

Towards Non-Black-Box Lower Bounds in Cryptography

Rafael Pass*, Wei-Lung Dustin Tseng, and Muthuramakrishnan Venkatasubramanian

Cornell University,
{rafael,wtdtseng,vmuthu}@cs.cornell.edu

Abstract. We consider average-case strengthenings of the traditional assumption that coNP is not contained in AM . Under these assumptions, we rule out generic and potentially *non-black-box* constructions of various cryptographic primitives (e.g., one-way permutations, collision-resistant hash-functions, constant-round statistically hiding commitments, and constant-round black-box zero-knowledge proofs for NP) from one-way functions, assuming the security reductions are *black-box*.

1 Introduction

In the past four decades, many cryptographic tasks have been put under rigorous treatment in an effort to realize these tasks under minimal assumptions. In particular, one-way functions are widely regarded as the most basic cryptographic primitive; their existence is implied by most other cryptographic tasks. Presently, one-way functions are known to imply schemes such as private-key encryption [GM84,GGM86,HILL99], pseudo-random generators [HILL99], statistically-binding commitments [Nao91], statistically-hiding commitments [NOVY98,HR07] and zero-knowledge proofs [GMW91]. At the same time, some other tasks still have no known constructions based on one-way functions (e.g., key agreement schemes or collision-resistant hash functions).

Following the seminal paper by Impagliazzo and Rudich [IR88], many works have addressed this phenomenon by demonstrating *black-box separations*, which rules out constructions of a cryptographic task using the underlying primitive as a *black-box*. For instance, Impagliazzo and Rudich rule out black-box constructions of key-agreement protocols (and thus also trapdoor predicates) from one-way functions; Simon [Sim98] rules out black-box constructions of collision-resistant hash functions from one-way functions. Furthermore, these impossibility results are *unconditional*.

Yet many classical cryptographic constructions (e.g., [FS90,DDN00,GMW91]) are non-black-box. This begs the question: to what extent does black-box separations give us insight into the actual separation of cryptographic primitives?

* Supported in part by a Microsoft New Faculty Fellowship, NSF CAREER Award CCF-0746990, AFOSR Award FA9550-08-1-0197 and BSF Grant 2006317

In this paper, we directly focus on providing lower bounds for *non-black-box* constructions of cryptographic primitives from one-way functions. We emphasize that although we consider non-black-box constructions, we still assume Turing (i.e., black-box) security reductions. For some of our results, we heavily leverage the existing literature on the impossibility of basing cryptography on NP hardness (these works also directly consider a Turing reduction of cryptographic primitives from NP). Perhaps surprisingly, we also make extensive use of known black-box separations. In other words, we demonstrate that some black-box separations *can* be modified to give further insight into the separation of cryptographic primitives.

Before stating our theorems, we first discuss our assumptions. Assumptions are necessary for non-black-box separations assuming black-box reductions; to show that a primitive P cannot be constructed using one-way functions, we must at least assume that a weak notion of so-called *somewhere-uninvertible* one-way functions exist—i.e. functions that cannot be inverted on *all* input lengths (as opposed to infinitely many lengths as in the traditional definition of one-way functions).¹ As one of the main contributions of the paper, we introduce general assumptions that we believe are reasonable, and are useful in establishing a variety of non-black-box separations.

1.1 Our Assumptions

Assumption 1. $\text{Dist}^{\text{1sided}}\text{-coNP} \not\subseteq \text{Heur}_{1/\text{poly}}\text{AM}$ is an average-case extension of the well-studied (and widely believed) classical assumption $\text{coNP} \not\subseteq \text{AM}$. Briefly, $\text{Dist}^{\text{1sided}}\text{-coNP}$ contains all coNP languages coupled with an efficiently samplable distribution over the no instances of the language. Such a language is considered to be in $\text{Heur}_{1/\text{poly}}\text{AM}$ if there exists an AM (constant-round) protocol that accepts the language, with the relaxation that soundness only needs to hold with high probability over the no instances, as measured by the given distribution. As we prove later, the assumption is equivalent to the existence of an efficiently computable function f that is not *heuristically co-range verifiable*—that is, there does not exist an AM protocol proving that an element is outside the range of f , where soundness holds with high probability for a random instance $f(x)$.² Assuming that there exists an efficiently computable function that is not heuristically co-range verifiable seems most reasonable (consider, for instance, proving that an element is not in the range of AES [DR02]). We additionally show that such a function is implied by the existence of pseudorandom generators³ secure against “promise-AM \cap coAM”.

Assumption 2. Our second assumption is of a different flavor: we assume the existence of one-way functions that are secure against $\text{PPT}^{\text{SAM}_d}$. Here SAM_d refers

¹ If Somewhere-Uninvertible OWFs do not exist, then every cryptographic primitive can be constructed from OWFs, because for every efficiently computable function, there would be a trivial reduction that inverts the function on all input lengths.

² See section 3 for a comparison with the literature of “average refutation” [FKO06].

³ Here we refer to BM-Y-type pseudo-random generators [BM84, Yao82].

to the depth- d collision finding oracle defined in [Sim98,HHRS07];⁴ $\text{PPT}^{\text{SAM}_d}$ refers to the class of probabilistic polynomial time machines with oracle access to SAM_d . This assumption is flexible since we can adjust the parameter d ; a larger d implies a stronger assumption (in fact, if $d = n/\log n$, the assumption is simply false since $\text{SAM}_{n/\log n}$ can in fact invert one-way functions [PV10]). In our work, we focus on the case $d = O(1)$ (constant depth), and refer to $\text{SAM}_{O(1)}$ simply as SAM .

Assumption 3. Our final and strongest assumption is $\text{Dist}^{\text{1sided}}\text{-coNP} \not\subseteq \text{Heur}_{1/\text{poly}} \text{IP}[\text{PPT}^{\text{NP}}]$ (heuristically verified by an interactive protocol where the prover is a probabilistic polynomial time machine with access to a NP oracle). It directly implies assumption 1, and relying on the work of Haitner, Mahmoody-Ghidary and Xiao [HMX10], we show that it implies assumption 2 as well in the case $d = O(1)$. Due to their similarity, $\text{Dist}^{\text{1sided}}\text{-coNP} \not\subseteq \text{Heur}_{1/\text{poly}} \text{IP}[\text{PPT}^{\text{NP}}]$ inherits many of the justifications as our first assumption in a weaker form (e.g., it is based on the classical assumption $\text{coNP} \not\subseteq \text{IP}[\text{PPT}^{\text{NP}}]$, and is equivalent to the existence of efficient functions whose co-range cannot be verified by $\text{IP}[\text{PPT}^{\text{NP}}]$ protocols). We treat assumption 3 as a unifying (and strongest) assumption that implies all of the results in our work.

Minimizing the assumption. It is natural to ask if the classical assumption $\text{coNP} \not\subseteq \text{AM}$, or perhaps the more standard average-case hardness assumption $\text{Dist-coNP} \not\subseteq \text{Heur}_{1/\text{poly}} \text{AM}$, are enough for our theorems (Dist-coNP consists of coNP languages coupled with efficiently samplable distributions that may range over all instances). We argue that it would be unlikely. In order to rule out constructions of cryptographic primitives based on OWFs, we first need to assume the existence of OWFs. But, it is unknown even if hard-on-the-average languages exist assuming only $\text{coNP} \not\subseteq \text{AM}$. Similarly, the stronger assumption $\text{Dist-coNP} \not\subseteq \text{Heur}_{1/\text{poly}} \text{AM}$ implies the existence of a hard-on-the-average language, but, as far as we know, does not imply the existence of OWFs (indeed, this is related to the question of whether one-way functions can be based on average-case hardness). Restricting to one-sided distributions (i.e., considering $\text{Dist}^{\text{1sided}}\text{-coNP}$ instead of Dist-coNP) is the next logical step, and this can be shown to imply a form of one-way functions (see full version).

1.2 Our results

As mentioned, we are able to prove many separation results by adapting numerous previous works to take advantage of our assumptions. We highlight the main separations here (grouped by their assumptions), and leave the numerous corollaries to the main text.

Based on the work of [Bra83], [AGGM06] and [Pas06], we have

⁴ Given an interactive Turing machine M and a transcript of $\leq d$ rounds, the SAM_d oracle samples uniformly from the set of random tapes on which the M would produce the given transcript.

Theorem 1 (Informal) *Assuming $\text{Dist}^{\text{1sided}}\text{-coNP} \not\subseteq \text{Heur}_{1/\text{poly}}\text{AM}$, one-way permutations and constant-round public-coin strongly witness-indistinguishable proofs for all of NP cannot be based on one-way functions with a Turing security reduction.*

Based on the work of [Sim98], [HHR07] and [PV10], we have

Theorem 2 (Informal) *Assuming the existence of one-way functions secure against $\text{PPT}^{\text{SAM}_{O(1)}}$ (implied by $\text{Dist}^{\text{1sided}}\text{-coNP} \not\subseteq \text{Heur}_{1/\text{poly}}\text{IP}[\text{PPT}^{\text{NP}}]$), collision-resistant hash functions, constant-round statistically hiding commitments, and constant-round black-box zero-knowledge proofs for all of NP cannot be based on one-way functions with a Turing security reduction.*

Remark 1. Based on the work of [HMX10], the results in Theorem 2 can be obtained under the weaker assumption of Theorem 1 if we restrict to security reductions that have constant adaptivity.

In addition to these theorems, we again stress the following philosophical contribution: with the right assumptions, not only are non-black-box separation results possible, many such separations can be based on existing techniques. For example, the black-box separation results of [Sim98], [HHR07] and [PV10] are essentially “upgraded” to non-black-box separations using our framework.

1.3 Our Techniques

Regarding the first assumption, $\text{Dist}^{\text{1sided}}\text{-coNP} \not\subseteq \text{Heur}_{1/\text{poly}}\text{AM}$, our separation results are largely based on previous works in the literature of separating cryptography from NP hardness, specifically ruling out constructions of one-way permutations [Bra83], size-verifiable one-way functions [AGGM06] and public-coin strongly witness-indistinguishable proofs [Pas06]. These works follow a common pattern: they take a (candidate) Turing security reduction of some cryptographic primitive P from NP, transform the reduction into an AM protocol, and conclude that $\text{coNP} \subseteq \text{AM}$, an unlikely consequence. By adapting their techniques, we show that a (candidate) Turing security reduction of the same primitive P from a one-way function can be transformed into an AM protocol that inverts the one-way function, and therefore the AM protocol may verify the co-range of f . This is a contradiction (not surprising since our assumption is an average case generalization of $\text{coNP} \not\subseteq \text{AM}$).

Our second assumption is used in a different fashion. Having justified the assumption that there exist one-way functions secure against $\text{SAM} = \text{SAM}_{O(1)}$, it follows that any cryptographic primitive P whose security can be broken using SAM cannot be based on one-way functions. This is because a Turing security reduction of primitive P from a one-way function f directly gives an algorithm that inverts f by using the SAM oracle, if $\text{SAM}_{O(1)}$ can be used to break the security of primitive P . The $\text{SAM}_{O(1)}$ oracle (as well as its variants) is particularly interesting in this aspect, since it is originally studied in the setting of black-box separations. Therefore, we know from previous works that in a relativized world

with the $\text{SAM}_{O(1)}$ oracle, there do not exist collision-resistant hash functions [Sim98], constant-round statistically hiding commitments [HHS07], and zero-knowledge proofs for all of NP [PV10]. In a similar spirit, other on black-box separations can also be extended also to non-black-box separations; the work then lies in justifying the resulting new assumption.

A note on Turing reductions. In this work, we only consider constructions with Turing security reductions; that is, reductions that use the adversary (supposedly breaking the security of the construction) as a black box. The non-black-box simulation technique of Barak [Bar01] demonstrates how the code of the adversary can be used in security proofs for certain interactive zero-knowledge protocols. Such non-black-box reductions might potentially also be useful in analyzing the security of other cryptographic tasks.

However, as we argue, in the context of basing cryptographic primitives on one another, Turing reductions provide a semantically stronger notion of security than non-black-box reductions. The existence of a Turing reduction from a primitive P to a primitive Q implies that any “physical device”—which might rely on physical phenomena—that breaks the security of primitive Q , can be used to break the security of primitive P . With a non-black-box security reduction, we would instead require an *explicit* description of the code of the attack on primitive Q . Such descriptions might be hard to find: consider, for instance, a “human-aided” computation, where a human is interacting with a computer program in order to break a crypto system;⁵ getting an explicit description of the attack would require providing an explicit (and “short”) description of the human brain.

2 Preliminaries

We assume familiarity with common complexity classes such as NP, AM, etc., as well as common cryptographic primitives such as one-way functions (OWF), collision-resistant hash-functions (CRH), zero-knowledge proofs (ZK), and witness-indistinguishable proofs (WI).

Let $[n]$ denotes the set $\{1, \dots, n\}$. Given an interactive protocol (P, V) (a pair of interactive Turing machines), let $\langle P, V \rangle(x)$ denote the output of V (the verifier) at the end of an execution with P (the prover), on common input x . Given a function $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$ and a polynomial $q(n)$, we say g is $q(n)$ **concatenations of f** to mean that for $x_1, \dots, x_{q(n)} \in \{0, 1\}^n$, $g(x_1, \dots, x_{q(n)}) = (f(x_1), \dots, f(x_{q(n)}))$ (on other input lengths, g considers part of the input to be padding appropriately).

2.1 Distributional Languages

Definition 3 (Distributional Languages). An **ensemble of distributions** is a collection $D = \{D_1, D_2, \dots\}$ where D_n is a distribution over $\{0, 1\}^n$. The en-

⁵ Practical attacks on crypto-systems are often not fully automatized, but do indeed rely on such interactions; see e.g., [AAG⁺00].

semble is **efficiently samplable** if there exists a probabilistic polynomial-time algorithm S that, on input 1^n , outputs a sample according to D_n . A **distributional language** is a pair (L, D) where L is a standard language and D is an ensemble of distributions.

A well known class of distributional languages is **Dist-coNP**; it contains the set of distributional languages (L, D) where $L \in \text{coNP}$ and D is efficiently samplable.

2.2 Hardness Amplification of One-Way Functions

The following lemma on hardness amplification of one-way functions is due to Yao [Yao82].

Lemma 4 ([Yao82]) *Let $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$ be an efficiently computable function. Given any polynomial $q(n)$, let g be $q(n)$ concatenations of f . Then there is a PPT oracle machine $\mathcal{A}^\mathcal{O}$ such that whenever \mathcal{O} is an oracle that inverts g with non-negligible probability, i.e., there exists some polynomial $p(n)$ such that for some set of n 's,*

$$\Pr_{x \leftarrow \{0, 1\}^{nq(n)}} [\mathcal{O}(g(x)) \in g^{-1}(g(x))] \geq 1/p(n)$$

then $\mathcal{A}^\mathcal{O}$ inverts f with probability $1 - 1/q(n)$, i.e., for the same set of n 's,

$$\Pr_{x \leftarrow \{0, 1\}^n} [\mathcal{A}^\mathcal{O}(f(x)) \in f^{-1}(f(x))] \geq 1 - 1/q(n)$$

3 On $\text{Dist}^{\text{1sided}}\text{-coNP} \not\subseteq \text{Heur}_{1/\text{poly}} \text{AM}$

In this section we discuss our first assumption, $\text{Dist}^{\text{1sided}}\text{-coNP} \not\subseteq \text{Heur}_{1/\text{poly}} \text{AM}$, starting with definitions, followed by its relation to other assumptions, and its implications on basing cryptography on one-way functions.

Definition 5. A distributional language (L, D) is in $\text{Dist}^{\text{1sided}}\text{-coNP}$ if and only if $L \in \text{coNP}$, D is efficiently samplable, and D only ranges over \bar{L} .

Remark 2. In other words, $(L, D) \in \text{Dist}^{\text{1sided}}\text{-coNP}$ if and only if $(L, D) \in \text{Dist-coNP}$ and D only sample instances *not* in L .

Definition 6. A distributional language (L, D) is in $\text{Heur}_{1/\text{poly}} \text{AM}$ if for every polynomial q , there exists an AM (i.e., constant-round public-coin) protocol (P, V) such that:

Completeness: If $x \in L$, $\Pr[\langle P, V \rangle(x) = 1] \geq 2/3$.

Soundness: For every $n \in \mathbb{N}$ and every machine P^* , with probability $1 - 1/q(n)$, an $x \in \{0, 1\}^n$ sampled from D_n conditioned on $x \notin L$ satisfies $\Pr[\langle P^*, V \rangle(x) = 1] \leq 1/3$.

Remark 3. As usual, the choice of $2/3$ and $1/3$ is arbitrary and can be amplified to $1 - 2^{-n}$ and 2^{-n} . Intuitively, the soundness condition means that L is almost in AM, except for a fraction of instances in \bar{L} that is sampled with (arbitrarily small) polynomial probability.

Remark 4. In a related work, Feige, Kim and Ofek give positive results in refuting restricted random coSAT instances on average [FKO06]. The main difference between the notion of average refutation and our definition of heuristic verifiability is in where errors are allowed. An average refutation algorithm may not refute a random unsatisfiable instance with small probability, but will never refute a satisfiable instance (i.e., perfect soundness). On a philosophical level, the work of [FKO06] gives a distribution of coSAT instances that may indeed be heuristically verifiable.

The complexity assumption we consider is $\text{Dist}^{\text{1sided}}\text{-coNP} \not\subseteq \text{Heur}_{1/\text{poly}}\text{AM}$, which is a strengthening of the more standard assumption that $\text{Dist-coNP} \not\subseteq \text{Heur}_{1/\text{poly}}\text{AM}$, which in turn is the heuristic analog of $\text{coNP} \not\subseteq \text{AM}$.

Relation to other assumptions. To get a more concrete handle on our assumption, we prove that $\text{Dist}^{\text{1sided}}\text{-coNP} \not\subseteq \text{Heur}_{1/\text{poly}}\text{AM}$ is equivalent to the existence of an efficiently computable function f that is not heuristically co-range verifiable, i.e., there does not exist an AM protocol proving that an instance is outside the range of f , where soundness holds only with high probability with respect to random instances of $f(x)$. We then present several candidates for such a function (such as AES [DR02] and Learning Parity with Noise [BFKL93]). Using this equivalence, we also show that $\text{Dist}^{\text{1sided}}\text{-coNP} \not\subseteq \text{Heur}_{1/\text{poly}}\text{AM}$ is implied by the existence of pseudorandom generators secure against $\text{BPP}(\text{Promise}(\text{AM} \cap \text{coAM}))$.⁶

3.1 Heuristic co-Range Verifiable Functions

Given a function f , consider the language $\text{Range}_f = \{f(x) \mid x \in \{0, 1\}^*\}$.

Definition 7. f is **heuristically co-range verifiable** if for any polynomial p , there exists an AM (i.e., constant-round public-coin) protocol (P, V) such that:

Completeness: For every $y \notin \text{Range}_f$, $\Pr[\langle P, V \rangle(y) = 1] \geq 2/3$.

Soundness: For every $n \in \mathbb{N}$ and every machine P^* , with probability $1 - 1/p(n)$ over $x \leftarrow \{0, 1\}^n$, $\Pr[\langle P^*, V \rangle(f(x)) = 1] \leq 1/3$.

Theorem 8 $\text{Dist}^{\text{1sided}}\text{-coNP} \not\subseteq \text{Heur}_{1/\text{poly}}\text{AM}$ if and only if there exists an efficiently computable function that is not heuristically co-range verifiable.

Proof. We show each direction separately.

⁶ Traditionally, NW-style [NW94] PRGs against AM have been considered in the literature (see e.g., [MV05]); in contrast, we require a BMY-style [BM84, Yao82] “cryptographic” PRG.

“if” part: Let f be a function that is not heuristically co-range verifiable. By padding the input/output of f , construct another efficiently computable function g that is length preserving (i.e., $|g(x)| = |x|$ for all x). It is easy to see that padding preserves heuristic co-range verifiability, and so g is also not heuristically co-range verifiable. Consider the $\text{Dist}^{\text{1sided}}\text{-coNP}$ distributional language (L, D) where $L = \overline{\text{Range}_g}$ and D_n is the distribution that results from computing g on a uniformly random $x \in \{0, 1\}^n$. Because g is not heuristically co-range verifiable, $(L, D) \notin \text{Heur}_{1/\text{poly}} \text{AM}$.

“only-if” part: Let (L, D) be a distributional language such that $(L, D) \in \text{Dist}^{\text{1sided}}\text{-coNP}$ and $(L, D) \notin \text{Heur}_{1/\text{poly}} \text{AM}$, and let $t(n)$ be a bound on the random bits required to efficiently sample from D_n . Define f on input $x \in \{0, 1\}^{t(n)}$ to be the result of sampling from D_n given randomness x (for other input lengths, f may treat part of the input as padding). f is an efficient function since D is efficiently samplable, and f is not heuristically co-range verifiable precisely because $(L, D) \notin \text{Heur}_{1/\text{poly}} \text{AM}$. \square

The statement “ f is heuristically co-range verifiable” can be viewed as an average-case (heuristic) variant of the statement “ $\text{Range}_f \in \text{coAM}$ ”. (Also observe that if f is efficiently computable then $\text{Range}_f \in \text{NP} \subseteq \text{AM}$.) We believe that the existence of such functions is a reasonable average-case generalization of $\text{SAT} \notin \text{coAM}$: Just as it seems “unlikely” that there exist AM proofs for proving that a string is outside an arbitrary NP set, it seems “unlikely” that there is a AM proof for proving that a string is outside the range an arbitrary efficiently computable function, even if we only require soundness to hold for a random string in the range of the function.

Candidate functions that are not heuristic co-range verifiable. Although many traditional one-way functions (based for example on the hardness of factoring, RSA, discrete log [Rab80], or lattice-based problems [GG00,AR05]) are co-range verifiable, there are also “natural” one-way functions for which we do not know of co-range verifiability protocols. We here briefly discuss a few functions that are not known to be heuristically co-range verifiable.

Generalized AES: AES is a permutation on 128 bits [DR02]; that is, for a 128-bit seeds, AES_s is a permutation on defined on $\{0, 1\}^{128}$. However, due to the algebraic nature of the construction of AES, it can easily be generalized to longer input lengths. Let AES^n denote this generalized version of AES to n -bit inputs. Now, consider the (one-way) function $f(x) = \text{AES}_x^{|x|}(0^{|x|})$. It would seem unlikely that this function is heuristically co-range verifiable.

Random Binary Linear Codes: A random binary linear code is obtained by encoding a message $x \in \{0, 1\}^n$ as Ax where A is a random $m \times n$ binary matrix. Given the matrix A and a codeword y , it is easy to find the corresponding message x when $m \geq n$. However, the problem of finding x becomes hard when only a “noisy” codeword is given. The *learning parity with noise* (LPN) problem requires finding a random secret x , given $(A, Ax + e)$ where e is a “short” (binary) error vector. The worst-case variant of the LPN problem (i.e. given a set of equations $Ax = s$ to find x that maximally satisfies

the equations) is known to be NP-hard even to approximate [Hås01]. The average-case version of LPN is also believed to be intractable: the $\text{LPN}_{p,m}$ assumption [BFKL93] states that for $p \in (0, \frac{1}{2})$ and polynomial m , there is no PPT algorithm that finds x with more than negligible probability given $(A, Ax + e \bmod 2)$ where A is a random $m \times n$ binary matrix and every component of e is set to 1 independently with probability p . It seems like a reasonable strengthening of the LPN assumption to say that the function $x \mapsto (A, Ax + e \bmod 2)$ is not heuristically co-range verifiable, for some choices of m and p . In other words, there is no AM-proof showing that a binary string y is “far” from Ax for any x , even if soundness only holds for randomly perturbed codewords.

Pseudo-random Generators secure against $\text{BPP}(\text{Promise}(\text{AM} \cap \text{coAM}))$: While not a specific function, we show that this class of PRGs are not heuristically co-range verifiable.

Definition 9. Let U_n denote the distribution of uniform bit-strings of length n . A collection of efficiently computable functions $\mathcal{G} = \{g_n : \{0,1\}^n \rightarrow \{0,1\}^{n+1}\}_{n \in \mathbb{N}}$ is a PRG secure against $\text{BPP}(\text{Promise}(\text{AM} \cap \text{coAM}))$ if no PPT adversary with a $\text{Promise}(\text{AM} \cap \text{coAM})$ oracle can distinguish the ensembles $\{g_n(U_n)\}_{n \in \mathbb{N}}$ and $\{U_{n+1}\}_{n \in \mathbb{N}}$ with non-negligible probability in n .

Claim 10 Let $g : \{0,1\}^n \rightarrow \{0,1\}^{n+1}$ be a PRG secure against $\text{BPP}(\text{Promise}(\text{AM} \cap \text{coAM}))$. Then g is not heuristically range verifiable.

Proof. Assume for contradiction that g is heuristically range verifiable. By the definition of heuristic range verifiability, there is a AM protocol (P, V) such that on input $g(x)$ for a uniformly random $x \in \{0,1\}^n$, V rejects $g(x)$ with probability at least $1 - 1/n$. Let $S = \{x \in \{0,1\}^n \mid \Pr[V \text{ rejects } g(x)] \leq 1/n\}$ (i.e., the set of x where V fails to reject $g(x)$). Then we must have

$$\Pr_{x \leftarrow \{0,1\}^n} [x \in S] \leq 2/n$$

Let $T = \{g(x) \mid x \in S\}$, i.e., the set of inputs where (P, V) has high soundness error. Now consider the promise problem $\Pi = (\Pi_Y, \Pi_N) = (\text{Range}_g - T, \overline{\text{Range}_g})$. Note that Π is trivially in $\text{NP} \subseteq \text{AM}$, and that $\Pi \in \text{coAM}$ by definition of T (via protocol (P, V)). Therefore $\Pi \in \text{AM} \cap \text{coAM}$.

We now describe a polynomial-time distinguisher \mathcal{D} that has oracle access to a decision procedure for the the promise problem Π . On input y , \mathcal{D} simply outputs $\Pi(y)$. To show that \mathcal{D} is a good distinguisher for g , observe that

$$\Pr_{x \leftarrow \{0,1\}^n} [\mathcal{D}(g(x)) = 1] \geq \Pr_x [g(x) \notin T] = \Pr_x [x \notin S] \geq 1 - \frac{2}{n}$$

On the other hand,

$$\Pr_{y \leftarrow \{0,1\}^{n+1}} [\mathcal{D}(y) = 1] \leq \Pr_y [y \notin \text{Range}_g] \leq \frac{1}{2} \quad \square$$

Claim 10 together with forthcoming theorems yields the following trade-off: if certain cryptographic primitives can be based on OWFs, then there does not exist PRGs secure against $\text{BPP}(\text{Promise}(\text{AM} \cap \text{coAM}))$.

3.2 Consequences of $\text{Dist}^{\text{1sided}}\text{-coNP} \not\subseteq \text{Heur}_{1/\text{poly}}\text{AM}$

The assumption $\text{Dist}^{\text{1sided}}\text{-coNP} \not\subseteq \text{Heur}_{1/\text{poly}}\text{AM}$ implies some impossibility results on basing cryptographic primitives on one-way functions. First, we provide an outline of our proof framework.

Recall that we consider arbitrary non-black-box (and even non explicit) constructions based on one-way functions, but restrict our attention to Turing (black-box) security reductions. This means a primitive P constructed from a one-way function f is accompanied by a PPT oracle reduction $R^{\mathcal{O}}$, such that whenever \mathcal{O} is an oracle that “breaks the security of P ”, $R^{\mathcal{O}}$ inverts the f with non-negligible probability. We will show that for certain primitives P and respective oracles \mathcal{O} that break the security of P , the reduction $R^{\mathcal{O}}$ can be emulated in an AM protocol, allowing the verifier of the AM protocol to invert the one-way function. Coupled with the Yao’s amplification lemma (Lemma 4), the verifier can actually invert f with very high probability, and therefore heuristically verify the co-range of f (by checking for a lack of inverses).

We present the lower-bound result for one-way permutations and Strong WI AM proofs based on OWFs below.

On Basing One-Way Permutations on One-Way Functions We first formalize the definition of basing one-way permutations (OWP) on one-way functions (OWF) with Turing (black-box) reductions, and show that such a construction is ruled out by the assumption $\text{Dist}^{\text{1sided}}\text{-coNP} \not\subseteq \text{Heur}_{1/\text{poly}}\text{AM}$.

Definition 11. We say that **OWPs can be based on OWFs** if:

Construction: There is a mapping that takes the description of any polynomial-time function f (candidate OWF) and outputs the description of a permutation $\phi = \phi_f$ (candidate OWP).

Reduction: For any polynomial-time function f , there is a PPT oracle algorithm R_f such that whenever \mathcal{O} inverts ϕ , i.e., there is a polynomial p such that $\Pr_{x \leftarrow \{0,1\}^n}[\mathcal{O}(\phi(x)) = x] \geq 1/p(n)$, $R_f^{\mathcal{O}}$ inverts f , i.e., there is some polynomial p' such that

$$\Pr_{x \leftarrow \{0,1\}^n}[R_f^{\mathcal{O}}(f(x)) \in f^{-1}(f(x))] \geq 1/p'(n)$$

The following theorem is proved using our framework combined with the work of [Bra83].

Theorem 12 *If OWPs can be based on OWFs, then $\text{Dist}^{\text{1sided}}\text{-coNP} \subseteq \text{Heur}_{1/\text{poly}}\text{AM}$ (contradicting our assumption).*

Proof. Suppose that OWPs can be based on OWFs. We will show that every efficiently computable function is heuristically co-range verifiable. Fix any efficient function f and polynomial $q(n)$ (as in the definition of heuristically co-range verifiability), and define g to be $q(n)$ concatenations of f . By assumption, there exists a permutation P_g and an efficient security reduction R_g such that, given

an oracle \mathcal{O} that inverts ϕ inverts g , $R_g^{\mathcal{O}}$ inverts g with non-negligible probability. Using Lemma 4, we can construct a new efficient reduction \tilde{R}_f that, given an oracle \mathcal{O} that inverts ϕ inverts g , $\tilde{R}_f^{\mathcal{O}}$ inverts f with probability $1 - 1/q(n)$.

Next we recall from [Bra83] an AM protocol that allows the verifier to run \tilde{R}_f without access to \mathcal{O} . The verifier starts by sending the prover a sufficiently long random string to act as the random tape of \tilde{R}_f . The prover then runs \tilde{R}_f with the given randomness, solving oracle queries as needed. When \tilde{R}_f terminates, the prover sends the output of \tilde{R}_f as well as any oracle query-answer pairs encountered in the execution of \tilde{R}_f to the verifier. The verifier can check the validity of the oracle query-answer pairs, and the validity of the execution using the given oracle query-answer pairs. On common input y , the verifier accepts if and only if $\tilde{R}_f(y)$ fails to find an inverse.

Completeness: If $y \notin \text{Range}_f$, and if the prover simulates $\tilde{R}_f(y)$ honestly, then the verifier will always accept the simulation, and of course \tilde{R}_f will never find an inverse to y under f . Hence we have completeness probability 1.

Soundness: We may assume that the verifier accepts the execution of $\tilde{R}_f(y)$ provided by the (possibly cheating) prover. In this case, the simulated execution of $\tilde{R}_f(y)$ is identical to a real execution of $\tilde{R}_f^{\mathcal{O}}(y)$ for a “perfect oracle” \mathcal{O} that answers all queries correctly; this is because every oracle has exactly one answer. Therefore:

$$\Pr_{x \leftarrow \{0,1\}^n} [\tilde{R}_f(f(x)) \in f^{-1}(f(x))] > 1 - 1/q(n)$$

By an averaging argument, we have that with probability at least $1 - 3/q(n)$ over a random $x \in \{0, 1\}^n$, $y = f(x)$,

$$\Pr[\tilde{R}_f(f(x)) \in f^{-1}(f(x))] > 2/3$$

in which case the verifier would reject.

This concludes that f is heuristically co-range verifiable.

Remark 5. The difficulty of extending Theorem 12 to other cryptographic primitives comes from constructing an AM protocol. For many primitives (e.g., collections of trapdoor one-way functions), an oracle that breaks the security of the primitive suffers from two caveats: some queries have no answers (which cannot be checked by the verifier), and some queries have multiple answers (which allow a cheating prover to adaptively select the answer). These difficulties are well known; see [BT03,AGGM06,HMX10].

Theorem 12 can be extended beyond one-way permutations. For example, it can rule out basing *certified collection of (trapdoor) permutations* on one-way functions [BY96]. In this case, an oracle query consists of a candidate permutation description and a candidate image. The verifier can check whether each description is indeed a valid permutation in the collection (certifiable), and if so

expect a unique inverse of the given image. (We may even extend the definition of “certified” to mean certifiable under an AM protocol.)

Another example is to rule out basing *size-verifiable, polynomial-sized pre-image one-way functions* on one-way functions [AGGM06]. In this case, size-verifiable one-way functions allow the verifier to check the pre-image size of any oracle query (in particular the verifier checks whether a pre-image exists). Then, the verifier may ask the prover to provide all polynomially many pre-images to force a unique answer.

On Basing Public-Coin Strongly Witness Indistinguishable Proofs on OWFs Using the same framework, we rule out the possibility of basing $O(1)$ -round public-coin strongly witness-indistinguishable proofs (Strong-WI AM) for languages in NP on OWFs. Below, we provide the result and brief overview of the proof. The complete proof will appear in the full version.

The definition of basing Strong-WI AM proofs on OWFs can be extended similarly to OWPs. Roughly speaking, for any language L , there exists a mapping from the description of any function f to a protocol $(P_f^{\text{SWI}}, V_f^{\text{SWI}})$ and a reduction R such that for any adversary O and pair of ensembles of distributions, $\{D_n^1\}_{n \in \mathbb{N}}$ and $\{D_n^2\}_{n \in \mathbb{N}}$, and D_n^1 and D_n^2 are distributions over $L \cap \{0, 1\}^n \times \{0, 1\}^*$, if O distinguishes proofs of statements using $(P_f^{\text{SWI}}, V_f^{\text{SWI}})$ sampled from the two distributions D_n^1 and D_n^2 , then R^O inverts f with non-negligible probability. The main result we obtain using the work of [Pas06] is

Theorem 13 *If there exists $O(1)$ -round Strong-WI AM proof systems with perfect completeness based on OWFs for all NP-languages, then $\text{Dist}^{\text{1sided}}\text{-coNP} \subseteq \text{Heur}_{1/\text{poly}}\text{AM}$*

On a high-level, [Pas06] shows how to construct a game G^f from any function f using a Strong-WI AM protocol for NP languages based on f such that there exists a reduction from breaking the game to inverting the function f . Additionally, he shows that a worst-case breaking oracle for G^f can be simulated using an AM protocol. We obtain our result using the same game G^f but instead of using any one-way function f , we use the function g obtained from any language $(L, D) \in \text{Dist}^{\text{1sided}}\text{-coNP}$ as in the proof for OWP. Since a worst-case breaker can be simulated using an AM protocol, following the proof technique from Theorem 12, it essentially follows that $(L, D) \in \text{Heur}_{1/\text{poly}}\text{AM}$.

4 On One-Way Functions Secure Against $\text{PPT}^{\text{SAM}_{O(1)}}$

In this section we explore our second assumption: the existence of one-way functions that cannot be inverted by $\text{PPT}^{\text{SAM}_{O(1)}}$: efficient algorithms that have access to a $\text{SAM}_{O(1)}$ oracle.

4.1 Definition of the SAM oracle

Let M be a probabilistic interactive Turing machine that runs a d -round protocol. Let $\text{TRANS}_i = (a_1, b_1, \dots, a_i, b_i)$ be a partial transcript of the messages exchange with $M(1^n)$ in an execution. We use $::$ to denote appending messages to a transcript. Define $R_{\text{TRANS}_i}(M)$ to be the set of all random tapes τ for which $M_\tau(1^n, a_1, b_1, \dots, b_{j-1}) = a_j$ for all $j < i$; we say that such a τ is *consistent* with respect to TRANS_i . Without loss of generality, we assume that M sends the first message (i.e., outputs a message on initiation). The oracle $\text{SAM}_{d(n)}$ takes inputs of the form $Q = (M(1^n), \text{TRANS}_i, r)$ where $\text{TRANS}_{i-1} = (a_1, b_1, \dots, b_{i-1})$ is a partial transcript and $r \in \{0, 1\}^*$. On input Q , $\text{SAM}_{d(n)}$ outputs $(\tau', \text{TRANS}_{i-1} :: a_i)$ such that $\tau' \in R_{\text{TRANS}_{i-1}}(M(1^n))$ and $M_{\tau'}(1^n, \text{TRANS}_i) = a_i$,⁷ with the following restrictions:

1. If $i > 1$, then $(a_1, b_1, \dots, a_{i-1})$ was the result of a previous query of the form $(M, (a_1, b_1, \dots, b_{i-2}), r')$ for some $r' \in \{0, 1\}^*$.
2. τ' is uniformly distributed in $R_{\text{TRANS}_{i-1}}(M)$ over the randomness of $\text{SAM}_{d(n)}$, independent of all other queries.
3. $\text{SAM}_{d(n)}$ answers queries only up to a depth $d(n)$, i.e. $i \leq d(n)$.

Otherwise, $\text{SAM}_{d(n)}$ outputs \perp . The role of r in the query is to obtain new and independent samples for each r and to allow a verifier to obtain the same sample query by querying on the same r .

Our above description of the $\text{SAM}_{d(n)}$ -oracle is a stateful instantiation of the oracle defined in [HHR07]. Just as in [HHR07], for our results, we need the oracle to be stateless; [HHR07] specify how to modify the oracle to achieve this (using “signatures”); we omit the details. When clear from context, we drop the input 1^n to M .

Definition 14. We say that a (one-way) function $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$ is **secure against** (or **hard to invert by**) $\text{PPT}^{\text{SAM}_d}$ if for every oracle PPT machine A there exists a negligible function $\nu(\cdot)$ such that

$$\Pr[x \leftarrow \{0, 1\}^n; y = f(x) : A^{\text{SAM}_d}(y) \in f^{-1}(y)] \leq \nu(n)$$

In this work, we focus on the $\text{SAM}_{O(1)}$ and in the rest of the paper, we refer to this oracle simply by **SAM**.

Definition 15. We say that a language L is in BPP^{SAM} if there exists an oracle PPT machine M such that the following holds:

Completeness: For every $x \in L$, $\Pr[M^{\text{SAM}}(x) = 1] \geq 2/3$

Soundness: For every $x \notin L$, $\Pr[M^{\text{SAM}}(x) = 1] \leq 1/2$

The second assumption that we consider to establish non black-box lower bounds is the existence of one-way functions that are secure against PPT^{SAM} . We justify our assumption in the next section.

⁷ It suffices to consider an oracle that merely outputs τ' , however, we consider **SAM** that additionally outputs $\text{TRANS}_{i-1} :: a_i$ for ease of exposition.

4.2 Relation to $\text{Dist}^{\text{1sided}}\text{-coNP} \not\subseteq \text{Heur}_{1/\text{poly}} \text{IP}[\text{PPT}^{\text{NP}}]$

Definition 16. A distributional language (L, D) is in $\text{Heur}_{1/\text{poly}} \text{IP}[\text{PPT}^{\text{NP}}]$ if for every polynomial q , there exists an interactive protocol (P, V) where $P \in \text{PPT}^{\text{NP}}$ (oracle PPT machine with oracle access to an NP oracle) such that:

Completeness: If $x \in L$, $\Pr[\langle P, V \rangle(x) = 1] \geq 2/3$.

Soundness: For every $n \in \mathbb{N}$ and every machine P^* , with probability $1 - 1/q(n)$, an $x \in \{0, 1\}^n$ sampled from D_n conditioned on $x \notin L$ satisfies $\Pr[\langle P^*, V \rangle(x) = 1] \leq 1/3$.

The assumption $\text{Dist}^{\text{1sided}}\text{-coNP} \not\subseteq \text{Heur}_{1/\text{poly}} \text{IP}[\text{PPT}^{\text{NP}}]$ is a heuristic extension of the worst case assumption $\text{coNP} \not\subseteq \text{IP}[\text{PPT}^{\text{NP}}]$, i.e., there are no interactive proofs for coSAT where the prover is efficient with a NP oracle. While $\text{coNP} \not\subseteq \text{IP}[\text{PPT}^{\text{NP}}]$ is not as well studied as more standard assumptions like $\text{coNP} \not\subseteq \text{AM}$, the search for the aforementioned interactive proof for coSAT has been open since the question was raised by Babai, Fortnow and Lund in 1991 [BFL91]. Next we show that $\text{Dist}^{\text{1sided}}\text{-coNP} \not\subseteq \text{Heur}_{1/\text{poly}} \text{IP}[\text{PPT}^{\text{NP}}]$ implies the existence of one-way functions secure against PPT^{SAM} ; the bulk of the technical content of the proof is taken from [HMX10].

Lemma 17 *If $\text{Dist}^{\text{1sided}}\text{-coNP} \not\subseteq \text{Heur}_{1/\text{poly}} \text{IP}[\text{PPT}^{\text{NP}}]$, then there exists a one-way function that is secure against PPT^{SAM} .*

Proof. We prove the contrapositive. Suppose all efficiently computable functions can be inverted by PPT^{SAM} . Fix any $(L, D) \in \text{Dist}^{\text{1sided}}\text{-coNP}$ and any polynomial q as in the definition of $\text{Heur}_{1/\text{poly}} \text{IP}[\text{PPT}^{\text{NP}}]$. We will show that $(L, D) \in \text{Heur}_{1/\text{poly}} \text{IP}[\text{PPT}^{\text{NP}}]$.

Let $t(n)$ be a bound on the randomness required to efficiently sample from D_n , define f on input $x \in \{0, 1\}^{t(n)}$ to be the result of sampling from D_n given randomness x , and let $g = g_q$ be $q(n)$ concatenations of f . By assumption, there is a PPT oracle algorithm R such that R^{SAM} inverts g with polynomial probability. By Lemma 4, we can further construct a PPT oracle algorithm \tilde{R} such that \tilde{R}^{SAM} inverts f with probability $1 - 1/q(n)$.

By the work of Haitner et. al [HMX10], the reduction \tilde{R} can be simulated in an interactive proof (P, V) where the P is an efficient algorithm with access to an NP oracle. Specifically, using Theorem 5.2 of [HMX10]⁸, with parameter $\delta = 1/q$, (P, V) has two properties:

Completeness: (P, V) has completeness error $1/q(n)$ (the probability that V aborts).

Soundness: For any (possibly cheating) prover P^* , if V does not abort, $\langle P^*, V \rangle(y)$ (the output of V) and the output of $\tilde{R}^{\text{SAM}}(y)$ has statistical difference at most $1/q(n)$.

⁸ The theorem number refers to the full version of [HMX10] on ECCC.

We modify the protocol so that V on input y accepts if and only if V does not abort during the simulation of \tilde{R} , and that \tilde{R} does not find an inverse of y under f . The resulting protocol shows that $(L, D) \in \text{Heur}_{1/\text{poly}} \text{IP}[\text{PPT}^{\text{NP}}]$:

Completeness: On input $y \in L$, i.e., $y \notin \text{Range}_f$, V only rejects during the simulation of \tilde{R} because \tilde{R} can never find an inverse to y . Therefore V rejects with probability at most $1/q(n)$.

Soundness: Let P^* be an arbitrary machine. On a random input $y \notin L$ distributed according to D_n , i.e., $y = f(x)$ for a random $x \in \{0, 1\}^{t(n)}$, $\tilde{R}^{\text{SAM}}(y)$ would find an inverse of y with probability $1 - 1/q(n)$. Therefore, if V does not reject the simulation of \tilde{R} provided by P^* , V would find an inverse of y with probability at least $1 - 2/q(n)$. By an averaging argument, with probability at least $1 - 3/q(n)$ over choosing y from D_n , $\Pr[\langle P^*, V \rangle(y) = 0] \geq 2/3$.

4.3 Consequences of the existence of one-way function secure w.r.t PPT^{SAM}

Assuming the existence of one-way function secure against PPT^{SAM} we show separation of collision-resistant hash-functions, $O(1)$ -round statistically-hiding commitments and $O(1)$ -round zero-knowledge proofs for NP from OWFs. On a high-level, for each of these primitives, we show that there exists an adversary that can break the security with oracle access to SAM. Therefore, if these primitives could be based on one-way functions, then we arrive at a contradiction under the assumption.

As with the case of one-way permutations, we consider arbitrary non-black-box (and even non explicit) constructions, but as before restrict attention to Turing (i.e., black-box) security reductions. The definitions of basing CRHs, statistically-hiding commitments and zero-knowledge proofs on one-way functions can be extended analogously from OWP. Below we discuss briefly how the SAM oracle can be used to break each primitive.

Collision-Resistant Hash-Functions: Recall that, the SAM oracle can sample uniform collisions for probabilistic interactive Turing machines. If we consider the efficient Turing machine that computes the CRH function, it follows that SAM can find a collision for a uniform input to the CRH if one exists. Since any length-compressing function with high-probability has collisions for uniformly chosen inputs, SAM breaks any CRH. We remark that it suffices to consider the potentially weaker SAM_1 -oracle to break CRHs. As a consequence, we obtain the following theorem.

Theorem 18 *Assuming the existence of one-way functions that are secure against PPT^{SAM} , we have that worst-case CRHs cannot be based on OWFs.*

As a corollary, we also obtain (a potentially weaker statement) that worst-case CRHs cannot be based on OWFs unless $\text{Dist}^{\text{1sided}}\text{-coNP} \subseteq \text{Heur}_{1/\text{poly}} \text{IP}[\text{PPT}^{\text{NP}}]$.

Statistically-Hiding Commitments: We show that, for every $O(1)$ -round statistically hiding commitment based on one-way functions, there exists a cheating sender who with oracle access to SAM violates the binding property of the commitment. Haitner, Hoch, Reingold and Segev [HHR07] prove that using the stronger SAM^π oracle (that finds collisions for PPT machines that access a random permutation oracle π), there is a cheating committer that can break the binding property of any fully black-box construction of a statistically-hiding commitment scheme based on one-way permutations. It essentially follows using the same proof that without access to any oracle π , SAM can break any statistically-hiding commitment scheme with a PPT committer. As a consequence, we obtain the following theorem.

Theorem 19 *Assuming the existence of one-way functions secure w.r.t. PPT^{SAM} , then there exists no $O(1)$ -round statistically-hiding bit-commitment scheme based on one-way function.*

As for the case of CRHs, we also have that there exists no $O(1)$ -round statistically-hiding bit-commitment scheme unless $\text{Dist}^{\text{1sided}}\text{-coNP} \subseteq \text{Heur}_{1/\text{poly}} \text{IP}[\text{PPT}^{\text{NP}}]$.

Zero-Knowledge Proofs: Using similar techniques we show how to extend to lower-bound of [PV10] on $O(1)$ -round zero-knowledge proofs based on one-way functions. Goldreich-Krawczyk [GK96b] showed that only languages in BPP have constant-round *public-coin* black-box zero-knowledge protocols. In [PV10], this lower bound was extended to “fully black-box” constructions of black-box zero-knowledge proofs (that could be *private-coin*) based on one-way functions. More precisely, they show that only languages decidable by oracle PPT machines with oracle access to SAM^π (for random permutation π) can have constant-round fully black-box zero-knowledge proofs. On a high-level, they establish this lower-bound, by providing a transformation that takes any private-coin zero-knowledge proof based on OWFs and produces a public-coin zero-knowledge proof in a SAM^π -relativized world and then concluding using the result of Goldreich-Krawczyk for public-coin protocols. Based on the result of [PV10], we obtain the following theorem.

Theorem 20 *Assume the existence of one-way functions that are secure w.r.t. PPT^{SAM} , there does not exist $O(1)$ -round computational zero-knowledge proofs for all of NP based on one-way functions.*

Following the proof of [PV10], we can show that only languages in PPT^{SAM} have $O(1)$ -round computational zero-knowledge proofs based on one-way functions. We complete the argument by noting that our assumption implies that $\text{NP} \not\subseteq \text{BPP}^{\text{SAM}}$, since otherwise, we can construct an oracle PPT machine that with oracle access to SAM inverts OWFs. We provide the formal proof in the full version.

Finally, we remark that Theorem 20 implies Theorem 19 relying on the result of Goldreich and Kahan [GK96a] and Theorem 19 implies Theorem 18 relying

on the result of Damgård, Pedersen and Pfitzmann [DPP98]. Nevertheless, the direct proofs are simpler and as mentioned before, it suffices to assume the weaker SAM₁-oracle for Theorem 18.

5 Acknowledgements

We are thankful to the anonymous TCC-reviewers for their insightful comments.

References

- [AAG⁺00] Fredrik Almgren, Gunnar Andersson, Torbjörn Granlund, Lars Ivansson, and Staffan Ulfberg. How we cracked the code book ciphers. Manuscript, 2000. http://codebook.org/codebook_solution.pdf.
- [AGGM06] Adi Akavia, Oded Goldreich, Shafi Goldwasser, and Dana Moshkovitz. On basing one-way functions on NP-hardness. In *STOC '06*, pages 701–710, 2006.
- [AR05] Dorit Aharonov and Oded Regev. Lattice problems in NP cap coNP. *J. ACM*, 52(5):749–765, 2005.
- [Bar01] Boaz Barak. How to go beyond the black-box simulation barrier. In *FOCS '01*, volume 0, pages 106–115, 2001.
- [BFKL93] Avrim Blum, Merrick Furst, Michael Kearns, and Richard Lipton. Cryptographic primitives based on hard learning problems. In *CRYPTO '93*, pages 278–291, 1993.
- [BFL91] László Babai, Lance Fortnow, and Carsten Lund. Non-deterministic exponential time has two-prover interactive protocols. *Computational Complexity*, 1:3–40, 1991.
- [BM84] Manuel Blum and Silvio Micali. How to generate cryptographically strong sequences of pseudo-random bits. *SIAM Journal on Computing*, 13(4):850–864, 1984.
- [Bra83] Gilles Brassard. Relativized cryptography. *IEEE Transactions on Information Theory*, 29(6):877–893, 1983.
- [BT03] Andrej Bogdanov and Luca Trevisan. On worst-case to average-case reductions for np problems. In *FOCS '03*, pages 308–317, 2003.
- [BY96] Mihir Bellare and Moti Yung. Certifying permutations: Noninteractive zero-knowledge based on any trapdoor permutation. *J. Cryptology*, 9(3):149–166, 1996.
- [DDN00] Danny Dolev, Cynthia Dwork, and Moni Naor. Nonmalleable cryptography. *SIAM Journal on Computing*, 30(2):391–437, 2000.
- [DPP98] Ivan Damgård, Torben P. Pedersen, and Birgit Pfitzmann. Statistical secrecy and multibit commitments. *IEEE Transactions on Information Theory*, 44(3):1143–1151, 1998.
- [DR02] Joan Daemen and Vincent Rijmen. *The Design of Rijndael*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2002.
- [FKO06] Uriel Feige, Jeong Han Kim, and Eran Ofek. Witnesses for non-satisfiability of dense random 3cnf formulas. In *FOCS '06*, pages 497–508, 2006.
- [FS90] Uriel Feige and Adi Shamir. Witness indistinguishable and witness hiding protocols. In *STOC '90*, pages 416–426, 1990.
- [GG00] Oded Goldreich and Shafi Goldwasser. On the limits of nonapproximability of lattice problems. *J. Comput. Syst. Sci.*, 60(3):540–563, 2000.

- [GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *Journal of ACM*, 33(4):792–807, 1986.
- [GK96a] Oded Goldreich and Ariel Kahan. How to construct constant-round zero-knowledge proof systems for NP. *Journal of Cryptology*, 9(3):167–190, 1996.
- [GK96b] Oded Goldreich and Hugo Krawczyk. On the composition of zero-knowledge proof systems. *SIAM Journal on Computing*, 25(1):169–192, 1996.
- [GM84] Shafi Goldwasser and Silvio Micali. Probabilistic encryption. *J. Comput. Syst. Sci.*, 28(2):270–299, 1984.
- [GMW91] Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity for all languages in np have zero-knowledge proof systems. *J. ACM*, 38(3):691–729, 1991.
- [Hås01] Johan Håstad. Some optimal inapproximability results. *J. ACM*, 48(4):798–859, 2001.
- [HHR07] 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 *FOCS*, pages 669–679, 2007.
- [HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28(4):1364–1396, 1999.
- [HMX10] Iftach Haitner, Mohammad Mahmoody, and David Xiao. A new sampling protocol and applications to basing cryptographic primitives on the hardness of NP. In *IEEE Conference on Computational Complexity*, pages 76–87, 2010.
- [HR07] Iftach Haitner and Omer Reingold. Statistically-hiding commitment from any one-way function. In *STOC '07*, pages 1–10, 2007.
- [IR88] Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In *CRYPTO '88*, pages 8–26, 1988.
- [MV05] Peter Bro Miltersen and N. V. Vinodchandran. Derandomizing arthur-merlin games using hitting sets. *Computational Complexity*, 14(3):256–279, 2005.
- [Nao91] Moni Naor. Bit commitment using pseudorandomness. *J. Cryptology*, 4(2):151–158, 1991.
- [NOVY98] Moni Naor, Rafail Ostrovsky, Ramarathnam Venkatesan, and Moti Yung. Perfect zero-knowledge arguments for p using any one-way permutation. *J. Cryptology*, 11(2):87–108, 1998.
- [NW94] Noam Nisan and Avi Wigderson. Hardness vs randomness. *J. Comput. Syst. Sci.*, 49(2):149–167, 1994.
- [Pas06] 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, 2006.
- [PV10] Rafael Pass and Muthuramakrishnan Venkatasubramanian. Private coins versus public coins in zero-knowledge proofs. To appear in TCC 2010, 2010.
- [Rab80] Michael O. Rabin. Probabilistic algorithm for testing primality. *Journal of Number Theory*, 12(1):128–138, 1980.
- [Sim98] Daniel R. Simon. Finding collisions on a one-way street: Can secure hash functions be based on general assumptions? In *EUROCRYPT '98*, pages 334–345, 1998.
- [Yao82] Andrew Chi-Chih Yao. Theory and applications of trapdoor functions (extended abstract). In *FOCS '82*, pages 80–91, 1982.