2013 IEEE 54th Annual Symposium on Foundations of Computer Science

From Unprovability to Environmentally Friendly Protocols

Ran Canetti
Boston University and
Tel Aviv University
Canetti@tau.ac.il

Huijia Lin
University of California, Santa Barbara
huijial@gmail.com

Rafael Pass
Cornell University
rafael@cs.cornell.edu

***Abstract*—An important security concern for cryptographic protocols is the extent to which they adversely affect the security of the systems in which they run. In particular, can we rule out the possibility that introducing a new protocol to a system might, as a "side effect", break the security of unsuspecting protocols in that system?**

Universally Composable (UC) security rules out such adverse side effects. However, many functionalities of interest provably cannot be realized with UC security unless the protocol participants are willing to put some trust in external computational entities.

We propose a notion of security that: (a) allows realizing practically any functionality by protocols in the plain model without putting trust in any external entity; (b) guarantees that secure protocols according to this notion have no adverse side-effects on existing protocols in the system — as long as the security of these existing protocols is proven via the traditional methodology of black box reduction to a game-based cryptographic hardness assumption with bounded number of rounds.

Our security notion builds on the angel-based security notion of Prabhakaran and Sahai. A key part in our analysis is to come up with a CCA-secure commitment scheme that (a) cannot be proven secure via a *black box reduction* to a game-based assumption, but (b) can be proven secure using a non-black-box reduction. To the best of our knowledge, this is the first time that the interplay between black-box provability and unprovability is used to demonstrate security properties of protocols.

I. Introduction

Traditionally, security of cryptographic protocols has been conceived by way of asserting properties that relate to the execution of the protocol itself, potentially within an adversarial execution environment. This approach is reflected in the basic notions of security in the literature, including [25], [6], [37], [19], [8], [24], [36], and even in notions of concurrent security such as [22], [39], [5], [26].

In contrast, when designing protocols for modern information systems, we would like to design the protocols in a way that allows asserting *overall security properties* of the information system within which these protocols are executed. This goal is made more challenging by the fact that current information systems are mostly an unstructured conglomerate of many loosely coordinated and dynamically changing components and protocols. Still, one wishes to identify those security properties of individual protocols that allow making global security guarantees, even in such unpredictable and complex systems.

The framework of universally composable (UC) security, with its associated universal composition theorem [9], [43], provide a general methodology for addressing that question. Indeed, within this framework, security properties of protocols are explicitly defined by way of the protocol's effect on the system it runs in (aka its "environment"). More specifically, one first formulates an "idealized protocol" whose interaction with its environment reflects the desired effect on the environment (in terms of the functionality and security properties), and then requires that the analyzed protocol ρ realizes the ideal protocol $\mathcal{G}$, in the sense that ρ has essentially the same effect on its environment as $\mathcal{G}$. The universal composition theorem essentially guarantees that if (a single instance of) the analyzed protocol ρ realizes the idealized protocol $\mathcal{G}$, then any system Π that potentially invokes (multiple concurrent instances of) $\mathcal{G}$ as part of its execution will continue to behave essentially the same when all instances of $\mathcal{G}$ are replaced by instances of ρ.

Protocols for realizing practically any idealized protocol (or, tasks) within the UC framework are known, e.g. [16], [30], [3], [10], [31], [34]. Indeed, these protocols can be used as sketched above to build systems with overall security properties. However, these protocols use in an essential way some global trust mechanisms. That is, the parties running these protocols need to put trust in the correct and non-malicious behavior of some external computational entities or network components. In fact

Ran Canetti is Supported by the Check Point Institute for Information Security, an ISF grant and NSF award No. 1218461.

Huijia Lin is Supported by DARPA grant FA8750-11-2-0225 and NSF grant CCF-1018064.

Pass is supported in part by a Alfred P. Sloan Fellowship, Microsoft New Faculty Fellowship, NSF Award CNS-1217821, NSF CAREER Award CCF-0746990, NSF Award CCF-1214844, AFOSR YIP Award FA9550-10-1-0093, and DARPA and AFRL under contract FA8750-11-2- 0211. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the US Government.

0272-5428/13 $26.00 © 2013 IEEE
DOI 10.1109/FOCS.2013.16

IEEE computer society

we know that many tasks of interest are impossible to realize in a UC manner without such underlying trust; impossibility holds even if ideally authenticated communication is guaranteed [11], [13], [44].

A natural and important question that arises from this fundamental impossibility is how to go about building secure systems that use cryptography *without making such trust assumptions.* One way to do that would be to resort to the traditional approach of treating the entire system as a single protocol and asserting its security using a notion such as one of those mentioned above. However, as argued, this approach is not helpful in the prevalent case where some of the components are unknown or dynamically changing.

To further exemplify this point, assume we have a system that includes, say, some encryption protocol Π that satisfies security property P. (Say, P is semantic security.) We add to the system some multi-party auction protocol ρ that was proven to satisfy property Q. (This protocol may be used by completely different parties than those that use the encryption protocol, and the protocols may have been designed without being aware of each other.) Can we be assured that the encryption protocol Π continues to satisfy property P and at the same time ρ continues to satisfy property Q? Note that here Π and ρ may run on inputs that are potentially adversely correlated. If the answer is positive for any ρ that has property Q, then we say that property Q is environmentally friendly to protocol Π and property P.

Now, if Q means "realizing some functionality (or, specification protocol) $\mathcal{G}$ with UC security", then the answer to the above question is positive for *any* property P and protocol Π, as long as Π is known to have property P in an execution environment that includes protocol $\mathcal{G}$. Similarly, if P were to mean "realizing some functionality with UC security" the answer would again be positive for any notion Q. In other words, UC security is environmentally friendly with respect to any protocol Π and security notion P, and any security notion Q is environmentally friendly with respect to UC security properties of Π. However, can we have environmental friendliness guarantees when both P and Q are more relaxed notions, and in particular are realizable without external trust assumptions?

THIS WORK. We make two main contributions: First, we motivate and formalize the concept of "environmental friendliness" of protocols. Next, we show how to realize any ideal functionality by protocols in the plain model[1], while guaranteeing environmental friendliness with respect to many protocols Π "in the environment" and corresponding security properties P. This allows us to regain much of the original appeal of UC security, while avoiding those trust assumptions that are inherent in full-fledged UC security.

Interestingly, whether our security definition is friendly to some protocol Π and property P depends not so much on Π and P, but rather on the way this property is proven. Specifically, our security definition is friendly to *any* Π and P, as long as: (a) P is a game-based property, and (b) the fact that Π satisfies P is proven via some (potentially non-explicit) black-box reduction to a game-based hardness assumption. We observe that this includes most cryptographic protocols known to date, and in particular all prevalent secure communication, identification, and electronic commerce protocols.

Our solution is based on the interplay between provability and black-box *un*provability of cryptographic schemes. Specifically, a key component in our construction is a new CCA-secure commitment scheme that cannot be proven secure via black box reduction to a game-based assumption, and at the same time can be proven secure using a non-black-box reduction.

A. Our Results in More Detail

To better present our results, let us first present in some more detail the general approach for defining security of protocols. One starts by considering an execution of the analyzed protocol with two adversarial entities: an *adversary,* that controls the communication, corrupts parties, and represents attacks on the protocol itself, and an *environment*, that controls the inputs and outputs of the parties and represents the "rest of the system". Protocol ρ is said to *emulate* protocol ϕ if for any adversary $\mathcal{A}$ there exists an adversary $\mathcal{S}$ such that no environment $\mathcal{E}$ can tell whether it is interacting with ρ and $\mathcal{A}$ or with ϕ and $\mathcal{S}$. Security requirements are then written by way of requiring that ρ emulate an ideal protocol where all parties hand their inputs to a trusted party and obtain their inputs from that party. The program run by the trusted party is called the *ideal functionality* that ρ realizes.

Different notions of security formalize the above definitional approach in different ways. While the differences span several aspects, here we concentrate on one central aspect, namely the complexities of the adversarial entities involved. UC security requires that all entities (namely, the environment $\mathcal{E}$ and adversaries

[1] In the plain model the parties have ideally authenticated communication channels and no other trusted set-up. Alternatively, using the results in [2], the results in this work can be translated to the *bare model* where the parties do not even have access to authenticated communication.

$\mathcal{A}$ and $\mathcal{S}$) are polytime. This is essential for proving the composition theorem, but also underlies the impossibility results in the plain model.

Super-polynomial security (SPS) [39], [5] relaxes UC security in that it allows the ideal-model adversary $\mathcal{S}$ (often called the *simulator*) to run in time T that is somewhat super-polynomial (say, $T(n) = O(n^{\log n})$). While for some functionalities this relaxation makes the guarantees provided by the ideal protocol meaningless, for many functionalities of interest this relaxation still provides meaningful guarantees with respect to the protocol execution itself. However, SPS provides much weaker environmental friendliness guarantees. To see this, consider a protocol ρ that realizes an ideal protocol $\mathcal{G}$ with SPS security, and a protocol Π that is guaranteed to satisfy property P in the presence of $\mathcal{G}$. Now, we would like to argue that Π continues to satisfy P even when the instances of $\mathcal{G}$ are replaced by instances of ρ. However, this is the case only if originally Π satisfied P against adversaries that run in time T! This limitation is especially bothersome since SPS security does not guarantee security against adversaries that run in time T.

Another relaxation of UC security is angel-based security [45]. Here, both the simulator and the adversary are given super-polynomial resources — however these resources are given as a function of the specific context in which these resources are needed. Specifically, the model of protocol execution is augmented with a computationally unbounded "helper" (or, "angel") entity that is aware of the execution and responds to queries of the adversary (or simulator) depending on the current state of the execution. (For instance, in [45] the angel essentially finds collisions in a collision resistant hash function, but makes sure that the collision point encodes the identity of a party that's under the control of the adversary.) In particular, a number of works present angels with respect to which one can realize general functionalities in the plain model [45], [35], [15]. Angel-based security has the potential to provide better environmental friendliness guarantees than SPS security, since the super-polynomial advice is restricted in scope and depends on the global state of the system.

It is stressed however that the meaningfulness of the notion is angel-specific. *In particular, whether a given protocol Π "in the environment" continues to satisfy security properties P when an angel-based protocol is added to the system critically depends upon the specific angel in use.* The only prior work that addresses this issue [15] argues (somewhat informally) that any protocol that's proven secure with respect to their specific angel is friendly to any set of protocols in the environment, as long as all of these protocols together complete within a constant number of rounds. This is arguably a rather restricted guarantee. Still, we use this notion of security as a basis for ours.

OUR RESULTS. We first make rigorous the notion of "environmental friendliness" of protocols. Next, we show a new angel, $\mathcal{H}$, such that: (a) It is possible to realize practically any ideal functionality with respect to angel $\mathcal{H}$, in the plain model. (b) There is a large class of protocols Π and security properties P such that any protocol ρ that securely realizes, with respect to angel $\mathcal{H}$, some functionality $\mathcal{F}$, is environmentally friendly to Π and P.

DEFINING ENVIRONMENTAL FRIENDLINESS. We focus on formalizing the notion of environmental friendliness with respect to the traditional notion of *game-based* security properties of [38], [46]. Using this approach, a security property P of a cryptographic scheme Π is defined via a game between an efficient challenger Chal (that depends on Π, usually by just running it) and an adversary $\mathcal{A}$, and the property P holds if any efficient adversary $\mathcal{A}$ (i.e., non-uniform polynomial time machines) wins the game with only negligible advantage over some threshold probability τ.

Given a game-based security property with a challenger Chal, consider the following transformation on P that is parameterized by two protocols ρ and $\mathcal{G}$. (We think of $\mathcal{G}$ as a "specification protocol" or ideal functionality, and ρ is a protocol that "emulates" $\mathcal{G}$ according to some yet-to-be-specified notion of security.) The transformation views the challenger Chal of the property as representing an interaction between Π and an adversary, in the presence of an environment that contains some running instances of the ideal protocol $\mathcal{G}$. It then transforms it into a new challenger $\mathsf{Chal}^{\rho/\mathcal{G}}$ that's identical to Chal except that each instance of $\mathcal{G}$ within Chal is replaced with an instance of ρ. (The exact mechanics of the replacement operation are detailed within.) We denote by $P^{\rho/\mathcal{G}}$ the resulting security property (with the same wining probability threshold as P).

Now, say that the pair $(\rho, \mathcal{G})$ is *environmentally friendly* to the pair (Π, P), if the transformed property $P^{\rho/\mathcal{G}}$ holds. The intuition behind this definition is that it directly formulates the following requirement: "If Π enjoyed security property P in systems that include instances of the ideal protocol $\mathcal{G}$ then Π will continue to enjoy security property P when each instance of $\mathcal{G}$ is replaced with an instance of ρ."

ACHIEVING ENVIRONMENTALLY FRIENDLINESS. In the specific context of angel-based security, a protocol ρ $\mathcal{H}$-emulates protocol $\mathcal{G}$ if it UC-emulates $\mathcal{G}$ in the presence of angel (or, helper) $\mathcal{H}$; therefore, an angel

defines a notion of secure computation. We say that an angel $\mathcal{H}$ is environmentally friendly to $\{(\Pi, P)\}$ if the security notion "UC-emulating $\mathcal{G}$ in the presence of $\mathcal{H}$" is environmentally friendly to $\{(\Pi, P)\}$.

To show that an angel $\mathcal{H}$ is environmentally friendly to a protocol Π with security property P, one has to show that giving the adversary $\mathcal{A}$ access to $\mathcal{H}$ when interacting with the challengers in P does not help $\mathcal{A}$ in winning the security game. We first observe that the angel proposed in [15] already provides some environmental friendliness guarantees. Specifically, the angel $\mathcal{H}$ used in that work is an instantiation of a CCA-secure commitment scheme with the following additional *robustness* property: Whenever the adversary has only a constant number of rounds of interaction with the external environment, it can be simulated by another adversary that does not interact with $\mathcal{H}$ at all. This can be used to argue that their angel $\mathcal{H}$ is environmentally friendly to any protocol Π that takes only a constant number of rounds (or, more precisely, whenever the security game used to define the security of Π takes only a constant number of rounds).

This is indeed a meaningful guarantee, which can be extended to any fixed polynomial number of rounds, at the price of making the CCA commitment scheme underlying $\mathcal{H}$ have even more rounds.

However, it is limited in that it cannot apply to those common protocols (such as, say, the IPSec, TLS, SSH, or HTTPS secure session protocols) that have an unbounded number of rounds, or are invoked an unbounded number of times.

To deal with such protocols and others, we take a different approach: Instead of bounding the actual communication of the protocol, we focus on the way in which security of the protocol is proven. That is, we show:

MAIN THEOREM (INFORMAL): Assume there exist trapdoor permutations and collision resistant hash functions. Then there exists an angel $\mathcal{H}$ such that: (a) practically any ideal functionality can be $\mathcal{H}$-realized in the plain model, and (b) $\mathcal{H}$ is environmentally friendly to any protocol Π with game-based property P *as long as the fact that* Π *satisfies* P *is proven via (potentially not explicitly specified) black box reduction to a game-based cryptographic hardness assumption* C *with a bounded polynomial number of rounds.*

Similarly to a game based primitive, a game based (or, falsifiable [38]) cryptographic hardness assumption is defined via a game between a polytime challenger and an adversary. The assumption states that no efficient adversary can win the game with noticeable advantage, where winning is an event recognizable by the challenger.

This shift of focus greatly extends the set of protocols and properties for which a notion of security can guarantee friendliness to, while preserving realizability in the plain model. Indeed, most proofs of security in the literature proceed via (potentially not explicitly specified) black box reduction to a hardness assumption that's formulated by way of a game with fixed polynomial (in fact, even constant) number of rounds. This holds even when the protocol itself has an unbounded number of communication rounds, and even when one asserts a security property that considers unboundedly many concurrently running instances of a simpler protocol.

To exemplify this point, consider the commonplace case of key exchange and secure communication channel protocols, such as the IPSec, TLS or SSH Internet standards. Security of such protocols is considered in a model where multiple communication sessions run concurrently [7], [12]. Consequently, from the point of view of environmental friendliness, these secure communication protocols "in the environment" have an unbounded number of rounds. (Indeed, the number of rounds depends on the number of communication sessions generated by the adversary.) In contrast, security proofs of prevalent protocols (e.g., [32], [29]) use black box reductions to assumptions that have only a small constant number of rounds. To the best of our knowledge this is the first time where there is a "tangible" security advantage to proving security via *black-box* reduction.

B. Our Techniques

At the heart of our techniques is a new connection between environmental friendliness and recent results on *black-box unprovability*. More specifically, we present a Chosen-Commitment-Attack (CCA) secure commitment scheme and show how the unprovability results in [40] can be extended to rule out black-box (but also non-uniform) proofs of security for the hiding property of this particular commitment scheme, based on any bounded-round game-based assumptions. Yet, we show—using *non-black-box techniques*, similar to those used in [1]—that the protocol indeed is hiding. We then show that any protocol that exhibits such a "gap" between non-black-box and black-box security proofs can be used to implement an angel that is environmental friendly to any protocol whose security is proven secure using a black-box reduction to any bounded-round game-based assumptions.

We remark that our results holds even with respect to black-box reductions that are not explicitly specified; at a technical level, this means the reductions may depend on some non-uniform advice (potentially not efficiently computable) about the adversary, for instance, the size of the adversary, or a sequence of random coins that

maximize the adversary's winning probability for every input length. This type of reductions can be viewed as "gray-box reductions"; in the rest of the paper, we refer to them as non-uniform black-box reductions, as opposed to uniform black-box reductions that are explicitly specified.

STEP 1: FROM UNPROVABILITY TO ENVIRONMENTAL FRIENDLINESS: Our starting point is the angel of [15]. At the basis of that angel lies a CCA secure commitment scheme $\langle C, R\rangle$. Roughly speaking, a tag-based commitment scheme (i.e., commitment scheme that takes an identifier—called the tag—as an additional input) is said to be *CCA-secure* if the value committed to with respect to the tag id remains hidden even if the receiver has access to a (super-polynomial time) oracle, called the *committed-value oracle* $\mathcal{O}$, that "breaks" commitments of the adversary's choice using any tag $\mathsf{id}' \neq \mathsf{id}$. More precisely, $\mathcal{O}$ interacts with the adversary playing the role of the receiver, in multiple and potentially concurrent commitments, and as soon as it completes an accepting commitment stage, it gives the committed value to the adversary. In [15] and its follow up work [33], the angel $\mathcal{H}$ (which is interactive and stateful) acts exactly the same as the *committed-value oracle* to the adversary (or simulator), *except that it only opens commitments where the tag encodes the identity of a corrupted party.* Intuitively, this angel allows the adversary or simulator to recover the committed values in all commitments generated corrupted parties, while the CCA property guarantees that the extra power provided by the angel does not help the adversary to compromise the security of commitments generated by honest parties.

Following the recent results on black-box unprovability [40], [23], [17], we say that a commitment scheme is unprovable from an assumption C if the existence of a black-box reduction basing its security on C implies that C is false. However, here we use the following stronger variant of this notion, called *strong unprovability:* Consider a (potentially computationally unbounded) "ideal" adversary $\mathcal{A}$ of a commitment scheme that acts as the honest receiver and returns the committed value with probability 1 at the end of the commit phase (using randomness generated by evaluating its internal random oracle RO). Then a commitment scheme is strongly unprovable if the existence of a black-box reduction that works only with the ideal adversary (rather than with any adversary violating hiding) already implies that underlying assumption is false.

Given a CCA secure commitment scheme that is strongly unprovable from an assumption C, we show that the corresponding angle $\mathcal{H}$ is environmentally friendly to any pair (Π, P) for which there is a black-box reduction basing the statement "scheme Π has property P" on assumption C. Intuitively, we need to show that giving an adversary B access to $\mathcal{H}$ does not help B in violating the security property P. Assume for contradiction that B with access to $\mathcal{H}$, denoted $B[\mathcal{H}]$, does break P. Since $\mathcal{H}$ (acting as the committed value oracle) can be perfectly emulated using oracle access to the ideal adversary $\mathcal{A}$ of the CCA commitment scheme, there is a B' that with oracle access to $\mathcal{A}$ violates P (by emulating $B[\mathcal{H}]$). Then, the black-box reduction R can "translate" this adversary $(B')^{\mathcal{A}}$ violating P into an adversary breaking C—namely, R with oracle access to $(B')^{\mathcal{A}}$ can break C. Therefore, there is another machine $R' = R^{B'}$ that with oracle access to only $\mathcal{A}$ breaks C; this contradicts the strongly unprovability of the CCA commitments, and thus $\mathcal{H}$ is environmentally friendly to any (Π, P) that admits a black-box security reduction to assumption C.

Now given an angel $\mathcal{H}$ that is environmental friendly to a pair (Π, P), we show that any secure computation protocol ρ that $\mathcal{H}$-emulates a functionality $\mathcal{G}$ is also environmental friendly to (Π, P). This amounts to showing that if the property $P = (\mathsf{Chal}, \tau)$ holds, then, the property $P^{\rho/\mathcal{G}} = (\mathsf{Chal}^{\rho/\mathcal{G}}, \tau)$ holds as well. Suppose not, and there is an adversary A that wins the challenger $\mathsf{Chal}^{\rho/\mathcal{G}}$ with non-negligible advantage, then it simply follows from the security of ρ that there is a simulator S that with access to $\mathcal{H}$ wins Chal with non-negligible advantage. However, this contradicts with the fact that having access to $\mathcal{H}$ does not help any adversary, including S, to violate property P. Therefore, we conclude that:

LEMMA (INFORMAL): Let $\langle C, R\rangle$ be a CCA-secure commitment scheme whose hiding property is strongly unprovable from a class $\mathcal{C}$ of game-based assumptions, and $\mathcal{H}$ the corresponding angel. Then assume the existence of trapdoor permutations, every functionality $\mathcal{G}$ can be $\mathcal{H}$-emulated by a protocol ρ that is environmental friendly to any protocol Π with properties P *provided that* Π *satisfies* P *is proven via a(potentially non-uniform) black box reduction to an assumption in* $\mathcal{C}$.

STEP 2: CONSTRUCTING STRONGLY UNPROVABLE CCA SECURE COMMITMENTS: To show our main theorem, it remains to construct a CCA-secure commitment scheme whose security (specifically, the hiding property) is strongly unprovable from any bounded-round assumptions. Our construction starts from the CCA commitment of [15]. However, that construction has a black-box security reduction to one way functions. We thus modify that construction to achieve strong unprovability, using techniques from non-black-box ZK protocols [1], [4], [41], [18], [27], [42], [14].

Let us first briefly review the CLP construction. The protocol follows the Feige-Shamir paradigm: The

receiver first sends a random n-bit string s which defines a "trapdoor" that is the pre-image t of s through a one-way permutation; the committer then sends a commitment c to the message m using a standard perfectly binding commitment scheme, followed by multiple 3-round witness indistinguishable special soundness proofs (WISSP) that it knows either a decommitment to c, or the trapdoor t. Messages in the WISSP proofs are scheduled in a special way (which builds upon the *message scheduling technique* of Dolev, Dwork and Naor [21]) so as to allow proving CCA security. The existence of the "trapdoor" is crucial for showing the hiding property: Note that a black-box *non-uniform* reduction can easily obtain a trapdoor, by receiving as an advice the pre-image r of the first message s from a (w.o.l.g. deterministic) malicious receiver R^*; this allows it to simulate the WISSP proofs to R^* without knowing the decommitment of the basic commitment c, and thus reducing the hiding of the protocol to the hiding of the basic commitment.

Therefore, to construct a strongly unprovable CCA commitment, we need (at least) a way to set up a trapdoor that prevents black-box reductions from obtaining a trapdoor. Following [1], we modify the CLP-protocol as follows: Instead of letting the receiver send a random string, the committer and receiver participates in a "trapdoor-setting sub-protocol" where the committer first sends a commitment c' to an all-zero string, obtains a random challenge r from the receiver, and then gives a proof (encrypted) that either it knows a decommitment of the main commitment c or that c' is a commitment to a program that generates r on input c'. Let $\langle C, R\rangle$ denote this new protocol. Using non-black-box simulation techniques, a "trapdoor" can be obtained by a reduction knowing the code of the cheating receiver, and thus hiding can be proven using a non-black-box reduction; but, any black-box reductions, even non-uniform one, cannot do so.

To show that the new construction is strongly unprovable, we build upon the techniques of developed by Pass [40] for showing that it is impossible to base the sequential witness hiding property of any constant-round (computational) special-sound proof/argument on any bounded-round game-based assumptions via black-box *uniform* reductions; in fact, a careful examination shows that the proof of [40] actually shows the *strong unprovability* of the sequential witness hiding property of (constant-round) special-sound proofs. It seems that we can directly apply this stronger result to show that $\langle C, R\rangle$ is strongly unprovable from bounded-round assumptions, as the hiding property of $\langle C, R\rangle$ implies the sequential witness hiding property of the 3-round special sound proofs contained in $\langle C, R\rangle$, which is strongly unprovable by [40].

However, this high-level approach does not go through easily for two reasons: First, the unprovability result of [40] only holds for black-box *uniform* reductions, whereas, we need to show that $\langle C, R\rangle$ is strongly unprovable even using black-box *non-uniform* reductions; second, the result of [40] only applies to special sound proofs for *unique witness language* (i.e., every true statement has a unique witness), this does not hold for the statements proven in Stage 3 which have at least two witnesses—the decommitment of c and a "trapdoor". Similar problems were encountered in a different context in [17]: This previous work first extends the unprovability result of [40] to handle black-box non-uniform reductions, and then separately, extends it to handle languages where true statements may have multiple witnesses, but only one of them is computationally easy to find (while all other witnesses are computationally hard to find as the "trapdoor"). However, their extended unprovability results are not sufficient for showing the unprovability of our CCA commitment scheme via non-uniform reductions, which requires handling non-uniform reduction and languages with multiple witnesses at the same time. Nevertheless, we show that building upon their techniques, we solve both of the above two challenges simultaneously.

II. Environmental Friendliness

In this section, we formalize the notion of environmental friendliness to game-based properties. We start with introducing some basic definitions.

Security Game and Game-Based Assumptions

We provide the definitions of *(canonical) security games* (or game for short) and *game-based assumptions*. Our definitions are almost identical to the notions of security games and falsifiable assumptions in the literature [40], [38], [20], [28], [47], [23], except that a game-based assumption is additionally associated with a *trivial strategy* that achieves certain threshold winning probability.

Definition 1. *A security game (or game) consists of an ITM* Chal*, called the challenger, that is polynomial-time in the length of the messages it receives, and a constant* $\tau_{\mathcal{C}}$*, called the threshold, in the interval* $[0, 1)$*. In an execution of a security game, the challenger* Chal *interacts with an adversary* A *on common input* 1^n *and outputs accept or reject at the end of the interaction.*

We say that A_n *breaks* Chal_n *with advantage* ε*, if* A_n *makes* Chal_n *accept with probability* $\tau_{\mathcal{C}} + \varepsilon$*. We say that* A *breaks* Chal*, or the game-based assumption* $\mathcal{C}$*, if* A_n *breaks* Chal_n *with advantage* $\varepsilon(n)$ *for infinitely many* $n \in N$ *for a non-negligible function* ε*.* ε *is the* advantage *of the adversary.*

Definition 2 (Game-Based Assumptions)**.** *A* game-based assumption *is simply a security game* $\mathcal{C} = (\mathsf{Chal}, \tau)$*, such that, there is a non-uniform PPT adversary* A*, called the* trivial strategy*, satisfying that* A_n *breaks* Chal_n *with probability at least* τ *(potentially without any advantage) for all* $n \in N$*. We say that assumption* $\mathcal{C}$ *holds if no non-uniform PPT adversary can break the game* (Chal, τ)*.*

We say that a game (and game-based assumption resp.) has $r(n)$ rounds, if for every n, the challenger Chal_n interacts in at most $r(n)$ rounds; and we say that a game has bounded rounds if the challenger interacts in a fixed bounded (polynomial) number of rounds.

GAME-BASED SECURITY PROPERTIES We say that a security property of a cryptographic scheme Π is game-based—called a game-based security property (or game-based property for short)—if it can be specified using a security game $P_\Pi = (\mathsf{Chal}, \tau)$ where Chal depends on the scheme Π.

Definition 3 (Game-Based Security Property)**.** *A* game-based security property *of a cryptographic scheme* Π *is simply a security game* $P_\Pi = (\mathsf{Chal}, \tau)$*. We say that the property* P_Π *holds if no non-uniform PPT adversary can break the game* (Chal, τ)*.* [2]

For many natural cryptographic primitives, their security requirements can be described as a game, for instance, the property of witness-indistinguishability of an interactive argument and the property of unforgibility of a signature scheme (which admits an unbounded-round security game). In fact, for most natural security properties, the challenger of the security game depend on the cryptographic scheme, but the threshold does not and is often either 0 (*e.g.* inverting a one-way function f) or $1/2$ (*e.g.* distinguishing a pseudo-random string generated by a PRG g from a uniform string).

We also note that not all security properties are game-based. For example, (adaptive) soundness of an interactive argument $\langle P, V\rangle$ and in general simulated-base security properties, such as the zero-knowledge property.

BLACK-BOX REDUCTIONS: Our definition of black-box reductions for game-based security properties is a special case of the definition of non-uniform black-box reductions for general cryptographic primitives in [17], when restricted to consider only security properties that are game-based.

Definition 4 (Black-Box Reductions for Game-Based Security Properties)**.** *Let* $P = (\mathsf{Chal}, \tau)$ *be a game-based security property of a cryptographic scheme* Π*. We say that a non-uniform PPT machine* R *is a* black-box security reduction *of* Π *from a game-based assumption* $C = (\mathsf{Chal}_C, \tau_C)$*, if there exists a function* $\mathsf{Z} : \{0,1\}^* \to \{0,1\}^*$ *and polynomials* s*,* a*,* m*, such that,* $m(n) = n^{\Theta(1)}$ *and for every family of* deterministic *circuits* A *that breaks* P*, that is, for infinitely many* $n \in N$*,* A_n *breaks* $(\mathsf{Chal})_n$ *with advantage* $1/p(n)$ *for a polynomial* p*, the following two conditions hold for every such* $n \in N$*:*
(i) $z = \mathsf{Z}(A_n)$ *has at most* $s(n \cdot p(n))$ *bits.*
(ii) $\Pr[R^{A_n}(1^{m(n)}, \mathsf{Z}(A_n))$ *breaks* $(\mathsf{Chal}_C)_{m(n)}] > \tau_C + 1/a(n \cdot p(n))$*.*

A. Formalize Environmental Friendliness

We say that a protocol ρ implementing a functionality $\mathcal{G}$ is *environmental friendly* to a cryptographic scheme Π with game-based property $P = (\mathsf{Chal}, \tau)$ if the following holds: If property P captures that executions of Π provide certain security guarantees in an (execution) environment where there coexists concurrent executions of (the ideal protocol $\pi_\mathcal{G}$ accessing) $\mathcal{G}$, *whose inputs are potentially correlated with the inputs and randomness of the executions of* Π, then, the same security guarantees hold even where executions of $\mathcal{G}$ are replaced with that of ρ in the environment. Security of Π in the latter scenario is captured via a new game $P^{\mathcal{G}/\rho} = (\mathsf{Chal}^{\mathcal{G}/\rho}, \tau)$, where the challenger $\mathsf{Chal}^{\mathcal{G}/\rho}$ proceeds identically to Chal except that every "sub-routine call" to $\mathcal{G}$ is replace with a "sub-routine call" to ρ (using the same inputs). Then, we say that ρ is environmental friendly to P if property $P^{\mathcal{G}/\rho}$ also holds. Below we formalize the notion of environmental friendliness; towards this, we first define the operation of making a "sub-routine call" to a protocol.

SUB-ROUTINE CALL TO A PROTOCOL: Let ρ be a m-party protocol for a set of players $\{P_1, \cdots, P_m\}$ in the $\mathcal{F}$-hybrid model, that is, it consists of m interactive algorithms $\rho = (A_1, \cdots, A_m)$ that accesses $\mathcal{F}$. We denote by $M = B^\rho$ an interactive machine that runs B who makes sub-routine calls to the protocol ρ. More precisely, for the i^{th} sub-routine call to ρ, B chooses a session id sid_i, a subset of players $S = \{P_{i_1}, \cdots, P_{i_k}\}$ and inputs $x_{i_1}, \cdots, x_{i_k}$; it then emulates $\mathcal{F}$ and honest players in S by spawning a sub-routine running $\mathcal{F}$ and a sub-routine running algorithm A_{i_j} with input $(\mathsf{sid}_i, x_{i_j})$ for each $j \in [k]$; messages from player P_{i_j} (emulated by the sub-routine running A_{i_j}) is delivered within M if its recipient is $\mathcal{F}$ or another player in S, and otherwise forwarded externally. Overall, messages from M consists of messages sent directly by B and that sent by its sub-routines; correspondingly, messages delivered to M are forwarded internally to B or one of its sub-

[2] Note that game-based security properties are not associated with any trivial strategies as game-based assumptions do.

routines appropriately; furthermore, the scheduling of the delivery of messages is controlled by B.

For an ideal functionality $\mathcal{G}$, we denote by $M = B^{\mathcal{G}}$ the interactive machine that makes sub-routine calls to the ideal protocol $\pi_{\mathcal{G}}$ of $\mathcal{G}$. Given an arbitrary machine M and protocol ρ (or functionality $\mathcal{G}$), one can always interpret M as B^{ρ} (or $B^{\pi_{\mathcal{G}}}$) by examining the sub-routine calls inside M. We denote by $M^{\rho/\phi}$ (or $M^{\mathcal{G}/\phi}$ resp.) the interactive machine that proceeds identically to M except that every sub-routine call to ρ (or $\pi_{\mathcal{G}}$ resp.) is replace with a sub-routine call to ϕ using the same input.

Definition 5 (Environmental Friendliness)**.** *Let $P = (\mathsf{Chal}, \tau)$ be a game-based security property of a cryptographic scheme Π, and ρ a protocol implementing a functionality $\mathcal{G}$. Then we say that ρ is* environmental friendly *to Π with property P, if the security property $P^{\mathcal{G}/\rho} = (\mathsf{Chal}^{\mathcal{G}/\rho}, \tau)$ holds.*

III. Achieving Friendliness

Traditionally, black-box unprovability results are often viewed as limitations towards achieving the desired security properties. In this work, we view the *unprovability via black-box reduction* of a cryptographic scheme as a feature. We show that given a robust CCA secure commitment scheme whose hiding property is *strongly unprovable* via black-box reductions from a class of assumptions, we can construct secure computation protocols that are "environmental friendly" to cryptographic schemes whose security is *proven* via black-box reductions from the same class of assumptions.

In this extended abstract, due to the lack of space, we skip reviewing the formal definitions of CCA-security and k-robustness of a commitment scheme, and the framework of angle-based security; we refer the reader to [15], [33] for formal definitions.

A. Strong Unprovability via BB Reductions

The hiding property of a commitment scheme Π is provable via black-box reductions based on an assumption C, if there is a machine R (a.k.a. the reduction) such that, given any adversary A that violates the hiding property of Π, R with black box access to A breaks the assumption C. Correspondingly, we say that (the hiding property of) a commitment scheme Π is *unprovable* via black-box reductions based on an assumption C, if the existence of a valid reduction R implies that the assumption C is false. In this work, we need a stronger notion of unprovability—called *strong unprovability*—which rules out even the possibility of having a reduction that works solely with a single stylized *ideal adversary*.

Ideal Adversary $\mathcal{A}$ of $\langle C, R\rangle$: The ideal adversary $\mathcal{A}$ participates in a single execution of the commit phase of $\langle C, R\rangle$, in which it follows the honest receiver strategy and at the end of the commit phase, it returns the unique committed value v if there is one, and $\perp$ otherwise; furthermore, $\mathcal{A}$ uses fresh randomness in answering any new query *even when rewound* by applying its internal random oracle $RO \leftarrow \mathbf{RO}_n$ on the partial transcript it has received so far. We note that, in a straight-line interaction (without rewindings) with a committer, the ideal adversary $\mathcal{A}$ emulates the committed-value oracle $\mathcal{O}$ of $\langle C, R\rangle$ perfectly, as the distribution of messages from $\mathcal{A}$ is identical to that from $\mathcal{O}$.

Definition 6 (Strong Unprovability of a Commitment Scheme)**.** *We say that a commitment scheme $\langle C, R\rangle$ is* strongly unprovable *via black-box reductions from a game-based assumptions $C = (\mathsf{Chal}_C, \tau_C)$, if the existence of a non-uniform PPT reduction R that with black-box access to the ideal adversary breaks the assumption C, implies the existence of a non-uniform PPT machine S that directly breaks C, where the former condition is formalized as follows:*

There is a non-uniform PPT machine R, a function Z, and polynomials s, m, a, such that (i) $m = \Theta(n)$, (ii) *for every $n \in N$ and $RO \in \mathbf{RO}_n$, $z = \mathsf{Z}(\mathcal{A}_n^{RO})$ has at most $1^{s(n)}$ bits, and* (iii) *for sufficiently large $n \in N$, the probability that $R^{\mathcal{A}_n^{RO}}(1^{m(n)}, z)$ breaks $(\mathsf{Chal}_C)_{m(n)}$ is at least $\tau_C + 1/a(n)$, where $RO \leftarrow \mathbf{RO}_n$ is a random oracle and $z = \mathsf{Z}(\mathcal{A}_n^{RO})$.*

B. Friendliness from Strong Unprovability

Consider any robust CCA secure commitment $\langle C, R\rangle$ that is *strongly unprovable* from a class of game-based assumptions $\mathcal{C}$; let $\mathcal{H}$ be the helper functionality that corresponds to the committed-value oracle of $\langle C, R\rangle$ as described in the introduction. It already follows from the main theorems in [15], [33] that every well-formed functionality can be $\mathcal{H}$-emulated.

Theorem 1 ([15], [33])**.** *Let $\langle C, R\rangle$ be a $T(n)$-round robust CCA secure commitment, and $\mathcal{H}$ the corresponding helper functionality. Assume the existence enhanced trapdoor permutations. Then, for every well-formed functionality $\mathcal{F}$, there exists a $O(T(n))$-round protocol ρ that $\mathcal{H}$-emulates $\mathcal{F}$.*

Furthermore, we show that these $\mathcal{H}$-secure protocols are also environmentally friendly.

Theorem 2. *Let $\langle C, R\rangle$ be a robust CCA secure commitment that is strongly unprovable from a set of game-based assumptions $\mathcal{C}$, and $\mathcal{H}$ the corresponding helper functionality. For any protocol ρ that $\mathcal{H}$-emulates a functionality $\mathcal{G}$, and every cryptographic scheme Π with*

a game-based security property P that is provable via black-box reductions from an assumption $C \in \mathcal{C}$, ρ is environmental friendly *to Π with property P, assuming that C is true.*

Proof: Let the property P be (Chal, τ) and the assumption C be $(\mathsf{Chal}_C, \tau_C)$, and R the black-box reduction of Π that base property P on C; interpret the challenger Chal as $B^{\mathcal{G}}$. Towards showing the theorem, we show that if ρ is not environmental friendly to Π, that is, the property $P^{\mathcal{G}/\rho} = (\mathsf{Chal}^{\mathcal{G}/\rho} = B^{\rho}, \tau)$ does not hold, then the assumption C is false. More precisely, assume that there is a non-uniform PPT adversary A and polynomial p, such that, A breaks the challenger B^{ρ} with advantage $1/p(n)$ (with probability $\tau_C + 1/p(n)$); then we construct another machine D that breaks C directly with a related inverse polynomial probability. Below we construct D in three steps.

STEP 1—BY THE $\mathcal{H}$-SECURITY OF ρ, SIMULATOR $S[\mathcal{H}]$ BREAKS PROPERTY P : As discussed in Section II-A, the interaction between B^{ρ} and the adversary A can be casted in the UC model and viewed as a concurrent-execution of the protocol ρ with adversary A and environment B. Since ρ is a $\mathcal{H}$-implementation of the ideal functionality $\mathcal{G}$, this concurrent-execution corresponds exactly to a real-world execution of ρ in the $\mathcal{H}$-model with adversary A and environment B[3]. Then, by our hypothesis, the environment B accepts with probability $\tau + 1/p(n)$ in the real-world.

It follows from the $\mathcal{H}$-security of ρ (and the UC composition theorem) that there exists a non-uniform PPT simulator S such that no environment can distinguish the real-world execution of ρ with A from an ideal-world execution of $\pi_{\mathcal{G}}$ with the simulator $S[\mathcal{H}]$ having access to the helper functionality $\mathcal{H}$; in particular, the probability that environment B accepts in the real-world must be almost the same as that in the idea-world except from a negligible amount. Therefore, the probability that B accepts in the ideal world is least $\tau + 1/2p(n)$. Note that the ideal-world execution corresponds exactly to the interaction between machine $B^{\mathcal{G}}$ and the adversary $S[\mathcal{H}]$; thus $S[\mathcal{H}]$ breaks the game $(\mathsf{Chal} = B^{\mathcal{G}}, \tau)$ with advantage at least $1/2p(n)$ for infinitely many $n \in N$.

STEP 2—BY THE BLACK-BOX REDUCTION R OF P, THERE IS $\overline{R}$ WITH ACCESS TO $S'^{\mathcal{A}^{RO}}$ BREAKS C: From Step 1, we have that $S[\mathcal{H}]$ breaks the game (Chal, τ), and thus violates the property P, with advantage $1/2p(n)$. We first show that there exists a polynomial-sized deterministic circuit S' when having oracle access to the ideal adversary $\mathcal{A}^{RO}$ of the commitment scheme $\langle C, R\rangle$, violates the property P with at least the same advantage. The circuit S' emulate the execution S with the best random coins (hardwired in) that maximizes the winning probability of S, and emulates the helper functionality $\mathcal{H}$ for S using the ideal adversary $\mathcal{A}^{RO}$ of $\langle C, R\rangle$; the latter can be done since by definition the helper functionality $\mathcal{H}$ can be emulated efficiently using a the committed-value oracle of $\langle C, R\rangle$, which in turn can be emulated efficiently using oracle access to the ideal adversary $\mathcal{A}^{RO}$ of $\langle C, R\rangle$. Therefore the relativized (deterministic) circuit $S'^{\mathcal{A}^{RO}}$ violates P with advantage at least $1/2p(n)$. Then, by the fact that the property P can be based on assumption C via a black-box reduction R, there exists a non-uniform PPT reduction $\overline{R}$ that with access to $S'^{\mathcal{A}^{RO}}$ breaks the assumption C with some inverse polynomial advantage $1/q(n)$.

STEP 3—BY THE STRONG UNPROVABILITY OF $\langle C, R\rangle$, THERE IS A MACHINE D THAT BREAKS C: As S' is efficient, we can construct another non-uniform PPT machine R' that with oracle access to merely $\mathcal{A}^{RO}$ emulates perfectly the execution of $\overline{R}$ by running S' internally and forwarding all its oracle queries to its own oracle; therefore R' with oracle access to the ideal adversary $\mathcal{A}^{RO}$ of $\langle C, R\rangle$ breaks the assumption C. Then, it follows from the strong unprovability of $\langle C, R\rangle$ that the existence of R' implies that there is another non-uniform PPT machine D that breaks the assumption C directly. This concludes the construction machine D. See the full version for a detailed description of R'. ■

C. Strongly Unprovable CCA Commitments

It remains to construct a robust CCA secure commitment scheme that is black-box unprovable. We show the following theorem.

Theorem 3. *Assume the existence of collision resistant hash functions and one-way permutations. For every polynomial $\lambda(n)$ and every constant $\delta > 0$, there exists a $O(n^{\delta} + \lambda(n))$-round robust CCA secure protocol that is strongly unprovable from any $\lambda(n)$-round game-based hardness assumptions.*

Our protocol is built upon previous constructions of CCA secure commitment schemes of [15] and non-black-box zero-knowledge protocol by Barak [1]. Furthermore, the proof of black-box unprovability relies on previous impossibility results in [40], [17]. Here Due to the lack of space, we defer the proof of the theorem to the full version.

Finally, Combining Theorem 1, 2 and 3, we conclude the main theorem.

[3]Formally, the $\mathcal{H}$-model considers only a single execution of the protocol under analysis. However, this is w.l.o.g. since, the UC composition theorem guarantees that security of a "stand-alone" execution implies that of concurrent executions.

REFERENCES

[1] B. Barak. How to go beyond the black-box simulation barrier. In *FOCS*, volume 0, pages 106–115, 2001.

[2] B. Barak, R. Canetti, Y. Lindell, R. Pass, and T. Rabin. Secure computation without authentication. In *CRYPTO*, pages 361–377, 2005.

[3] B. Barak, R. Canetti, J. B. Nielsen, and R. Pass. Universally composable protocols with relaxed set-up assumptions. In *FOCS*, pages 186–195, 2004.

[4] B. Barak, O. Goldreich, S. Goldwasser, and Y. Lindell. Resettably-sound zero-knowledge and its applications. In *FOCS*, pages 116–125, 2001.

[5] B. Barak and A. Sahai. How to play almost any mental game over the net - concurrent composition via super-polynomial simulation. In *FOCS*, pages 543–552, 2005.

[6] D. Beaver. Foundations of secure interactive computing. In *CRYPTO*, pages 377–391, 1991.

[7] M. Bellare and P. Rogaway. Entity authentication and key distribution. In *CRYPTO*, pages 232–249, 1993.

[8] R. Canetti. Security and composition of multiparty cryptographic protocols. *Journal of Cryptology*, pages 143–202, 2000.

[9] R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *FOCS*, pages 136–145, 2001.

[10] R. Canetti, Y. Dodis, R. Pass, and S. Walfish. Universally composable security with global setup. In *TCC*, pages 61–85, 2007.

[11] R. Canetti and M. Fischlin. Universally composable commitments. In *CRYPTO*, pages 19–40, 2001.

[12] R. Canetti and H. Krawczyk. Analysis of key-exchange protocols and their use for building secure channels. In *EUROCRYPT*, pages 453–474, 2001.

[13] R. Canetti, E. Kushilevitz, and Y. Lindell. On the limitations of universally composable two-party computation without set-up assumptions. In *EUROCRYPT*, pages 68–86, 2003.

[14] R. Canetti, H. Lin, and O. Paneth. Public-coin concurrent zero-knowledge in the global hash model. In *TCC*, pages 80–99, 2013.

[15] R. Canetti, H. Lin, and R. Pass. Adaptive hardness and composable security in the plain model from standard assumptions. In *FOCS*, pages 541–550, 2010.

[16] R. Canetti, Y. Lindell, R. Ostrovsky, and A. Sahai. Universally composable two-party and multi-party secure computation. In *STOC*, pages 494–503, 2002.

[17] K.-M. Chung, H. Lin, M. Mahmoody, and R. Pass. On the power of nonuniformity in proofs of security. To Appear ITCS 2013, 2012.

[18] Y. Deng, V. Goyal, and A. Sahai. Resolving the simultaneous resettability conjecture and a new non-black-box simulation strategy. In *FOCS*, pages 251–260, 2009.

[19] Y. Dodis and S. Micali. Parallel reducibility for information-theoretically secure computation. In *CRYPTO*, pages 74–92, 2000.

[20] Y. Dodis, R. Oliveira, and K. Pietrzak. On the generic insecurity of the full domain hash. In *CRYPTO*, pages 449–466, 2005.

[21] D. Dolev, C. Dwork, and M. Naor. Nonmalleable cryptography. *SIAM Journal on Computing*, 30(2):391–437, 2000.

[22] C. Dwork, M. Naor, and A. Sahai. Concurrent zero-knowledge. *J. ACM*, 51(6):851–898, 2004.

[23] C. Gentry and D. Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In *STOC*, pages 99–108, 2011.

[24] O. Goldreich. *Foundations of Cryptography — Basic Applications*. Cambridge University Press, 2004.

[25] S. Goldwasser and L. A. Levin. Fair computation of general functions in presence of immoral majority. In *CRYPTO*, pages 77–93, 1990.

[26] V. Goyal. Constant round non-malleable protocols using one way functions. In *STOC*, pages 695–704, 2011.

[27] V. Goyal and A. Jain. On the round complexity of covert computation. In *STOC*, pages 191–200, 2010.

[28] I. Haitner and T. Holenstein. On the (im)possibility of key dependent encryption. In *TCC*, pages 202–219, 2009.

[29] T. Jager, F. Kohlar, S. Schäge, and J. Schwenk. On the security of tls-dhe in the standard model. In *CRYPTO*, pages 273–293, 2012.

[30] Y. T. Kalai, Y. Lindell, and M. Prabhakaran. Concurrent composition of secure protocols in the timing model. *J. Cryptology*, 20(4):431–492, 2007.

[31] J. Katz. Universally composable multi-party computation using tamper-proof hardware. In *EUROCRYPT*, pages 115–128, 2007.

[32] H. Krawczyk. Sigma: The 'sign-and-mac' approach to authenticated diffie-hellman and its use in the ike-protocols. In *CRYPTO*, pages 400–425, 2003.

[33] H. Lin and R. Pass. Black-box constructions of composable protocols without set-up. In *CRYPTO*, pages 461–478, 2012.

[34] H. Lin, R. Pass, and M. Venkitasubramaniam. A unified framework for concurrent security: universal composability from stand-alone non-malleability. In *STOC*, pages 179–188, 2009.

[35] T. Malkin, R. Moriarty, and N. Yakovenko. Generalized environmental security from number theoretic assumptions. In *TCC*, pages 343–359, 2006.

[36] S. Micali, R. Pass, and A. Rosen. Input-indistinguishable computation. In *FOCS*, pages 367–378, 2006.

[37] S. Micali and P. Rogaway. Secure computation (abstract). In *CRYPTO*, pages 392–404, 1991.

[38] M. Naor. On cryptographic assumptions and challenges. In *CRYPTO*, pages 96–109, 2003.

[39] R. Pass. Simulation in quasi-polynomial time, and its application to protocol composition. In *EUROCRYPT*, pages 160–176, 2003.

[40] R. Pass. Limits of provable security from standard assumptions. In L. Fortnow and S. P. Vadhan, editors, *STOC*, pages 109–118. ACM, 2011.

[41] R. Pass and A. Rosen. Concurrent non-malleable commitments. In *FOCS*, pages 563–572, 2005.

[42] R. Pass, A. Rosen, and W.-L. D. Tseng. Public-coin parallel zero-knowledge for np. *J. Cryptology*, 2011.

[43] B. Pfitzmann and M. Waidner. Composition and integrity preservation of secure reactive systems. In *ACM Conference on Computer and Communications Security*, pages 245–254, 2000.

[44] M. Prabhakaran and M. Rosulek. Cryptographic complexity of multi-party computation problems: Classifications and separations. In *CRYPTO*, pages 262–279, 2008.

[45] M. Prabhakaran and A. Sahai. New notions of security: achieving universal composability without trusted setup. In *STOC*, pages 242–251, 2004.

[46] O. Reingold, L. Trevisan, and S. P. Vadhan. Notions of reducibility between cryptographic primitives. In *TCC*, pages 1–20, 2004.

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