Adaptive One-way Functions and Applications

Omkant Pandey[1], Rafael Pass[2*], and Vinod Vaikuntanathan[3**]

[1] UCLA (omkant@cs.ucla.edu)
[2] Cornell University (rafael@cs.cornell.edu)
[3] MIT (vinodv@mit.edu)

Abstract. We introduce new and general complexity theoretic hardness assumptions. These assumptions abstract out concrete properties of a random oracle and are significantly stronger than traditional cryptographic hardness assumptions; however, assuming their validity we can resolve a number of long-standing open problems in cryptography.

Keywords. Cryptographic Assumptions, Non-malleable Commitment, Non-malleable Zero-knowledge

1 Introduction

The state-of-the-art in complexity theory forces cryptographers to base their schemes on unproven hardness assumptions. Such assumptions can be general (e.g., the existence of one-way functions) or specific (e.g., the hardness of RSA or the Discrete logarithm problem). Specific hardness assumptions are usually stronger than their general counterparts; however, as such assumptions consider primitives with more structure, they lend themselves to constructions of more efficient protocols, and sometimes even to the constructions of objects that are not known to exist when this extra structure is not present. Indeed, in recent years, several new and more exotic specific hardness assumptions have been introduced (e.g., [12, 4, 11]) leading to, among other things, signatures schemes with improved efficiency, but also the first provably secure construction of identity-based encryption.

In this paper, we introduce a new class of strong but *general* hardness assumptions, and show how these assumptions can be used to resolve certain long-standing open problems in cryptography. Our assumptions are all abstractions of concrete properties of a random oracle. As such, our results show that for the problems we consider, random oracles are not necessary; rather, provably secure constructions can be based on concrete hardness assumptions.

* Supported by NSF CAREER Grant No. CCF-0746990, AFOSR Award No. FA9550-08-1-0197, BSF Grant No. 2006317.
** Supported in part by NSF Grant CNS-0430450.

1.1 Adaptive Hardness Assumptions

We consider *adaptive* strengthenings of standard general hardness assumptions, such as the existence of one-way functions and pseudorandom generators. More specifically, we introduce the notion of collections of adaptive 1-1 one-way functions and collections of adaptive pseudorandom generators. Intuitively,

- A *collection of adaptively 1-1 one-way functions* is a family of 1-1 functions $\mathcal{F}_n = \{f_{\mathsf{tag}} : \{0,1\}^n \mapsto \{0,1\}^n\}$ such that for every tag, it is hard to invert $f_{\mathsf{tag}}(r)$ for a random r, even for an adversary that is granted access to an "inversion oracle" for $f_{\mathsf{tag}'}$ for every $\mathsf{tag} \neq \mathsf{tag}'$. In other words, the function f_{tag} is one-way, even with access to an oracle that invert all the other functions in the family.
- A collection of *adaptive pseudo-random generators* is a family of functions $\mathcal{G}_n = G_{\mathsf{tag}} : \{0,1\}^n \mapsto \{0,1\}^m$ such that for every tag, G_{tag} is a pseudorandom even if given access to an oracle that decides whether given y is in the range of $G_{\mathsf{tag}'}$ for $\mathsf{tag}' \neq \mathsf{tag}$.

Both the above assumptions are strong, but arguably not "unrealistically" strong. Indeed, both these assumptions are satisfied by a (sufficiently) length-extending random oracle.[4] As such, they provide concrete mathematical assumptions that can be used to instantiate random oracles in certain applications. We also present some concrete candidate instantiations of these assumptions. For the case of adaptive 1-1 one-way functions, we provide construction based on the the "adaptive security" of Factoring, or the Discrete Log problem. For the case of adaptive PRGs, we provide a candidate construction based on a generalization of the advanced encryption standard (AES).

Related Assumptions in the Literature. Assumptions of a related flavor have appeared in a number of works. The class of "one-more" assumptions introduced by Bellare, Namprempre, Pointcheval and Semanko [4] are similar in flavor. Informally, the setting of the one-more RSA-inversion problem is the following: The adversary is given values $z_1, z_2, \ldots, z_k \in \mathbb{Z}_N^*$ (for a composite $N = pq$, a product of two primes) and is given access to an oracle that computes RSA inverses. The adversary wins if the number of values that it computes an RSA inverse of, exceeds the number of calls it makes to the oracle. They prove the security of Chaum's blind-signature scheme under this assumption. This flavor of assumptions has been used in numerous other subsequent works [5, 6].

Even more closely related, Prabhakaran and Sahai [31] consider an assumption of the form that there are collision-resistant hash functions that are secure even if the adversary has access to a "collision-sampler". In a related work, Malkin, Moriarty and Yakovenko [24] assume that the discrete logarithm problem in $\mathbb{Z}_p^*$ (where p is a k-bit prime) is hard even for an adversary that has access to an oracle that computes discrete logarithms in $\mathbb{Z}_q^*$ for any k-bit prime

[4] Note that a random function over, say, $\{0,1\}^n \to \{0,1\}^{4n}$ is 1-1 except with exponentially small probability.

$q \neq p$. Both these works use the assumption to achieve secure computation in a relaxation of the universal composability framework. (In a sense, their work couples the relaxed security notion to the hardness assumption. In contrast, we use adaptive hardness assumptions to obtain protocols that satisfy the traditional notion of security.)

1.2 Our Results

Non-Interactive Concurrently Non-Malleable Commitment Schemes. Non-malleable commitment schemes were first defined and constructed in the seminal paper of Dolev, Dwork and Naor [17]. Informally, a commitment scheme is non-malleable if no adversary can, upon seeing a commitment to a value v, produce a commitment to a related value (say $v-1$). Indeed, non-malleability is crucial to applications which rely on the *independence* of the committed values. A stronger notion—called concurrent non-malleability–requires that no adversary, after receiving commitments of $v_1, \ldots, v_m$, can produce commitments to related values $\tilde{v}_1, \ldots, \tilde{v}_m$; see [28, 23] for a formal definition.

The first non-malleable commitment scheme of [17] was interactive, and required $O(\log n)$ rounds of interaction, where n is a security parameter. Barak [1] and subsequently, Pass and Rosen [29, 28] presented constant-round non-malleable commitment schemes; the protocols of [29, 28] are the most round-efficient (requiring 12 rounds) and the one of [28] is additionally concurrently non-malleable. We note that of the above commitment schemes, [17] is the only one with a black-box proof of security, whereas the schemes of [1, 29, 28] rely on the non-black-box proof technique introduced by Barak [1].[5]

Our first result is a construction of a *non-interactive, concurrently non-malleable* string commitment scheme, from a family of adaptive one-way permutations; additionally our construction only requires a black-box proof of security.

Theorem 1 (Informal). *Assume the existence of collections of adaptive 1-1 permutations. Then, there exists a non-interactive concurrently non-malleable string commitment scheme with a black-box proof of security.*

If instead assuming the existence of adaptive PRGs, we show the existence of 2-round concurrent non-malleable commitment with a black-box proof of security.

Theorem 2 (Informal). *Assume the existence of collections of adaptive PRGS. Then, there exists a 2-round concurrently non-malleable string commitment scheme with a black-box proof of security.*

[5] Subsequent to this work, Lin, Pass and Venkitasubramaniam [23] have presented constructions of concurrent non-malleable commitments using a black-box security proof, based on only one-way functions. Their construction, however, uses $O(n)$ communication rounds.

Round-optimal Black-box Non-malleable Zero-knowledge. Intuitively, a zero-knowledge proof is *non-malleable* if a man-in-the-middle adversay, receiving a proof of a statement x, will not be able to provide a proof of a statement $x' \neq x$ unless he could have done so without hearing the proof of x. Dolev, Dwork and Naor [17] defined non-malleable zero-knowledge ($\mathcal{ZK}$) and presented an $O(\log n)$-round $\mathcal{ZK}$ proof system. Barak [1] and subsequently, Pass and Rosen [29] presented constant-round non-malleable $\mathcal{ZK}$ argument system. Again, the protocol of [17] is the only one with a black-box proof of security.

We construct a 4-*round non-malleable* $\mathcal{ZK}$ *argument* system with a black-box proof of security (that is, a black-box simulator). Four rounds is known to be optimal for black-box $\mathcal{ZK}$ [20] (even if the protocol is not required to be non-malleable) and for non-malleable protocols (even if they are not required to be $\mathcal{ZK}$) [22].

Theorem 3 (Informal). *Assume the existence of collections of adaptive 1-1 one-way function. Then, there exists a* 4*-round non-malleable zero-knowledge argument system with a black-box proof of security. Assume, instead, the existence of collections of adaptive one-way permutations. Then, there exists a* 5*-round non-malleable zero-knowledge argument system with a black-box proof of security.*

It is interesting to note that the (seemingly) related notion of concurrent zero-knowledge cannot be achieved in $o(\log n)$ rounds with a black-box proof of security. Thus, our result shows that (under our new assumptions), the notion of non-malleability and concurrency in the context of $\mathcal{ZK}$ are quantitatively different.

Efficient Chosen-Ciphertext Secure Encryption. Chosen ciphertext (CCA) security was introduced in the works of [26, 32] and has since been recognized as a *sine-qua-non* for secure encryption. Dolev, Dwork and Naor [17] gave the first construction of a CCA-secure encryption scheme based on general assumptions. Their construction, and the subsequent construction of Sahai [33], uses the machinery of non-interactive zero-knowledge proofs, which renders them less efficient than one would like. In contrast, the constructions of Cramer and Shoup [15, 16] are efficent, but are based on specific number-theoretic assumptions.

Bellare and Rogaway [7] proposed an encryption scheme that is CCA-secure in the random oracle model (see below for more details about the random oracle model). We show complexity-theoretic assumptions that are sufficient to replace the random oracle in this construction. We mention that, previously, Canetti [13] showed how to replace random oracles in a related construction to get a semantically secure encryption scheme, but without CCA security. In a more recent work, Boldyreva and Fischlin [10] also show how to obtain a weakened notion of non-malleability, but still without CCA security.

Interactive Arguments for which Parallel-repetition does not reduce the soundness error. A basic question regarding interactive proofs is whether parallel repetition of such protocols reduces the soundness error. Bellare, Impagliazzo and Naor [3] show that there are interactive *arguments* (i.e., computationally-sound) proofs in the Common Reference String (CRS) model, for which parallel-repetition does not reduce the soundness error. Their construction relies on non-malleable encryption, and makes use of the CRS to select the public-key for this encryption scheme. However, if instead relying on a non-interactive concurrent non-malleable commitment scheme in their construction, we can dispense of the CRS altogether. Thus, by Theorem 1, assuming the existence of collections of adaptive 1-1 one-way functions, we show that there exists an interactive argument for which parallel repetition does not reduce the soundness error. We also mention that the same technique can be applied also to the strengthened construction of [30].

Our Techniques. Our constructions are simple and efficient. In particular, for the case of non-malleable commitment schemes, we show that appropriate instantiations of the Blum-Micali [9] or Naor [25] commitment schemes in fact are non-malleable. The proof of these schemes are also "relatively straight-forward" and follow nicely from the adaptive property of the underlying primitives.

Next, we show that by appropriately using our non-malleable commitment protocols in the Feige-Shamir [18] $\mathcal{ZK}$ argument for $\mathcal{NP}$, we can also get a round-optimal black-box non-malleable $\mathcal{ZK}$ proof for $\mathcal{NP}$. Although the construction here is straight-forward, its proof of correctness is less so. In particular, to show that our protocol is non-malleable, we rely on a techniques that are quite different from traditional proofs of non-malleability: in particular, the power of the "adaptive" oracle will only be used inside hybrid experiments; the simulation, on the other hand, will proceed by traditional rewinding. Interestingly, to get a round-optimal solution, our proof inherently relies on the actual Feige-Shamir protocol and high-lights some novel features of this protocol.

Interpreting Our Results. We offer two interpretations of our results:

- The *optimistic* interpretation: Although our assumptions are strong, they nonetheless do not (a priori) seem infeasible. Thus, if we believe that e.g., AES behaves as an adaptively secure PRG, we show efficient solutions to important open questions.
- The *conservative* interpretation: As mentioned, our constructions are black-box; namely, both the construction of the cryptographic objects and the associated security proof utilize the underlying primitive—adaptive one-way permutations or adaptive PRGs—as a black-box, and in particular, do not refer to a specific implementation of these primitives. Thus, a conservative way to view our results is that to show even black-box lower-bounds and impossibility results for non-interactive concurrent non-malleable commitments and non-malleable zero-knowledge proofs, one first needs to to refute our assumptions. Analogously, it means that breaking our CCA-secure encryptions scheme, or proving a general parallel-repetition theorem for interactive arguments, first requires refuting our assumptions.

A cryptographer could choose to make "mild" assumptions such as $\mathcal{P} \neq \mathcal{NP}$, "relatively mild" ones such as the existence of one-way functions, secure encryption schemes or trapdoor permutations, or "preposterous" ones such as "this scheme is secure". Whereas preposterous assumptions clearly are undesirable, mild assumptions are—given the state-of-the-art in complexity theory—too weak for cryptographic constructions of non-trivial tasks. Relatively mild assumptions, on the other hand, are sufficient for showing the feasibility of essentially all known cryptographic primitives.

Yet, to obtain efficient constructions, such assumptions are—given the current-state-of-art—not sufficient. In fact, it is a priori not even clear that although feasibility of a cryptographic task can be based on a relatively mild assumptions, that an "efficient" construction of the primitive is possible (at all!). One approach to overcome this gap is the random oracle paradigm, introduced in the current form by Bellare and Rogaway [7]: the proposed paradigm is to prove the security of a cryptographic scheme in the random-oracle model—where all parties have access to a truly random function—and next instantiate the random oracle with a concrete function "with appropriate properties". Nevertheless, as pointed out in [14] (see also [21, 2]) there are (pathological) schemes that can be proven secure in the random oracle model, but are rendered insecure when the random oracle is replaced by any concrete function (or family of functions).

In this work we, instead, investigate a different avenue for overcoming this gap between theory and practice, by introducing strong, but general, hardness assumption. When doing so, we, of course, need to be careful to make sure that our assumptions (although potentially "funky") are not preposterous. One criterion in determining the acceptability of a cryptographic assumption A is to consider (1) what the assumption is used for (for instance, to construct a primitive P, say) and (2) how much more "complex" the primitive P is, compared to A. For example, a construction of a pseudorandom generator assuming a one-way function is non-trivial, whereas the reverse direction is not nearly as interesting. Unfortunately, the notion of "complexity" of an assumption is hard to define. We here offer a simple interpretation: view complexity as "succinctness". General assumption are usually more succinct than specific assumptions, one-way functions are "easier" to define than, say, pseudorandom functions. Given this point of view, it seems that our assumptions are not significantly more complex than traditional hardness assumption; yet they allow us to construct considerably more complex objects (e.g., non-malleable zero-knowledge proofs).

On Falsifiability/Refutability of Our Assumptions. Note that the notions of non-malleable commitment and non-malleable zero-knowledge both are defined using simulation-based definitions. As such, simply assuming that a scheme is, say, non-malleable zero-knowledge, seems like a very strong assumption, which is hard to falsify[6]—in fact, to falsify it one needs to show (using a mathematical proof) that no Turning machine is a good simulator. In contrast, to falsify our

[6] Recall that falsifiability is Popper's classical criterion for distinguishing scientific and "pseudo-scientific" statements.

assumptions it is sufficient to exhibit an attacker (just as with the traditional cryptographic hardness assumptions).

To make such "qualitative" differences more precise, Naor [27] introduced a framework for classifying assumptions, based on how "practically" an assumption can refuted. Whereas non-malleability, a priori, seems impossible to falsify (as there a-priori is not a simple way to showing that no simulator exists). In contrast, traditional assumptions such as "factoring is hard" can be easily refuted simply by publishing challenges that a "falsifier" is required to solve. Our assumptions cannot be as easily refuted, as even if a falsifier exhibits an attack against a candidate adaptive OWF, it is unclear how to check that this attack works. However, the same can be said also for relatively mild (and commonly used) assumptions, such as "factoring is hard for subexponential-time".[7]

Additionally, we would like to argue that our assumptions enjoy a similar "win/win" situation as traditional cryptographic hardness assumptions. The adaptive security of the factoring or discrete logarithm problems seem like natural computational number theoretic questions. A refutation of our assumptions (and its implication to factoring and discrete logarithm problem) would thus be interesting in its own right. Taken to its extreme, this approach suggest that we might even consider assumptions that most probably are *false*, such as e.g., assuming that AES is an (adaptive one-way) *permutation*, as long as we believe that it might be hard to *prove* that the assumption is false.

2 New Assumptions and Definitions

The following sections introduce our definitions of adaptively secure objects—one-way functions, pseudorandom generators and commitment schemes—and posit candidate constructions for adaptively secure one-way functions and pseudorandom generators.

2.1 Adaptive One-Way Functions

In this paper, we define a *family* of adaptively secure injective one-way functions, where each function in the family is specified by an index $\mathsf{tag} \in \{0,1\}^n$. The adaptive security requirement says the following: consider an adversary that picks an index tag^* and is given $y^* = f_{\mathsf{tag}^*}(x^*)$ for a random x^* in the domain of f_{tag^*}, and the adversary is supposed to compute x^*. The adversary, in addition, has access to a "magic oracle" that on input (tag, y) where $\mathsf{tag} \neq \mathsf{tag}^*$, and get back $f^{-1}_{\mathsf{tag}}(y)$. In other words, the magic oracle helps invert all functions f_{tag} different from the "target function" f_{tag^*}. The security requirement is that the

[7] Note that the assumption that factoring is hard for subexponential-time can be falsified by considering a publishing a very "short" challenge (of length polylog$n$). However, in the same vein, our assumption can be falsified by considering challenges of length $\log n$; then it is easy to check if someone can exhibit an efficient attack on the adaptive security of an assumed one-way function, since the inverting oracle can also be efficiently implemented.

adversary have at most a negligible chance of computing x^*, even with this added ability. Note that the magic oracle is just a fictitious entity, which possibly does not have an efficient implementation (as opposed to the decryption oracle in the definition of CCA-security for encryption schemes which can be implemented efficiently given the secret-key). More formally,

Definition 1 (Family of Adaptive One-to-one One-way Functions). *A family of* injective *one-way functions* $\mathcal{F} = \{f_{\mathsf{tag}} : D_{\mathsf{tag}} \mapsto \{0,1\}^*\}_{\mathsf{tag} \in \{0,1\}^n}$ *is called adaptively secure if,*

- (EASY TO SAMPLE AND COMPUTE.) *There is an efficient randomized* domain-sampler D*, which on input* $\mathsf{tag} \in \{0,1\}^n$*, outputs a random element in* D_{tag}*. There is a deterministic polynomial algorithm* M *such that for all* $\mathsf{tag} \in \{0,1\}^n$ *and for all* $x \in D_{\mathsf{tag}}$, $M(\mathsf{tag}, x) = f_{\mathsf{tag}}(x)$.
- (ADAPTIVE ONE-WAYNESS.) *Let* $\mathcal{O}(\mathsf{tag}, \cdot, \cdot)$ *denote an oracle that, on input* tag' *and* y *outputs* $f^{-1}_{\mathsf{tag}'}(y)$ *if* $\mathsf{tag}' \neq \mathsf{tag}$, $|\mathsf{tag}'| = |\mathsf{tag}|$ *and* $\perp$ *otherwise.*
  The family $\mathcal{F}$ *is adaptively secure if, for any probabilistic polynomial-time adversary* A*, there exists a negligible function* μ *such that for all* n*, and for all tags* $\mathsf{tag} \in \{0,1\}^n$,
  $$\Pr[x \leftarrow D_{\mathsf{tag}} : A^{\mathcal{O}(\mathsf{tag},\cdot,\cdot)}(\mathsf{tag}, f_{\mathsf{tag}}(x)) = x] \leq \mu(n)$$
  where the probability is over the random choice of x *and the coin-tosses of* A*.*

A potentially incomparable assumption is that of an adaptively secure injective one-way function (as opposed to a family of functions); here the adversary gets access to an oracle that inverts the function on any y' that is different from the challenge y (that the adversay is supposed to invert). However, it is easy to see that an adaptively secure one-way function with subexponential security and a dense domain implies a family of adaptively secure one-way functions, as defined above. In fact, our construction of a family of adaptively secure one-way functions based on factoring goes through this construction.

Hardness Amplification. A strong adaptively secure one-way function is one where no adversary can invert the function with probability better than some negligible function in k (even with access to the inversion oracle). A weak one, on the other hand, only requires that the adversary not be able to invert the function with a probability better than $1 - 1/\text{poly}(k)$ (even with access to the inversion oracle).

We remark that we can construct a collection of strong adaptively secure one-way function from a collection of weak adaptively secure one-way function. The construction is the same as Yao's hardness amplification lemma. We defer the details to the full version.

Candidates We now present candidates for adaptively secure one-way functions, based on assumptions related to discrete-log and factoring.

Factoring. First, we show how to build an adaptively secure one-way function (not a family of functions) from the factoring assumption. Then, we show how to turn it into a family of functions, assuming, in addition, that factoring is subexponentially-hard.

The domain of the function f is $\{(p, q) \mid p, q \in \mathcal{P}_n, p < q\}$, where $\mathcal{P}_n$ is the set of all n-bit primes. Given this notation, $f(p, q)$ is defined to be pq. Assuming that it is hard to factor a number N that is a product of primes, even with access to an oracle that factors all other products of two primes, this function is adaptively secure.

We now show how to turn this into a family of adaptively secure one-way functions. The index is simply an $n' = n^{1/\epsilon}$-bit string (for some $\epsilon > 0$) $i = (i_1, i_2)$. The domain is the set of all strings (j_1, j_2) such that $p = i_1 \circ j_1$ and $q = i_2 \circ j_2$ are both n-bit primes. The function then outputs pq. Since we reveal the first $n' = n^{1/\epsilon}$ bits of the factors of $N = pq$, we need to assume that factoring is subexponentially hard (even with access to an oracle that factors other products of two primes). The function is clearly injective since factoring forms an injective function. In the full version, we additionally provide candidates for adaptive one-way functions based on the RSA and Rabin functions.

Discrete Logarithms. The family of adaptive OWFs $\mathcal{F}_{DL}$ is defined as follows: The domain of the function is a tuple (p, g, x) such that p is a $2n$-bit prime p whose first n bits equal the index i, g is a generator for $\mathbb{Z}_p^*$ and x is a $2n-1$-bit number. The domain is easy to sample–the sampler picks a "long-enough" random string r and a $2n-1$-bit number x. The function f_i uses r to sample a $2n$-bit prime p whose first n bits equal i (this can be done by repeated sampling, and runs in polynomial time assuming a uniformness conjecture on the density of primes in large intervals) and a generator $g \in \mathbb{Z}_p^*$. The output of the function on input (p, g, x) is $(p, g, g^x \bmod p)$. f_i is injective since the output determines p and g; given p and g, $g^x \bmod p$ next determines x uniquely since $x < 2^{2n-1}$ and p, being a $2n$-bit prime, is larger than 2^{2n-1}.

We also mention that the adaptive security of this family can be based on the subexponential adaptive security of the one-way function (as opposed to family) obtained by simply sampling random p, g, x (or even random p being a safe prime) and outputting p, g, g^x. (In the full version of the paper, we additionally show how to obtain our results under a different variant of *polynomial-time* adaptive hardness of the above one-way function; roughly speaking, the variant we require here is that the adversary gets access to an oracle that inverts the function on any input length.)

2.2 Adaptive Pseudorandom Generator

A family of adaptively secure pseudorandom generators $\mathcal{G} = \{G_{\mathsf{tag}}\}_{\mathsf{tag} \in \{0,1\}^*}$ is defined in a similar way to an adaptive one-way function. We require that the output of the generator G, on a random input x and an adversarially chosen tag be indistinguishable from uniform, even for an adversary that can query a magic oracle with a value (tag', y) (where $\mathsf{tag}' \neq \mathsf{tag}$) and get back 0 or 1 depending on whether y is in the range of $G_{\mathsf{tag}'}$ or not.

Definition 2 (Adaptive PRG). *A family of functions* $\mathcal{G} = \{G_{\mathsf{tag}} : \{0,1\}^n \mapsto \{0,1\}^{s(n)}\}_{\mathsf{tag}\in\{0,1\}^n}$ *is an adaptively secure pseudorandom generator (PRG) if* $|G_{\mathsf{tag}}(x)| = s(|x|)$ *for some function* s *such that* $s(n) \geq n$ *for all* n *and,*

– *(*EFFICIENT COMPUTABILITY.*) There is a deterministic polynomial-time algorithm* M_G *such that* $M_G(x, \mathsf{tag}) = G_{\mathsf{tag}}(x)$*.*
– *(*ADAPTIVE PSEUDORANDOMNESS.*) Let* $\mathcal{O}(\mathsf{tag}, \cdot, \cdot)$ *denote an oracle that, on input* (tag', y) *such that* $\mathsf{tag}' \neq \mathsf{tag}$*,* $|\mathsf{tag}'| = |\mathsf{tag}|$*, outputs* 1 *if* y *is in the range of* $G_{\mathsf{tag}'}$ *and* 0 *otherwise.*
The PRG G *is adaptively secure if, for any probabilistic polynomial-time adversary* A*, there exists a negligible function* μ *such that for all* n *and for all tags* $\mathsf{tag} \in \{0,1\}^n$*,*

$$\left| \Pr[y \leftarrow G_{\mathsf{tag}}(U_n) : A^{\mathcal{O}(\mathsf{tag},\cdot,\cdot)}(y) = 1] - \Pr[y \leftarrow U_m : A^{\mathcal{O}(\mathsf{tag},\cdot,\cdot)}(y) = 1] \right| \leq \mu(n)$$

where the probability is over the random choice of y *and the coin-tosses of* A*.*

Candidates For the case of adaptive PRGs, we provide a candidate construction based on the advanced encryption standard (AES). AES is a permutation on 128 bits; that is, for a 128-bit seed s, AES_s is a permutation defined on $\{0,1\}^{128}$. However, due to the algebraic nature of the construction of AES, it can easily be generalized to longer input length. Let AES_n denote this generalized version of AES to n-bit inputs. Our candidate adaptive pseudorandom generator AESG_{tag} is simply $\mathsf{AESG}_{tag}(s) = \mathsf{AES}_s(tag \circ 0) \circ \mathsf{AES}_s(tag \circ 1)$.

2.3 Adaptively Secure Commitment Schemes

In this subsection, we define adaptively secure commitment schemes. Let $\{\textsc{Com}_{\mathsf{tag}} = \langle S_{\mathsf{tag}}, R_{\mathsf{tag}} \rangle\}_{\mathsf{tag}\in\{0,1\}^*}$ denote a family of commitment protocols, indexed by a string tag. We require that the commitment scheme be secure, even against an adversary that can query a magic oracle on the transcript of a commitment interaction and get back a message that was committed to in the transcript. More precisely, the adversary picks an index tag and two equal-length strings x_0 and x_1 and gets a value $y_b = \textsc{Com}_{\mathsf{tag}}(x_b; r)$, where b is a random bit and r is random. The adversary can, in addition, query a magic oracle on (y', tag') where $\mathsf{tag}' \neq \mathsf{tag}$ and get back the some x' such that $y' \in \textsc{Com}_{\mathsf{tag}'}(x'; r')$ (if y' is a legal commitment) and $\perp$ otherwise. [8] The security requirement is that the adversary cannot distinguish whether y_b was a commitment to x_0 or x_1, even with this extra power.

[8] In case the transcript corresponds to the commitment of multiple messages, the oracle returns a canonical one of them. In fact, one of our commitment schemes is perfectly binding and thus, does not encounter this problem.

Definition 3 (Adaptively-Secure Commitment). *A family of functions* $\{\textsc{Com}_{\mathsf{tag}}\}_{\mathsf{tag}\in\{0,1\}^*}$ *is called an adaptively secure commitment scheme if* S_{tag} *and* R_{tag} *are polynomial-time and*

- Statistical Binding: *For any* tag*, over the coin-tosses of the receiver* R*, the probability that a transcript* $\langle S^*, R_{\mathsf{tag}}\rangle$ *has two valid openings is negligible.*
- Adaptive Security: *Let* $\mathcal{O}(\mathsf{tag},\cdot,\cdot)$ *denote the oracle that, on input* $\mathsf{tag}' \neq \mathsf{tag}$, $|\mathsf{tag}'| = |\mathsf{tag}|$ *and* c*, returns an* $x \in \{0,1\}^{\ell(n)}$ *if there exists strings* r_S *and* r_R*, such that* c *is the transcript of the interaction between* S *with input* x *and random coins* r_S *and* R *with random coins* r_R*, and* $\perp$ *otherwise.*
  For any probabilistic polynomial-time oracle TM A*, there exists a negligible function* $\mu(\cdot)$ *such that for all* n*, for all* $\mathsf{tag} \in \{0,1\}^n$ *and for all* $x, y \in \{0,1\}^{\ell(n)}$*,*

$$\Big| \Pr[c \leftarrow \langle S_{\mathsf{tag}}(x), R_{\mathsf{tag}}\rangle; A^{\mathcal{O}(\mathsf{tag},\cdot,\cdot)}(c,\mathsf{tag}) = 1] - \Pr[c \leftarrow \langle S_{\mathsf{tag}}(y), R_{\mathsf{tag}}\rangle; A^{\mathcal{O}(\mathsf{tag},\cdot)}(c,\mathsf{tag}) = 1] \Big| \leq \mu(n)$$

3 Non-Malleable Commitment Schemes

In this section, we construct non-malleable string-commitment schemes. We first construct adaptively-secure bit-commitment schemes based on an adaptively secure injective OWF and an adaptively secure PRG—the first of these constructions is non-interactive and the second is a 2-round commitment scheme. We then show a simple "concatenation lemma", that constructs an adaptively secure string commitment scheme from an adaptively-secure bit-commitment scheme. Finally, we show that an adaptively secure commitment scheme are also concurrently non-malleable. The complete proofs are deferred to the full version.

Lemma 1. *Assume that there exists a family of adaptively secure injective one-way functions. Then, there exists an adaptively secure bit-commitment scheme. Furthermore, the commitment scheme is non-interactive.*

Further, assuming the existence of a family of adaptively secure pseudorandom generators, there exists a 2*-round adaptively secure bit-commitment scheme.*

The first of these constructions follows by replacing the injective one-way function in the Blum-Micali [9] commitment scheme, with an adaptively secure one, and the second follows from the Naor commitment scheme [25] in an analogous way.

Lemma 2 (Concatenation Lemma). *If there is an adaptively secure family of bit-commitment schemes, then there is an adaptively secure family of string-commitment schemes.*

The concatenation lemma follows by simply committing to each bit of the message independently using a single-bit commitment scheme $\textsc{Com}_{\mathsf{tag}}$.

Finally, in the full version we show that any adaptively secure commitment scheme is concurrenly non-malleable according to the definition of [23]. The proof is essentially identical to the proof of [17] that any CCA-secure encryption scheme is also non-malleable.

Lemma 3. *If* $\{\textsc{Com}_{\mathsf{tag}}\}_{\mathsf{tag}\in\{0,1\}^n}$ *is a tag-based adaptively secure commitment scheme, then it is also concurrently non-malleable.*

4 Four-Round Non-Malleable Zero-Knowledge

In this section, we present a 4-round non-malleable zero-knowledge argument system. We start by reviewing the notion of non-malleable zero-knowledge [17] and refer the reader to [29] for a formal definition of the notion we consider in this work.

Non-malleable $\mathcal{ZK}$ proofs: An informal definition. Let Π_{tag} be a tag-based family of $\mathcal{ZK}$ proofs. Consider a man-in-the-middle adversary that participates in two interactions: in the left interaction the adversary A is verifying the validity of a statement x by interacting with an honest prover P using tag. In the right interaction A proves the validity of a statement x' to the honest verifier V using $\mathsf{tag}' \neq \mathsf{tag}$. The objective of the adversary is to convince the verifier in the right interaction. Π_{tag} is, roughly speaking, non-malleable, if for any man-in-the-middle adversary A, there exists a stand-alone prover S that manages to convince the verifier with essentially the same probability as A (without receiving a proof on the left).

Our protocol. The argument system is the Feige-Shamir protocol [18], compiled with an adaptively secure commitment scheme. In our analysis we rely on the following properties of the Feige-Shamir protocol:

- The first prover message is (perfectly) independent of the witness used by the prover (and even the statement). This property has previously been used to simplify analysis, but here we inherently rely on this property to *enable* our analysis.
- Given a random accepting transcript, and the *openings* of the commitments in the first message, it is possible to "extract a witness". In other words, any transcript implicitly defines a witness; additionally, given a random transcript, this witness will be valid with a high probability (if the transcript is accepting).

In what follows, we present a sketch of the protocol and the proof. The complete proof is deferred to the full version.

4.1 An Adaptively Secure WI Proof of Knowledge

The main component in the NMZK protocol is a three-round witness-indistinguishable (WI) proof of knowledge (POK); see [19] for a definition of witness indistinguishability and proof of knowledge. The protocol is simply a parallelization of the 3-round $\mathcal{ZK}$ proof $\tilde{\Pi}$ for the $\mathcal{NP}$-complete language of Hamiltonicity [8, 18], with the only change that the commitment scheme used in the proof is adaptively secure. Let Π_{tag} denote this family of protocols; it is a family which is parameterized by the tag of the adaptively secure commitment.

We show that this family of protocols satisfy two properties:

- it has an "adaptive WI" property which, roughly stated, means that the transcripts of the protocol when the prover uses two different witnesses w_1 and w_2 are computationally indistinguishable, even if the distinguisher has access to a magic oracle that inverts all commitments $\text{Com}_{\mathsf{tag}'}$, where $\mathsf{tag}' \neq \mathsf{tag}$.
- a random transcript of $\tilde{\Pi}_{\mathsf{tag}}$ uniquely defines a witness (even though not it is not computable in polynomial-time). We define this to be the *witness implicit in the transcript* in an instance of Π_{tag}. Furthermore, we show that the implicit witness in Π_{tag} is computable given access to $\mathcal{O}(\mathsf{tag}', \cdot, \cdot)$ for any $\mathsf{tag}' \neq \mathsf{tag}$.

4.2 The Non-Malleable Zero-Knowledge Argument System

The non-malleable ZK protocol consists of two instances of the protocol Π_{tag} running in conjunction, one of them initiated by the verifier and the other initiated by the prover. We will denote the copy of Π_{tag} initiated by the verifier as Π^V_{tag} and the one initiated by the prover as Π^P_{tag}.

Recall that Π_{tag} is a parallelized version of a 3-round protocol $\tilde{\Pi}_{\mathsf{tag}}$; let A_i, C_i and Z_i denote the messages in the i'th repetion in these three rounds. In the description of the protocol, we let messages in the protocol Π^V_{tag} (resp. Π^P_{tag}) appear with a superscript of V (resp. P).

Theorem 4. *Assume that* Com *is a non-interactive adaptively secure commitment scheme. Then, the protocol in Figure 1 is a* 4*-round non-malleable zero-knowledge argument system.*

Proof (Sketch). Completeness, soundness and zero-knowledge properties of the protocol follow directly from the corresponding properties of the Feige-Shamir protocol. In Lemma 4, we show that the protocol non-malleable.

In other words, for every man-in-the-middle adversary A that interacts with the prover P_{tag} on a statement x and convinces the verifier $V_{\mathsf{tag}'}$ (for a $\mathsf{tag}' \neq \mathsf{tag}$) in a right-interaction on a statement x' (possibly the same as x), we construct a stand-alone prover that convinces the verifier on x' with the same probability as A, but *without access to the left-interaction.* The construction of the stand-alone prover in the proof of non-malleability (see Lemma 4) relies on the adaptive security of the commitment scheme $\text{Com}_{\mathsf{tag}}$. It is important to note that the stand-alone prover itself runs in classical polynomial-time, and in particular does not use any oracles. Access to the commitment-inversion oracle is used only to show that the stand-alone prover works as expected (and in particular, that it convinces the verifier with the same probability as does the MIM adversary).

Lemma 4. *The protocol* $\mathrm{NM}_{\mathsf{tag}}$ *in Figure 1 is non-malleable.*

Proof (Sketch). For every man-in-the-middle adversary A, we construct a stand-alone prover S: the construction of the stand-alone prover S proceeds in three steps.

Non-Malleable Zero-Knowledge Argument $\mathrm{NM}_{\mathsf{tag}}$

COMMON INPUT: An instance $x \in \{0,1\}^n$, presumably in the language L.
PROVER INPUT: A witness w such that $(x, w) \in R_L$.
ROUND 1: **(Verifier)** Pick w_1 and w_2 at random and compute $x_i = f(w_i)$ for $i \in \{1,2\}$.
Let the $\mathcal{NP}$-relation $R_V = \{((x_1, x_2), w') \mid$ either $f(w') = x_1$ or $f(w') = x_2\}$.
Initiate the WI protocol Π^V_{tag} with the statement $(x_1, x_2) \in L_V$. In particular,
$\mathbf{V} \to \mathbf{P}$: Send (x_1, x_2) to P. Send $A^V_1, A^V_2, \ldots, A^V_n$ to P.
ROUND 2: **(Prover)** Let the $\mathcal{NP}$-relation R_P be

$$\{((x, x_1, x_2), w) \mid \text{either } (x, w) \in R_L \text{ or } f(w) = x_1 \text{ or } f(w) = x_2\}$$

Initiate a WI protocol Π^P_{tag} with common input (x, x_1, x_2). Also, send the second-round messages of the protocol Π^V_{tag}. In particular,
(2a) $\mathbf{P} \to \mathbf{V}$: Send $A^P_1, A^P_2, \ldots, A^P_n$ to V.
(2b) $\mathbf{P} \to \mathbf{V}$: Send $C^V_1, C^V_2, \ldots, C^V_n$ to V.
ROUND 3: **(Verifier)** Send round-2 challenges of the protocol Π^P_{tag} and round-3 responses of Π^V_{tag}.
(3a) $\mathbf{V} \to \mathbf{P}$: Send $C^P_1, \ldots, C^P_n$ to P.
(3b) $\mathbf{V} \to \mathbf{P}$: Send $Z^V_1, \ldots, Z^V_n$ to P.
ROUND 4: **(Prover)** P verifies that the transcript $\{(A^V_i, C^V_i, Z^V_i)\}_{i \in [n]}$ is accepting for the subprotocol Π^V_{tag}. If not, abort and send nothing to V. Else,
$\mathbf{P} \to \mathbf{V}$: Send $Z^P_1, \ldots, Z^P_n$ to V.
V accepts iff the transcript $\{(A^P_i, C^P_i, Z^P_i)\}_{i \in [n]}$ is accepting for the subprotocol Π^P_{tag}.

Fig. 1. NON-MALLEABLE ZERO-KNOWLEDGE PROTOCOL $\mathrm{NM}_{\mathsf{tag}}$ FOR A LANGUAGE L

1. Run the adversary A with "honestly generated" verifier-messages on the right interaction, and extract the witness for the WIPOK Π^V_{tag} that the adversary initiates on the left interaction.
2. Use the witness thus obtained to simulate the left-interaction of the adversary A and rewind the WI proof of knowledge $\Pi^P_{\mathsf{tag}'}$ it initiates on the right interaction to extract the witness w' for the statement x'.
3. Finally provide an honest proof to the outside verifier of the statement x' using the tag tag' and witness w'.

Carrying out this agenda involves a number of difficulties. We first describe how to accomplish Step 1. This is done by invoking the simulator for the Feige-Shamir protocol, and is described below. Informally, S extracts the witness w' that the MIM A uses in the subprotocol Π^V_{tag} in the left-interaction. Then, S acts as the honest prover using the witness w' in the protocol Π^P_{tag}.

We now describe how to carry out Step 2 of the agenda, and show that at the end of Step 2, S extracts a witness for the statement x' that the MIM adversary A uses in the right-interaction with essentially the same probability that A convinces the verifier on the right-interaction. S starts by running the protocol in the left-interaction using the witness w' it extracted using the strategy in Step 1. Consider the moment when A outputs the first message on the left (that is, the first message in the subprotocol Π^V_{tag}). Consider two cases.

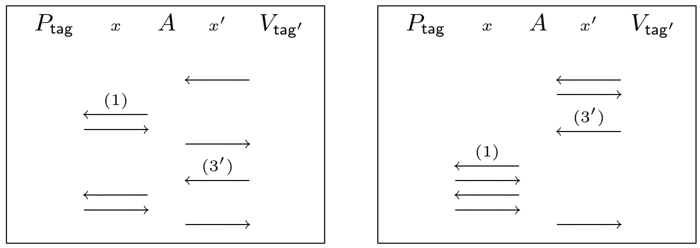

Fig. 2. Two scheduling strategies (i) on the left and (ii) on the right

Case One: In the first case, A has not yet received the round-3 messages in the right interaction (that is, the challenges in the subprotocol $\Pi^P_{\mathsf{tag}'}$) (See Figure 2(i)). In this case, the Round-1 message that A sends on the left interaction is independent of the Round-3 message in the right interaction. Now, S proceeds as follows: S runs the left-interaction as a normal prover P_{tag} would with the fake-witness w', and rewinds the protocol $\Pi^P_{\mathsf{tag}'}$ on the right-interaction to extract a witness for the statement x'. Since the rewinding process does not change the messages in the right-interaction before round 3, S can use w' to produce the left-interaction just as an honest prover with witness w' would; note that we here rely on the property of the Feige-Shamir protocol that the first message sent by the prover (i.e., round 2) is independent of the statement and witness used.

Case Two: In the second case, A has already received the challenges in the subprotocol $\Pi^P_{\mathsf{tag}'}$ in the right interaction (See Figure 2(ii)). In this case, trying to rewind in the WIPOK $\Pi^P_{\mathsf{tag}'}$ on the right is problematic, since A could change the first message on the left, every time it is fed with a different challenge in round-3 on the right-interaction. In this case, S proceeds as follows: Every time the extractor for the WIPOK $\Pi^P_{\mathsf{tag}'}$ in the right-interaction rewinds, S repeats the entire procedure in Step 1 of the agenda to extract a witness w' corresponding to the (potentially new) Round-1 message in the left interaction. S then simulates the left-interaction with the witness thus extracted. Note that due to the particular scheduling, the extraction procedure on the right-interaction is unaffected by the rewinding on the left.

To analyze the correctness of the above simulator, we first show that the view generated by S following Step 1 of the agenda is indistinguishable from the view of A in a real interaction, even to a distinguisher that has access to the oracle $\mathcal{O}(\mathsf{tag}, \cdot, \cdot)$ that inverts $\textsc{Com}_{\mathsf{tag}'}$ for any $\mathsf{tag}' \neq \mathsf{tag}$. Then, we use this to show that the *implicit witness* in the transcript of the subprotocol $\Pi^P_{\mathsf{tag}'}$ in the right-interaction is indistinguishable between the simulated and the real execution. This means that the witness that S extracts from the right interaction of A is computationally indistinguishable from the witness that A uses in the real interaction. We defer an analysis of the running-time to the full version; intuitively it follows that the running-time in expectation is polynomial since when performing rewinding on the right, we can *perfectly* emulate the messages on the left with the same distribution as when generating the initial view in Stage 1.

5 CCA2-Secure Encryption Scheme

Bellare and Rogaway [7] showed how to construct an efficient encryption scheme that is CCA2-secure in the random oracle model, starting from any trapdoor permutation. We show that the same scheme is CCA2-secure in the standard model (that is, without assuming random oracles) by instantiation their scheme with adaptively secure primitives.

To prove security of the construction, we assume an adaptively secure variant of perfectly one-way hash functions (defined by Canetti [13]), and a family of trapdoor permutations that is hard to invert even with access to an oracle that inverts the perfectly one-way hash function. We note that Canetti [13] (define and) use perfectly one-way hashing with auxiliary input to prove *IND-CPA* security (semantic security) of the [7] construction.

We sketch the notion of adaptively secure perfectly one-way hashing w.r.t auxiliary information and give a high-level intuition of the security proof; the complete definition and proof is deferred to the full version. Consider a family of functions $\mathcal{H}$ such that for a random function $H \leftarrow \mathcal{H}$, it is computationally infeasible to distinguish between $h \leftarrow H(r; s)$ (for random r, s) and a random value, even if the adversary is given (1) $g(r)$, where g is an uninvertible function evaluated on the input r, and (2) access to an oracle that inverts every $h' \neq h$ (namely, the oracle, given any $h' \neq h$, computes (r', s') such that $h' = H(r'; s')$).

Theorem 5. *Let* TDPGen *be a family of trapdoor permutations that are uninvertible with access to the $\mathcal{H}$-inverting oracle, and let $\mathcal{H}$ be an adaptively secure perfectly one-way hash family with auxiliary information. Then, the scheme in Figure 3 is an IND-CCA2-secure encryption scheme.*

Proof (Idea). The proof is analogous to that for the two-key paradigm of Naor and Yung [26]. The main idea of the proof is that there are two ways to decrypt a ciphertext – the first is using the trapdoor f^{-1} (as the legal decryption algorithm Dec does), and the second is using an oracle that inverts H. Given a ciphertext $c = (c_0, c_1, c_2, s_1, s_2)$ and access to such an oracle, we first compute r' such that $H((r', s_1, c_1); s_2) = c_2$, and check that $c_0 = f(r')$. If the check passes, output $m = c_1 \oplus H(r'; s_1)$, otherwise output $\perp$. This allows the simulator to answer the decryption queries of the adversary, given access to an oracle that inverts H. Even with access to such an oracle, the adversary can neither (1) invert $f(r)$ on a new value r (since f is uninvertible even with access to the H-inverting

$\mathsf{Gen}(1^n)$: Run $\mathsf{TDPGen}(1^n)$ and get a pair (f, f^{-1}). Run $\mathsf{PHGen}(1^k)$ to get a perfectly one-way hash function H. Let $\text{PK} = (f, H)$ and $\text{SK} = f^{-1}$.

$\mathsf{Enc}(\text{PK}, m)$:

1. Pick random $r \leftarrow \{0,1\}^n$. Compute $c_0 = f(r)$ and $c_1 = m \oplus H(r; s_1)$ for random s_1.
2. Let $c' = (r, s_1, c_1)$. Compute $c_2 = H(c'; s_2)$ for random s_2.

Output the ciphertext $\mathbf{c} = (c_0, c_1, c_2, s_1, s_2)$.

$\mathsf{Dec}(\text{SK}, c)$: Parse $\mathbf{c}$ as $(c_0, c_1, c_2, s_1, s_2)$.

1. Compute $r' = f^{-1}(c_0)$, and $m' = c_1 \oplus H(r'; s_1)$.
2. Let $c' = (r', s_1, c_1)$. Output m' if $H(c'; s_2) = c_2$. Otherwise output $\perp$.

Fig. 3. An IND-CCA2-secure encryption scheme.

oracle), nor (2) distinguish $H(r; \cdot)$ from random (since $\mathcal{H}$ is an adaptively secure perfectly one-way hash family). Thus, even with access to the decryption oracle, the scheme is semantically secure; that is to say that the scheme itself is IND-CCA2-secure.

Acknowledgements

We are very grateful to Yuval Ishai for illuminating discussions. First author thanks Vipul Goyal, Ivan Visconti, and Darrel Carbajal for their ideas/comments.

References

1. Boaz Barak. Constant-round coin-tossing with a man in the middle or realizing the shared random string model. In *FOCS*, pages 345–355, 2002.
2. Mihir Bellare, Alexandra Boldyreva, and Adriana Palacio. An uninstantiable random-oracle-model scheme for a hybrid-encryption problem. In *EUROCRYPT*, pages 171–188, 2004.
3. Mihir Bellare, Russell Impagliazzo, and Moni Naor. Does parallel repetition lower the error in computationally sound protocols? In *FOCS*, pages 374–383, 1997.
4. Mihir Bellare, Chanathip Namprempre, David Pointcheval, and Michael Semanko. The power of rsa inversion oracles and the security of chaum's rsa-based blind signature scheme. In *Financial Cryptography*, pages 319–338, 2001.
5. Mihir Bellare and Gregory Neven. Transitive signatures based on factoring and rsa. In *ASIACRYPT*, pages 397–414, 2002.
6. Mihir Bellare and Adriana Palacio. Gq and schnorr identification schemes: Proofs of security against impersonation under active and concurrent attacks. In *CRYPTO*, pages 162–177, 2002.
7. Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *First ACM Conference on Computer and Communications Security*, pages 62–73, Fairfax, 1993. ACM.
8. Manuel Blum. How to prove a theorem so no one can claim it. In *Proc. of The International Congress of Mathematicians*, pages 1444–1451, 1986.
9. Manuel Blum and Silvio Micali. How to generate cryptographically strong sequences of pseudo random bits. In *FOCS*, pages 112–117, 1982.
10. Alexandra Boldyreva and Marc Fischlin. On the security of oaep. In *ASIACRYPT*, pages 210–225, 2006.

11. Dan Boneh. The decision diffie-hellman problem. In *ANTS*, pages 48–63, 1998.
12. Dan Boneh and Matthew K. Franklin. Identity-based encryption from the weil pairing. In *CRYPTO*, pages 213–229, 2001.
13. Ran Canetti. Towards realizing random oracles: Hash functions that hide all partial information". In Burt Kaliski, editor, *Proceedings CRYPTO '97*, pages 455–469. Springer-Verlag, 1997. Lectures Notes in Computer Science No. 1294.
14. Ran Canetti, Oded Goldreich, and Shai Halevi. The random oracle methodology, revisited. *J. ACM*, 51(4):557–594, 2004.
15. Ronald Cramer and Victor Shoup. A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In *CRYPTO*, pages 13–25, 1998.
16. Ronald Cramer and Victor Shoup. Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In *EUROCRYPT*, pages 45–64, 2002.
17. Danny Dolev, Cynthia Dwork, and Moni Naor. Nonmalleable cryptography. *SIAM J. Comput.*, 30(2):391–437, 2000.
18. Uriel Feige and Adi Shamir. Witness indistinguishable and witness hiding protocols. In *STOC*, pages 416–426, 1990.
19. Oded Goldreich. *Foundations of Cryptography: Basic Tools.* Cambridge University Press, 2001. Earlier version available on http://www.wisdom.weizmann.ac.il/ oded/frag.html .
20. Oded Goldreich and Hugo Krawczyk. On the composition of zero-knowledge proof systems. *SIAM J. Comput.*, 25(1):169–192, 1996.
21. Shafi Goldwasser and Yael Tauman Kalai. On the (in)security of the fiat-shamir paradigm. In *FOCS*, pages 102–, 2003.
22. Jonathan Katz and Hoeteck Wee. Black-box lower bounds for non-malleable protocols. 2007.
23. Huijia Lin, Rafael Pass, and Muthuramakrishnan Venkitasubramaniam. Concurrent non-malleable commitments from any one-way function. In *TCC*, pages 571–588, 2008.
24. Tal Malkin, Ryan Moriarty, and Nikolai Yakovenko. Generalized environmental security from number theoretic assumptions. In *TCC*, pages 343–359, 2006.
25. Naor. Bit commitment using pseudorandomness. *J. of Cryptology*, 4, 1991.
26. M. Naor and M. Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *STOC '90: Proceedings of the twenty-second annual ACM symposium on Theory of computing*, pages 427–437, New York, NY, USA, 1990. ACM Press.
27. Moni Naor. On cryptographic assumptions and challenges. In *CRYPTO*, pages 96–109, 2003.
28. Rafael Pass and Alon Rosen. Concurrent non-malleable commitments. In *FOCS*, pages 563–572, 2005.
29. Rafael Pass and Alon Rosen. New and improved constructions of non-malleable cryptographic protocols. In *STOC*, pages 533–542, 2005.
30. Krzysztof Pietrzak and Douglas Wikström. Parallel repetition of computationally sound protocols revisited. In *TCC*, pages 86–102, 2007.
31. Manoj Prabhakaran and Amit Sahai. New notions of security: achieving universal composability without trusted setup. In *STOC*, pages 242–251, 2004.
32. Charles Rackoff and Daniel R. Simon. Cryptographic defense against traffic analysis. In *STOC '93: Proceedings of the twenty-fifth annual ACM symposium on Theory of computing*, pages 672–681, New York, NY, USA, 1993. ACM Press.
33. Amit Sahai. Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In *FOCS*, pages 543–553, 1999.