *uniform* proofs of security—that is, the "security reduction" is a uniform polynomial-time algorithm. In contrast, some quite commonly used techniques in cryptography make use of *nonuniformity* in the proof of the security (see e.g., [26–28, 33, 39]). For example, a common technique is to use a hybrid argument that involves nonuniformly fixing some "prefix" of the experiment to be the "best possible" value for the reduction. This may be thought of as a mild form of non-black-box access to the code (of the primitive and the adversary) by the security reduction. In some cases, initial nonuniform proofs of security were eventually made uniform (a celebrated example is the result of of [32], making the security reduction of the pseudorandom generator construction of [33] uniform), but in general it is not clear whether nonuniformity makes proofs of security more powerful or not, leaving open the following question:

Are nonuniform proofs of security inherently more powerful than uniform ones, or can any nonuniform proof of security be made uniform?

A closely related question is whether (known) barriers for uniform security extend to the nonuniform regime:

Do (known) black-box separations handling uniform security reductions extend to handle also nonuniform security reductions?

In this work we address the above two questions. We answer the first question affirmatively—showing that a known protocol (specifically, the non-malleable commitment of [39], which is based on the existence of one-way permutations) that uses a nonuniform black-box proof of security cannot be proven secure using a uniform black-box proof of security. Thus, assuming the existence of one-way permutations, nonuniform proofs of security are provably more powerful than uniform ones.

Regarding the second question, we show that, although in general black-box separations for the uniform regime do not extend to uniform regime, many *known* black-box separations for the uniform regime actually do extend to the nonuniform regime. We prove our results by providing general techniques for extending certain types of black-box separations to handle nonuniformity. But before further explaining our result, we need to formally define nonuniform proofs of security.

1.1 Formalizing Nonuniform Proofs of Security

Let us start by recalling the formalization of a black-box proofs of security from [46]. In a black-box proof of security for a cryptographic scheme Q based on some cryptographic assumption P, a security proof S^A is an efficient oracle Turing machine that uses any adversary A who breaks Q as an oracle and breaks the security of P; we use the terms "security proof" or "security reduction" interchangeably. In a *black-box construction* [46] the implementation of the new primitive Q uses a implementation of the primitive P as an oracle, and the security reduction also may only access P as an oracle.

In order to allow the security reduction to be nonuniform, an initial approach would be to allow the security reduction S to be a nonuniform circuit (rather than an efficient Turing machine). More formally, we may extend the definition of a black-box security reduction to the nonuniform setting by requiring the existence of an efficient *circuit* S that for every adversary A breaking the security of Q, S^A breaks P. This "naive" way of defining a nonuniform proof of security, however, does not capture known nonuniform proof techniques where, for instance, the security reduction fixes some "prefix" of a computation to its "best value".[1] The reason is that for such nonuniform proofs of security, the nonuniform advice (i.e., the "best value" fixed by the reduction) can depend on the adversary A. Thus, such use of nonuniform advice may be viewed as a limited form of non-black-box use of the adversary. To also capture such techniques, we instead allow the nonuniform advice to be selected as a function of the adversary A and P.

[1] In fact, known black-box separations for the uniform regime easily extend also to this "naive" nonuniform setting.

DEFINITION 1.1 (GENERAL CONSTRUCTIONS). *A (general)* construction *of the primitive $\mathcal{Q}$ from another primitive $\mathcal{P}$ is consists of the following:*

- **Construction Mapping Q.** *There exists a mapping $Q(\cdot)$ such that: if P is an efficient implementation for $\mathcal{P}$, $Q(P)$ is an efficient implementation for $\mathcal{Q}$.*
- **Nonuniform Proof of Security.** *For every (computationally unbounded) adversary A breaking the security of the efficient construction $Q(P)$, there is a polynomial-size circuit S (which may depend on A and Q) such that S^A breaks the security of P.*

DEFINITION 1.2 (FULLY BLACK-BOX CONSTRUCTIONS). *A* fully black-box construction *of the primitive $\mathcal{Q}$ from another primitive $\mathcal{P}$ consists of:*

- **Black-box Construction Q.** *For every (computationally unbounded) oracle P implementing $\mathcal{P}$, Q^P implements $\mathcal{Q}$.*
- **Nonuniform Proof of Security.** *For every (computationally unbounded) oracle P implementing $\mathcal{P}$ and every (computationally unbounded) adversary A breaking the security of Q^P (as an implementation of $\mathcal{Q}$), there exists a polynomial-size security reduction S (which may depending on A and P) such that $S^{P,A}$ breaks the security of P as an implementation of $\mathcal{P}$.*

The above way of treating nonuniform proofs of security is closely related to a notion considered in [3,49] in the context of hardness amplification; however, in this work our focus is on non-uniform security reductions between *cryptographic primitives*.

Before describing our results we mention the work of Koblitz an Menezes [37] in which they argue that it is "unnatural and undesirable" to use non-uniform security reductions (and non-uniform security in general). We take no stand on this general claim, however, due to extensive use of nonuniform techniques in cryptographic constructions, we believe it is crucial to understand the power and limits of such techniques in a precise and formal way.

1.2 Our Results

1.2.1 Separating the power of uniform and nonuniform security reductions

Our first results shows a separation between the power of uniform and nonuniform proofs of security.

THEOREM 1.3 (INFORMALLY STATED). *There exists a construction C of commitments from one-way permutations with a nonuniform security reduction such that: the existence of a uniform security reduction for C means that the one-way permutation in use is efficiently invertible. Thus, nonuniform proofs of security are provably more powerful than uniform ones.*

As mentioned earlier, the commitment scheme we consider is not an artificial one—it is the actual non-malleable commitment scheme constructed in [39].[2] To prove the above

[2] [39] also provide a construction of a non-malleable commitment based on one-way functions; our separation does not seem to apply to this protocol.

theorem, we show that the recent framework of [44]—which proves barriers to cryptographic constructions using uniform security reductions—can be extended to rule security proofs for commitment schemes of the above type. Interestingly, as we shall shortly see, the actual result proven in [44] indeed extends to the nonuniform regime, but we can still use this framework for our separation.

Related Work. The work of Backes and Unruh [4] can be interpreted as comparing the power of uniform and nonuniform security reductions (the former are called *constructive* proofs by [4]). They present a cryptographic protocol whose security can be based on collision-resistant hash functions only through a non-uniform security reduction, but it is proved under the assumption: there are one-way permutations secure against uniform adversaries and *insecure* against nonuniform adversaries. However, using an assumption of the same flavor as the conclusion as such makes the result of [4] provide perhaps little insight into whether such a separation can be obtained from scratch. Our Theorem 1.3, on the other hand, establishes such separation *unconditionally*. Also, Lu [40] gives two examples (strong vs. weak one-way functions and pseudorandom generators vs. one-way functions) in which the security reduction (with general constructions) would need to be nonuniform if it is "too efficient" (i.e. asks "a few" oracle queries to the adversary).

1.2.2 Handling nonuniform reductions in fully black-box constructions

As mentioned, following the seminal work of Impagliazzo and Rudich [34] there has been a large body of work proving barriers to providing a "black-box construction" of various tasks from various primitives. Ten years after the work of [34], using a "reconstruction technique," Gennaro and Trevisan [21] (in the context of studying the efficiency of the black-box constructions) showed that a random permutation $P\colon \{0,1\}^n \mapsto \{0,1\}^n$ with overwhelming probability is *nonuniformly-hard* to invert. We first observe that the result of [21] can be used to directly extend the original result of [34] to rule out black-box constructions of key-agreement from one-way permutations also with respect to nonuniform proofs of security. This follows since: **(1)** [34] show how to break all key-agreement protocols through a *single* inefficient but poly(n)-query attack, and **(2)** [21] show that any fixed computationally unbounded algorithm that gets poly(n) bits of advice about the random permutation P and may ask poly(n) queries to P, has negligible probability of inverting P; thus, if a black-box construction with a nonuniform security proof had existed, we could use the poly(n)-query attacker A^P of [34] in conjunction with the nonuniform security reduction S to obtain a fixed computationally unbounded algorithm S' that inverts a random permutation P by getting poly(n) bits of advice about P and asking only poly(n)-queries to it (contradicting [21]).

The "reconstruction technique" of [21] has subsequently been employed in several other black-box separation results [19, 29, 30] and by the same argument these results also extend to the nonuniform setting.

In this work, we also establish new black-box separations results in the non-uniform setting, and a primitive that we focus on is the families of collision-resistant hash functions. We note that although nonuniform hardness results have been proved also for other (idealized forms of) cryptographic primitives [14], as far as we know, no nonuniform hardness result have been proved for families of collision-resistant hash functions, leaving open the question of whether there exists a black-box construction of a key-agreement protocol from families of collision-resistant hash functions with a nonuniform proof of security. (It is well-known that in the uniform setting, the separation of [34] easily extends to hash functions.) We start by proving such a separation also in the nonuniform setting.

THEOREM 1.4. *There is no fully black-box construction of key-agreement protocols from families of collision-resistant hash functions even with a nonuniform proof of security.*

Since public-key encryption and oblivious transfer both imply key-agreement in a black-box way [23], Theorem 1.4 extends to separate families of collision-resistant hash functions from those primitives as well. Our proof proceeds by proving a nonuniform hardness lower bound for families of collision-resistant hash functions, and next relying on the above proof template sketched for the case of one-way permutations.

THEOREM 1.5 (NONUNIFORM COLLISION RESISTANCE). *Let the function $h\colon \{0,1\}^k \times \{0,1\}^m \mapsto \{0,1\}^n$ be chosen uniformly at random and $k \geq 4n, m > n$. Then with $1 - \mathrm{negl}(n)$ probability h will be $2^{n/10}$-secure as a collision-resistant hash function family $\{h_K : K \in \{0,1\}^k\}$. Namely, any circuit of size $2^{n/10}$ with h gates can find collision in h for at most $2^{n/10}$ fraction of the keys $K \in \{0,1\}^k$.*

The proof of Theorem 1.5 follows by a simple counting argument and an application of a lemma due to Unruh [51] (see Lemma 4.2). The proof extends to any *family* version of natural cryptographic primitives.

By applying Theorem 1.5 we also extend some earlier lower bounds [5, 20] on the *efficiency* of black-box constructions from families of collision-resistant hash functions to the nonuniform regime.

THEOREM 1.6. (Efficiency of constructions using FCRHs.)

1. *Any fully black-box construction $G\colon \{0,1\}^\ell \mapsto \{0,1\}^{\ell+k}$ of PRGs from families of collision-resistant hash functions with a nonuniform security reduction needs $\Omega(k/\log n)$ oracle calls to FCRHs.*
2. *Any fully black-box construction of digital signatures from FCRHs with a nonuniform security reduction and for messages of length at least n bits and needs $\Omega(n)$ oracle calls to FCRHs.*

1.2.3 Handling nonuniform reductions in general constructions

In recent years, new types of black-box separations have emerged. These types of separation apply even to non-black-box constructions, but still only rule out black-box proofs of security: Following the works of Brassard [11] and Akavia et al [2], demonstrating limitations of "NP-hard Cryptography",[3] Pass [43] and Pass, Tseng and Venkitasubramaniam [45] demonstrate that under certain (new) complexity theoretic assumptions, various cryptographic tasks cannot

[3]See also the results of Feigenbaum and Fortnow [17] and the result of Bogdanov and Trevisan [8] that demonstrate limitations of NP-hard cryptography for *restricted* types of reductions.

be based on *one-way functions* using a black-box security reduction, even if the protocol uses the one-way function in a non-black-box way. Very recently, two independent works demonstrate similar types of lower bounds, but this time ruling our security reductions to a *general* set of intractability assumptions: Pass [44] demonstrates unconditional lower bounds on the possibility of using black-box reductions to prove the security of several primitives (e.g., Schnorr's identification scheme, commitment scheme secure under weak notions of selective opening, Chaum Blind signatures, etc) based on any "bounded-round" intractability assumption (where the challenger uses an a-priori bounded number of rounds, but is otherwise unbounded). Gentry and Wichs [22] provides a lower bound (assuming the existence of strong pseudorandom generators) on the possibility of using black-box security reductions to prove soundness of "succinct non-interactive arguments" (SNARGs) based on any "falsifiable" assumption (where the challenger is computationally bounded). Both of the above-mentioned work fall into the "meta-reduction" paradigm of Boneh and Venkatesan [9], which was previously used to prove separations for restricted types of reductions (see e.g., [12, 18, 31]).

As with the literature on fully black-box separations, the results of [22, 44] are only proved in the context of uniform security reductions. Here, we show that these results actually do extend to the nonuniform regime as well.

THEOREM 1.7 (INFORMALLY STATED). *The separations results of [22, 44] hold also when considering nonuniform black-box proofs of security.*

2. PRELIMINARIES

2.1 Notation

For any Boolean string x, by $|x|$ we denote the length of x. By $[k]$ we denote the set $\{1, \ldots, k\}$. We use bold letters (e.g., $\mathbf{x}$) when referring to random variables. By $x \xleftarrow{\$} \mathbf{x}$ we mean that x is sampled according to the distribution of the random variable $\mathbf{x}$. We use calligraphic letters (e.g., $\mathcal{S}$) to denote sets (e.g., events over random variables) and cryptographic primitives (e.g., one-way function). We use sans-serif letters (e.g., NP) to denote complexity classes. For a set $\mathcal{S}$, by $\mathbf{U}_{\mathcal{S}}$ we mean the random variable with uniform distribution over $\mathcal{S}$, and by $x \xleftarrow{\$} \mathcal{S}$ we mean $x \xleftarrow{\$} \mathbf{U}_S$. By the *support* of the random variable $\mathbf{y}$, represented by $\mathrm{Supp}(\mathbf{y})$, we mean $\{y \mid \Pr[\mathbf{y} = y] > 0\}$. For jointly distributed random variables $(\mathbf{x}, \mathbf{y})$, and for any $y \in \mathrm{Supp}(\mathbf{y})$, the conditional distribution $(\mathbf{x} \mid y)$ is the random variable $\mathbf{x}$ conditioned on $\mathbf{y} = y$.

When we say that an event parameterized by n occurs with negligible probability, denoted by $\mathrm{negl}(n)$, we mean that it occurs with probability $n^{-\omega(1)}$, and we say it happens with overwhelming probability if it happens with probability $1 - \mathrm{negl}(n)$. We call two random variables $\mathbf{x}, \mathbf{y}$ (or their corresponding distributions) over the support set $\mathcal{S}$ ε-close if their statistical distance, defined as $\Delta(\mathbf{x}, \mathbf{y}) = \frac{1}{2} \cdot \sum_{s \in \mathcal{S}} |\Pr[\mathbf{x} = s] - \Pr[\mathbf{y} = s]|$, is at most ε. We call an algorithm D an ε-distinguisher between the random variables $\mathbf{x}$ and $\mathbf{y}$ if $|\Pr[D(\mathbf{x}) = 1] - \Pr[D(\mathbf{y}) = 1]| \geq \varepsilon$. It is easy to see that if there is any D that ε-distinguishes between $\mathbf{x}$ and $\mathbf{y}$, then $\Delta(\mathbf{x}, \mathbf{y}) \geq \varepsilon$.

We use the term *efficient* for any *probabilistic* algorithm that runs in polynomial time over its input length. By the *size* of a circuit we refer to the number of bits that is required to describe it. So the number of circuits of size k will be at most 2^k.

2.2 Intractability Assumptions and Security Reductions

We recall the definition of an intractability assumption from [44] (see also [15, 22, 30, 42, 47]).

DEFINITION 2.1 (INTRACTABILITY ASSUMPTIONS). *A* intractability assumption $\mathcal{C}$ *is a two party game between a* challenger Chal *and an adversary* Adv *where both parties get* 1^n *as common input and* Chal *at the end outputs accept or reject. Any intractability assumption* $\mathcal{C}$ *has a* security threshold $\tau_{\mathcal{C}}$ *assigned to it which is a constant in the interval* $[0, 1)$. *We say that an interactive algorithm* Adv breaks $\mathcal{C}$, *if* Adv *(over the common input* 1^n*) can make* Chal *accept with probability* $\tau_{\mathcal{C}} + \varepsilon(n)$ *for a nonnegligible function* $\varepsilon(n)$. *When the adversary wins with probability* $\tau_{\mathcal{C}} + \varepsilon$ *we say that he has won the game with* advantage ε. *We say that* $\mathcal{C}$ *can be* uniformly (resp. nonuniformly) broken *if there is a PPT (resp.* poly(n)*-sized circuit)* Adv *that breaks* $\mathcal{C}$. *A* falsifiable *assumption is an intractability assumption where* Chal *is polynomial-time in the length of the messages it receives. A* bounded-round *assumption is an intractability assumption where the game between* $\mathcal{C}$ *and* Adv *has a fixed* poly(n) *number of rounds.*

Cryptographic Primitives. A cryptographic primitive is a syntactical requirement over a set of algorithms performing some cryptographic task. For example if $\mathcal{P}$ denotes the primitive one-way permutation, then it simply requires some algorithm P that computes some (hopefully one-way) permutation. We call (a computationally unbounded) oracle P an *implementation* of the primitive $\mathcal{P}$, if P satisfies the syntactical requirements of $\mathcal{P}$ (when all composed in one algorithm). We say an algorithm P *efficiently* implements $\mathcal{P}$ if it implements $\mathcal{P}$ and runs in polynomial time. We always assume that the algorithm P implementing the primitive $\mathcal{P}$ takes as its first input 1^n (and if it is efficient, it will run in time $\mathrm{poly}(n, |x|)$ where x is the main input). In most primitives the security parameter n can be related to the input length. For example in case of one-way function, n is usually taken to be equal to the input length, or for the case of PKE, it could be the length of the random seed used in the key-generation.

The security of (the efficient implementations of) almost all cryptographic primitives can be modeled as an intractability assumption. For instance, the security game of signature schemes is falsifiable, but it is not bounded-round, since the adversary is allowed to choose the number of received signatures before trying to forge one; soundness of a constant-round interactive argument, on the other hand, is a bounded-round assumption, but not falsifiable since checking whether an attacker succeeds is not efficient. The security threshold is usually either 0 (e.g., inverting a one-way function) or 1/2 (e.g., distinguishing a PRG from a uniform string).

DEFINITION 2.2 (CRYPTOGRAPHIC PRIMITIVES). *A cryptographic primitive* $\mathcal{P}$ *is a tuple* $(\mathcal{F}_{\mathcal{P}}, \mathcal{R}_{\mathcal{P}})$ *where every* $P \in \mathcal{F}_{\mathcal{P}}$ *is a function implementing* P, *and* $\mathcal{R}_{\mathcal{P}}$ *is a relation whose first component is always in* $\mathcal{F}_{\mathcal{P}}$. *When* $(P, A) \in \mathcal{R}_{\mathcal{P}}$,

we say that the "adversary" A breaks P (as an implementation of $\mathcal{P}$). By P_n we denote the implementation P restricted to the security parameter n. We restrict ourself to "natural" primitives where the security of any implementation P of $\mathcal{P}$ is defined through an intractability game with security threshold $\tau_{\mathcal{P}}$ that only depends on $\mathcal{P}$.

In the following, we formalize the definition of cryptographic constructions (fully black-box and general) and security reductions (uniform and non-uniform) separately. Formal variants of Definitions 1.2 and 1.1 can be obtained directly from these definitions.

DEFINITION 2.3 (CRYPTOGRAPHIC CONSTRUCTIONS). *A (general)* construction *of the primitive $\mathcal{Q}$ from another primitive $\mathcal{P}$ is a mapping $Q(\cdot)$ such that: if P is an efficient implementation for $\mathcal{P}$, $Q(P)$ is an efficient implementation for $\mathcal{Q}$. A* black-box construction *is a particular form of construction such that: $Q(P)$ is defined through an efficient (uniform) oracle machine Q accessing P only as an oracle, and Q^P is an implementation for $\mathcal{Q}$ for* any *(possibly inefficient) implementation P of $\mathcal{P}$.*

We might use the terms "non-black-box" and "general" (constructions) interchangeably.

DEFINITION 2.4 (SECURITY REDUCTIONS). *We say that a (general or black-box) construction Q of a primitive $\mathcal{Q}$ from another primitive $\mathcal{P}$ has a* nonuniform *(black-box) security reduction, if for every implementations P for $\mathcal{P}$ and $Q(P)$ for $\mathcal{Q}$, and every adversary A such that A_n breaks $Q_n(P)$ (over security parameter n) with advantage $\varepsilon \geq 1/\operatorname{poly}(n)$, there is a $\operatorname{poly}(n/\varepsilon)$-sized* oracle circuit *$S$ (whose code might depend on (P, A)) such that S^{P,A_n} breaks the security of P_m for some polynomially related security parameter $m = n^{\Theta(1)}$. A* uniform *security reduction is defined similarly by requiring the code of $S^{P,A_n}(1^n, 1^{1/\varepsilon})$ to be uniform and independent of the choices of the oracles (P, A).*

We might use the terms "security reduction" and "proof of security" interchangeably.

One can potentially define non-black-box proofs of security as well, but throughout this paper security reductions are always black-box.

REMARK 2.5. *The reason that we allow the security reduction S to call A only over a single security parameter n (which is polynomially related to $m = n^{\Theta(1)}$) is that here we work with the "almost-everywhere" notion of security (where any infinite–but arbitrarily sparse–sequence of security parameters $\{n_1 < n_2 \ldots\}$ over which A "wins" shall be transformed into a sequence of security parameters $\{m_1 < m_2 \ldots\}$ over which P is broken).*

The formal definition of a general construction of a primitive $\mathcal{Q}$ from another primitive together $\mathcal{P}$ with a non-uniform security reduction (i.e. the formalized form of Definition 1.1) can be obtained directly from Definitions 2.3 and 2.4. For the case of black-box constructions, we employ the following terminology whose uniform variant is due to [46].

DEFINITION 2.6 (FULLY BLACK-BOX CONSTRUCTIONS). *A* fully *black-box construction of a primitive $\mathcal{Q}$ from another primitive $\mathcal{P}$ consists of a black-box construction Q of $\mathcal{Q}$ from $\mathcal{P}$ together with a (uniform or nonuniform) black-box security reduction for this construction.*

For every fixed efficient implementation P of a primitive $\mathcal{P}$, and every (black-box or non-black-box) construction Q of primitive $\mathcal{Q}$ from $\mathcal{P}$, a black-box reduction of the security of $Q(P)$ to that of P can usually be modeled as an intractability assumption. Thus, intractability assumptions allow us to model the (uniform or nonuniform) proofs of security for cryptographic constructions directly (regardless of whether the construction is black-box or not).

Note that any intractably assumption $\mathcal{C} = (\mathsf{Chal}, \mathsf{Adv})$ can be considered as a cryptographic primitive $\widetilde{\mathcal{C}}$ such that: the set of implementations of $\widetilde{\mathcal{C}}_n$ (over security parameter n) only contains the empty function, and the security of $\widetilde{\mathcal{C}}_n$ is defined based on the interactive game $(\mathsf{Chal}, \mathsf{Adv})$ over security parameter n. Based on this perspective, the following definition formalizes what it means to base a cryptographic primitive on such a primitive.

DEFINITION 2.7 (REDUCTIONS TO ASSUMPTIONS). *We say that a cryptographic primitive $\mathcal{Q}$ can be based on a intractability assumption $\mathcal{C} = (\mathsf{Chal}, \mathsf{Adv})$ through a uniform (resp. non-uniform) black-box reduction iff there exists a construction of $\mathcal{Q}$ from $\mathcal{P} = \widetilde{\mathcal{C}}$ with a uniform (resp. non-uniform) security reduction.*

Refuting the possibility of basing a cryptographic primitive $\mathcal{Q}$ on any (falsifiable/bounded round/general) intractability assumption through a uniform (resp. non-uniform) black-box reduction immediately rules out the possibility of basing $\mathcal{Q}$ on a large class of natural cryptographic primitives (whose security for efficient implementations are of the form of intractability assumptions) for uniform (resp. nonuniform) black-box security reductions.

2.3 Special Soundness and Witness Hiding

We assume the reader is familiar with the notions Witness Indistinguishability and Commitment schemes. We refer the reader to [25] for formal definitions.

Special Soundness. Recall that a three-round public-coin interactive proof is said to be *special-sound*, if a valid witness to the statement x can be efficiently computed from any two accepting proof-transcripts of x which have the same first message but different second messages. [44] considers a relaxation of this notion—referred to as *computational special-soundness*—where **(a)** the number of communication rounds is any constant (instead of just three), **(b)** the extractor may need a polynomial number of accepting transcripts (instead of just two), and **(c)** extraction need only succeed if the transcripts are generated by communicating with a computationally-bounded prover (see Section 4 for a formal definition). All traditional constant-round public-coin proofs of knowledge protocols (such as [7, 26, 28, 48], as well as instantiations of [7, 26] using statistically-hiding commitments) satisfy this property, and continue to do so also under parallel repetition. We say that a computationally special-sound protocol has a *large challenge space* if the length of the verifier challenge is $\omega(\log n)$ on common inputs of length n.

DEFINITION 2.8 (COMPUTATIONAL SPECIAL-SOUNDNESS). *Let (P, V) be a k-round (where k is a constant) public-coin interactive argument for the language $L \in \mathsf{NP}$ with witness relation R_L. (P, V) is said to be* computationally special-sound *if there exists a polynomial $m(\cdot)$, and a polynomial-*

time extractor machine X, such that for every polynomial-time deterministic machine P^*, and every polynomial $p(\cdot)$, there exists a negligible function $\mu(\cdot)$ such that the following holds for every $x \in L$ and every auxiliary input z for P^*. Let $\vec{T} = (T_1, T_2, \ldots T_{p(|x|)})$ denote transcripts in $p(|x|)$ random executions between $P^*(x, z)$ and $V(x)$ where V uses the same randomness for the first $k-2$ messages (thus, the first $k-1$ messages are the same in all transcripts). Then, the probability (over the randomness used to generate $\vec{T}$) that:

1. *$\vec{T}$ contains a set of $m(|x|)$ accepting transcripts with different round $k-1$ messages; and*
2. *$X(\vec{T})$ does not output a witness $w \in R_L(x)$*

is smaller than $\mu(|x|)$. We say that a computationally special-sound protocol has a large challenge space *if the length of the verifier challenge is $\omega(\log n)$ on common inputs of length n.*

Witness Hiding. A desirable property of interactive proofs is that they "hide" the witness used by the prover. We will consider a very weak notion of sequential witness hiding: roughly speaking, a protocol is said to be weakly sequential witness hiding if no polynomial time attacker can always recover the witness for any statement that it hear polynomially many sequential proofs of.

Definition 2.9 (Breaking Weak Witness Hiding). *Let (P, V) be an argument for the language L with witness relation R_L. We say that (a potentially unbounded) A* breaks weak $\ell(\cdot)$-sequential witness hiding of *(P, V) with respect to R_L if for every $n \in \mathbb{N}$, $x \in L \cap \{0,1\}^n$, and $w \in R_L(x)$, A wins in the following experiment with probability 1: Let $A(x)$ sequentially communicate with $P(x, w)$ $\ell(n)$ times; A is said to win if it outputs a witness w' such that $w' \in R_L(x)$. (P, V) is called* weakly $\ell(\cdot)$-sequentially witness hiding *w.r.t R_L if no polynomial time algorithm A breaks weak $\ell(\cdot)$-sequential witness hiding of (P, V) w.r.t R_L.*

Let us now state the uniform separation result of [44]:

Theorem 2.10 (Main Result of [44]). *Let (P, V) be a computationally-special-sound argument with large challenge space for the language L with a unique witness relation R_L, and let $\mathcal{C}$ be an $r(\cdot)$-round assumption where $r(\cdot)$ is a polynomial. If for every polynomial $\ell(\cdot)$ there exists a black-box security reduction S for basing weak $\ell(\cdot)$-sequential witness hiding of (P, V) w.r.t R_L on intractability assumption $\mathcal{C}$ w.r.t threshold $\tau_{\mathcal{C}}$, then there exists an efficient algorithm B such that $B(1^n)$ breaks $\mathcal{C}$ w.r.t. $\tau_{\mathcal{C}}$ with advantage $\varepsilon(n)$ for a nonnegligible $\varepsilon(n)$.*

Pass [44] shows how Theorem 2.10 can be used to separate several well-known cryptographic protocols/primitive (e.g., Schnorr's identification scheme, commitment scheme secure under weak notions of selective opening) from any intractability assumption w.r.t. uniform black-box reductions. It directly follows that if Theorem 2.10 is extended to the nonuniform setting, then these corollaries also extend to the nonuniform setting. We refer the reader to [44] for more details on these corollaries.

3. SEPARATING UNIFORM AND NONUNIFORM SECURITY REDUCTIONS

In this section, we demonstrate that a commitment scheme from [39] (LPV), which is proven secure based on the existence of one-way permutations through a *nonuniform* black-box proof of security, cannot be based on any bounded-round falsifiable assumptions through a uniform black-box proof of security; in particular, it cannot be based on the existence of one-way permutations using a uniform black-box proof of security. As such, we separate the power of nonuniform and uniform black-box proofs of security (assuming the existence of one-way permutations). Let us start by reviewing the LPV protocol.

The LPV protocol is a non-malleable commitment scheme [16], where the committer and the receiver receive as common input a security parameter n and an identity id of some length $\ell(n)$. To commit to a value v, committer and the receiver proceeds in three stages. In Stage 1, the receiver sends a random string $s \in \{0,1\}^n$. Then in Stage 2, the committer commits to v by sending string c using a non-interactive perfectly binding commitment scheme com. Finally, in Stage 3, the committer proves that either c is a valid commitment to v or it knows the pre-image of s through a one-way permutation f (that is, that it knows a string r such that $f(r) = s$). More specifically, the witness relation used is defined as follows: witness relation $R_L = \{((s, c), y) \mid f(y) = s \text{ or } c = \mathsf{com}(y; r) \text{ for some } r\}$. This is proved using $4\ell(n)$ invocations of a 3-round public-coin WI special-sound (WISSP) argument. Messages in these WISSP arguments are scheduled according to a special scheduling based on id that guarantees that there exist at least $2\ell(n)$ sequential WISSP arguments in the protocol. Our result will only rely on this fact, and thus for simplicity of exposition, below we outline our result w.r.t. a simplified protocol consisting of 2ℓ sequential WISSP arguments in Stage 3.[4] We refer to this protocol as $(C, R)_\ell$.

The following lemma is shown in [39].

Lemma 3.1 ([39]). *Let ℓ be any polynomial. Assuming the existence of one-way permutations, $(C, R)_\ell$ is perfectly binding and computationally hiding with a nonuniform black-box proof of security.*

To highlight the power of nonuniformity in security reductions, let us briefly sketch a proof of this lemma. Fix a polynomial ℓ. Assume that there is an adversary A (w.l.o.g. deterministic) that breaks the hiding property of the LPV protocol $(C, R)_\ell$; that is, for two values v_0, v_1, the adversary A after receiving a commitment to one of the values chosen at random, can guess with inverse polynomial probability which value it has received a commitment to. We now demonstrate the existence of a nonuniform reduction R that with black-box access to A breaks the computational hiding property of the perfectly binding commitment com. Since A is deterministic, the first message s that it sends is fixed. Thus the reduction R can receive as a nonuniform advice the pre-image r of s through the OWP f (i.e, a string r such that $f(r) = s$). Then, after receiving a com commitment c to one of the two values v_0, v_1 chosen at random, it can use the adversary A to guess which value it receives a commitment to by forwarding c to A and simulating all the WISSP

[4] It is easy to see that the same argument applies also to the original LPV protocol.

arguments using r as a fake witness; finally, it outputs A's guess. It follows from the witness indistinguishability property of the WISSP argument that R^A has inverse polynomial advantage in guessing whether it received a commitment to v_0 or v_1.

Let us now turn to showing the impossibility of basing the hiding property of the LPV protocol on any bounded-round falsifiable assumptions through a *uniform* black-box proof of security.

THEOREM 3.2. *Let $\mathcal{C}$ be an $r(\cdot)$-round falsifiable assumption where $r(\cdot)$ is a polynomial. If for every polynomial $\ell(\cdot)$ there exists an efficient black-box security reduction S for basing the hiding property of $\langle C, R\rangle_\ell$ on assumption $\mathcal{C}$ w.r.t threshold $\tau_{\mathcal{C}}$, then there exists an efficient algorithm B and a polynomial $p(\cdot)$ such that for infinitely many $n \in N$, $B(1^n)$ breaks $\mathcal{C}$ with advantage $1/p(n)$.*

Here, we provide an outline of the proof. See the full version for a full proof.

3.1 Outline of the Proof of Theorem 3.2

At first sight, it would seem that Theorem 3.2 directly follows from Theorem 2.10: At a very high-level, the LPV protocol simply consists of many sequentially repeated special-sound arguments (that are also constant-round and public coin) for the statement (s, c) defined by the messages in Stage 1 and 2. Additionally, demonstrating hiding, at the very least implies that these protocols are weakly $l(n)$-sequentially witness hiding (or else, the committed value can be completely recovered!).

However, this approach does not go through as Theorem 2.10 only provides a separation in the case of *unique witness languages*, whereas R_L admits two witnesses for every statement. To get around this problem, we consider a unique witness relation $R_{L'}$ for which every statement (s, c) only has one unique witness that is the unique committed value in c.

But if we change the language, the proofs in Stage 3 of the protocol are no longer special-sound for the witness relation $R_{L'}$, since for a specific instance (s, c), it might be easy to invert s and therefore the value extracted from a WISSP argument (in Stage 3 of the LPV protocol) might be the pre-image r instead of the committed value, violating the (computational) special soundness property for $R_{L'}$. We resolve the problem by observing that, in fact, the proof of the Theorem 2.10, when restricted to falsifiable bounded-round assumption (as opposed to general bounded-round assumption) in [44] itself is black-box and uniform[5]. Roughly speaking, theorem 2.10 states that if (P, V) is computationally special-sound for a unique witness language L, then it is impossible to base the sequential witness hiding property on any bound-round falsifiable assumptions through uniform black-box security reduction. This is proven by showing that for any (public-coin, constant-round) interactive protocol (P, V), if there is a uniform black-box security reduction S for basing its sequential witness hiding property for $R_{L'}$ on any bounded-round falsifiable assumption, then there is a *uniform meta reduction M* that with black-box access to S can violate the computational special-soundness of (P, V) w.r.t. $R_{L'}$. In our context, since (P, V) is a special-sound proof for R_L (as opposed to $R_{L'}$), this meta reduction may either violate computational special-soundness of (P, V) w.r.t. R_L, or may output a pre-image r to s. However, as long as s is chosen at random by M, we can thus use such an M to invert a random string s through f (assuming that computational special-soundness of (P, V) for R_L holds). See Section 3 for the detailed proof.

Let us briefly comment on the why the above proof sketch does not extend to nonuniform proofs of security (whereas as we shall see in Section 4.1, the main theorem of [44] does). The problem is that if the reduction S is nonuniform, it may get as nonuniform advice the string s (or some function of it); this means that in the above proof, M can no longer choose the string s uniformly at random (since S would notice this). Indeed, in the actual nonuniform proof of security of LPV, the reduction does get the pre-image of s as nonuniform advice.

[5]In fact, the proof of Theorem 2.10 as stated in [44] in actually *nonuniform*, but for the special case that the intractability assumption is falsifiable, the proof is uniform.

4. REFUTING NONUNIFORM REDUCTIONS TO INTRACTABILITY ASSUMPTIONS

In this section we prove Theorem 1.7.

4.1 Extending Theorem 2.10 to Nonuniform Security Reductions

Here we show how to extend Theorem 2.10 to handle nonuniform proofs of security. Pass showed how Theorem 2.10 can be used to prove that certain well known cryptographic protocols/primitive (e.g., Schnorr's identification scheme, commitment scheme secure under weak notions of selective opening) can not be based on any bounded-round intractability assumption through a black-box proof of security. It follows that all corollaries of [44] also extend to the nonuniform regime.

THEOREM 4.1. *Let (P, V) be a computationally-special-sound argument with large challenge space for the language L with a unique witness relation R_L, and let $\mathcal{C}$ be an $r(\cdot)$-round assumption where $r(\cdot)$ is a polynomial. If for every polynomial $\ell(\cdot)$ there exists a nonuniform black-box security reduction S for basing weak $\ell(\cdot)$-sequential witness hiding of (P, V) w.r.t R_L on intractability assumption $\mathcal{C}$ w.r.t threshold $\tau_{\mathcal{C}}$, then there exists an nonuniform polynomial-time algorithm B and a polynomial $p(\cdot)$ such that for infinitely many $n \in N$, $B(1^n)$ breaks $\mathcal{C}$ with advantage $1/p(n)$.*

Note that in Theorem 4.1 we consider also nonuniform security reductions S, but the conclusion is slightly weaker than the conclusion of Theorem 2.10: the machine B breaking the assumption $\mathcal{C}$ is no longer uniform, but it now is a nonuniform polynomial-time algorithms.

In the following we outline the proof of Theorem 4.1. See the full version of the paper for the full proof.

4.1.1 Outline of the Proof of Theorem 4.1

Let us start by briefly outlining the high-level approach used in the result of [44] and explain why handling nonuniform proofs of security becomes an issue. Assume there exists a security reduction S (and for now, assume that S is uniform) such that S^A breaks the assumption $\mathcal{C}$ whenever A breaks weak sequential witness hiding of a (computationally) special-sound argument (P, V) for a language with unique witnesses. We want to use S to directly break $\mathcal{C}$ *without the*

help of A. So, (following the paradigm of [9]), the goal will be to efficiently emulate A for S (i.e., we will construct a "meta-reduction" M which uses the underlying reduction S to break $\mathcal{C}$). We consider a particular computationally unbounded oracle A that after hearing an appropriate number of proofs using (P, V) (acting as a verifier) simply outputs a witness to the statement proved. The idea is to "extract" out the witness that A is supposed to provide S by "rewinding" S—since (P, V) is computationally special-sound, S, intuitively, must know a witness for all statements x that it proves to A. (There are several obstacles in formalizing this approach. The main one is that the reduction S is not a "stand-alone" prover—it might *rewind and reset* the oracle A, so it is no longer clear that it needs to "know" a witness for x in order to convince A of x. It is here that the proof of [44] relies on the fact that there are multiple proofs being provided by S; this gives the meta-reduction more opportunities to rewind S, which enables extraction even if S "nests" its queries to A in an arbitrary way. We refer the reader to [44] for further details.

Let us point out a crucial component in the above proof: To succeed in its emulation of S, it is imperative that whenever A is acting as a verifier, it chooses fresh random coins to generate its messages, even in case A is rewound (technically, this is achieved by letting A generate its messages by applying a random function to its queries). This is needed to ensure that M can "rewind" S (in order to extract out a witness), sending it new verifier messages, while ensuring that S provides an answer back with the same probability as if S communicated with A. In other words, to ensure that M succeeds in extracting witnesses from S, we require A's verification message to have essentially "full entropy", or else S may be able to notice that it is being rewound, and may abort its computation. In the context of nonuniform reductions, we can no longer guarantee that A's answers have high entropy: S gets a nonuniform advice string as a function of A, and thus the conditional entropy of the answers of A drops. Our approach for getting around this problem is that: although the conditional entropy of A's answers drops (once S gets its nonuniform advice), for all but a polynomial number of "bad" queries to A, the answers to the remaining "good" queries will still have high enough entropy. In fact, by an argument due to Unruh [51], it can be shown that the conditional distribution of answers to "good" queries is statistically close to the original distribution of A.

LEMMA 4.2 (INFORMAL VARIANT OF [51]). *Suppose A is a randomized oracle, and suppose S is an oracle algorithm that gets as input a non-uniform (polynomial-size) advice z as function of (the description of) A, and then asks polynomially many queries to A. Then there is a "pre-sampling" algorithm* Samp *that given z samples s query-answers of A (according to their true distribution based on A), and the view of $S^A(z(A))$ is $1/\operatorname{poly}(n)$-close in another experiment in which S is given z, then $\mathsf{Samp}(z)$ samples a partial domain of A, then A gets* resampled *on every other point other than the outputs of $\mathsf{Samp}(z)$, and finally $S(z)$ gets executed with oracle access to the (newly sampled) A.*

If the reduction S had only queried these "good" queries (that are *not* presampled by Samp), we would already be done. But, S may of course ask also "bad" queries (i.e. the ones presampled by $\mathsf{Samp}(\cdot)$). To deal with this, we present a *nonuniform meta reduction M—M* receives as a nonuniform advice the set of "bad" queries (which may depend on the nonuniform advice of S), and for each of these queries, the answer that A actually would provide. M can then perfectly emulate "bad" queries (using its nonuniform advice), and as before emulate good queries (in a statistically close manner) by using rewindings. As a conclusion we get that the existence of a security reduction S can be used to break the intractability assumption $\mathcal{C}$ in *nonuniform polynomial time*.

4.2 Succinct Non-Interactive Arguments

In this section we show how to extend a result of Gentry and Wichs [22] on the impossibility of basing succinct non-interactive arguments (SNARGs) for NP on any falsifiable assumption to the nonuniform regime. We start by recalling the formal definition of SNARGs.

A succinct non-interactive argument system Π consists of three efficient algorithms (G, P, V): The generation algorithm G on input security parameter 1^n outputs a common reference string crs and a private verification state priv. The prover algorithm P on input crs, a statement $x \in \{0,1\}^n$ and a witness w outputs a proof π. The verifier algorithm V on input priv, x, and π outputs a bit $b \in \{0,1\}$ represents whether V accepts or rejects the proof π for x.

DEFINITION 4.3 (SUCCINCT NONINTERACTIVE ARGUMENTS). *$\Pi = (G, P, V)$ is a succinct non-interactive argument (SNARG) for an NP language L with relation R_L if following holds. (Below $|x| = n$ and by negligible we mean $\operatorname{negl}(n)$.)*

- **Completeness:** *For every $(x, w) \in R_L$ the probability that the verifier V rejects in the following experiment is negligible: (i) $(\mathsf{crs}, \mathsf{priv}) \leftarrow G(1^n)$, (ii) $\pi \leftarrow P(\mathsf{crs}, x, w)$, and (iii) $b \leftarrow V(\mathsf{priv}, x, \pi)$.*
- **(Adaptive) Soundness:** *For every efficient cheating prover P^*, the probability that the verifier V accepts in the following experiment is negligible: (i) $(\mathsf{crs}, \mathsf{priv}) \leftarrow G(1^n)$, (ii) P^* on input crs outputs both a statement x and a proof π, and (iii) $b \leftarrow V(\mathsf{priv}, x, \pi)$.*
- **Succinctness:** *The length of a proof $\pi \leftarrow P(\mathsf{crs}, x, w)$ generated by the prover is $n^{o(1)}$.*

We mention that for the simplicity of exposition, in the above definition, we restrict the length of statements to be the same as the security parameter n, and define the succinctness property by $|\pi| \leq n^{o(1)}$. Our extension of the results of [22] holds for the general definition in [22] as well. We now recall the formal result of [22].

THEOREM 4.4 ([22]). *Assuming the existence of sub-exponentially hard one-way functions, for any SNARG $\Pi = (G, P, V)$ that satisfies the completeness and succinctness properties, the adaptive soundness of Π can not be based on any falsifiable assumption $\mathcal{C}$ through a* uniform *black-box reduction, unless $\mathcal{C}$ can already be broken (nonuniformly).*

Note that Theorem 4.4 only rules out *uniform* black-box security proof of adaptive soundness. Our goal is to extend the theorem to also rules out *nonuniform* black-box security proof. Towards this goal, we note that core of the proof of Theorem 4.4 in [22] is the construction of two adversarial provers, and we identify the key properties of the two adversarial provers that are needed by us.[6]

[6]Lemma 4.1 in [22] is stated only w.r.t. *efficient* distinguishers (indicating them to be uniform), but since it assumes the

LEMMA 4.5 (LEMMA 4.1 IN [22]). *Assuming the existence of sub-exponentially hard pseudorandom generators, for any SNARG system $\Pi = (G, P, V)$ that satisfies the completeness and succinctness properties, there exist two adversarial provers A and B that satisfy the following properties.*

- *A breaks the adaptive soundness of Π but is inefficient, whereas B is efficient.*
- *A and B are computationally indistinguishable in the following sense: for any polynomial-sized circuit family $\{D_n\}$ it holds that*

$$\left|\Pr[D_n^A = 1] - \Pr[D_n^B = 1]\right| = \text{negl}(n).$$

- *Both A and B are randomized and* stateless *and use fresh independent randomness to generate answers to distinct queries (which are the* crs*'s).*

The third property allows us to view A (resp., B) as a deterministic algorithm that given any query $q = \mathsf{crs}$, access fresh randomness $RO(q)$ and returns $A(q, RO(q))$ (resp., $B(q, RO(q))$). We can think of $RO(\cdot)$ as a random oracle with long enough outputs length[7]. More explicitly, by $A[RO]$ we emphasize on the fact that A is using the sampled random oracle $RO \xleftarrow{\$} \mathbf{RO}$.

We are ready to extend Theorem 4.4 to rules out *nonuniform* black-box security proof.

THEOREM 4.6. *Assuming the existence of sub-exponentially hard pseudorandom generators, for any SNARG system $\Pi = (G, P, V)$ that satisfies the completeness and succinctness properties, the adaptive soundness of Π can not be based on any falsifiable assumption $\mathcal{C}$ through a* nonuniform *black-box reduction, unless $\mathcal{C}$ can already be broken (nonuniformly).*

Recall that $\mathbf{RO}$ denotes the distribution of random oracles, and for any partial length preserving function F suppose $\mathbf{RO}[F]$ denotes the distribution of random oracles with pre-sampled part F.

To prove Theorem 4.6 we again use Lemma 4.2 an a careful hybrid argument similar to the proof of Theorem 4.1. See the full version of the paper for the full proof.

5. REFUTING NONUNIFORM REDUCTIONS FOR FULLY-BLACK-BOX CONSTRUCTIONS

In this section we prove our results of Theorems 1.4, 1.5, and 1.6. We start by proving Theorem 1.5 and then will prove Theorems 1.4 and 1.6 based on that.

We also show how the techniques used in the proofs of these theorems can be used to obtain different proofs of nonuniform hardness for other primitives such as one-way permutations as well as extending known black-box separations in the uniform regime to the nonuniform regime.

subset-membership problem to be non-uniformly hard, the very same proof given for Lemma 4.1 handles nonuniform distinguishers as well, so we state and use this lemma in this form.

[7]By padding the queries q appropriately, we can assume w.l.o.g that $RO(\cdot)$ is length preserving.

Nonuniform Collision-Resistance of Random Hashing and Beyond.

In this section we prove the Theorem 1.5 about the nonuniform hardness of random functions as families of collision-resistant hash functions. In fact we prove the following stronger theorem:

THEOREM 5.1 (NONUNIFORM COLLISION RESISTANCE). *Let the function $\mathbf{h}\colon [K] \times [M] \mapsto [N]$ be chosen uniformly at random. If for $d, t \in \mathbb{N}$ it holds that $dN^3 < 2t^5K$, then for every t-query adversary A with d bits of advice z that may depend on h: $A^h(i, z)$ can find a collision pair for the hash function $h_i(\cdot) = h(i, \cdot)$ with probability at most $3t^2/N$ over the choice of $h \xleftarrow{\$} \mathbf{h}, i \xleftarrow{\$} [K]$ and the randomness of A.*

Instead of proving Theorem 5.1, in fact we prove a more general lemma about the non-uniform hardness of *families* of ideal primitives, if the ideal primitive is already (uniformly) secure against bounded-query adversaries.

LEMMA 5.2 (HARDNESS OF FAMILIES OF PRIMITIVES). *Suppose $\mathbf{O}$ is a randomized oracle such that any computationally unbounded t-query adversary A "wins" in the "game" $A^{\mathbf{O}}$ only with probability at most ε where the notion of winning only depends on the transcript of $A^{\mathbf{O}}$. Suppose FO is another randomized oracle which consists of $K = 2^k$ independent samples $O_1, \ldots, O_K$ from $\mathbf{O}$ accessed through a k-bit prefix to the queries. We define $A^{\mathsf{FO}}(i)$ wins iff the transcript of the interaction of A with O_i indicates a win for A. Suppose A also receives d bits of advice z about the oracle FO. Then for any such advice function, and any $s \le K$:*

$$\Pr_{i \xleftarrow{\$} [K], FO \xleftarrow{\$} \mathsf{FO}} [A^{FO}(i, z(FO)) \text{ wins}] \le \varepsilon + s/K + \sqrt{td/2s}.$$

The high level idea behind the proof of Lemma 5.2 is to use Lemma 4.2 as follows. Each sampled oracle O_i could be considered as an answer returned by a random oracle over domain $[K]$. By Lemma 5.2 the oracle answers to most of the indexes in $[K]$ remain statistically close to uniform, even given a "small" advice about all of $\mathsf{FO} = [O_1, \ldots, O_K]$. So for a "typical" $i \xleftarrow{\$} [K]$, the job of the adversary to win against O_i is just as hard as the case there were no advice.

See the full version of the paper for a formal proof of Lemma 5.2 and how to derive Lemma 5.1 from it.

A Lemma for Ruling Out Nonuniform Security Reductions.

Variants of the following lemma in the *uniform* regime has been used in some previous work [5, 13, 41].

LEMMA 5.3. *Let $\mathcal{P}$ and $\mathcal{Q}$ be two cryptographic primitives and $\mathcal{P}$ has security threshold zero and let $\mathbf{O}$ be a randomized oracle. Suppose the following two holds:*

1. *For* any *black-box construction $Q^{\mathbf{O}}$ of $\mathcal{Q}$ from $\mathbf{O}$ there is an adversary Adv who asks $q(n)$ oracle queries to $\mathbf{O}$ and breaks the security of $Q_n^{\mathbf{O}}$ (over security parameter n) with non-negligible probability $\varepsilon > 1/\text{poly}(n)$.*
2. *There exists a black-box construction $P^{\mathbf{O}}$ of $\mathcal{P}$ from $\mathbf{O}$ such that for any $m = n^{\Theta(1)}$ any $\text{poly}(n)(q(n)+1)$-query adversary T who receives $\text{poly}(n)$ bits of advice about $\mathbf{O}$ can break $P_m^{\mathbf{O}}$ only with negligible probability $\text{negl}(n)$.*

Then there is no black-box construction of $\mathcal{Q}$ from $\mathcal{P}$ even with a nonuniform proof of security. Moreover, to rule out constructions in which Q_n only calls P_n, we only need to consider $m = n$ in second condition above.

See the full version of the paper for a proof of Lemma 5.3.

Proving Theorems 1.4 and 1.6.

Here we prove Theorems 1.4 and 1.6.

Roughly speaking Theorem 1.4 can be concluded from Theorem 5.1 similar to the way we used the result of [21] to extend the result of [34] to the nonuniform regime. Formal proof is as follows. In the rest of this section we prove Theorem 1.4. In the following we prove Theorem 1.4.

PROOF OF THEOREM 1.4. We employ Lemma 5.3 as follows. We let the randomized oracle $\mathbf{O}$ be the random oracle. Theorem 5.1 shows that there is a construction of FCRHs relative to $\mathbf{O}$ such that any adversary T who gets $2^{o(n)}$ bits of advice about $\mathbf{O}$ and asks $2^{o(n)}$ queries to it can break it only with advantage $2^{-\Omega(n)}$. This would imply the second requirement of Lemma 5.3 with $q(n) = \text{poly}(n) \leq 2^{n/10}$.

We also use the following result which is the main step toward separating one-way functions from key-agreement proved in [34] and provides us with the adversary Adv as needed for the first condition of Lemma 5.3.

LEMMA 5.4. *Suppose Π is a key-agreement protocol in which Alice and Bob each ask n oracle queries to the random oracle $O \xleftarrow{\$} \mathbf{RO}$, and they agree on a key with probability $1 - \text{negl}(n)$. There is a computationally unbounded adversary Adv who only accesses the public messages sent between Alice and Bob, asks at most $\text{poly}(n)$ oracle queries to O, and finds the key with probability $1 - 1/n^2$.*

Since both conditions of Lemma 5.3 are satisfied, Theorem 1.4 follows immediately.

□

Now we prove Theorem 1.6.

PROOF OF THEOREM 1.6. We will first prove Theorem 1.6 for PRGs and then will prove it for the case o digital signatures.

Part 1: Pseudorandom Generators. Suppose on the contrary that $G\colon \{0,1\}^{\ell} \mapsto \{0,1\}^{\ell+k}$ is a black-box construction of PRGs that stretches its input by k and $p = \lambda(n) \cdot (k/\log n)$ oracle queries while $1/\text{poly}(n) \leq \lambda(n) \leq o(1)$. In this case we again use $\mathbf{O}$ to denote a random oracle, but we use a slightly different construction of FCRHs h (to be used as $\mathcal{H}$) relative to $\mathbf{O}$. Given any query (z, x) where $z \in \{0,1\}^{4n}$ and $x \in \{0,1\}^{2n}$ to choose the n-bit output we choose the last $r = \min \log^2 n, \frac{\log n}{10\lambda(n)} = \omega(\log n)$ bits of y uniformly at random, and we copy the $n - \frac{\log n}{\lambda(n)}$ first bits of x as the $n - \frac{\log n}{\lambda(n)}$ first bits of the output y.

Note that any adversary breaking collision resistance of h has to find a collision for the last r bits of h as well, so as a mental experiment we pretend that the output length of h is $r \leq \log^2 n$ Since $\text{poly}(n) \cdot 2^{\log^2 n} \leq \text{poly}(n) \cdot 2^{4n}$, Theorem 5.1 proves that the construction $\mathbf{h} = h^{\mathbf{O}}$ is non-uniformly secure in the sense that any $\text{poly}(n)$ query attacker with $\text{poly}(n)$ bits of advice about $\mathbf{O}$ can find a collision in $h^{\mathbf{O}}$ only with negligible probability. This shows that the second requirement of Lemma 5.3 holds for $q(n) = \text{poly}(n)$.

LEMMA 5.5 ([20]). *There is an adversary Adv who asks* no *queries to $\mathbf{O}$ but is able to distinguish the output of $G(U_\ell)$ from $U_{\ell+k}$ with advantage $1/4$.*

We prove a proof for sake of completeness.

PROOF. Then the total number of random bits used in the computation of $G^h(\mathbf{U}_\ell)$ is $\ell + p \cdot \frac{\log n}{10\lambda(n)} < \ell + k$ (where $\mathbf{U}_\ell$ is a random string of length ℓ). Therefore the support set of the function $G^h(\mathbf{U}_\ell)$ has size at most $2^{\ell+k-1}$. The latter implies that a computationally unbounded adversary is able to distinguish between $G^h(\mathbf{U}_\ell)$ and $\mathbf{U}_{\ell+k}$ with advantage at least $1/4$ (by outputting 1 whenever the given y is in the support of $G^h(\mathbf{U}_\ell)$ and outputting 0 elsewhere). □

The attacker Adv of Lemma 5.5 satisfies the first requirement of Lemma 5.3. Therefore, the first part of Theorem 1.6 follows directly from Lemma 5.3.

Part 2: Signature Schemes. Now we prove the second part of Theorem 1.6. Suppose for sake of contradiction that D is a black-box construction of digital signatures for message space $\{0,1\}^n$ using $o(n)$ queries to $\mathcal{H}_n$. We use the same construction h for $\mathcal{H}_n$ as that of the proof of Theorem 1.4; namely a random FCRHs with key length $4n$, input length $2n$, and output length $2n$. By Theorem 5.1, $h^{\mathbf{O}}$ is $2^{n/10}$ secure, even if the adversary gets $2^{n/10}$ bits of advice about $\mathbf{O}$. Thus we obtain the first requirement of Lemma 5.3 for any $q(n) = 2^{o(n)}$.

LEMMA 5.6 ([5]). *For any construction of digital signatures D for message space $\{0,1\}^n$ in the random oracle model in which the key-generation, signing, and verification algorithms ask p oracle queries can be broken with advantage $1/\text{poly}(n)$ by an adversary Adv who asks $\text{poly}(n)2^{O(p)}$ oracle queries.*

Lemma 5.6 shows that if D asks only $o(n)$ many queries to $\mathcal{H}_n$, then Adv can break it with $q(n) = 2^{o(n)}$ queries. This implies that the first requirement of Lemma 5.3 holds as well and so Lemma 5.3 implies the second part of Theorem 1.6 directly.

6. REFERENCES

[1] Takahiro Matsuda 0002 and Kanta Matsuura. On black-box separations among injective one-way functions. In Yuval Ishai, editor, *TCC*, volume 6597 of *Lecture Notes in Computer Science*, pages 597–614. Springer, 2011.

[2] Adi Akavia, Oded Goldreich, Shafi Goldwasser, and Dana Moshkovitz. On basing one-way functions on np-hardness. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*, pages 701–710, 2006.

[3] Sergei Artemenko and Ronen Shaltiel. Lower bounds on the query complexity of nonuniform and adaptive reductions showing hardness amplification. *Electronic Colloquium on Computational Complexity (ECCC)*, 18:16, 2011.

[4] Michael Backes and Dominique Unruh. Limits of constructive security proofs. In Josef Pieprzyk, editor, *ASIACRYPT*, volume 5350 of *Lecture Notes in Computer Science*, pages 290–307. Springer, 2008.

[5] Boaz Barak and Mohammad Mahmoody-Ghidary. Lower bounds on signatures from symmetric primitives. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, 2007.

[6] Boaz Barak and Mohammad Mahmoody-Mahmoody. Merkle puzzles are optimal - an $O(n^2)$-query attack on any key exchange from a random oracle. In Shai Halevi, editor, *Advances in Cryptology - CRYPTO 2009, 29th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2009. Proceedings*, volume 5677 of *Lecture Notes in Computer Science*, pages 374–390. Springer, 2009.

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

[8] Andrej Bogdanov and Luca Trevisan. On worst-case to average-case reductions for np problems. *SIAM Journal on Computing*, 36(4):1119–1159, 2006.

[9] Boneh and Venkatesan. Breaking RSA may not be equivalent to factoring. In *EUROCRYPT: Advances in Cryptology: Proceedings of EUROCRYPT*, 1998.

[10] Dan Boneh, Periklis A. Papakonstantinou, Charles Rackoff, Yevgeniy Vahlis, and Brent Waters. On the Impossibility of Basing Identity Based Encryption on Trapdoor Permutations. In *FOCS*, pages 283–292, 2008.

[11] Gilles Brassard. Relativized cryptography. In *Proceedings of the 20th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 383–391. IEEE Computer Society, 1979.

[12] Emmanuel Bresson, Jean Monnerat, and Damien Vergnaud. Separation results on the "one-more" computational problems. In Tal Malkin, editor, *CT-RSA*, volume 4964 of *Lecture Notes in Computer Science*, pages 71–87. Springer, 2008.

[13] Dana Dachman-Soled, Yehuda Lindell, Mohammad Mahmoody, and Tal Malkin. On black-box complexity of optimally-fair coin-tossing. In *Theory of Cryptography Conference - TCC 2011*, 2011.

[14] Anindya De, Luca Trevisan, and Madhur Tulsiani. Time space tradeoffs for attacks against one-way functions and PRGs. In Tal Rabin, editor, *Advances in Cryptology - CRYPTO 2010, 30th Annual Cryptology Conference, Santa Barbara, CA, USA, August 15-19, 2010. Proceedings*, volume 6223 of *Lecture Notes in Computer Science*, pages 649–665. Springer, 2010.

[15] Yevgeniy Dodis, Roberto Oliveira, and Krzysztof Pietrzak. On the generic insecurity of the full domain hash. *Lecture Notes in Computer Science*, pages 449–??, 2005.

[16] Danny Dolev, Cynthia Dwork, and Moni Naor. Nonmalleable cryptography. *SIAM Journal on Computing*, 30(2):391–437 (electronic), 2000. Preliminary version in STOC 1991.

[17] Joan Feigenbaum and Lance Fortnow. Random-self-reducibility of complete sets. *SIAM Journal on Computing*, 22(5):994–1005, 1993.

[18] Marc Fischlin and Dominique Schröder. On the impossibility of three-move blind signature schemes. In Henri Gilbert, editor, *EUROCRYPT*, volume 6110 of *Lecture Notes in Computer Science*, pages 197–215. Springer, 2010.

[19] Gennaro, Gertner, and Katz. Lower bounds on the efficiency of encryption and digital signature schemes. In *STOC: ACM Symposium on Theory of Computing (STOC)*, 2003.

[20] Rosario Gennaro, Yael Gertner, Jonathan Katz, and Luca Trevisan. Bounds on the efficiency of generic cryptographic constructions. *SIAM J. Comput.*, 35(1):217–246, 2005.

[21] Rosario Gennaro and Luca Trevisan. Lower Bounds on the Efficiency of Generic Cryptographic constructions. In *FOCS*, pages 305–313, 2000.

[22] Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In Lance Fortnow and Salil P. Vadhan, editors, *STOC*, pages 99–108. ACM, 2011.

[23] Yael Gertner, Sampath Kannan, Tal Malkin, Omer Reingold, and Mahesh Viswanathan. The Relationship between Public Key Encryption and Oblivious transfer. In *FOCS*, pages 325–335, 2000.

[24] Yael Gertner, Tal Malkin, and Omer Reingold. On the Impossibility of Basing Trapdoor Functions on Trapdoor Predicates. In *FOCS*, pages 126–135, 2001.

[25] Oded Goldreich. *Foundations of Cryptography: Basic Applications*. Cambridge University Press, 2004.

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

[27] Oded Goldreich and Yair Oren. Definitions and properties of zero-knowledge proof systems. *Journal of Cryptology*, 7(1):1–32, 1994.

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

[29] Iftach Haitner, Jonathan J. Hoch, Omer Reingold, and Gil Segev. Finding collisions in interactive protocols – A tight lower bound on the round complexity of statistically-hiding commitments. In *Proceedings of the 47th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE Computer Society, 2007.

[30] Iftach Haitner and Thomas Holenstein. On the (im)possibility of key dependent encryption. In Omer Reingold, editor, *Theory of Cryptography, 6th Theory of Cryptography Conference, TCC 2009, San Francisco, CA, USA, March 15-17, 2009. Proceedings*, volume 5444 of *Lecture Notes in Computer Science*, pages 202–219. Springer, 2009.

[31] Iftach Haitner, Alon Rosen, and Ronen Shaltiel. On the (im)possibility of arthur-merlin witness hiding protocols. In *Theory of Cryptography, Fourth Theory of Cryptography Conference, TCC 2009*, 2009.

[32] Johan Håstad. Pseudo-random generators under uniform assumptions. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 387–394, 1990.

[33] Russell Impagliazzo, Leonid A. Levin, and Michael Luby. Pseudo-random generation from one-way functions. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing (STOC)*, pages 12–24, 1989.

[34] Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing (STOC)*, pages 44–61. ACM Press, 1989.

[35] Jonathan Katz, Dominique Schröder, and Arkady Yerukhimovich. Impossibility of blind signatures from one-way permutations. In Yuval Ishai, editor, *TCC*, volume 6597 of *Lecture Notes in Computer Science*, pages 615–629. Springer, 2011.

[36] Jeong Han Kim, Daniel R. Simon, and Prasad Tetali. Limits on the efficiency of one-way permutation-based hash functions. In *Proceedings of the 40th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 535–542, 1999.

[37] Neal Koblitz and Alfred Menezes. Another look at non-uniformity. *IACR Cryptology ePrint Archive*, 2012:359, 2012. informal publication.

[38] Lin, Trevisan, and Wee. On hardness amplification of one-way functions. In *Theory of Cryptography Conference (TCC), LNCS*, volume 2, 2005.

[39] Huijia Lin, Rafael Pass, and Muthuramakrishnan Venkitasubramaniam. Concurrent non-malleable commitments from any one-way function. In Ran Canetti, editor, *TCC*, volume 4948 of *Lecture Notes in Computer Science*, pages 571–588. Springer, 2008.

[40] Chi-Jen Lu. On the security loss in cryptographic reductions. In Antoine Joux, editor, *EUROCRYPT*, volume 5479 of *Lecture Notes in Computer Science*, pages 72–87. Springer, 2009.

[41] Mohammad Mahmoody and Rafael Pass. The curious case of non-interactive commitments - on the power of black-box vs. non-black-box use of primitives. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO*, volume 7417 of *Lecture Notes in Computer Science*, pages 701–718. Springer, 2012.

[42] Moni Naor. On cryptographic assumptions and challenges. *Lecture Notes in Computer Science*, 2729:96–109, 2003.

[43] Rafael Pass. Parallel repetition of zero-knowledge proofs and the possibility of basing cryptography on NP-hardness. In *IEEE Conference on Computational Complexity*, pages 96–110. IEEE Computer Society, 2006.

[44] Rafael Pass. Limits of provable security from standard assumptions. In Lance Fortnow and Salil P. Vadhan, editors, *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pages 109–118. ACM, 2011.

[45] Rafael Pass, Wei-Lung Dustin Tseng, and Muthuramakrishnan Venkitasubramaniam. Towards non-black-box lower bounds in cryptography. In Yuval Ishai, editor, *Theory of Cryptography - 8th Theory of Cryptography Conference, TCC 2011, Providence, RI, USA, March 28-30, 2011. Proceedings*, volume 6597 of *Lecture Notes in Computer Science*, pages 579–596. Springer, 2011.

[46] Omer Reingold, Luca Trevisan, and Salil P. Vadhan. Notions of reducibility between cryptographic primitives. In *Theory of Cryptography, First Theory of Cryptography Conference, TCC 2004*, volume 2951 of *Lecture Notes in Computer Science*, pages 1–20. Springer, 2004.

[47] Guy N. Rothblum and Salil P. Vadhan. Are PCPs inherent in efficient arguments? *Computational Complexity*, 19(2):265–304, 2010.

[48] Claus-Peter Schnorr. Efficient signature generation by smart cards. *J. Cryptology*, 4(3):161–174, 1991.

[49] Ronen Shaltiel and Emanuele Viola. Hardness amplification proofs require majority. *SIAM J. Comput*, 39(7):3122–3154, 2010.

[50] Daniel R. Simon. Finding Collisions on a One-Way Street: Can Secure Hash Functions Be Based on General Assumptions? In *EUROCRYPT*, pages 334–345, 1998.

[51] Dominique Unruh. Random oracles and auxiliary input. In Alfred Menezes, editor, *CRYPTO*, volume 4622 of *Lecture Notes in Computer Science*, pages 205–223. Springer, 2007.

[52] Yevgeniy Vahlis. Two Is a Crowd? A Black-Box Separation of One-Wayness and Security under Correlated Inputs. In *TCC*, pages 165–182, 2010.