From Unprovability to Environementally Friendly Protocols

Ran Canetti* Huijia Lin† Rafael Pass‡

Abstract

An important property of cryptographic protocols is to what extent they adversely affect the security of the systems in which they run, and in particular whether introducing a new protocol to a system might 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 pariticipants 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 notions 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 provably cannot be proven secure via a *black box reduction* to a game-based assumption, but that nonetheless 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.

*Boston University and Tel Aviv University, Email: Canetti@tau.ac.il. Supported by the Check Point Institute for Information Security, an ISF grant and NSF award No. 1218461.

†MIT and Boston University, Email: huijia@csail.mit.edu.

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

Contents

1 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 [GL90, Bea91, MR91, DM00, Can00, Gol04, MPR06], and even in notions of concurrent security such as [DNS04, Pas03a, BS05].

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 [Can01, PW00], 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 protcol" whose interaction with its environment reflects the desired security properties, and then requires that the analyzed protocol has essentially the same effect on its environment as the idealized protocol. The universal composition theorem essentially guarantees that if a single instance of the analyzed protocol has the same effect on its environment as a single instance of the ideal protocol, then multiple concurrently running instances of the analyzed protocol have the same effect on the environment as multiple concurrent instances of the ideal protocol.

Protocols for realizing practically any idealized protocol (or, functionality) within the UC framework are known, e.g. [CLOS02, KLP07, BCNP04, CDPW07, Kat07, LPV09]. Indeed, these protocols can be used as sketched above to build systems with overall security propeties. 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 we know that many functionalities of interest are impossible to realize in a UC manner without such underlying trust; impossibility holds even if ideally authenticated communication is guaranteed [CF01, CKL03, PR08].

A natural and important question that arises from this fundamental impossibility is how to go about building secure systems that use cryptography — without making 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 the ones mentioned above. However, as argued above, this approach is not helpful for analyzing security of current information systems where some of the components are unknown.

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 secure according to security notion 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? We will use this scenario as a running example throughout the introduction.

When P means "UC security" the answer is positive, regardless of notion Q. But we'd like to make sure that introducing ρ does not adversely affect even security properties that are weaker than

UC. When Q means "UC security" the answer is positive for *any* property P (as long as protocol π is known to have property P alongside the ideal trusted-party based protocol). However, no other notion of security, and in particular no notion that is generally realizable given only authenticated communication, provides such a general guarantee.

An alternative approach, taken in this work, is to find meaningful notions of security that (a) are realizable by protocols that don't use any external trust mechanisms, and (b) guarantee that an admitted protocol adversely affects the external system in the least possible way, or in other words preserves the security of *as many as possible* "protocols in the environment". Previous work that has addressed this issue (either implicitly or explicitly) provides only relatively weak and informal guarantees regarding the "environmental friendliness" of protocols. The present work makes two main contributions: First, we clarify 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 a large class of "protocols in the environment". 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.

Our Idea in a Nutshel At the heart of our work is a new connection between enviromental friendliness and recent results on *black-box unprovability*. More specifically, we present a concurrently secure implementation of a commitment scheme and show how the unprovability results in [Pas11] can be extended to rule out black-box (but also non-unform) proofs of security for the hiding property of this particular commitment scheme, based on a general class of "game-based assumption". Yet, we show—using *non-black-box techniques*, similar to those used in [Bar01]—that the protocol indeed is hiding. (As an independent contribution, as far as we know, this yields the first example a security property that can be proven using a non-black-box security reduction, but for which even *non-uniform* black-box reductions do not exists.) Roughly speaking, we next show that any protocol that exhibits such a "gap" between non-black-box and black-box security proofs cannot harm the security of any protocol whose security is proven secure using a black-box reduction (potenially a non-uniform one) to same class of game-based assumptions.

1.1 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 π is said to *emulate* protocol ϕ is for any adversary $\mathcal{A}$ there exists an adversary $\mathcal{S}$ such that no environment $\mathcal{E}$ can tell whether it is interacting with π and $\mathcal{A}$ or with ϕ and $\mathcal{S}$. Security requirements are then written by way of requiring that π 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 π realizes.

Notions of security differ in the complexity of the different components, and in the way by which the different components interact. We briefly review three such notions, leading to our new notion. Basic security restricts the communication between the environment and the adversary.

[1] In the plain model the parties have ideally authenticated communication channels and no other trusted set-up. Alternatively, using the results in [BCL+05], the results in this work can be translated to the *bare model* where the parties do not even have access to authenticated communication.

Specifically, it only allows the environment to give information to the adversary once at the onset of the protocol and receive information from the adversary once at the end of the interaction. This essentially implies that the only ("protocols in the environment" with respect to which basic security guarantees friendliness are those protocols where there is *no flow of information* between the protocol execution and the rest of the system during the execution of the protocol. (In terms of our running example, if the security property Q that protocol ρ realizes is basic security, then the only protocols π whose security properties are quaranteed to be preserved are those protocols that remain frozen throughout the execution of ρ.) This is indeed a strong restriction.

Super-polynomial security (SPS) [Pas03a, BS05] relaxes UC security in that it allows the adversary $\mathcal{S}$ (often called the *simulator*) to run in time T that is somewhat super-polynomial (say, $T(n) = O(n^{\log n})$). For many functionalities this relaxation still provides meaningful guarantees with respect to the protocol execution itself. However, as far as environmental friendliness is concerned, SPS security is only guaranteed to preserve those properties that are already guaranteed to hold against adversaries that run in time T. Using our running example, this means that if the security notion Q is taken to be SPS security, then protocol π is guaranteed to continue satisfying property $\mathbf{P}$ in the presence of protocol ρ only if protocol π originally satisfies property $\mathbf{P}$ even against adversaries that run in time T. This limitation is especially bothersome since, in order to allow proofs to go through, SPS protocols are often designed so that they are breakable by adversaries whose complexity is T' which is not that much larger than T. This requires making rather precise determination of the hardness of the underlying primitives — so as to make them not too easy and at the same time not too hard.

Another relaxation of UC security is angel-based security [PS04]. 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 [PS04] 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 [PS04, MMY06, CLP10]. Angel-based security has the potential to provide better environmental friendliness guarantees than SPS security, since the super-polynomial advice is restricted in scope. 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 $\mathbf{P}$ 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 [CLP10] 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 $\mathbf{P}$ such that any protocol ρ that securely realizes, with respect to angel $\mathcal{H}$, some functionality $\mathcal{F}$, is environmentally friendly to Π and $\mathbf{P}$.

Defining environmental friendliness. At high level, environmental friendliness is defined as follows. We start by somewhat generalizing the traditional notion of a *cryptographic primitive* [Nao03, RTV04], or more precisely of a *game based* primitive. Originally, the security of a protocol Π (e.g., Blum's 3-round WI protocol) implementing a game based primitive (e.g., witness indistinguishable protocols) is defined via a game between an efficient challenger Chal_Π (that depends on Π, usually by just running it) and an adversary $\mathcal{A}$. Intuitively, a game defines a security property, and a protocol Π satisfies this security property if any efficient adversary $\mathcal{A}$ (i.e., non-uniform polynomial time machines) wins the game with only negligible advantage over some threshold probability. We extend the definition of a game based primitive to allow a primitive to require multiple security properties $\mathbf{P}$—defined via a *set of games* with challengers $\{\mathsf{Chal}^i_\Pi\}_{i\in I}$—and requires that no efficient adversary can win any of the games with any noticeable advantage; we say that such a primitive is *extended-game-based*. The idea here is that the security of a protocol Π in a certain execution environment can be viewed as a security property, and thus represented using a game. A protocol satisfies the overall requirement if it withstands all environments under consideration, represented as a set of games.

Next, we define the following transformation on extended game based primitives. The transformation is parameterized by two protocols ρ and $\pi_\mathcal{G}$, where $\pi_\mathcal{G}$ is thought of as an "ideal protocol" that manifests some ideal functionality $\mathcal{G}$, and ρ is a protocol that "emulates" $\pi_\mathcal{G}$ according to some yet-to-be-specified notion of security. Given a primitive $\mathcal{P}$ with properties $\mathbf{P} = \{\mathsf{Chal}^i_\Pi\}_{i\in I}$, the transformation then views each challenger Chal^i_Π as representing an interaction between Π and an adversary, in the presence of an environment that contains some running instances of the ideal protocol $\pi_\mathcal{G}$. It then adds to $\mathbf{P}$ a new challenger $\mathsf{Chal}^{i,\rho/\pi_\mathcal{G}}_\Pi$ that's identical to Chal^i_Π except that each instance of $\pi_\mathcal{G}$ within Chal^i_Π is replaced with an instance of ρ. We let $\mathbf{P}_{\rho/\pi_\mathcal{G}}$ denote the transformed (or, enhanced) properties. (The exact mechanics of the replacement are detailed within.)

Now, say that the pair $(\rho, \pi_\mathcal{G})$ is *environmentally friendly* to a protocol Π implementing a extended-game-based primitive $\mathcal{P}$ with properties $\mathbf{P}$ (or simply friendly to a pair $(\Pi, \mathbf{P})$), if Π satisfies also the enhanced properties $\mathbf{P}_{\rho/\pi_\mathcal{G}}$. The intuition behind this definition is straightforward: It directly formulates the guarantee that if Π enjoyed security property $\mathbf{P}$ in systems that include instances of the ideal protocol $\pi_\mathcal{G}$ then Π will continue to enjoy security property $\mathbf{P}$ when each instance of $\pi_\mathcal{G}$ is replaced with an instance of ρ.

In the concrete context of angel-based security, a protocol ρ $\mathcal{H}$-emulates protocol $\pi_\mathcal{G}$ if it UC-emulates $\pi_\mathcal{G}$ in the presence of angel (or, helper) $\mathcal{H}$; therefore, an angel defines a notion of secure computation. As discussed before, the goal here is to find a notion, or equivalently an angel $\mathcal{H}$, such that all protocols that $\mathcal{H}$-realizes some ideal functionality is enviromental friendly to a large set of pairs $\{(\Pi, \mathbf{P})\}$; in this case we say that $\mathcal{H}$ is environmentally friendly to $\{(\Pi, \mathbf{P})\}$.

Constructing environmentally friendly angels. Intuitively, in order to show that an angel oracle $\mathcal{H}$ is environmentally friendly to a protocol Π with security property $\mathbf{P}$, one has to show that giving the adversary $\mathcal{A}$ access to $\mathcal{H}$ when interacting with the challengers in $\mathbf{P}$ does not help $\mathcal{A}$ in winning the security game. One way to do that was informally proposed in [CLP10]: The angel $\mathcal{H}$ used in that work was an instantiation of a CCA-secure commitment scheme with the following additional *robustness* property: Any adversary that interacts with $\mathcal{H}$ using a constant number of rounds can be simulated by another adversary that does not interact with $\mathcal{H}$ at all. This was used by [CLP10] 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 we formalize in this work and extend to any

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

This work takes 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 properties $\mathbf{P}$ *as long as the fact that* Π *satisfies* $\mathbf{P}$ *is proven via (potentially non-uniform) black box reduction to a game-based cryptographic hardness assumption* C *with a bounded polynomial number of rounds.*
Here similarly to a game based primitive, a game based (or, falsifiable [Nao03]) cryptographic hardness assumption is defined via a game between a challenger and an adversary. The assumption states that no efficient adversary can win the game with noticiable advantage.

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 non-uniform) black box reduction to hardness assumption that's formulated via 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.

We remark that, 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.

Our techniques. We give a very high level overview of the technical challenges and our methods for solving them. Our starting point is the angel of [CLP10]. At the basis of that angel lies a CCA-commitment scheme S, which is a commitment scheme that remains secure even against an adversary that has access to an oracle D that "breaks" commitments of the adversary's choice by finding the committed values for the adversary. That is, D interacts with the adversary playing the role of the receiver, in multiple and potentially concurrent commitments, and as soon as it completes a successful commitment stage, it gives the committed value to the adversary. Similarly to D, the angel $\mathcal{H}_S$ of [CLP10] and its follow up work [LP12] plays the receiver in commitments, and as soon as it completes a successful commitment it gives the commit value of the commitment to the adversary — as long as the committed string specifies an identity of a party and that party is corrupted (i.e., under the control of the adversary). Using this angel, they construct a commitment scheme which is extractable by a straight line simulator (using the angel), and from this, a protocol that $\mathcal{H}_S$-realizes the ideal commitment functionality $\mathcal{F}_{\mathsf{com}}$. We note that their protocol for $\mathcal{H}_S$-realizing $\mathcal{F}_{\mathsf{com}}$ works generically for any CCA commitment S.

In the CCA commitment of [CLP10], the committer provides the receiver with a commitment c to the message m using a standard perfectly binding commitment scheme, along with $y = f(x)$ for a random x where f is a one way permutation. It then proves to the receiver via multiple witness indistinguishable proofs with special soundness (WISSP) that it holds either x or a decommitment to c. The WISSP proofs are scheduled in a special way (which builds upon the *message scheduling technique* of Dolev, Dwork and Naor [DDN00]) so as to allow proving CCA security.

We modify the [CLP10] construction of CCA commitments so as to obtain a different CCA commitment S' that allows us to demonstrate the environmental friendliness of the resulting angel

$\mathcal{H}_{S'}$. To motivate the modification, we first review our method of demonstrating environmental friendliness.

Consider a protocol Π that instantiates a primitive with properties $\mathbf{P} = \{\mathsf{Chal}^i_\Pi\}_{i\in I}$ via a black box reduction $\mathcal{R}$ to a cryptographic assumption that involves a challenger C. That is, for any $\mathsf{Chal} = \mathsf{Chal}^i_\Pi$ and any polytime adversary $\mathcal{A}$ that wins the game with Chal, the interaction $C \leftrightarrow \mathcal{R}^{\mathcal{A}}$ results in breaking the assumption. To demonstrate environmental friendliness of $\mathcal{H}$, consider protocols ρ and $\pi_{\mathcal{G}}$ such that ρ $\mathcal{H}$-emulates $\pi_{\mathcal{G}}$. We would like to show that the interaction $C \leftrightarrow \mathcal{R}^{\mathcal{A}}$ results in breaking the assumption even when we're only guaranteed that $\mathcal{A}$ wins the game with the modified challenger $\mathsf{Chal}^{\rho/\pi_{\mathcal{G}}}$. (Recall that a priori it might indeed be easier to break the challenger $\mathsf{Chal}^{\rho/\pi_{\mathcal{G}}}$, since it involves replacing the ideal $\pi_{\mathcal{G}}$ with the real ρ, thus potentially giving $\mathcal{A}$ more opportunities to break the security of Π.)

To our help comes the fact that ρ $\mathcal{H}$-emulates $\pi_{\mathcal{G}}$, thus there exists a simulator $\mathcal{S}$ that can essentially translate an $\mathcal{A}$ that breaks $\mathsf{Chal}^{\rho/\pi_{\mathcal{G}}}$ into an $\mathcal{A}'$ that breaks Chal. The properties of $\mathcal{R}$ now imply that the interaction $C \leftrightarrow \mathcal{R}^{\mathcal{A}'}$ results in breaking the assumption. However, we're still not done since the new adversary $\mathcal{A}'$ has access to the super-polynomial $\mathcal{H}$, and so having $\mathcal{A}'$ win the game with C does not imply breaking the underlying hardness assumption.

We bridge this gap by showing that if there exists a reduction $\mathcal{R}$ such that $C \leftrightarrow \mathcal{R}^{\mathcal{A}^{\mathcal{H}}}$ wins the game with C then there is another polytime machine M such that $C \leftrightarrow M$ wins the game as well. In other words, we show that the security of $\mathcal{H}$ cannot be proven via black box reduction to a game with C. To do that, we build on the techniques of [Pas11] that shows a reminiscent result with respect to the impossibility of proving security of sequential witness hiding for unique witness languages via black box reduction to a game based assumption. The idea there is to show that if the reduction $\mathcal{R}$ (playing the prover in a witness hiding protocol) manages to convince the adversary (playing the verifier) in multiple sequential interactions then it is possible to extract the witness from $\mathcal{R}$. We essentially use the same idea here. Indeed, recall that also in the [CLP10] CCA commitment there are multiple proofs happening sequentially.

Here, however, the setting is more complex than in [Pas11] in a number of ways: First, here the proofs all have multiple witnesses (and inherently so, due to the use of witness indistinguishability). Second, here we need to consider also non-uniform reductions, whereas [Pas11] only deals with uniform reductions. To get around this problem, we use techniques developed in a follow-up work [CLMP12] that extends the result of [Pas11] in two separate directions: 1) It extends the uniform separation result for witness hiding protocols to the non-uniform setting and 2) it shows an analogue of the uniform separation result for a commitment scheme of [LPV08]. Building upon their techniques, we solve both of the above two challenges simultaneously.

In more details, we modify the underlying CCA commitment scheme as follow: First the "trapdoor generation mechanism" in [CLP10] (namely having the committer send $y = f(x)$) does not withstand non-uniform reductions—indeed, such reductions might have x as an advice and so resist attempts to extract the committed value from them. We handle this by using a different "trapdoor generation mechanism" that withstands non-uniform reductions. Specifically, we use a mechanism a la [Bar01] where the committer first commits c' to an all-zero string, obtains a random challenge r from the receiver, and then gives a committed proof 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'. 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, and allows for showing that the modified commitment scheme is indeed unprovable via any non-uniform black-box reductions.

Organization. Section 2 contains the basic notations and definitions. In Section 3 we review the definition of UC security with super-polynomial time helpers and previous feasiblity results in this model. In Section 4, we formalize the notion of environmental friendliness. We show in Section 5 that robust CCA secure commitments that have a strong unprovability property leads to environmentally friendly secure computation protocols. Finally, in Section 6, we provide a construction of such strongly unprovable robust CCA secure commitments, which then gives a construction of secure computation protocols satisfying UC security with super-polynomial time helpers that are environmental friendly.

2 Preliminaries

2.1 Notations

Let N denote the set of all positive integers. For any integer $n \in N$, let $[n]$ denote the set $\{1, 2, \ldots, n\}$, and let $\{0,1\}^n$ denote the set of n-bit long strings. We assume familiarity with the basic notions of *Interactive Turing Machines* [GMR89] (ITM for brevity) and interactive *protocols*. Given a pair of ITMs, A and B, we denote by $\langle A(x), B(y)\rangle(z)$ the random variable representing the (joint) outputs of A and B, on common input z and private inputs x and y respectively, and when the random input to each machine is uniformly and independently chosen. We denote by $A[B]$ the interactive Turing machine composed of the ITM A that sends messages externally, while interacting with the ITM B, and A^B the ITM composed of the ITM A while having oracle access to B. We denote by $\mathcal{PPT}$ probabilistic polynomial time Turing machines. Adversaries are modeled as non-uniform machines. All ITMs take (implicitly) as the first input a security parameter 1^n in unary; we use P_n to denote the algorithm P restricted to the security parameter 1^n.

2.2 Witness Relations

We recall the definition of a witness relation for a $\mathcal{NP}$ language [Gol01].

Definition 1 (Witness relation)**.** *A* witness relation *for a language $L \in \mathcal{NP}$ is a binary relation R_L that is polynomially bounded, polynomial time recognizable and characterizes L by $L = \{x : \exists y \, s.t. \, (x,y) \in R_L\}$. We say that a language L has unique witness if for every $x \in L$ there is a unique w such that $(x,w) \in R_L$.*

We say that y is a witness for the membership $x \in L$ if $(x,y) \in R_L$. We will also let $R_L(x)$ denote the set of witnesses for the membership $x \in L$, i.e., $R_L(x) = \{y : (x,y) \in L\}$. In the following, we assume a fixed witness relation R_L for each language $L \in \mathcal{NP}$.

2.3 Indistinguishability

Definition 2 (Computational Indistinguishability)**.** *Let Y be a countable set. Two ensembles $\{A_{n,y}\}_{n\in N, y\in Y}$ and $\{B_{n,y}\}_{n\in N, y\in Y}$ are said to be* computationally indistinguishable *(denoted by $\{A_{n,y}\}_{n\in N, y\in Y} \approx \{B_{n,y}\}_{n\in N, y\in Y}$), if for every $\mathcal{PPT}$ "distinguishing" machine D, there exists a negligible function $\nu(\cdot)$ so that for every $n \in N, y \in Y$:*

$$|\Pr[a \leftarrow A_{n,y} \ : \ D(1^n, y, a) = 1] - \Pr[b \leftarrow B_{n,y} \ : \ D(1^n, y, b) = 1]| < \nu(n)$$

2.4 Interactive Proofs

We use the standard definitions of interactive proofs (and interactive Turing machines) [GMR89] and arguments (a.k.a computationally-sound proofs) [BCC88]. Given a pair of interactive Turing machines, P and V, we denote by $\langle P(w), V\rangle(x)$ the random variable representing the (local) output of V, on common input x, when interacting with machine P with private input w, when the random input to each machine is uniformly and independently chosen.

Definition 3 (Interactive Proof System)**.** *A pair of interactive machines $\langle P, V\rangle$ is called an* interactive proof system *for a language L if there is a negligible function $\nu(\cdot)$ such that the following two conditions hold:*

- Completeness: *For every $x \in L$, and every $w \in R_L(x)$,* $\Pr[\langle P(w), V\rangle(x) = 1] = 1$
- Soundness: *For every $x \notin L$, and every interactive machine B,* $\Pr[\langle B, V\rangle(x) = 1] \leq \nu(|x|)$

In case that the soundness condition is required to hold only with respect to a computationally bounded prover, the pair $\langle P, V\rangle$ is called an interactive argument *system.*

2.5 Witness Indistinguishable Proofs

The notion of *witness indistinguishability* ($\mathcal{WI}$) was introduced by Feige and Shamir in [FS90]. Roughly speaking, an interactive proof is said to be $\mathcal{WI}$ if the verifier's output is "computationally independent" of the witness used by the prover for proving the statement. In this context, we focus on languages $L \in \mathcal{NP}$ with a corresponding witness relation R_L. Namely, we consider interactions in which, on common input x, the prover is given a witness in $R_L(x)$. By saying that the output is computationally independent of the witness, we mean that for any two possible $\mathcal{NP}$-witnesses that could be used by the prover to prove the statement $x \in L$, the corresponding outputs are computationally indistinguishable.

Definition 4 (Witness-indistinguishability)**.** *Let $\langle P, V\rangle$ be an interactive proof system for a language $L \in \mathcal{NP}$. We say that $\langle P, V\rangle$ is* witness-indistinguishable *for R_L, if for every $\mathcal{PPT}$ ITM V^* and for every two sequences $\{w^1_{n,x}\}_{n\in N, x\in L\cap\{0,1\}^n}$ and $\{w^2_{n,x}\}_{n\in N, x\in L\cap\{0,1\}^n}$, such that $w^1_{n,x}, w^2_{n,x} \in R_L(x)$ for every x, the following probability ensembles are computationally indistinguishable.*

- $\{\langle P(w^1_{n,x}), V^*(z)\rangle(x)\}_{n\in N, x\in L\cap\{0,1\}^n, z\in\{0,1\}^*}$
- $\{\langle P(w^2_{n,x}), V^*(z)\rangle(x)\}_{n\in N, x\in L\cap\{0,1\}^n, z\in\{0,1\}^*}$

2.6 Special-sound $\mathcal{WI}$ proofs

A 3-round public-coin interactive proof for the language $L \in \mathcal{NP}$ with witness relation R_L is special-sound with respect to R_L, if for any two accepting transcripts (α, β, γ) and $(\alpha', \beta', \gamma')$ such that the first messages, α, α', are the same but the challenges β, β' are different, there is a deterministic procedure to extract the witness from the two transcripts and runs in polynomial time. In this paper, we rely on 3-round special sound proofs that are also witness indistinguishable 3-round special-sound $\mathcal{WI}$ proofs for languages in $\mathcal{NP}$ can be based on the existence of non-interactive statistically binding commitment schemes, which in turn can be based on one-way permutations [GMW91, FS90, HILL99, Nao91].

2.7 Game-Based Assumptions

We provide the definition of game-based assumptions; our definition is almost identical to the notion of falsifiable assumptions in the literature [Pas11, Nao03, DOP05, HH09, RV10, GW11]), except that we additionally require that there is always a *trivial strategy* that achieve some basic winning probability threshold. Towards this, we first define the notion of a (canonical) security game (or game for short), which will be used in other definition later as well.

Definition 5. *[Security Game] 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 *can make* 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 said to be the* advantage *of the adversary.*

Definition 6. *[Game-Based Assumptions] A* game-based assumption *is simply a security game* $\mathcal{C} = (\mathsf{Chal}, \tau_{\mathcal{C}})$*. We say that assumption* $\mathcal{C}$ *holds if no non-uniform PPT adversary can break* $\mathcal{C}$*.*

Furthermore, we assume that there is always a non-uniform PPT adversary A*, called the* trivial strategy*, such that* A_n *breaks* Chal_n *with probability at least* τ *(without any advantage) for every* $n \in N$*.*

We remark that the requirement that for every game-based assumption, there is always an adversary that achieves the basic winning probability threshold τ is crucial later for showing that non-uniform and uniform black-box security reductions from such an assumption can be "lifted" to work with even relativized adversaries; see Lemma 1. For most falsifiable assumptions in the literature, this is requirement hold: For instance, in bit-guessing games where the threshold is $1/2$, the trivial strategy is to always guess a random bit, and for games with threshold 0, the trivial strategy is simply to abort.

2.8 Cryptographic Primitives

A cryptographic primitive puts forward a syntactical requirement and a security requirement over a set of algorithms performing some task. For example if $\mathcal{P}$ denotes the primitive one-way permutation, then the syntactical requirement requires any algorithm Π instantiating the one-way permutation to compute a permutation, while the security requirement requires the computed permutation to be one-way. We call a (potentially computationally unbounded) oracle Π an *implementation* of the primitive $\mathcal{P}$, if Π satisfies the syntactical requirements of the components of $\mathcal{P}$ (when all composed in one algorithm). We say an algorithm Π (potentially interactive) *efficiently* implements $\mathcal{P}$ if it implements $\mathcal{P}$ and runs in polynomial time. Formally, following the definition in [RTV04], the syntactical requirement is modeled as the set $\mathcal{F}_{\mathcal{P}}$ that consists of all algorithms satisfying the syntax required, and the security requirement is defined as a relation $\mathcal{R}_{\mathcal{P}}$ that, on input an implementation of the primitive and an adversary, decides whether that adversary breaks the implementation.

Definition 7. *[Cryptographic Primitives] A cryptographic primitive $\mathcal{P}$ is a tuple $(\mathcal{F}_\mathcal{P}, \mathcal{R}_\mathcal{P})$ where every $\Pi \in \mathcal{F}_\mathcal{P}$ is a function implementing $\mathcal{P}$, and $\mathcal{R}_\mathcal{P}$ is a relation whose first component is always in $\mathcal{F}_\mathcal{P}$. When $(\Pi, A) \in \mathcal{R}_\mathcal{P}$, we say that the "adversary" A breaks Π (as an implementation of $\mathcal{P}$). We say that an implementation $\Pi \in \mathcal{F}_\mathcal{P}$ is computationally secure if for all non-uniform PPT adversary A, $\mathcal{R}_\mathcal{P}(\Pi, A) = 0$. In this case, we say Π has security property defined through $\mathcal{R}_\mathcal{P}(P, \cdot)$.*

Statistical security can be defined similarly by requiring that the relation holds false for all adversaries. In this work we focus on computational security; in the rest of the paper, by a secure implementation we implicitly refer to a computationally secure implementation.

Remark 1. *We note that the separation between syntactical and security requirement is not inherent, since the former can always be modeled as a part of the latter, by simply considering every algorithm that does not satisfy the syntactical requirement as insecure. However, as discussed in [RTV04], this separation provides the convenience of separating the structural properties that are "easy" to achieve (e.g., the algorithm needs to compute a permutation) from the hardness properties that are "hard" to achieve (e.g., the one-wayness). Here we follow the convention.*

Game-Based Primitives For many natural cryptographic primitives, the security of the primitive can be modeled as game. That is, an implementation Π of such a primitive $\mathcal{P}$ is secure if no non-uniform $\mathcal{PPT}$ adversary can break an game-based assumption $(\mathsf{Chal}_\Pi, \tau_\Pi)$ associated with this implementation Π. For instance, the security of a witness-indistinguishable protocol is defined through a game-based assumption that is bounded-round (the same as the number of round of the WI protocol); the security of signature schemes is defined through an game-based assumption that is not bounded-round, since the adversary is allowed to choose the number of received signatures before trying to forge one. For many natural cryptographic primitives, the game threshold τ_Π only depends on the primitive $\mathcal{P}$ but not the implementation Π; and the threshold usually is either 0 (*e.g.* inverting a one-way function) or 1/2 (*e.g.* distinguishing a PRG from a uniform string).

Definition 8. *[Game-Based Primitives] A cryptographic primitive $\mathcal{P} = (\mathcal{F}_\mathcal{P}, \mathcal{R}_\mathcal{P})$ is game-based if for every $\Pi \in \mathcal{F}_\mathcal{P}$, there is a game-based assumption $(\mathsf{Chal}_\Pi, \tau_\Pi)$ such that $\mathcal{R}_\mathcal{P}(P, A) = 1$ if and only if A breaks the game-based assumption $(\mathsf{Chal}_\Pi, \tau_\Pi)$.*

We note that not all cryptographic primitives are game-based. For example, soundness of a constant-round interactive argument is defined through a game where the challenger is not efficient, and the zero-knowledge property of a protocol is not even defined through a game.

2.9 Commitment Schemes

A commitment scheme $\langle C, R \rangle$ consists of a pair of $\mathcal{PPT}$ ITMs C and R that interact in a commit stage and a reveal stage. In this work, we consider commitment schemes that are perfectly binding and computationally hiding. Furthermore, we restrict our attention to commitment schemes where the reveal phase is non-interactive—the committer decommits to value v by simply sending a decommitment pair (v, d). In fact, in this paper, we will need commitment schemes that have unique decommitment. Such a perfectly binding commitment scheme exists assuming the existence of one-way permutations.

Additionally, we consider tag-based commitment schemes [PR05, DDN00]: A *tag-based commitment scheme* with $l(n)$-bit identities [PR05, DDN00] is a scheme where, in addition to the security parameter 1^n, the committer and the receiver also receive a "tag"—a.k.a. the identity—id of length $l(n)$ as common input.

2.10 CCA-Secure Commitments

Security under chosen-ciphertext-attacks (CCA security) [RS91] has been studied extensively in the context of encryption schemes, where the confidentiality of encrypted messages is guaranteed even in the presence of a decryption oracle. The analogy of CCA security for commitment schemes is the notion of chosen-commitment-attack (CCA) secure commitment schemes introduced by [CLP10]. Roughly speaking, a commitment scheme is CCA if the commitment scheme retains its hiding property even if the receiver has access to an *oracle* that "breaks the security" of other commitments of the receiver's choice. Below we recall the definition of CCA secure commitment scheme in [LP12], which is almost identical to that of [CLP10] except that in the latter the oracle "breaks" other commitments and sends back the corresponding decommitment information, but the definition in the former considers an oracle that returns only the committed value.

Let $\langle C, R\rangle$ be a tag-based commitment scheme with n-bit identities. A committed-value oracle $\mathcal{O}$ of $\langle C, R\rangle$ acts as follows in interaction with an adversary A: it participates with A in many sessions of the commit phase of $\langle C, R\rangle$ as an honest receiver, using identities of length n, chosen adaptively by A. At the end of each session, if the session is *valid*, it reveals the unique committed value of that session to A; otherwise, it sends $\perp$. (If a session has multiple committed values, the decommitment oracle also returns $\perp$. The statistically binding property guarantees that this happens with only negligible probability.)Loosely speaking, a tag-based commitment scheme $\langle C, R\rangle$ is said to be CCA-secure w.r.t. the committed-value oracle, if the hiding property of the commitment holds even with respect to adversaries with access to the committed-value oracle $\mathcal{O}$. More precisely, denote by $A[\mathcal{O}]$ the adversary A with access to the committed-value oracle $\mathcal{O}$. Let $\mathsf{IND}_b(\langle C, R\rangle, A, n, z)$, where $b \in \{0,1\}$, denote the output of the following probabilistic experiment: on common input 1^n and auxiliary input z, $A[\mathcal{O}]$ (adaptively) chooses a pair of challenge values $(v_0,\ v_1) \in \{0,1\}^n$—the values to be committed to—and an identity $\mathsf{id} \in \{0,1\}^n$, and receives a commitment to v_b using identity id. Finally, the experiment outputs the output y of $A[\mathcal{O}]$; the output y is replaced by $\perp$ if during the execution A sends $\mathcal{O}$ any commitment using identity id. (That is, any execution where the adversary queries the decommitment oracle on a commitment using the same identity as the commitment it receives, is considered invalid).

Definition 9. *[CCA-secure Commitments.] Let $\langle C, R\rangle$ be a tag-based commitment scheme. We say that $\langle C, R\rangle$ is* CCA-secure w.r.t. the committed-value oracle*, if for every $\mathcal{PPT}$ ITM A, the following ensembles are computationally indistinguishable:*

- $\{\mathsf{IND}_0(\langle C, R\rangle, A, n, z)\}_{n\in N, z\in\{0,1\}^*}$
- $\{\mathsf{IND}_1(\langle C, R\rangle, A, n, z)\}_{n\in N, z\in\{0,1\}^*}$

2.10.1 k-Robust Commitments

Consider a man-in-the-middle adversary that participates in an *arbitrary* left interaction with a *limited number of rounds*, while having access to a committed oracle.

Definition 10. *Let $\langle C, R\rangle$ be a tag-based commitment scheme. We say that $\langle C, R\rangle$ is* k-robust*, if there exists a $\mathcal{PPT}$ simulator S, such that, if for every $\mathcal{PPT}$ adversary A, and every $\mathcal{PPT}$ k-round ITM B, the following two ensembles are computationally indistinguishable.*

- $\{\mathsf{out}_{B,A[\mathcal{O}]}[\langle B(y), A[\mathcal{O}](z)\rangle(1^n, x)]\}_{n\in N, x,y,z\in\{0,1\}^*}$
- $\{\mathsf{out}_{B,S^A}[\langle B(y), S^{A(z)}\rangle(1^n, x)]\}_{n\in N, x,y,z\in\{0,1\}^*}$

where $\mathsf{out}_{A,B}[\langle B(y), A(z)\rangle(x)]$ *denote the joint output of* A *and* B *in an interaction between them, on common input* x *and private inputs* z *to* A *and* y *to* B *respectively, with uniformly and independently chosen random inputs to each machine.*

Thus, roughly speaking, $\langle C, R\rangle$ is k-robust if the (joint) output of every k-round interaction, with an adversary having access to the oracle $\mathcal{O}$, can be simulated without the oracle. In other words, having access to the oracle does not help the adversary in participating in any k-round protocols much.

We say that a tag-based commitment $\langle C, R\rangle$ is robust if it is k-robust w.r.t. the committed-value oracle for every constant k; we say that $\langle C, R\rangle$ is k-robust CCA secure (or robust CCA secure resp.) if it is both k-robust (or robust resp.) and CCA secure w.r.t. the committed-value oracle.

On the identity length: Recall that in the definition of robust CCA-security, the adversary can pick arbitrary identities id of length n for both the left and the right interaction. We may also consider a restricted notion of robust CCA-security where the adversary is restricted to use identities of some bounded length. As shown in [CLP10, LP12], standard techniques [DDN00] can be used to show that any *robust* CCA-secure commitment that is secure for identities of length $\ell(n) = n^{\varepsilon}$ can be turned into a robust CCA-secure commitment (that is secure for identities of length n) by adding one extra message at the beginning of the protocol: Simply add a stage at the beginning of the protocol where the committer first generates a key pair (sk, vk) for a signature scheme, such that the verification-key vk is of length n^{ε}[2]; the sender then sends vk, and a signature of the n-bit identity (using sk) to the receiver and next runs $\langle C, R\rangle$ with identity vk. Robustness is here needed to argue that the signature scheme is secure even for attackers that have access to the CCA-oracle. We refer the reader to [CLP10, LP12] for a formal proof and omit the proof here.

Proposition 1. *Let* ε *be any constant such that* $0 < \varepsilon < 1$, ℓ *a polynomial such that* $\ell(n) = n^{\varepsilon}$, *and* $\langle C, R\rangle$ *a* k*-round robust CCA-secure commitment scheme with* ℓ*-bit identities. Then assuming the existence of one-way functions, there exists a robust* $k+1$*-round CCA-secure commitment scheme* $\langle \hat{C}, \hat{R}\rangle$ *with* n*-bit identities.*

3 UC with Super-Polynomial Time Helpers

3.1 UC and Global UC security

We briefly review UC and externalized UC (EUC) security. For full details see [Can00, CDPW07]. The original motivation to define EUC security was to capture settings where all ITMs in the system have access to some global, potentially trusted information (such as a globally available public key infrastructure or a bulletin board) [CDPW07]. Here however we use the EUC formalism to capture the notion of global helper functionalities that are available only to the corrupted parties.

We first review the model of computation, ideal protocols, and the general definition of securely realizing an ideal functionality. Next we present hybrid protocols and the composition theorem.

The basic model of execution. Following [GMR89, Gol01], a protocol is represented as an interactive Turing machine (ITM), which represents the program to be run within each participant. Specifically, an ITM has three tapes that can be written to by other ITMs: the input and subroutine

[2]The existence of signature schemes is implied by the existence of one-way functions [Rom90]; to get a signature scheme with a "short" verification-key, simply "scale-down" the security parameter.

output tapes model the inputs from and the outputs to other programs running within the same "entity" (say, the same physical computer), and the incoming communication tapes and outgoing communication tapes model messages received from and to be sent to the network. It also has an identity tape that cannot be written to by the ITM itself. The identity tape contains the program of the ITM (in some standard encoding) plus additional identifying information specified below. Adversarial entities are also modeled as ITMs.

We distinguish between ITMs (which represent static objects, or programs) and *instances of ITMs*, or *ITI*s, that represent interacting processes in a running system. Specifically, an ITI is an ITM along with an identifier that distinguishes it from other ITIs in the same system. The identifier consists of two parts: A session-identifier (SID) which identifies which protocol instance the ITI belongs to, and a party identifier (PID) that distinguishes among the parties in a protocol instance. Typically the PID is also used to associate ITIs with "parties", or clusters, that represent some administrative domains or physical computers.

The model of computation consists of a number of ITIs that can write on each other's tapes in certain ways (specified in the model). The pair (SID,PID) is a unique identifier of the ITI in the system. With one exception (discussed within) we assume that all ITMs are probabilistic polynomial time.[3]

Security of protocols. Protocols that securely carry out a given task (or, protocol problem) are defined in three steps, as follows. First, the process of executing a protocol in an adversarial environment is formalized. Next, an "ideal process" for carrying out the task at hand is formalized. In the ideal process the parties do not communicate with each other. Instead they have access to an "ideal functionality," which is essentially an incorruptible "trusted party" that is programmed to capture the desired functionality of the task at hand. A protocol is said to securely realize an ideal functionality if the process of running the protocol amounts to "emulating" the ideal process for that ideal functionality. Below we overview the model of protocol execution (called the *real-life model*), the ideal process, and the notion of protocol emulation.

The model for protocol execution. The model of computation consists of the parties running an instance of a protocol π, an adversary A that controls the communication among the parties, and an environment Z that controls the inputs to the parties and sees their outputs. We assume that all parties have a security parameter $k \in \mathbf{N}$. (We remark that this is done merely for convenience and is not essential for the model to make sense). The execution consists of a sequence of *activations,* where in each activation a single participant (either Z, A, or some other ITM) is activated, and may write on a tape of at most *one* other participant, subject to the rules below. Once the activation of a participant is complete (i.e., once it enters a special waiting state), the participant whose tape was written on is activated next. (If no such party exists then the environment is activated next.)

The environment is given an external input z and is the first to be activated. In its first activation, the environment invokes the adversary A, providing it with some arbitrary input. In the context of UC security, the environment can from now on invoke (namely, provide input to) only ITMs that consist of a single instance of protocol π. That is, all the ITMs invoked by the environment must have the same SID and the code of π. In the context of EUC security the

[3] An ITM is $\mathcal{PPT}$ if there exists a constant $c > 0$ such that, at any point during its run, the overall number of steps taken by M is at most n^c, where n is the overall number of bits written on the *input tape* of M in this run. In fact, in order to guarantee that the overall protocol execution process is bounded by a polynomial, we define n as the total number of bits written to the input tape of M, *minus the overall number of bits written by M to input tapes of other ITMs*; see [Can01].

environment can in addition invoke an additional ITI that interacts with all parties. We call this ITI the helper functionality, denoted $\mathcal{H}$.

Once the adversary is activated, it may read its own tapes and the outgoing communication tapes of all parties. It may either deliver a message to some party by writing this message on the party's incoming communication tape or report information to Z by writing this information on the subroutine output tape of Z. For simplicity of exposition, in the rest of this paper we assume authenticated communication; that is, the adversary may deliver only messages that were actually sent. (This is however not essential since authentication can be realized via a protocol, given standard authentication infrastructure [Can04].)

Once a protocol party (i.e., an ITI running π) is activated, either due to an input given by the environment or due to a message delivered by the adversary, it follows its code and possibly writes a local output on the subroutine output tape of the environment, or an outgoing message on the adversary's incoming communication tape.

The protocol execution ends when the environment halts. The output of the protocol execution is the output of the environment. Without loss of generality we assume that this output consists of only a single bit.

Let $\text{EXEC}_{\pi,A,Z}(k,z,r)$ denote the output of the environment Z when interacting with parties running protocol π on security parameter k, input z and random input $r = r_Z, r_A, r_1, r_2, ...$ as described above (z and r_Z for Z; r_A for A, r_i for party P_i). Let $\text{EXEC}_{\pi,A,Z}(k,z)$ denote the random variable describing $\text{EXEC}_{\pi,A,Z}(k,z,r)$ when r is uniformly chosen. Let $\text{EXEC}_{\pi,A,Z}$ denote the ensemble $\{\text{EXEC}_{\pi,A,Z}(k,z)\}_{k\in N, z\in\{0,1\}^*}$.

Ideal functionalities and ideal protocols. Security of protocols is defined via comparing the protocol execution to an *ideal protocol* for carrying out the task at hand. A key ingredient in the ideal protocol is the *ideal functionality* that captures the desired functionality, or the specification, of that task. The ideal functionality is modeled as another ITM (representing a "trusted party") that interacts with the parties and the adversary. More specifically, in the ideal protocol for functionality $\mathcal{F}$ all parties simply hand their inputs to an ITI running $\mathcal{F}$. (We will simply call this ITI $\mathcal{F}$. The SID of $\mathcal{F}$ is the same as the SID of the ITIs running the ideal protocol. (the PID of $\mathcal{F}$ is null.)) In addition, $\mathcal{F}$ can interact with the adversary according to its code. Whenever $\mathcal{F}$ outputs a value to a party, the party immediately copies this value to its own output tape. We call the parties in the ideal protocol dummy parties. Let $\pi(\mathcal{F})$ denote the ideal protocol for functionality $\mathcal{F}$.

Securely realizing an ideal functionality. We say that a protocol π *emulates* protocol ϕ if for any adversary A there exists an adversary $\mathcal{S}$ such that no environment Z, on any input, can tell with non-negligible probability whether it is interacting with A and parties running π, or it is interacting with $\mathcal{S}$ and parties running ϕ. This means that, from the point of view of the environment, running protocol π is 'just as good' as interacting with ϕ. We say that π *securely realizes* an ideal functionality $\mathcal{F}$ if it emulates the ideal protocol $\pi(\mathcal{F})$. More precise definitions follow. A distribution ensemble is called *binary* if it consists of distributions over $\{0,1\}$.

Definition 11. *Let π and ϕ be protocols. We say that π* UC-emulates *(resp.,* EUC-emulates*) ϕ if for any adversary A there exists an adversary $\mathcal{S}$ such that for any environment Z that obeys the rules of interaction for UC (resp., EUC) security we have $\text{EXEC}_{\phi,\mathcal{S},Z} \approx \text{EXEC}_{\pi,A,Z}$.*

Definition 12. *Let $\mathcal{F}$ be an ideal functionality and let π be a protocol. We say that π* UC-realizes *(resp.,* EUC-realizes*) $\mathcal{F}$ if π* UC-emulates *(resp.,* EUC-emulates*) the ideal protocol $\pi(\mathcal{F})$.*

Security with dummy adversaries. Consider the adversary $\mathcal{D}$ that simply follows the instructions of the environment. That is, any message coming from one of the ITIs running the protocol is forwarded to the environment, and any input coming from the environment is interpreted as a message to be delivered to the ITI specified in the input. We call this adversary the dummy adversary. A convenient lemma is that UC security with respect to the dummy adversary is equivalent to standard UC security. That is:

Definition 13. *Let π and ϕ be protocols. We say that π* UC-emulates *(resp.,* EUC-emulates*) ϕ* w.r.t the dummy adversary *$\mathcal{D}$ if there exists an adversary $\mathcal{S}$ such that for any environment Z that obeys the rules of interaction for UC (resp., EUC) security we have* $\text{EXEC}_{\phi,\mathcal{S},Z} \approx \text{EXEC}_{\pi,\mathcal{D},Z}$.

Theorem 1. *Let π and ϕ be protocols. Then π* UC-emulates *(resp.,* EUC-emulates*) ϕ if and only if π* UC-emulates *(resp.,* EUC-emulates*) ϕ* with respect to the dummy adversary.

Hybrid protocols. Hybrid protocols are protocols where, in addition to communicating as usual as in the standard model of execution, the parties also have access to (multiple copies of) an ideal functionality. Hybrid protocols represent protocols that use idealizations of underlying primitives, or alternatively make *trust assumptions* on the underlying network. They are also instrumental in stating the universal composition theorem. Specifically, in an $\mathcal{F}$-hybrid protocol (i.e., in a hybrid protocol with access to an ideal functionality $\mathcal{F}$), the parties may give inputs to and receive outputs from an unbounded number of copies of $\mathcal{F}$.

The communication between the parties and each one of the copies of $\mathcal{F}$ mimics the ideal process. That is, giving input to a copy of $\mathcal{F}$ is done by writing the input value on the input tape of that copy. Similarly, each copy of $\mathcal{F}$ writes the output values to the subroutine output tape of the corresponding party. It is stressed that the adversary does not see the interaction between the copies of $\mathcal{F}$ and the honest parties.

The copies of $\mathcal{F}$ are differentiated using their SIDs. All inputs to each copy and all outputs from each copy carry the corresponding SID. The model does not specify how the SIDs are generated, nor does it specify how parties "agree" on the SID of a certain protocol copy that is to be run by them. These tasks are left to the protocol. This convention seems to simplify formulating ideal functionalities, and designing protocols that securely realize them, by freeing the functionality from the need to choose the SIDs and guarantee their uniqueness. In addition, it seems to reflect common practice of protocol design in existing networks.

The definition of a protocol securely realizing an ideal functionality is extended to hybrid protocols in the natural way.

The universal composition operation. We define the universal composition operation and state the universal composition theorem. Let ρ be an $\mathcal{F}$-hybrid protocol, and let π be a protocol that securely realizes $\mathcal{F}$. The composed protocol ρ^π is constructed by modifying the code of each ITM in ρ so that the first message sent to each copy of $\mathcal{F}$ is replaced with an invocation of a new copy of π with fresh random input, with the same SID, and with the contents of that message as input. Each subsequent message to that copy of $\mathcal{F}$ is replaced with an activation of the corresponding copy of π, with the contents of that message given to π as new input. Each output value generated by a copy of π is treated as a message received from the corresponding copy of $\mathcal{F}$. The copy of π will start sending and receiving messages as specified in its code. Notice that if π is a $\mathcal{G}$-hybrid protocol (i.e., ρ uses ideal evaluation calls to some functionality $\mathcal{G}$) then so is ρ^π.

The universal composition theorem. Let $\mathcal{F}$ be an ideal functionality. In its general form, the composition theorem basically says that if π is a protocol that UC-realizes $\mathcal{F}$ (resp., *EUC*-realizes $\mathcal{F}$) then, for any $\mathcal{F}$-hybrid protocol ρ, we have that an execution of the composed protocol ρ^π "emulates" an execution of protocol ρ. That is, for any adversary A there exists a simulator $\mathcal{S}$ such that no environment machine Z can tell with non-negligible probability whether it is interacting with A and protocol ρ^π or with $\mathcal{S}$ and protocol ρ, in a UC (resp., EUC) interaction. As a corollary, we get that if protocol ρ UC-realizes $\mathcal{F}$ (resp., EUC-realizes $\mathcal{F}$), then so does protocol ρ^π.[4]

Theorem 2 (Universal Composition [Can01, CDPW07])**.** *Let $\mathcal{F}$ be an ideal functionality. Let ρ be a $\mathcal{F}$-hybrid protocol, and let π be a protocol that UC-realizes $\mathcal{F}$ (resp., EUC-realizes $\mathcal{F}$). Then protocol ρ^π UC-emulates ρ (resp., EUC-emulates ρ).*

An immediate corollary of this theorem is that if the protocol ρ UC-realizes (resp., EUC-realizes) some functionality $\mathcal{G}$, then so does ρ^π.

3.2 UC Security with Super-polynomial Helpers

We modify the definitions of UC security by giving the corrupted parties access to an external "helper" entity, in a conceptually similar way to [PS04]. This entity, denoted $\mathcal{H}$, is computationally unbounded, and can be thought of as providing the corrupted parties with some judicious help. (As we'll see, this help will be used to assist the simulator to "reverse engineering" the adversary in order to extract relevant information hidden in its communication.)

The definition uses the formalism of EUC security [CDPW07]. Specifically, we model the helper entity as an ITM that is invoked directly by the environment, and that interacts with the environment and the corrupted parties. More formally, let $\mathcal{H}$ be an ITM. An environment Z is called *aided by* $\mathcal{H}$ if: (a) Z invokes a single instance $\mathcal{H}$ immediately after invoking the adversary; (b) As soon as a party (i.e., an ITI) P is corrupted (i.e., P receives a `corrupted` message), Z lets $\mathcal{H}$ know of this fact; (c) $\mathcal{H}$ interacts only with the corrupted parties. Then:

Definition 14. *Let π and ϕ be protocols, and let $\mathcal{H}$ be a helper functionality (i.e., an ITM). We say that π $\mathcal{H}$-EUC-emulates ϕ if for any adversary A there exists an adversary $\mathcal{S}$ such that for any environment Z that's aided by $\mathcal{H}$ we have $\text{EXEC}_{\phi,\mathcal{S},Z} \approx \text{EXEC}_{\pi,A,Z}$.*

The meaningfulness of relativized UC security of course depends on the particular helper ITM in use. Still, it is easy to see that if protocol π $\mathcal{H}$-EUC-emulates protocol ϕ where $\mathcal{H}$ obeys the above rules and runs in time $T(n)$, then π UC-emulates ϕ according to a relaxed notion where the adversary $\mathcal{S}$ can run in time $poly(T(n))$. As noted in the past, for many protocols and ideal functionalities, this relaxed notion of security suffices even when $T(n) = exp(n)$ [Pas03b, PS04, BS05, MMY06].

Universal Composition with super-polynomial helpers. The universal composition theorem generalizes naturally to the case of EUC, even with super-polynomial helper functionalities:

Theorem 3 (universal composition for relativized UC)**.** *Let $\mathcal{F}$ be an ideal functionality, let $\mathcal{H}$ be a helper functionality, let π be an $\mathcal{F}$-hybrid protocol, and let ρ be a protocol that $\mathcal{H}$-EUC-realizes $\mathcal{F}$. Then protocol π^ρ $\mathcal{H}$-EUC-emulates π.*

[4] The universal composition theorem in [Can01] applies only to "subroutine respecting protocols", namely protocols that do not share subroutines with any other protocol in the system. In [CDPW07] the theorem is extended to protocols that share subroutines with arbitrary other protocols, as long as the composed protocol, ρ^π, realizes $\mathcal{F}$ with EUC security.

Proof. The proof of Theorem 3 follows the same steps as the proof of Theorem 2 (see e.g. the proof in [Can00]). The only difference is in the construction of the distinguishing environment Z_π (see there). Recall that Z_π takes an environment Z that distinguishes between an execution of π and an execution of π^ρ, and uses it to distinguish between an execution of ρ and an ideal evaluation of $\mathcal{F}$. For this purpose, Z_π emulates for Z an execution of π^ρ.

Now, in the presence of the helper $\mathcal{H}$, Z_ρ must emulate for Z also the interaction with $\mathcal{H}$. Note that Z_π cannot run $\mathcal{H}$ on its own, since $\mathcal{H}$ may well be super-polynomial in complexity. Instead, Z_π will forward to the external instance of $\mathcal{H}$ each message sent to $\mathcal{H}$ by Z. Similarly, whenever any of the corrupted parties that Z_π locally runs sends a message to $\mathcal{H}$, Z_π externally invokes a party with the same ID and code, corrupts it, and instructs it to send the query to the external instance of $\mathcal{H}$. The responses of $\mathcal{H}$ are handled analogously.

Note that the proof uses the fact that the helper functionality $\mathcal{H}$ does not take messages directly from the adversary. Indeed, Z_π cannot emulate for the external instance of $\mathcal{H}$ messages coming from the adversary. □

3.3 Previous Feasibility Results for UC with Super-Poly Helpers

As shown in [CLP10, LP12], assuming the existence of a robust CCA secure commitment scheme $\langle C, R\rangle$ and the existence a stand-alone semi-honest oblivious transfer (OT) protocol, there exists a super-polynomial time helper functionality $\mathcal{H}$, such that almost all functionalities —more precisely, all well-formed functionalities as defined in [CLOS02]— can be $\mathcal{H}$-EUC-emulated. In complement, in [CLP10] a $O(n^\varepsilon)$-round robust CCA commitment scheme is constructed from any one-way functions using a *non-uniform* security reduction[5]. Roughly speaking, the super-poly time helper $\mathcal{H}$ they considered simply "breaks" commitments of $\langle C, R\rangle$ in the same way as its committed-value oracle $\mathcal{O}$ does, subject to the condition that player P_i in a protocol instance *sid* can only query the functionality on commitments that uses identity (P_i, sid). More precisely, every party P_i in a secure computation can simultaneously engage with $\mathcal{H}$ in multiple sessions of the commit phase of $\langle C, R\rangle$ as a committer using identity P_i, where the functionality simply forwards all the messages internally to the committed-value oracle $\mathcal{O}$, and forwards P_i the committed value returned from $\mathcal{O}$ at the end of each session. See figure 1 for a formal description of the functionality. Clearly this functionality can also be implemented in sub-exponential time. In summary, the results in [CLP10, LP12] showed the following two theorems.

Theorem 4 ([CLP10, LP12])**.** *Let $\langle C, R\rangle$ be a $T(\cdot)$-round $t(\cdot)$-robust CCA secure commitment and $\mathcal{H}$ the helper functionality corresponding to $\langle C, R\rangle$. Assume the existence of a $t(\cdot)^\alpha$-round stand-alone semi-honest secure oblivious transfer protocol, where α is some universal constant in the interval $(0, 1)$. Then, for every well-formed functionality $\mathcal{F}$, there exists a $O(T(\cdot) + t(\cdot))$-round protocol ρ that $\mathcal{H}$-EUC-emulates $\mathcal{F}$.*

Theorem 5 ([CLP10])**.** *Let ℓ be a polynomial. Assume the existence of one-way functions. Then, for every $\delta > 0$, there exists an $O(\max(n^\delta, \ell(n)))$-round $\ell(n)$-robust CCA-secure commitment scheme (where n is the security parameter), with a non-uniform security reduction.*

[5] In fact, the commitment scheme constructed in [CLP10] satisfies the stronger notion of robust CCA security w.r.t. decommitment oracle; since a committed-value oracle can be emulated trivially using a decommitment oracle, the CLP construction is also robust CCA secure w.r.t. the committed-value oracle, which is the notion considered in this paper.

Functionality $\mathcal{H}$

Corrupted Parties: Upon receiving an input (`Corrupt`, P_i, sid) from the environment, record (`Corrupt`, P_i, sid).

Initialization: Upon receiving an input (`Init`,P_i, sid, k) from party P_i in the protocol instance sid, if there is no previously recorded tuple (`Corrupt`, P_i, sid) or there is a previously recorded session (P_i, sid, k), ignore this message; otherwise, initialize a session of $\langle C, R\rangle$ with $\mathcal{O}$ using identity (P_i, sid), and record session (P_i, sid, k).

Accessing $\mathcal{O}$: Upon receiving an input (`Mesg`,P_i, sid, k, m) from party P_i in the protocol instance sid, if there is no previously recorded session (P_i, sid, k), ignore the message; otherwise, forward m to $\mathcal{O}$ in the k^{th} session that uses identity (P_i, sid), obtain a reply m', and return (`Mesg`,P_i, sid, k, m') to P_i.

Figure 1: The ideal functionality $\mathcal{H}$

4 Formalizing Environmental Friendliness

In this section, we define the notion of environmental friendliness, which considers the relation between a secure multi-party computation protocol and a cryptographic implementation—whether the secure multi-party computation protocol is as "friendly" to (or, does not hurt the security of) the cryptographic implementation as the ideal functionality it realizes is (or, does). To this end, we do not consider all cryptographic implementations, instead, only consider implementations of a particular type of primitives—called extended-game-based primitives —which is a natural generalization of game-based primitives.

4.1 Extended-Game-Based Primitives

Game-Based primitive defines the security of an implementation Π through a security game $(\mathsf{Chal}_\Pi, \tau_\Pi)$. A natural generalization is to allow specifying security using a set of security games $\left\{\mathsf{Chal}^i_\Pi, \tau^i_\Pi\right\}_{i\in I_\Pi}$, leading to the notion of extended-game-based primitives.

Definition 15. *[Extended-Game-Based Primitives] A extended-game-based primitive $\mathcal{P} = (\mathcal{F}_\mathcal{P}, \mathcal{R}_\mathcal{P})$ satisfies that for every implementation $\Pi \in \mathcal{F}_\mathcal{P}$, there is a set of games $\left\{\mathsf{Chal}^i_\Pi, \tau^i_\Pi\right\}_{i\in I_\Pi}$, such that, $\mathcal{R}_\mathcal{P}(\Pi, A) = 1$ if and only if there is an $i \in I_\Pi$ such that A breaks the game $(\mathsf{Chal}^i_\Pi, \tau^i_\Pi)$.*

Clearly, the traditional game-based primitives are special cases of extended-game-based primitives. In general, extended-game-based primitive provides more flexibility in defining the security of cryptographic implementations. For instance, it is unclear how to model the witness hiding property of a protocol using a single game, as the hiding property is required to hold for all ensembles $\{\mathcal{D}_n\}_{n\in N}$ of distributions over statements (of a certain length) that are hard on average to find a valid witness of. (If a game allows an adversary to select the distribution it wants to attack at, the challenger can not decided efficiently whether the distribution is hard on average or not.) But, using a family of games, where each game has a challenger that samples from a particular hard-on-average distribution, we can easily define the witness hiding property. Another example is the adaptive soundness property of delegation schemes, where soundness must hold against adversaries proving statements of its choice. It seems that since delegation only considers statements corresponding to deterministic polynomial time computation, its security can be defined using a single game where the challenger first determines on its own if the statement chosen by the adversary is

false or not and then verifies if the adversary succeeds in "cheating". The problem is that since the adversary may choose statement of any polynomial-time complexity, this challenger does not run in any bounded polynomial time. But using a family of games, where each challenger verifies only statements of a bounded polynomial time complexity, the soundness can be defined.

Black-Box Reductions for Extended-Game-Based Primitives We define uniform and non-uniform security reductions for extended-game-based primitives; our definition extends the definition of uniform and non-uniform reductions for game-based primitives in [CLMP12] in a straightforward way.

Definition 16. *[Non-Uniform Security Reductions for Extended-Game-Based Primitives] Let* $\mathcal{P} = (\mathcal{F}_{\mathcal{P}}, \mathcal{R}_{\mathcal{P}})$ *be an extended-game-based primitive, and* Π *an implementation of* $\mathcal{P}$ *whose security is defined via a set of games* $\{(\mathsf{Chal}_i, \tau_i)\}_{i \in I}$.

We say that Π *has a* non-uniform *(black-box) security reduction* R *from a game-based assumption* (Chal, τ)*, if* R *is a non-uniform PPT machine, and there exists a function* $\mathsf{Z} : \{0,1\}^* \to \{0,1\}^*$*, and polynomials* s*,* a*,* m *such that* $m(n) = n^{\Theta(1)}$ *satisfying the following: For every family of* deterministic *circuits* A *that breaks the security of* Π*, that is, there is an* $i \in I$*, and a polynomial* p *such that, for infinitely many* $n \in N$*,*

$$\Pr[A_n \text{ breaks } (\mathsf{Chal}_i)_n] > \tau_i + 1/p(n)$$

the following two conditions hold for every $n \in N$ *for which the above holds:*

1. $z = \mathsf{Z}(A_n)$ *has at most* $s(n \cdot p(n))$ *bits.*
2. $\Pr[z = \mathsf{Z}(A_n) : R^{A_n}(1^{m(n)}, z) \text{ breaks } \mathsf{Chal}_{m(n)}] > \tau + 1/a(n \cdot p(n))$.

Furthermore, a *uniform* security reduction is defined similarly except that $R^{A_n}(1^{m(n)}, 1^{p(n)})$ does not receive any non-uniform advice, but receives $1^{p(n)}$.

We note that in the above definition, the non-uniform reduction only works with deterministic circuits. However, since it is black-box and its behavior is almost independent of the adversary (except from that its advantage and the size of the non-uniform advice it may take depend on the advantage of the adversary), plus the fact that every game-based assumption has a trivial strategy, we can show that even for families of deterministic circuits with access to *randomly chosen oracles*, for instance, circuits with access to the random oracle, there exist a non-uniform reduction $\overline{R}$, (function $\overline{\mathsf{Z}}$ and polynomials $\overline{m}$, $\overline{s}$, $\overline{a}$) that with access to the deterministic circuit aided by a randomized oracle achieves the same as the non-uniform reduction does for deterministic circuits. See the full version for a formal statement and proof. This lemma will be instrumental in the proof later.

Lemma 1. *Let* $\mathcal{P} = (\mathcal{F}_{\mathcal{P}}, \mathcal{R}_{\mathcal{P}})$ *be a extended-game-based primitive, and* Π *an implementation of* $\mathcal{P}$ *whose security is defined through a set of games* $\{(\mathsf{Chal}_i, \tau_i)\}_{i \in I}$*. Assume that* Π *has a non-uniform reduction* R *from* (Chal, τ)*. Then the following holds.*

For every ensemble $\{\mathcal{D}_n\}_n$ *of distributions over all functions that maps* n *bits to* $\mathrm{poly}(n)$ *bits for some polynomial, and every deterministic families of* oracle *circuits* A *satisfying that there is an* $i \in I$ *and a polynomial* p*, such that, for infinitely many* $n \in N$*,*

$$\Pr[f \leftarrow \mathcal{D}_n \ : \ A_n^f \text{ breaks } (\mathsf{Chal}_i)_n] > \tau_i + 1/p(n) \ , \qquad (*)$$

then, there exists a non-uniform PPT machine $\overline{R}$*, a function* $\overline{\mathsf{Z}}$ *and polynomials* $\overline{s}$*,* $\overline{a}$ *and* $\overline{m}$ *where* $\overline{m}(n) = n^{\Theta(1)}$*, such that, the following two conditions hold for every* $n \in N$ *for which the above holds,*

1. *For every* $f_n \in \mathcal{D}_n$, $z = \overline{\mathsf{Z}}(A_n^{f_n})$ *has at most* $\overline{s}(n \cdot p(n))$ *bits, and*

2. $\Pr[f \leftarrow \mathcal{D}_n,\ z = \overline{\mathsf{Z}}(A_n^f) :\ \overline{R}^{A_n^f}(1^{\overline{m}(n)}, z)$ *breaks* $(\mathsf{Chal})_{\overline{m}(n)}] > \tau + 1/\overline{a}(n \cdot p(n))$

Proof. Let R be associated with function Z and polynomials s, a and m as defined in Definition 16. For every $n \in N$ such that (*) holds, it follows from Markov inequality that, for every constant c there is a $(1 - 1/c) \cdot 1/p(n)$ fraction of $f_n \leftarrow \mathcal{D}_n$ satisfying that the probability that $A_n^{f_n}$ breaks $(\mathsf{Chal}_i)_n$ is $\tau_i + 1/cp(n)$; we say such a function f_n is $\mathsf{good}_{1/c}$. Thus, we can consider families of deterministic circuit $\{A_n^{f_n}\}_{n \in N}$ where f_n is $\mathsf{good}_{1/4}$ for every $n \in N$ such that (*) holds; by the definition of non-uniform reduction, we have that

1. $z = \mathsf{Z}(A_n^{f_n})$ has $s(n \cdot 4p(n))$ bits, and

2. $\Pr[z = \mathsf{Z}(A_n) :\ R^{A_n^{f_n}}(1^{m(n)}, z)$ breaks $(\mathsf{Chal})_{m(n)}] > \tau + 1/a(n \cdot 4p(n))$

In other words, the above two conditions hold for every $\mathsf{good}_{1/4}$ f_n. Next we construct $\overline{\mathsf{Z}}$, $\overline{s}$, $\overline{a}$, $\overline{m}$ and $\overline{R}$ as follows:

- Let $\overline{s}(m) = s(4m)$, $\overline{a}(m) = 4m \cdot a(4m)$ and $\overline{m}(n) = m(n)$.
- Let $\overline{\mathsf{Z}}$ be a function such that $\overline{\mathsf{Z}}(A_n^f)$ outputs the first $\overline{s}(n \cdot p(n)) = s(n \cdot 2p(n))$ bits of $\mathsf{Z}(A_n^f)$.
- Let $\overline{R}$ be a machine that with oracle access to $A_n^{f_n}$ and on input $(1^{\overline{m}(n)}, \overline{\mathsf{Z}}(A_n^{f_n}))$, interacts with $\mathsf{Chal}_{\overline{m}(n)}$ as follows: It first estimates the advantage of $A_n^{f_n}$ with error at most $1/100p(n)$, by emulating a sufficiently many executions between $A_n^{f_n}$ and $(\mathsf{Chal}_i)_n$ and taking average of the winning probability of $A_n^{f_n}$ in these executions; this can be done in time polynomial in the running time of $(\mathsf{Chal}_i)_n$ and $p(n)$. Then, if the estimated advantage $p_n^{f_n}$ is larger than $1/3p(n)$, $\overline{R}$ internally runs R by forwarding all its oracle queries to $A_n^{f_n}$ and messages to $\mathsf{Chal}_{\overline{m}(n)}$. Otherwise, $\overline{R}$ runs the trivial strategy that achieves winning probability τ against $\mathsf{Chal}_{\overline{m}(n)}$. (Recall that for every game with threshold τ there is a trivial strategy that wins with that probability; see definition 6.)

To analyze the advantage of $\overline{R}(1^{\overline{m}}, \overline{\mathsf{Z}}(A_n^{f_n}))$ with access to $A_n^{f_n}$ for a randomly sampled $f_n \leftarrow \mathcal{D}_n$, consider the following two cases: First, when the estimated advantage of $A_n^{f_n}$ is larger than $1/3p(n)$, it holds that except with negligible probability, that the actual advantage of $A_n^{f_n}$ is at least $1/4p(n)$; in this case, $\overline{R}$ emulates the execution of R with oracle $A_n^{f_n}$ interacting with $\mathsf{Chal}_{m(n)}$ perfectly, and thus achieves at least advantage $1/a(n \cdot 4p(n))$ (against $\mathsf{Chal}_{\overline{m}(n)}$); therefore,

$$\Pr[f_n \leftarrow \mathcal{D}_n\ :\ \overline{R}^{A_n^f}(1^{\overline{m}(n)}, 1^{p(n)}) \text{ breaks } (\mathsf{Chal})_{\overline{m}(n)} \mid \text{Case 1 occurs }] > \tau + 1/a(n \cdot 4p(n)) - \text{negl}(n)$$

Second, in the case when the estimated advantage is smaller than $1/4p(n)$, $\overline{R}$ does not achieve any advantage, but breaks the challenger with probability at least τ. Then by the fact that there are at least a $1/2p(n)$ fraction of f_n's that are $\mathsf{good}_{1/2}$, and conditioning on a $\mathsf{good}_{1/2}$ f_n is sampled, the probability that case 1 occurs in the execution is overwhelming, we have that the probability that Case 1 occur is at least $1/2p(n) - \text{negl}(n)$. Thus, the overall advantage of $\overline{R}$ is at least $1/2p(n) \cdot 1/a(n \cdot 4p(n)) - \text{negl}(n) > 1/\overline{a}(n \cdot p(n))$.

□

4.2 Implementation of a Functionality

We provide a simple definition for what it means that a multi-party protocol implements a functionality. In the literature of *secure* multi-party computation, there has proposed many different notions of secure computation for what secure implementation of a functionality means, like stand-alone security, UC security, super-polynomial time simulation security, UC with super-polynomial time helper, to name a few. Every notion of secure computation essentially defines a cryptographic primitive for every functionality, as it specifies both the syntactical and security requirements for an implementation of that functionality. Usually, the syntactical requirement are the same across different notions of secure computation, namely, an honest execution of the protocol should compute the functionality correctly. However, the security requirements can be very different at different aspects, for instance, stand-alone security v.s. concurrent security, or polynomial time simulation v.s. super-polynomial time simulation etc. Following the convention of defining the implementation of a cryptographic primitive irrespective of its security, we say that a protocol is an implementation of a functionality if it satisfies only the syntactical requirements, more precisely,

Definition 17. *We say that a protocol ρ is an implementation of a functionality $\mathcal{G}$ if for every possible input, an honest execution of ρ with that input computes $\mathcal{G}$ correctly with probability 1, where the probability is over the randomness in the protocol execution.*

Remark 2. *Beyond following the convention, the reason that we define an implementation of a functionality irrespective of any security requirement is that it provides a "minimal object" as the basis for defining different properties. In some sense, the security requirements (as defined by different notions of secure computation) are different properties of implementations. Similarly, as we will see shortly below, environmental friendliness is also defined as a property of an implementation. The key difference between them is that the former is concerned of how well an implementation "mimics" a functionality in terms of correctness and privacy guarantees, whereas the latter is concerned of "friendliness" towards other cryptographic implementations. Some notion of secure computation, namely, UC security, already includes "friendliness" as a part of its security requirement, in other words, UC security implies environmental friendliness; but, many notions do not, including stand-alone security, concurrent self/general composition, or provide only limited "friendliness", like super-polynomial time simulation, angel-based security and UC with super-polynomial time helpers. Thus we think it is meaningful to defining "friendliness" as a separate property and can then potentially study its relation with different security properties.*

4.3 Formalize Environmental Friendliness

Roughly speaking, we say that a protocol ρ implementing a functionality $\mathcal{G}$ is *environmental friendly* to an implementation Π of a extended-game-based primitive $\mathcal{P}$ if the following holds: Suppose that Π is secure in an *ideal world* where the adversary attacking Π simultaneously participates in many executions of the ideal protocol $\pi_{\mathcal{G}}$ of $\mathcal{G}$ (the inputs in executions of $\pi_{\mathcal{G}}$ and Π may be correlated), then it remains secure even in the *real world* where the ideal protocol is replaced with ρ. Towards formalizing this intuitive requirement, we first define the ideal and real worlds, and the meaning of security in these two worlds, as follows.

Ideal and Real Worlds: Both the ideal and the real worlds follow the same model of protocol execution as defined in UC security except from a few important differences. As in UC, an execution consists of honest parties running an instance of a protocol π, an adversary A that controls the communications between parties, and an environment Z that controls the inputs to the parties and

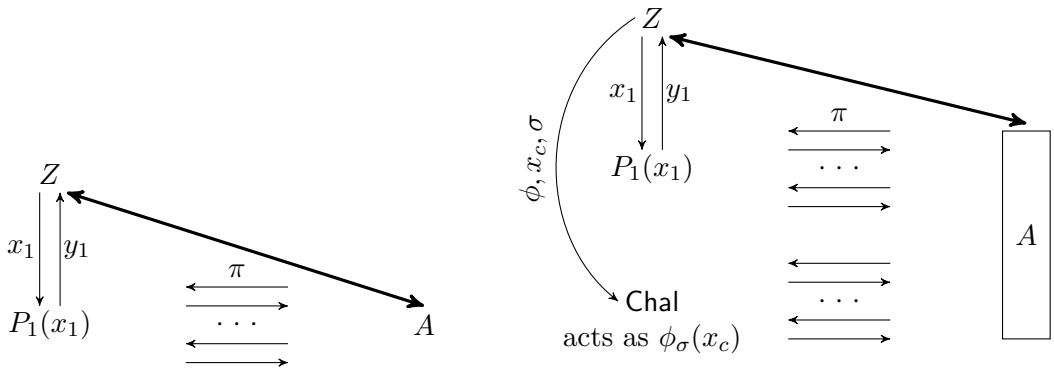

Figure 2: Left: A UC execution of a two-party protocol $\pi = \langle P_1, P_2 \rangle$, where A controls P_2. Right: An execution of experiment $\mathsf{Exp}(\pi, A, Z)$ where the adversary A attacks an implementation Π of a extended-game-based primitive $\mathcal{P}$ while participating in an execution of π; the upper protocol instance is an execution of π between A and the honest player P_1, while the lower instance is the interaction between the adversary and the challenger Chal.

sees their outputs, while communicating with A in an arbitrary way; these entities are invoked in turns and interacts with each other in the same way as in UC. (See Figure 2 on the left for a depiction of the UC model and we refer the reader to Appendix 3.1 for more details on the UC model.) Different from UC, to capture the scenario where the adversary (while participating in an execution of π) simultaneously attacks an implementation Π of $\mathcal{P}$, we introduce an additional entity—a universal challenger Chal—in the model; the environment Z initiates the universal challenger Chal by sending it the code ϕ of an interactive machine, the input x_c and the randomness σ at the beginning of the execution, which later interacts with A acting as challenger $\phi_\sigma(x_c)$. We remark that the fact that Z knows all information (ϕ, x_c, σ) about the challenger means it may leak information about Chal to A through the execution of π by setting the inputs of the honest parties depending on (ϕ, x_c, σ), or via its communication with A. See Figure 2 on the right for depiction of our model of protocol execution. Let $\mathsf{Exp}(\pi, A, Z)$ denote the experiment specified above. Then the ideal world simply consists of an execution of $\mathsf{Exp}(\pi_\mathcal{G}, A, Z)$ running the ideal protocol $\pi_\mathcal{G}$ of the functionality $\mathcal{G}$, and the real world an execution of $\mathsf{Exp}(\rho, A, Z)$ running the protocol ρ implementing $\mathcal{G}$. See Figure 3.

Remark 3. *A drawback of the above definition is that it only allows for a single session of execution of $\pi_\mathcal{G}$ and ρ in the ideal and real worlds respectively, (as the environment only opens a single session of the protocol execution). For some notions of secure computation, like stand-alone security, this definition suffices. For some other notions of secure computation, like UC security, we would like to extend the definition to allow for multiple sessions of execution of $\pi_\mathcal{G}$ and ρ. This can be easily dealt with by considering the multi-session extension $\hat{\mathcal{G}}$ of $\mathcal{G}$, which allows for an arbitrary number of invocations of $\mathcal{G}$. Then, for every notion of secure computation satisfying that the concurrent execution of a protocol ρ that securely implements $\mathcal{G}$, gives a protocol $\hat{\rho}$ that securely implements $\hat{\mathcal{G}}$, we can define the ideal and real worlds to be respectively executions of $\mathsf{Exp}(\pi_{\hat{\mathcal{G}}}, A, Z)$ and $\mathsf{Exp}(\hat{\rho}, A, Z)$. This allows for considering environmental friendliness with respect to the concurrent execution of a secure computation protocol.*

Security in the Ideal and Real Worlds: The security of an implementation Π in experiment $\mathsf{Exp}(\pi, A, Z)$ can be formalized as a game (Chal', τ), where the challenger encompasses all the entities

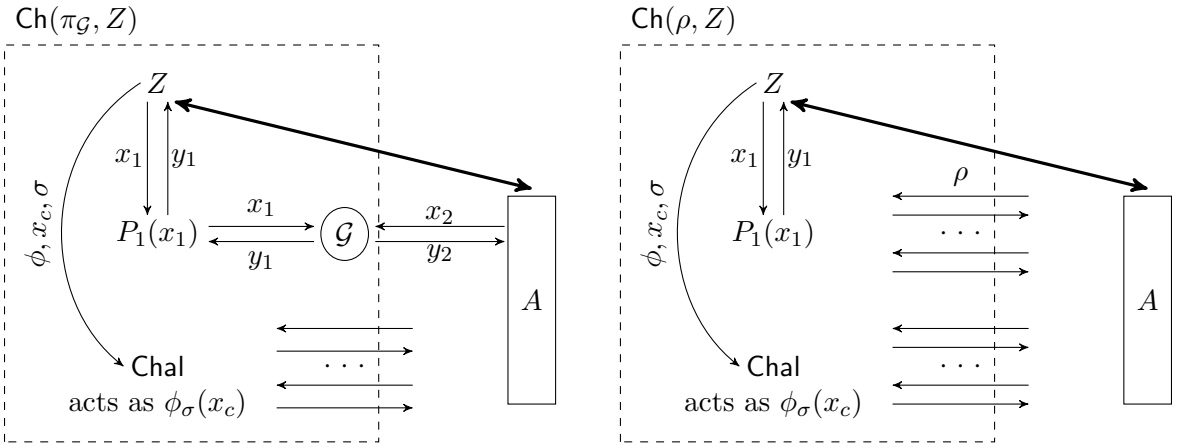

Figure 3: Left: The ideal world execution $\mathsf{Exp}(\pi_{\mathcal{G}}, A, Z)$. Right: The real world execution $\mathsf{Exp}(\rho, A, Z)$. For simplicity, $\mathcal{G}$ is a functionality involving only two players P_1 and P_2. In the pictures, the upper protocol instance is either an execution of $\pi_{\mathcal{G}}$ (on the left) or ρ (on the right) between A and the honest player P_1, while the lower instance is always the interaction between the adversary and the challenger Chal. Additionally, both executions can be formulated as games, by viewing entities in the dashed rectangles as the challenger in the game.

in the experiment except from the adversary A (including Z and honest players of an instance π and the challenger Chal). Formally, let $\mathsf{Ch}(\pi, Z)$ denote the machine that internally emulates the execution of all entities in $\mathsf{Exp}(\pi, A, Z)$ except A; externally it interacts A and forwards messages to/from A from/to the appropriate entities it emulates internally. In Figure 3, we use the dashed rectangles to mark the challengers $\mathsf{Ch}(\pi_{\mathcal{G}}, Z)$ and $\mathsf{Ch}(\rho, Z)$ in the ideal and real worlds respectively. We say that a protocol ρ implementing $\mathcal{G}$ is environmental friendly to Π if for every Z such that Π is secure in the ideal world with Z—that is, the security of Π implies that no adversary can break the game $(\mathsf{Ch}(\pi_{\mathcal{G}}, Z), \tau)$ for some threshold τ—then Π remains secure in real world with Z—that is, no adversary can break the game $(\mathsf{Ch}(\rho, Z), \tau)$ for the same τ, either.

Definition 18. *[Environmental Friendliness]Let* Π *be a secure implementation of a extended-game-based primitive* $\mathcal{P}$ *that has security property defined through a set of games* Δ_Π*. We say that a protocol* ρ *implementing a functionality* $\mathcal{G}$ *is* environmental friendly *to* Π*, if* Π *has (potentially stronger) security property defined through* Δ'_Π *as follows:*

$$\Delta'_\Pi = \Delta_\Pi \cup \{(\mathsf{Ch}(\rho, Z), \tau) \ : \ \forall \ (\mathsf{Ch}(\pi_{\mathcal{G}}, Z), \tau) \in \Delta_\Pi\}$$

5 From Unprovability to Environmental Friendliness

In the literature, a rich body of work has been devoted to investigating the limits of the power of (uniform and non-uniform) black-box security reductions, one of the most commonly used proof techniques. Previous negative results showed that many cryptographic implementations cannot be proven secure using black-box reductions from certain class of assumptions. For instance, it is impossible to prove, using a black-box security reduction, the security of a black-box key-agreement protocol constructed using some OWF from the one-wayness of the same OWF [IR88], nor the soundness of a succinct non-interactive argument system from any game-based assumptions [GW11]. In this work, we change perspective and view unprovability via black-box reduction as a feature of an implementation. We show that in the case of robust CCA secure commitment

schemes, a strong notion of unprovability of the hiding property via black-box reduction actually leads to the construction of secure computation protocols that are environmental friendly to a large class of cryptographic implementations.

5.1 Unprovability of a Commitment Scheme via Black-Box Reductions

The hiding property of a perfectly binding commitment scheme $\langle C, R\rangle$ is unprovable via a uniform or non-uniform reduction from a game-based assumption if the existence of such a reduction implies that the assumption is false. We will consider a very weak notion of hiding property: Roughly speaking, a protocol is said to be weakly hiding if no polynomial time attacker can always recover the committed value after receiving a commitment of $\langle C, R\rangle$.

Breaking Weak Hiding: We say that (a potentially unbounded) A breaks the weak hiding property of a perfectly binding commitment scheme $\langle C, R\rangle$, if for all $n \in N$, A wins the following game $(\mathsf{Chal}, 0)$ with advantage 1: $A(1^n)$ after receiving a commitment of $\langle C, R\rangle$ to a randomly chosen $v \leftarrow \{0,1\}^n$ from Chal_n outputs the value v with probability 1.

Definition 19 (Unprovability via Non-Uniform (Black-Box) Reductions)**.** *We say that a commitment scheme $\langle C, R\rangle$ is* unprovable *via non-uniform reductions from a game-based assumption $C = (\mathsf{Chal}, \tau)$, if the existence of a non-uniform PPT reduction R for basing the weak hiding property of $\langle C, R\rangle$ on C implies the existence of a non-uniform PPT machine S that breaks* Chal.

In this work, we will in fact need a stronger notion of unprovability, which rules out even the possiblity of having a reduction that works solely with a single stylized *ideal adversary* defined below:

Ideal Adversary $\mathcal{A}$ of $\langle C, R\rangle$: In a straight-line interaction, the ideal adversary $\mathcal{A}$ behaves like the committed-value oracle of $\langle C, R\rangle$: It follows the honest receiver strategy and at the end of the commit phase, $\mathcal{A}$ returns the unique committed value v if there is one, otherwise, it returns $\perp$; furthermore, $\mathcal{A}$ also uses fresh randomness in answering any new query *even when rewound* by using its internal random oracle. More formally, $\mathcal{A}_n$ will internally access its own random oracle $RO \leftarrow \mathbf{RO}_n$, and upon any query q as a partial transcript t of a commitment, it does the following:

- If t is a full transcript of a commitment, it returns the unique committed value v if t can only be decommitted to one value; otherwise, it returns $\perp$ if the commitment is invalid or if t can be decommitted to multiple values.
- If t is not a full transcript of a commitment, $\mathcal{A}$ applies RO to t and uses $RO(t)$ as the randomness for generating the next message in response to t according to $\langle C, R\rangle$.

Definition 20. *[Strong Unprovability via Non-Uniform (Black-Box) Reductions] We say that a commitment scheme $\langle C, R\rangle$ is* strongly unprovable *via non-uniform reductions from a set $\mathcal{C}$ of game-based assumptions, if the first condition in the following implies the second condition:*

1. *There exists a non-uniform PPT reduction R that with black-box access to the ideal adversary breaks an assumption $C = (\mathsf{Chal}, \tau) \in \mathcal{C}$: That is, there exists a function Z, and polynomials s, m, and a, such that, 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*
$$\Pr[RO \leftarrow \mathbf{RO}_n,\ z = \mathsf{Z}(A_n^{RO})\ :\ R^{\mathcal{A}_n^{RO}}(1^{m(n)}, z) \text{ breaks } \mathsf{Chal}_{m(n)}] > 1/a(n).$$
2. *There exists a non-uniform PPT machine S that breaks C.*

We say that $\langle C, R\rangle$ is strongly unprovable via non-uniform reductions from *a set* of game-based assumptions $\mathcal{C}$, if it is strongly unprovable from any of the assumption $C \in \mathcal{C}$.

5.2 Environmental Friendly $\mathcal{H}$-EUC-Secure Protocols

Next we show that when a robust CCA secure commitment $\langle C, R\rangle$ is *strongly unprovable* using non-uniform reductions from a class of game-based assumptions $\mathcal{C}$, then, protocols satisfying $\mathcal{H}$-EUC-security with the helper corresponding to the committed-value oracle of $\langle C, R\rangle$ is environmental friendly to a large class of cryptographic implementations: It is environmental friendly to every implementation of extended-game-based primitives whose security is *provable* using a non-uniform reduction from an assumption $C \in \mathcal{C}$, *assuming that C is true*. (Note that in the case that the assumption C is false, even if a non-uniform reduction exists, the implementation may still be insecure, and environmental friendliness to an insecure implementation is not meaningful.) Formally,

Theorem 6. *Let $\mathcal{C}$ be a set of game-based assumptions, $\langle C, R\rangle$ a robust CCA secure commitment that is strongly unprovable using non-uniform security reductions from any assumption in $\mathcal{C}$, and $\mathcal{H}$ the helper functionality corresponding to $\langle C, R\rangle$. Consider any $\mathcal{H}$-EUC-secure protocol ρ realizing some functionality $\mathcal{G}$. For every implementation Π of an extended-game-based primitive that has a non-uniform reduction from an assumption $C \in \mathcal{C}$, the following holds*

- Either*, the assumption C is false (and thus the implementation Π may be insure),*
- Or *ρ (implementing $\mathcal{G}$) is* environmental friendly *to Π.*

Proof. Fix any ρ, $\mathcal{G}$, Π and C as described in the theorem statement; furthermore let Δ be the set of games that defines the security property of Π and R the non-uniform reduction that base the security property (defined by Δ) of Π on the assumption $C = (\mathsf{Chal}', \tau')$. We need to show that assuming that C holds, ρ (implementing $\mathcal{G}$) is environmental friendly to Π, that is, Π actually has the (potentially stronger) security property defined through Δ' as follows:

$$\Delta' = \Delta \cup \{(\mathsf{Ch}(\rho, Z), \tau) \ : \ \forall \ (\mathsf{Ch}(\pi_{\mathcal{G}}, Z), \tau) \in \Delta\} \qquad (*)$$

Let $\mathcal{Z}$ denote the set of PPT machines Z such that $(\mathsf{Ch}(\pi_{\mathcal{G}}, Z), \tau) \in \Delta$. Towards showing (*), it boils down to prove that, given that no efficient adversary can break any game in $\{(\mathsf{Ch}(\pi_{\mathcal{G}}, Z), \tau)\}_{Z\in\mathcal{Z}}$ (due to the security property of Π defined by Δ), no efficient adversary can break any *modified game* in $\{(\mathsf{Ch}(\rho, Z), \tau)\}_{Z\in\mathcal{Z}}$ either. Recall by definition, the game between $\mathsf{Ch}(\pi, Z)$ and an adversary A emulates perfectly an execution of the experiment $\mathsf{Exp}(\pi, Z, A)$. Therefore, it is equivalent to prove that for every $Z \in \mathcal{Z}$, given that Π is secure in the ideal world $\mathsf{Exp}(\pi_{\mathcal{G}}, Z, A)$ (against all efficient adversaries A), Π remains secure even in the real world $\mathsf{Exp}(\rho, Z, A)$. Assume for contradiction that there is a $Z \in \mathcal{Z}$, a polynomial-size adversary A and a polynomial p, such that, for infinitely many security parameters $n \in N$, A_n makes the challenger Chal_n in the real world execution $\mathsf{Exp}(\rho, Z, A)$ (with common input 1^n) accept with probability $\tau + 1/p(n)$ for some polynomial p. We derive a contradiction in the following three steps.

Step 1—By the $\mathcal{H}$-EUC-security of ρ, simulator $S[\mathcal{H}]$ breaks Chal in the ideal world: It follows from the $\mathcal{H}$-EUC-security of ρ that for this adversary A, there is a polynomial-size simulator S such that no environment can distinguish an $\mathcal{H}$-EUC-execution of ρ with A from an execution of $\pi_{\mathcal{G}}$ with the simulator $S[\mathcal{H}]$ having access to the helper functionality $\mathcal{H}$; Therefore for every sufficiently large $n \in N$, the probability that A makes Chal accept in the real world $\mathsf{Exp}(\rho, Z, A)$ is almost the same as that the simulator $S[\mathcal{H}]$ does in the ideal world $\mathsf{Exp}(\rho, Z, S[\mathcal{H}])$, except from a

negligible difference. In other words, $S[\mathcal{H}]$ breaks the game $(\mathsf{Ch}(\pi_{\mathcal{G}}, Z), \tau)$ with advantage at least $1/2p(n)$—for infinitely many $n \in N$, $S_n[\mathcal{H}_n]$ breaks $\mathsf{Ch}(\pi_{\mathcal{G}}, Z)_n$ with advantage at least $1/2p(n)$. [6]

Step 2—By the n.u. reduction $\mathcal{R}$ of Π, there is $\overline{R}$ with access to $S'[\mathcal{A}^{RO}]$ breaks C: We first note that by definition the helper functionality $\mathcal{H}$ can be emulated efficiently using a the committed-value oracle of $\langle C, R\rangle$ (see figure 1) which in turn can be emulated efficiently using the ideal adversary $\mathcal{A}^{RO}$ of $\langle C, R\rangle$. [7]Therefore, there is a polynomial-size simulator S' such that, for infinitely many $n \in N$, $S'_n[\mathcal{A}_n^{RO}]$ breaks $\mathsf{Ch}(\pi_{\mathcal{G}}, Z)_n$ with advantage at least $1/2p(n)$. Furthermore, since S' is a non-uniform PPT machine, it is without loss of generality to assume that it is deterministic (since it can receive as non-uniform advice the best random coins that maximizes its winning probability). Then, since $(\mathsf{Ch}(\pi_{\mathcal{G}}, Z), \tau) \in \Delta$ and Π has a non-uniform PPT reduction R for basing its security on assumption $C = (\mathsf{Chal}', \tau')$, by Lemma 1, there exists a non-uniform PPT machine $\overline{R}$ that with oracle access to $S'[\mathcal{A}^{RO}]$ breaks the underlying assumption C.

Step 3—By the strong unprovability of $\langle C, R\rangle$, there is B breaking C directly: 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 underlying assumption C. Finally, it follows directly from the strong unprovability of $\langle C, R\rangle$ via non-uniform that the existence of R' implies that there is another non-uniform PPT machine S'' that breaks the assumption C directly. This contradicts with the hypothesis that the assumption C holds. For completeness, we provide the construction of R' below.

Construction of R': By Lemma 1, $\overline{R}$ with oracle access to $S'[\mathcal{A}^{RO}]$ breaks the underlying assumption C. More precisely, there is a function $\overline{\mathsf{Z}}$, and polynomials $\overline{s}$, $\overline{a}$ and $\overline{m}$ (s.t., $\overline{m}(n) = n^{\Theta(1)}$) such that, the following two conditions are true.

1. For every $RO \in \mathbf{RO}_n$, $z = \overline{\mathsf{Z}}(S'[\mathcal{A}_n^{RO}])$ has at most $\overline{s}(n \cdot 2p(n))$ bits, and
2. $\Pr[RO \leftarrow \mathbf{RO}_n,\ S'' = S'[\mathcal{A}_n^{RO}],\ z = \overline{\mathsf{Z}}(S'') :\ \overline{R}^{S''}(1^{\overline{m}(n)}, z)$ breaks $(\mathsf{Chal})_{\overline{m}(n)}] > \tau + 1/\overline{a}(n \cdot 2p(n))$

Using $\overline{R}$ we can construct a non-uniform PPT reduction R' for the weak hiding property of $\langle C, R\rangle$ that works the ideal adversary $\mathcal{A}$. We construct R' and its associated function Z', s', m' and a' as follows:

- Let $s'(n) = \overline{s}(n \cdot p(n))$, $m'(n) = \overline{m}(n)$ and $a'(n) = \overline{a}(n \cdot 2p(n))$.
- Let $\mathsf{Z}'(\mathcal{A}_n^{RO}) = \overline{\mathsf{Z}}(S'[\mathcal{A}_n^{RO}])$.
- Let R' be such that, $R'(1^{m'(n)}, \mathsf{Z}'(\mathcal{A}_n^{RO}))$ with oracle access to $\mathcal{A}_n^{RO}$, emulates the execution of $\overline{R}(1^{\overline{m}(n)}, \overline{\mathsf{Z}}(S'[\mathcal{A}_n^{RO}]))$ with oracle access to $S'[\mathcal{A}_n^{RO}]$ by emulating S' internally and forwarding its oracle queries to its own oracle.

[6]Recall that the helper functionality is stateful and its security parameter is always the same to the security parameter used in the whole $\mathcal{H}$-EUC-execution, which is also used by the simulator.

[7]Note that the committed-value oracle can simultaneously accept many sessions of $\langle C, R\rangle$ and uses independent random coins in every session. This behavior can be emulated using $\mathcal{A}^{RO}$, by assuming w.l.o.g. that the commitment scheme $\langle C, R\rangle$ has its first message sent from the committer and includes a session id as a part of its first messages: In an stand alone execution of $\langle C, R\rangle$, the honest committer always sends 0 as the session id; but, when emulating the $committed - value$ oracle using the ideal adversary $\mathcal{A}^{RO}$, different session id's can be used.

Since R' emulates the execution of $\overline{R}$ perfectly, we have that

1. For every $RO \in \mathbf{RO}_n$, $z = \mathsf{Z}'(\mathcal{A}_n^{RO})$ has at most $s'(n)$ bits, and
2. $\Pr[RO \leftarrow \mathbf{RO}_n,\ z = \mathsf{Z}'(\mathcal{A}_n^{RO}) :\ (R')^{\mathcal{A}_n^{RO}}(1^{m'(n)}, z)$ breaks $(\mathsf{Chal})_{m'(n)}] > \tau + 1/a'(n)$

Thus, R' is indeed a non-uniform machine that with access to the ideal adversary $\mathcal{A}^{RO}$ breaks the underlying assumption C. □

6 Achieving Environmental Friendliness

6.1 Warm-Up: From Robustness to Environmental Friendliness

Already in [CLP10] the notion of environmental friendliness was discussed informally: They suggested intuitively that any $\mathcal{H}$-EUC-secure protocol with helper functionality corresponding to a robust CCA-secure commitment scheme is "environmental friendly" to any cryptographic implementation whose security is defined using constant-round games. However, the discussion stayed at the intuition level without any formal definitions or proofs. Using our definition of environmental friendliness, we now formalize and prove their intuition, which follows easily from a simplified version of the proof of Theorem 1.

Theorem 7. *Let λ be any super-constant polynomial, and $\langle C, R\rangle$ a $\lambda(\cdot)$-robust CCA secure commitment and $\mathcal{H}$ the helper functionality corresponding to $\langle C, R\rangle$. Then every $\mathcal{H}$-EUC-secure protocol ρ realizing some functionality $\mathcal{G}$ is environmental friendly to every secure implementation Π of a extended-game-based primitive $\mathcal{P}$ whose security is defined through a set Δ of games of $\lambda(\cdot)$ rounds.*

Proof. The proof follows essentially from Step 1 of the proof of Theorem 1. Consider the above proof at the end of Step 1, which establishes that: Assuming for contradiction there is an adversary A that makes the challenger Chal in the real world accept with probability $\tau + 1/p(n)$, then by applying the $\mathcal{H}$-EUC-security of ρ, there is a simulator $S[\mathcal{H}]$ that breaks challenger $\mathsf{Ch}(\pi_{\mathcal{G}}, Z)$ with advantage at least $1/2p(n)$. Since the helper functionality $\mathcal{H}$ can be emulated perfectly using the committed-value oracle $\mathcal{O}$ of $\langle C, R\rangle$, we have that there is another machine S' that with access to $\mathcal{O}$, that is $S'[\mathcal{O}]$, breaks $\mathsf{Ch}(\pi_{\mathcal{G}}, Z)$ with the same advantage $1/2p(n)$. Furthermore, since $(\mathsf{Ch}(\pi_{\mathcal{G}}, Z)_n, \tau) \in \Delta$ and $\mathsf{Ch}(\pi_{\mathcal{G}}, Z)_n$ interacts in only $\lambda(\cdot)$ rounds, by the $\lambda(\cdot)$-robustness of $\langle C, R\rangle$, there exists machine S'' that without access to $\mathcal{O}$ breaks $\mathsf{Ch}(\pi_{\mathcal{G}}, Z)$ with at least advantage $1/4p(n)$. This contradicts with the security of Π and concludes the proof. □

Then, combining with the construction of $\lambda(\cdot)$-robust CCA secure commitment scheme in [CLP10] (Theorem 5) we have,

Corollary 1. *Let δ be any positive constant, λ a super-constant function, and $\mathcal{H}$ the helper functionality corresponding to $(C, R)_{\lambda,\delta}$. Assume the existence of a $\lambda(\cdot)^{\alpha}$-round stand-alone semi-honest secure oblivious transfer protocol where α is some universal constant in the interval $(0, 1)$, and the existence of one-way permutations.*

Then, for every well-formed functionality $\mathcal{G}$, there exists a $O(n^{\delta} + \lambda(\cdot))$-round protocol ρ that $\mathcal{H}$-EUC-emulates $\mathcal{G}$. Furthermore, ρ is environmental friendly to every secure implementation Π of a extended-game-based primitive $\mathcal{P}$ whose security is defined through a set Δ of games of $\lambda(\cdot)$ rounds.

6.2 Friendliness to Implementations with Non-Uniform Reductions

Towards achieving environmental friendliness, we present a new $\lambda(\cdot)$-robust CCA secure commitment scheme that is strongly unprovable even via *non-uniform* reductions from any $\lambda(\cdot)$-round game-based assmuption; but, in contrast, can be proven hiding using *non-black-box techniques* based on the existence of families of collision resistant hash functions and one-way permutations. The new scheme when combined with the feasiblity result of of [CLP10] (Theorem 4) gives new constructions of $\mathcal{H}$-EUC-secure protocols for almost all functionalities; furthermore, by our main theorem (Theorem 1), the strong unprobability property of the new scheme implies that the new $\mathcal{H}$-EUC-secure protocols are environemental friendly to any implementation that can be proven secure using even non-uniform reductions from any $\lambda(\cdot)$-round game-based assumption.

6.2.1 A New Robust CCA Secure Commitment Scheme $(\tilde{C}, \tilde{R})_{\lambda,\delta}$

As in [CLP10], the new scheme $(\tilde{C}, \tilde{R})_{\lambda,\delta}$ is constructed in two steps: First, a $O(n^\delta + \lambda(n))$-round $\lambda(\cdot)$-robust CCA secure commitment scheme $(\tilde{C}, \tilde{R})^\ell_{\lambda,\delta}$ for $\ell(\cdot)$-bit identities with $\ell = n^\varepsilon$ for some $\delta > \varepsilon > 0$ is constructed, and then, a commitment scheme $(\tilde{C}, \tilde{R})_{\lambda,\delta}$ with the same $\lambda(\cdot)$-robustness and CCA security, but for n-bit identities, is derived by applying Proposition 1.

High-Level Description of $(\tilde{C}, \tilde{R})^\ell_{\lambda,\delta}$: The construction of the scheme $(\tilde{C}, \tilde{R})^\ell_{\lambda,\delta}$ for $\ell(\cdot)$-bit identities is similar to the (slightly modified) CLP protocol. [8] The protocol follows the Feige-Shamir's zero-knowledge protocol paradigm [FS87] for achieving the hiding property, while relying on the *message scheduling technique* of Dolev, Dwork and Naor [DDN00] for achieving CCA security. Following the Feige-Shamir paradigm, a commitment of $(\tilde{C}, \tilde{R})^\ell_{\lambda,\delta}$ has a "trapdoor" embedded inside, which is computationally hidden for a cheating committer, but can be obtained by the security reduction for the hiding property to allow for completing a partial transcript of a commitment without knowing the committed value. The key difference between the CLP protocol and the new protocol lies in how a "trapdoor" is set up in the protocol.

- In the former, the receiver simply sends a random n-bit string at the *beginning* of the protocol, which decides a "trapdoor" that is the pre-image of the string through a one-way permutation. A non-uniform reduction can fix the best coins for a cheating receiver and learn the "trapdoor" determined by the, now fixed, n-bit string from the cheating receiver. However, no a uniform reduction can obtain a "trapdoor". This gap is essential for showing that the CLP protocol is provable via non-uniform reductions, but strongly unprovable via uniform reductions.

- In contrast, our new scheme sets-up a "trapdoor" using standard techniques in the literature of non-black-box ZK protocols [Bar01, BGGL01, PR05, DGS09, GJ10, PRT11, CLP12], and not at the beginning of the protocol but *after* the committer has committed to a value (using a basic commitment scheme). Using non-black-box simulation techniques, a "trapdoor" can be obtained by a reduction knowing the code of the cheating receiver, but not by any, even

[8] We clarify that the original CLP protocol is in fact slightly different from the one reviewed in Section 6.2: The original protocol relies on primitives including one-way functions, a 2-round interactive statistically binding commitment and a 4-round WISSP protocol, whereas the protocol reviewed in Section 6.2 instantiates these primitives with one-way permutations, non-interactive perfectly binding commitment and 3-round WISSP protocol. The reason that we consider this variant of the CLP protocol is that as shown in Section 6.2 this protocol is already strongly unprovable via uniform reductions from bounded-round game-based assumptions, whereas the original CLP protocol based on one-way functions is provable via a uniform reduction.

non-uniform, black-box reductions. This allows for showing that the new scheme is provable via a non-black-box reduction, but strongly unprovable via black-box reductions.

Formal Description of $(\tilde{C}, \tilde{R})^{\ell}_{\lambda,\delta}$: The protocol $(\tilde{C}, \tilde{R})^{\ell}_{\lambda,\delta}$ relies on the following primitives:

- A non-interactive perfectly binding commitment scheme com *with unique decommitment* (that is not only the committed value is unique, so is the decommitment) and a 3-round public-coin WISSP protocol $\langle P_s, V_s \rangle$, both of which exist assuming the existence of one-way permutations.
- Additionally, it also relies on a family of hash functions $\{\mathcal{H}_n\}_n$: to simplify the exposition, we here assume that both com and $\{\mathcal{H}_n\}_n$ are collision resistant against circuits of size $T'(\cdot)$, where $T'(\cdot)$ is "nice" super-polynomial function. As in [BG02], this assumption can be weakened to just collision resistance against polynomial-size circuits by modifying the protocol to use a "good" error-correcting code $\mathtt{ECC}$ (i.e., with constant distance and with polynomial-time encoding and decoding), and replace commitments $\mathsf{com}(h(\cdot))$ with $\mathsf{com}(h(\mathtt{ECC}(\cdot)))$ in the following protocol.
- A 4-round public coin UA argument system as constructed in [BG08], where the length of the messages is upper bounded by n.

An important building block of $(\tilde{C}, \tilde{R})^{\ell}_{\lambda,\delta}$ is a sub-protocol for setting up a "trapdoor" that can be obtained using non-black-box techniques. We present the trapdoor-setting sub-protocol $\langle \mathsf{Sen}, \mathsf{Rec} \rangle$ in Figure 4.

This sub-protocol and its variants have been used extensively in constructions of non-black-box ZK protocols [PR05, PRT11, CLP12], it follows from standard techniques that the protocol has the following two properties: First, it is *sound* in the sense that for every polynomial size adversary A, the probability that A after an interaction with Rec outputs a valid trapdoor is negligible; second, it is *non-black-box simulatable* in the sense that, for every polynomial size cheating receiver A, there is a polynomial size simulator S that outputs a simulated view of A that is indistinguishable from its view in an interaction with Sen, together with a trapdoor. The simulator S uses the following cheating strategy as in [Bar01]: In the first phase, after receiving h from A, it commits to the hash of the code of A and obtains a challenge r; in the second phase, it commits to valid UA prover's messages β, θ (using randomness s_1, s_2) showing that $(h, c, r) \in \Phi$, using the code of A, c, and the appropriate randomness used in generating c as the witness; since indeed $A(c) = r$, it creates an accepting underlying UA transcript, and thus can output $(\beta, \theta, s_1, s_2)$ as a trapdoor. (It follows directly from the hiding property of com that the simulated view of A is indistinguishable from its real view.)

Using the trapdoor-setting sub-protocol $\langle \mathsf{Sen}, \mathsf{Rec} \rangle$, we now turn to describing the protocol $(\tilde{C}, \tilde{R})^{\ell}_{\lambda,\delta}$. Let η be a polynomial such that $\eta(n) = n^{\delta-\varepsilon}$. To commit to a value v, the Committer $\tilde{C}$ and the Receiver $\tilde{R}$, on common input 1^n and the identity $\mathsf{id} \in \{0,1\}^{\ell(n)}$, proceed in the following three stages in the commit phase.

- *Stage 1 (Commitment Message):* The Committer sends a commitment c' to v using com.
- *Stage 2 (Trapdoor-Setting Sub-Protocol):* The Committer and the Receiver interacts using the trapdoor-setting sub-protocol, where the Committer acts as Sen and the Receiver acts as Rec. Let $\mathcal{T}$ be the transcript generated.
- *Stage 3 (Body of Proof):* The Committer proves that *either* c' is a valid commitment to v or $\mathcal{T} \in \Lambda$. This is proved using $4\ell(n)\eta(n) + 8\lambda(n)$ invocations of a 3-round public-coin WISSP

Trapdoor-Setting Sub-Protocol $\langle \mathsf{Sen}, \mathsf{Rec} \rangle$

On common input 1^n, the Sender Sen and the Receiver Rec proceed in two phases:

First phase (A Slot): Sen and Rec exchange the following three messages.

1. Rec chooses a randomly sampled hash function $h \leftarrow \mathbf{H_n}$.
2. Sen sends a commitment c to 0^n using com.
3. Rec replies with a random "challenge" r of length $n + |c|$.

We refer to the last two messages (c, r) a *slot*.

Second Phase (Encrypted UA*):* Sen and Rec exchange the following four messages.

1. Rec sends a randomly sampled n-bit string α.
2. Sen sends a commitment $\hat{\beta}$ to 0^n using com.
3. Rec sends a randomly sampled n-bit string γ.
4. Sen sends a commitment $\hat{\theta}$ to 0^n using com.

We refer to the transcript $(\alpha, \hat{\beta}, \gamma, \hat{\theta})$ an *encrypted* UA. Let β and θ be the values committed to in $\hat{\beta}$ and $\hat{\theta}$; then we refer to $(\alpha, \beta, \gamma, \theta)$ the *underlying* UA.

The Trapdoor Relation: Let $\mathcal{T}$ be a transcript of the sub-protocol. A trapdoor of $\mathcal{T}$ is a witness w of $\mathcal{T}$ for the following language Λ: $\mathcal{T} \in \Lambda$ if there is $w = (\beta, \theta, s_1, s_2)$ such that,

1. $\hat{\beta} = \mathsf{com}(\beta, s_1)$, and $\hat{\theta} = \mathsf{com}(\theta, s_2)$, and
2. $(\alpha, \beta, \gamma, \theta)$ is an accepting UA transcript for showing the membership of (h, c, r) in language $\Phi = \{(h, c, r) \ : \ \exists \ (\Pi, \sigma, s) \text{ s.t. } c = \mathsf{com}(h(\Pi), s) \ \wedge \ |\sigma| < |r| - n \ \wedge \ \Pi(\sigma) = r\}$

Note that since the verification of a UA transcript takes a fixed polynomial time, $\Lambda \in \mathbf{NP}$.

Figure 4: Trapdoor-Setting Sub-Protocol $\langle \mathsf{Sen}, \mathsf{Rec} \rangle$

proof system $\langle P_s, V_s \rangle$ where the verifier challenge has length $3n$[9]. The messages in the first $4\ell(n)\eta(n)$ proofs are scheduled based on the identity id and relies on scheduling pairs of proofs according to schedules design_0 and design_1 depicted in Figure 5. More precisely, the proof stage consist of $\ell(n)+2\lambda(n)$ phases. In phase $i \leq \ell(n)$, the committer provides $\eta(n)$ sequential $\mathsf{design}_{\mathsf{id}_i}$ pairs of proofs, followed by $\eta(n)$ sequential $\mathsf{design}_{1-\mathsf{id}_i}$ pairs of proofs; and in phase $\ell(n) < j \leq \ell(n) + 2\lambda(n)$, the committer provided one design_0 followed by one design_1.

In the reveal phase, the Committer simply decommits to the commitment c'. The Receiver accepts if the decommitment is valid and rejects otherwise. The round complexity of $\langle C, R \rangle$ is $O(\ell(n)\eta(n)+\tau(n))$. Since $\ell(n) = n^{\varepsilon}$ and $\eta(n) = n^{\delta-\varepsilon}$, the protocol has $O(n^{\delta} + \tau(n))$ rounds.

Obtaining $\lambda(\cdot)$-Robust CCA-Secure Commitment $(\tilde{C}, \tilde{R})_{\lambda,\delta}$ for n-bit Identities: Next by Proposition 1, the protocol $(\tilde{C}, \tilde{R})^{\ell}_{\lambda,\delta}$ for identities of length $\ell(n) = n^{\varepsilon}$ can be transformed into a full-fledged $\lambda(\cdot)$-robust CCA secure commitment scheme for identities of length n. The transformation simply adds one-round at the beginning of the protocol where the committer sends a ℓ-bit verification key vk of a signature scheme chosen at random, together with a signature of the

[9] In [CLP10], the length restriction helps the proof of CCA security. Furthermore, as we shall see later on, the length restriction also helps the proof of strong unprovability.

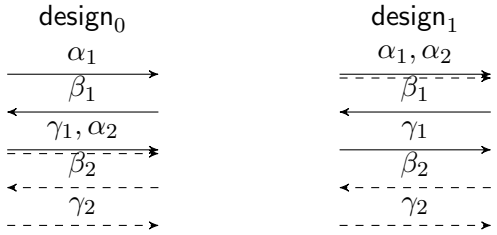

Figure 5: Description of the schedules used in Stage 3 of the protocol. $(\alpha_1, \beta_1, \gamma_1)$ and $(\alpha_2, \beta_2, \gamma_2)$ are respectively the transcripts of a pair of 3-round special-sound proofs.

n-bit identity it receives (using the corresponding signing key), it then commits to a value v using $(\tilde{C}, \tilde{R})^{\ell}_{\lambda,\delta}$ using vk as the identity; we denote the transformed protocol by $(\tilde{C}, \tilde{R})_{\lambda,\delta}$.

Theorem 8. *Fix any constant $\delta > 0$ and polynomial λ. Assuming the existence of families of collision resistant hash functions and one-way permutations, $(\tilde{C}, \tilde{R})_{\lambda,\delta}$ is a $O(n^{\delta} + \lambda(n))$-round perfectly binding $\lambda(\cdot)$-robust CCA secure commitment scheme, with a non-black-box security proof for the hiding property.*

The perfect binding property of the scheme follows directly from perfect binding property of the basic scheme com. Below we provide proofs of the hiding, CCA security and $\lambda(\cdot)$-robustness properties. Note that by Proposition 1, it suffices to prove these properties for the scheme $(\tilde{C}, \tilde{R})^{\ell}_{\lambda,\delta}$.

Proof of Hiding: Assume for contradiction that there is a polynomial-size adversary A that breaks the hiding property $(\tilde{C}, \tilde{R})^{\ell}_{\lambda,\delta}$; that is, for two values v_0, v_1, the adversary A after receiving a commitment to one of the values chosen at random, can guess with inverse polynomial probability which value it has received a commitment to. We demonstrate the existence of a non-black-box reduction R that with access to the code of A breaks the computational hiding property of com. The reduction R, on input 1^n and after receiving a com commitment c' to one of the two values v_b chosen at random, tries to use the adversary A to guess which value it receives a commitment to by forwarding c' to A as Stage 1 of a $(\tilde{C}, \tilde{R})^{\ell}_{\lambda,\delta}$ commitment to v_b and simulating the rest of the commitment as follows: It simulates Stage 2 by using the simulator of the trapdoor-setting sub-protocol $\langle \mathsf{Sen}, \mathsf{Rec} \rangle$ for the adversary A with c' hard-wired in—denote it as $A_{c'}$—obtaining a simulated view *view* together with a trapdoor; Then, it completes the simulation from *view* by cheating in the $\mathcal{WISSP}$ proofs in Stage 3 using the trapdoor as a "fake" witness. It follows from the witness indistinguishability property of the WISSP argument that R has inverse polynomial advantage in guessing whether it received a commitment to v_0 or v_1.

Proof of CCA Security: The CCA security w.r.t. the committed-value oracle $\mathcal{O}$ of $(\tilde{C}, \tilde{R})^{\ell}_{\lambda,\delta}$ follows essentially from the same proof of that of the CLP protocol. Recall that the new scheme differs from the CLP protocol only at the construction of the trapdoor-setting sub-protocol, using protocol $\langle \mathsf{Sen}, \mathsf{Rec} \rangle$ v.s. a random image of a one-way permutation, and the order in which the sub-protocol is executed, before v.s. after the commitment message from the committer. Below we first briefly review the CLP proof; as we will see, it turns out that the proof is completely oblivious of the order in which the trapdoor-setup sub-protocol is executed, and goes through w.r.t. any canonical trapdoor-setup sub-protocol that is *constant-round* and *public-coin*, which is satisfied by $\langle \mathsf{Sen}, \mathsf{Rec} \rangle$. Thus the robust CCA security of $(\tilde{C}, \tilde{R})^{\ell}_{\lambda,\delta}$ follows.

Sketch of the CLP Proof: Recall that proving CCA-security w.r.t. $\mathcal{O}$ amounts to showing that the views of A in experiments IND_0 and IND_1 are indistinguishable (when A has oracle access to $\mathcal{O}$).

Let us refer to the adversary's interaction with C as the *left interaction*, and its interactions with $\mathcal{O}$ as the *right interactions*. The main hurdle in showing the indistinguishability of IND_0 and IND_1 is that the oracle $\mathcal{O}$ is not efficiently computable; if it were, indistinguishability would directly follow from the hiding property of the left interaction. The main idea of the security proof of [CLP10] is then to implement the oracle $\mathcal{O}$ by following the honest receiver strategy to emulate messages belonging to the right commitments and extracting the committed values from the adversary, via "rewinding" the special-sound proofs in the right interactions. The two main technical challenges in simulating the oracle $\mathcal{O}$ are:

- First, once the simulation starts rewinding the right interactions, A might send new messages also in the left interaction. So, if done naively, this would rewind the left interaction, which could violate its hiding property. This problem is solved by relying on the special message scheduling in Stage 3 of the protocol: The message scheduling ensures that for every accepting right interaction with an identity that is different from the left interaction, there exists many points—called $\mathsf{safe\text{-}points}$—in the interaction, from which one can rewind the right interaction without requesting any *new* message in the left interaction.

- Second, in the experiment IND_b, the adversary A expects to receive the committed value at the very moment it completes a commitment to its oracle. If the adversary "nests" its oracle calls, these rewindings become recursive and the running-time of the extraction quickly becomes exponential. To avoid the extraction time from exploding, the simulation strategy in [CLP10] rewinds from $\mathsf{safe\text{-}points}$ using a concurrent extraction strategy that is similar to that used in the context of concurrent ZK by Richardson and Killian [RK99]; it guarantees that for every accepting right interaction (with an identity different from that of the left interaction), one $\mathcal{WISSP}$ proof in that interaction is rewound and a valid witness is extracted.

To complete the proof, it only remains to argue that for every accepting right interaction, the witness extracted from a $\mathcal{WISSP}$ proof is indeed the same committed value that $\mathcal{O}$ would return (except with negligible probability)—in other words, it is not a "trapdoor". For the CLP protocol, this holds, since a "trapdoor" is the pre-image of a random n-bit string through a OWP, and if a simulation strategy were able to output a "trapdoor", one could use it to construct a machine that violates the one-wayness of the OWP.

Handling Canonical Trapdoor-Setting Sub-Protocol: For the new scheme $(\tilde{C}, \tilde{R})^{\ell}_{\lambda,\delta}$, to show that the simulation strategy never outputs a "trapdoor", we rely on the fact that $\langle \mathsf{Sen}, \mathsf{Rec} \rangle$ is sound and canonical (i.e., public-coin and constant-round). First consider the hypothetical scenario that the simulation strategy were straight-line, that is, it does not use rewindings to extract the decommitment of right interactions. Then, since the simulation strategy emulates messages in the right interactions honestly, the probability that it extracts a trapdoor must be negligible, as otherwise, we can construct a machine that violates the soundness of $\langle \mathsf{Sen}, \mathsf{Rec} \rangle$ by simply guessing and forwarding externally the execution of $\langle \mathsf{Sen}, \mathsf{Rec} \rangle$ for which a trapdoor would be extracted; the probability of guess correctly is inverse polynomial since the total number of executions of $\langle \mathsf{Sen}, \mathsf{Rec} \rangle$ is upper bounded by the running time of the simulation, which is a polynomial. Unfortunately, the simulation strategy is not straight-line and heavily relies on rewindings. However, this is not a problem when the sub-protocol is public-coin and constant-round: If the simulation strategy extracts out a trapdoor for one of the executions of $\langle \mathsf{Sen}, \mathsf{Rec} \rangle$ with non-negligible probability, the probability of guessing correctly all messages in that execution is inverse polynomial as it consists of only a constant number of rounds, (and there are at most a polynomial number of messages generated in the whole simulation). Furthermore, as the sub-protocol is public-coin, messages in

the guessed execution can be forwarded out externally even amid rewindings. Therefore, the same argument as in the case of straight-line simulation goes through, and it follows from the soundness of $\langle \mathsf{Sen}, \mathsf{Rec} \rangle$ that simulation strategy never outputs a "trapdoor" except with negligible probability.

Proof of $\lambda(\cdot)$-Robustness: As in [CLP10], the proof of $\lambda(\cdot)$-robustness proceeds almost identically as that for CCA-security, except that, the first challenge discussed above becomes much easier to so solve. This is because now the left interaction has only $\lambda(\cdot)$ rounds, and as long as each right interaction has more than $n^\delta + \lambda(\cdot)$ WISSP proofs, there exists sufficiently many "safe points" from which one can rewind the right interaction without rewinding the left; we omit the details here.

6.2.2 Strong Unprovability of $(\tilde{C}, \tilde{R})_{\lambda,\delta}$

In this section, we show that the protocol $(\tilde{C}, \tilde{R})_{\lambda,\delta}$ is in fact strongly unprovable via any *non-uniform* black-box reductions from any $\lambda(\cdot)$ round game-based assumptions.

Theorem 9. *Fix any constant $\delta > 0$ and polynomial λ. Assume the existence of one-way permutations and families of collision resistant hash functions. $(C, R)_{\lambda,\delta}$ is strongly unprovable via* non-uniform *black-box reductions from any $\lambda(\cdot)$-round game-based assumptions.*

The proof of the theorem relies on the recent uniform separation result of [Pas11]; let us first briefly review this result.

The Uniform Separation Result of [Pas11]: Recently, Pass showed that it is impossible to base the *sequential (weak) witness hiding* property of a *computationally special-sound* interactive argument system for **NP** (that is additionally public-coin and constant-round) on any bounded-round assumption, via *uniform* reductions, assuming that there is a unique witness **NP** language. Recall that our construction of robust CCA secure commitment scheme $(\tilde{C}, \tilde{R})_{\lambda,\delta}$ contains many sequential invocations of a 3-round public-coin (statistical) special-sound ($\mathcal{WI}$) interactive proof system $\langle P_s, V_s \rangle$ for **NP** with $3n$-bit challenge messages. It is easy to see that Pass's result applies to $\langle P_s, V_s \rangle$ when it is used to prove **NP** language L with a unique witness relation R_L; below we formally state the result restricted to the special case of the protocol $\langle P_s, V_s \rangle$ for a unique witness **NP** language L, and bounded-round game-based assumptions [10].

Breaking Weak Witness Hiding of $\langle P_s, V_s \rangle$ w.r.t. R_L: We say that (a potentially unbounded) *A breaks weak $l(\cdot)$-sequential witness hiding of $\langle P_s, V_s \rangle$* for language L, if for every $n \in \mathcal{N}$, $x \in L \cap \{0,1\}^n$, and $w = R_L(x)$, A wins the following game with probability 1: Let $A(x)$ sequentially communicate with $P_s(x, w)$ $l(n)$ times; A is said to win if it outputs the unique witness w after the interaction. $\langle P_s, V_s \rangle$ is called *weakly $l(\cdot)$-sequentially witness hiding* w.r.t. R_L if no polynomial time machine A breaks weak $l(\cdot)$-sequential witness hiding of $\langle P_s, V_s \rangle$ w.r.t R_L.

Theorem 10 (Special Case of [Pas11])**.** *Let L be a* **NP** *language with a unique witness relation R_L, and C a $\lambda(\cdot)$-round game-based assumption where $\lambda(\cdot)$ is a polynomial. If for every polynomial $l(\cdot)$ there exists a* uniform *black-box security reduction R for basing the weak $l(\cdot)$-sequential witness hiding of $\langle P_s, V_s \rangle$ w.r.t R_L on C, then there exists a uniform PPT machine B that breaks C directly.*

The above theorem shows merely the unprovability (via uniform reductions) of the $l(\cdot)$-sequential weak witness hiding property of $\langle P_s, V_s \rangle$ for *all* l from any $\lambda(\cdot)$-round assumption for an arbitrary λ. In fact, the proof of this theorem in [Pas11] established a even stronger result:

[10] The bound-round assumptions considered in [Pas11] includes all assumptions defined through a bound-round game where the challenger may even be inefficient.

1. First, there is a more precise account on the relation between λ and l, in particular, the unprovability from any λ-round assumption holds for *every* sufficiently large l, such that, $l(n) > \lambda(n) + n^\delta$ for some constant $\delta > 0$.

2. Second, the proof showed an analogue of the *strong unprovability* w.r.t. the $l(\cdot)$-sequential witness hiding property, namely, it is impossible to even have a uniform reduction that works solely with an "ideal adversary" for breaking the $l(\cdot)$-sequential witness hiding property of $\langle P_s, V_s \rangle$ w.r.t. R_L (analogous to the ideal adversary for breaking the hiding property of a commitment scheme).

3. Third, the strong unprovability property itself is proved using a uniform black-box reduction M from the special soundness property of $\langle P, V \rangle$ w.r.t. the unique witness relation R_L.

4. Finally, due to the black-box nature of the proof, 1 and 2 holds even with respect to a "benign" type of *non-uniform* reductions that receive an advice that is *independent* of the "ideal adversary" (in particular, independent of the internal random oracle of the "ideal adversary"). This type of non-uniform reduction can be handled since it acts essentially as a uniform one (without any knowledge about the ideal adversary) and is not differentiated by the black-box security reduction M.

IDEAL ADVERSARY $\mathcal{A}$ FOR BREAKING $\ell(\cdot)$-SEQUENTIAL WITNESS HIDING OF $\langle P_s, V_s \rangle$ W.R.T. R_L: The ideal adversary $\mathcal{A}$ is defined similar to the ideal adversary for the hiding property of a commitment scheme. In a straight-line execution, the ideal adversary $\mathcal{A}_n$ first receives $l(n)$ sequential proofs of $\langle P_s, V_s \rangle$ of a statement x; in each proof of $\langle P_s, V_s \rangle$, it generates the random challenge by evaluating its internal random oracle on the partial transcripts of messages it has received so far; at the end of the $l(n)$ proofs, $\mathcal{A}$ returns the unique witness w of x; it returns $\perp$ if x is false. Since $\mathcal{A}$ generates its challenges using the internal random oracle, it always sends independent and random challenges in every proof of $\langle P_s, V_s \rangle$ *even when it is rewound.*

Theorem 11 (Implicit in [Pas11])**.** *Let L be a* **NP** *language with a unique witness relation R_L, and $C = (\mathsf{Chal}, \tau)$ a $\lambda(\cdot)$-round game-based assumption where $\lambda(\cdot)$ is a polynomial. For* every *polynomial $l(\cdot)$ such that $l(n) > \lambda(n) + n^\delta$ for some constant δ, Then, there exists a uniform PPT machine B and a polynomial p, such that, the following holds:*

- *For every non-uniform deterministic machine R and polynomials s, m, a satisfying that for infinitely many $n \in N$ there is a non-uniform advice $z^* \in \{0,1\}^{s(n)}$, such that,*

$$\Pr[RO \leftarrow \mathbf{RO}_n, \ : \ R^{\mathcal{A}_n^{RO}}(1^{m(n)}, z^*) \text{ breaks } \mathsf{Chal}_{m(n)}] > \tau + 1/a(n) \tag{1}$$

  where $\mathcal{A}$ is the ideal adversary for breaking $\ell(\cdot)$-sequential witness hiding of $\langle P_s, V_s \rangle$ w.r.t. R_L,

- *it holds that for infinitely many $n \in N$ for which the above holds,*

$$\Pr[B^{R(1^{m(n)}, z^*)}(1^{m(n)}, 1^{a(n)}) \text{ breaks } \mathsf{Chal}_{m(n)}] > \tau + 1/p(m(n)a(n)) \tag{2}$$

Furthermore, the above statement is proven using a uniform black-box security reduction M from the special-soundness property of $\langle P_s, V_s \rangle$ w.r.t. R_L: That is, there is a polynomial q, such that, For every R, s, m and a w.r.t. which the first condition above holds but not the second, it holds that for every sufficiently large $n \in N$ for which inequality (1) holds but not (2), $M^{R(1^{m(n)}, z^*)}(1^{m(n)}, 1^{a(n)})$ outputs a statement $x \in L \cap \{0,1\}^n$ and two accepting transcripts (α, β, γ) and $(\alpha, \beta', \gamma')$ of $\langle P_s, V_s \rangle$

proof of x, satisfying that $\beta \neq \beta'$, but, the deterministic extractor of $\langle P_s, V_s \rangle$ fails to extract the unique witness of $w = R_L(x)$, with probability $1/q(m \cdot a(n))$. Moreover, the statement x that M outputs is contained in one of the query answers returned from R. [11]

Proof Overview of Strong Unprovability: At first sight, it would seem that the strong unprovability property of $(\tilde{C}, \tilde{R})_{\lambda,\delta}$ directly follows from Theorem 11: At a very high-level, the protocol $(\tilde{C}, \tilde{R})_{\lambda,\delta}$ consists of at least $l(n) = \lambda(n) + n^\delta$ sequentially repeated invocations of $\langle P_s, V_s \rangle$ for the statement that either the message c' of Stage 1 is a valid commitment to v, or a trapdoor is embedded in the transcript $\mathcal{T}$ of Stage 2 (i.e., $\mathcal{T} \in \Lambda$). Additionally, demonstrating hiding, at the very least implies that $\langle P_s, V_s \rangle$ is weakly $l(n)$-sequentially witness hiding (or else, the committed value can be completely recovered!). However, this approach does not go through for two reasons:

- First, Theorem 11 only considers *non-uniform* reductions whose advice is independent of the ideal adversary $\mathcal{A}^{RO}$, whereas we want to show strong unprovability of $(\tilde{C}, \tilde{R})_{\lambda,\delta}$ even for non-uniform reductions who may receive an advice about $\mathcal{A}^{RO}$.
- Second, Theorem 11 only provides a separation in the case where $\langle P_s, V_s \rangle$ is used to prove statements of a *unique witness language*, however, in our protocol $(\tilde{C}, \tilde{R})_{\lambda,\delta}$, the statements proved using $\langle P_s, V_s \rangle$ may admit two witnesses—a valid witness w_1 that is the unique decommitment of c' and a "fake" witness w_2 that is a "trapdoor" embedded in the Stage 2. More precisely, the language under consideration is the following:

$$L' = \left\{ (c', \mathcal{T}) \ : \ \exists w_1 = (v, s) \text{ s.t. } c' = \mathsf{com}(v, s) \ \vee \ \exists w_2 = w' \text{ s.t. } (\mathcal{T}, w') \in R_\Lambda \right\}$$

To get around this problem, we use techniques developed in a follow-up work [CLMP12] that extends the result of [Pas11] in two separate directions: 1) It extends the uniform separation result for witness hiding protocols to the non-uniform setting and 2) it shows an analogue of the uniform separation result for the of non-malleable commitment scheme of [LPV08]. Building upon their techniques, we solve both of the above two challenges simultaneously.

Towards overcoming the first challenge, roughly speaking, we show that any reduction R that receives an advice z *depending* on the ideal adversary $\mathcal{A}^{RO}$ it has oracle access to, can be emulated using a reduction $\overline{R}$ that receives an advice z^* *independent* of $\mathcal{A}^{RO}$. The main difficulty lies in that when R receives z, the conditional entropy of the answers from RO drops, whereas given z^*, all answers from RO are still uniformly random. As in [CLMP12], our approach for getting around this problem is that: although the conditional entropy of RO's answers drops (given z), this happens only for a polynomial number of "bad" queries to RO; the answers to the remaining "good" queries will still have high enough entropy. In fact, by an argument due to Unruh [Unr07], it can be shown that the conditional distribution of answers to "good" queries is statistically close to the original distribution of RO. Thus, if the reduction R together with the ideal adversary $\mathcal{A}$, $R^{\mathcal{A}}$, had only queried these "good" queries to RO, we would already be done (since the answers from RO are information theoretically independent to z, and can be emulated using an another oracle RO' completely independent of z). However, $R^{\mathcal{A}}$ may of course ask also "bad" queries. To deal with this, we construct another reduction $\overline{R}$—$\overline{R}$ receives as a nonuniform advice z^* the set of queries Q from R (to $\mathcal{A}$) that may lead $\mathcal{A}$ to ask a "bad" query to its internal random oracle RO (Q may depend on z), and for each query in Q, $\overline{R}$ additionally receives (as a part of z^*) the answer that $\mathcal{A}^{RO}$ would provide for that query; then, $\overline{R}$ can perfectly emulate the execution of R

[11] As we shall see later in the proof of Claim 1, this property is important in the proof of strong unprovability of $(\tilde{C}, \tilde{R})_{\lambda,\delta}$.

internally, by emulating the answers to queries in Q using its nonuniform advice, and forwarding other queries not in Q to $\mathcal{A}^{RO}$; since $\overline{R}$ never asks any query in Q to $\mathcal{A}$, $\mathcal{A}$ in turn never asks any "bad queries" to RO; therefore the answers from RO are statistically close to random even conditioned on z^*.

Towards overcoming the second challenge, we consider a *unique* witness language L defined over pairs $(c', \mathcal{T})$, such that the only valid witness is the decommitment of c' through com.

$$L = \left\{(c', \mathcal{T}) \ : \ \exists w_1 = (v, s) \text{ s.t. } c' = \mathsf{com}(v, s)\right\}$$

Had the proofs of $\langle P_s, V_s\rangle$ in Stage 3 of the commitment scheme are special-sound w.r.t. R_L, we could have directly applied the technique discussed above and the result of [Pas11] (as stated in Theorem 11) to derive that for every (non-uniform) reduction R that with oracle access to the ideal adversary $\mathcal{A}$ can break the underlying assumption C, there is a meta-reduction B that with black box access to $\overline{R}$ as constructed above can directly break C. Unfortunately, proofs of $\langle P_s, V_s\rangle$ in Stage 3 are not special-sound w.r.t. R_L: For a specific instance $(c', \mathcal{T})$, it might be easy to obtain a "trapdoor" for $\mathcal{T}$, and the value extracted from a $\langle P_s, V_s\rangle$ proof (in Stage 3 of the commitment) might be a trapdoor instead of the unique decommitment—violating the special soundness for R_L. We resolve the problem by relying on the fact that the result of [Pas11] itself is proven using a uniform black-box security reduction M, meaning that if the meta-reduction B with oracle access to $\overline{R}$ does not break C, then M with access to $\overline{R}$ can violate the special-soundness of the $\langle P_s, V_s\rangle$ proofs w.r.t. R_L: That is, it can output a statement $x = (c', \mathcal{T})$ and two accepting transcripts (α, β, γ), $(\alpha, \beta', \gamma')$ for x such that $\beta \neq \beta'$, and the value extracted from the two transcripts is not the unique decommitment of c'. Then, since proofs of $\langle P_s, V_s\rangle$ are indeed special-sound for $R_{L'}$ (as opposed to R_L), the value extracted must be a "trapdoor" for $\mathcal{T}$. Therefore as long as the transcript $\mathcal{T}$ is generated at random by $M^{\overline{R}}$, we can use such an $M^{\overline{R}}$ to violate the soundness of $\langle \mathsf{Sen}, \mathsf{Rec}\rangle$. See the proof of Claim 1 for a detailed proof.

Formal Proof of Theorem 9: To prove the theorem formally, we will rely on the following lemma due to Unruh.

Lemma 2 ([Unr07])**.** *There is an (inefficient) algorithm* Samp *that gets as input some* (z, ξ) *for* $z \in \{0,1\}^*, \xi \in \mathcal{N}$ *and outputs a partial function* F *with* ξ *defined points such that the following holds. If* D *is a (computationally unbounded) oracle algorithm that receives an auxiliary input* z *of length* $|z| = d$ *and asks* t *queries to its oracle, then for any function* $Z\colon \{0,1\}^* \to \{0,1\}^d$ *the view of* D *in the following two experiments is* $\sqrt{dt/2\xi}$*-close in statistical distance:*

1. **(1)** $RO \xleftarrow{\$} \mathbf{RO}_n$. **(2)** $z = z(RO)$. **(3)** *Execute* $D^{RO}(z)$.
2. **(1)** $RO \xleftarrow{\$} \mathbf{RO}_n$. **(2)** $z = z(RO)$. **(3)** $F = \mathsf{Samp}(z, \xi)$. **(4)** $RO' \xleftarrow{\$} \mathbf{RO}_n[F]$ **(5)** *Execute* $D^{RO'}(z)$.

Formal Proof of Theorem 9. Towards showing the strongly unprovability of $(\tilde{C}, \tilde{R})_{\lambda,\delta}$, fix any non-uniform reduction R that with black-box access to the ideal adversary $\mathcal{A}$ of $(\tilde{C}, \tilde{R})_{\lambda,\delta}$ breaks a game-based assumption $C = (\mathsf{Chal}, \tau)$, that is, there is a function Z, and polynomials s, m, a such that for infinitely many $n \in N$, $\mathsf{Z}(\mathcal{A}_n^{RO}) \in \{0,1\}^{s(n)}$ for every $RO \in \mathbf{RO}$, and the following holds:

$$\Pr[RO \xleftarrow{\$} \mathbf{RO}_n \ : \ R^{\mathcal{A}_n^{RO}}(1^{m(n)}, z) \text{ breaks } \mathsf{Chal}_{m(n)}] > 1/a(n)$$

Given R we construct a machine E that breaks C directly. Towards this, we consider a sequence of experiments $\mathcal{E}_0$ to $\mathcal{E}_6$ running different machines. We show that in all experiments the machines

under consideration breaks Chal with non-negligible advantage for infinitely many $n \in N$ for which the above condition holds, and the machine in the last experiment is efficient which gives us the construction of E. Below for convenience, when we say for every (resp. sufficiently large, or infinitely many) $n \in N$, we mean every (resp. sufficiently large, or infinitely many) $n \in N$ for which the above condition holds.

Experiment $\mathcal{E}_0$: $\mathcal{E}_0$ runs the reduction R honestly in the following steps:

(1) $RO \xleftarrow{\$} \mathbf{RO}_n$, $z = \mathsf{Z}(\mathcal{A}_n^{RO})$,

(2) Run $R^{\mathcal{A}_n^{RO}}(1^{m(n)}, z)$ with $\mathsf{Chal}_{m(n)}$.

By our hypothesis, for every $n \in N$, R breaks $\mathsf{Chal}_{m(n)}$ with advantage $1/a(n)$.

Experiment $\mathcal{E}_1$: $\mathcal{E}_1$ proceeds identically to $\mathcal{E}_0$ except that after the non-uniform advice z is computed, the random oracle is re-sampled as in Lemma 2. More precisely, let $t(n)$ be an upper bound on the number of queries that $R^{\mathcal{A}_n^{RO}}$ makes to the random oracle RO in $\mathcal{E}_1$, and $\xi(n)$ be such that $\sqrt{s(n)t(n)/2\xi(n)} \leq 1/2a(n)$. Then the experiemnt takes the following steps:

(1) $RO \xleftarrow{\$} \mathbf{RO}_n$, $z = \mathsf{Z}(\mathcal{A}_n^{RO})$, $F = \mathsf{Samp}(z, \xi(n))$.

(2) $RO_1 \xleftarrow{\$} \mathbf{RO}_n[F]$.

(3) Run $R^{\mathcal{A}_n^{RO_1}}(1^{m(n)}, z)$ with $\mathsf{Chal}_{m(n)}$.

It follows directly from Lemma 2, that for every $n \in N$, R breaks $\mathsf{Chal}_{m(n)}$ with advantage $1/2a(n)$.

Experiment $\mathcal{E}_2$: $\mathcal{E}_2$ proceeds identically to $\mathcal{E}_1$ except that instead of running the reduction R that requires oracle access to the ideal adversary $\mathcal{A}$ of the commitment scheme $(\tilde{C}, \tilde{R})_{\lambda,\delta}$, it runs a reduction S that works with the ideal adversary $\tilde{\mathcal{A}}$ for breaking the l-sequential witness hiding property of $\langle P_s, V_s \rangle$ w.r.t. the unique witness relation R_L (Recall that $w = R_L((c', \mathcal{T}))$ is the unique decommitment of c') for $l(n) = n^\delta + \lambda(n)$; the reduction S takes as a randomized advice $y = (z, F, W)$, where z and F are sampled from the same distribution as in experiment $\mathcal{E}_1$, and the set W contains all tuples (c', v) such that v is the value committed to in c' using com and there is a partial transcript $q = (m_1, m_2, \cdots, m_i)$ satisfying that c' is contained in m_1 and q is answered in F (i.e., $F(q)$ is defined); let W be the function that computes W from F. Then the experiment $\mathcal{E}_2$ proceeds as follows:

(1) $RO \xleftarrow{\$} \mathbf{RO}_n$, $z = \mathsf{Z}(\mathcal{A}_n^{RO})$, $F = \mathsf{Samp}(z, \xi(n))$, $W = \mathsf{W}(F)$, $y = (z, F, W)$.

(2) $RO_1 \xleftarrow{\$} \mathbf{RO}_n[F]$.

(3) Run $S^{\tilde{\mathcal{A}}_n^{RO_1}}(1^{m(n)}, y)$ with $\mathsf{Chal}_{m(n)}$.

The reduction $S(1^{m(n)}, y)$, with oracle access to $\tilde{\mathcal{A}}_n^{RO_1}$, internally emulates the execution of $R(1^{m(n)}, z)$ by forwarding all its messages to $\mathsf{Chal}_{m(n)}$ externally to the challenger it is interacting with and emulating the ideal adversary $\mathcal{A}_n^{RO_1}$ of the commitment scheme $(\tilde{C}, \tilde{R})_{\lambda,\delta}$ for R as follows: Whenever R sends an oracle query q, it parses q as a partial transcript $q = (m_1, m_2, \cdots m_i)$ of the committer's messages in a commitment of $(\tilde{C}, \tilde{R})_{\lambda,\delta}$. It is without loss of generality to assume that before querying q, R must have queried all prefixes of q (as otherwise, one can always construct another reduction R' that observes this restriction). Then,

1. If q is already answered by F, return $F(q)$ as the answer to R.
2. Else if some prefix of q is already answered by F, that is, there is $j < i$ such that $(m_1, \cdots, m_j)$ is answered by F, emulate the answer internally as follows: If after q, a receiver's message is expected, toss fresh random coins and use it as an answer; otherwise, if after q, the committed value is expected, return the committed value v corresponding to the commitment c' contained in m_1 found in W.
3. Otherwise, if neither q itself nor any of its prefixes are answered by F,
   (a) If in the partial transcript corresponding to q, Stage 2 of the commitment of $(\tilde{C}, \tilde{R})_{\lambda,\delta}$ is not completed yet, toss fresh random coins and use it as an answer, (that is, receiver's messages in Stage 1 and 2 of $(\tilde{C}, \tilde{R})_{\lambda,\delta}$ are internally emulated);
   (b) Else if Stage 2 is completed, then let c' be the com commitment in Stage 1 and $\mathcal{T}$ the transcript of Stage 2 corresponding to q[12], and $m'_1, \cdots, m'_j$ the rest of the messages in q, which belongs to $\langle P_s, V_s \rangle$ proofs of the statement $x = (c', \mathcal{T})$ in Stage 3 of a commitment of $(\tilde{C}, \tilde{R})_{\lambda,\delta}$. If m'_j belongs to the first $l(n)$ *sequential* proofs of x in the commitment, forward $q' = (x, m'_1, \cdots, m'_j)$ to the ideal adversary $\tilde{\mathcal{A}}$ (that is, messages in the first $l(n)$ sequential proofs of $\langle P_s, V_s \rangle$ are forwarded out to the ideal adversary $\tilde{\mathcal{A}}$); if the answer from $\tilde{\mathcal{A}}$ is a random string, use it as an answer to R; otherwise if the answer is the unique witness w of x according to R_L, which by definition is a decommiment of c' to some value v, record v. In the other case where m'_j does not belong to one of the first $l(n)$ sequential proofs, simply return fresh random coins.
   (c) Otherwise, if a committed value is expected, previously there must have been $l(n)$ sequential proofs forwarded to $\tilde{\mathcal{A}}$ and a committed value v obtained; then use v as the answer.

We note that S emulates the ideal adversary $\mathcal{A}^{RO_1}$ of the commitment scheme perfectly for R: If a query q is not determined by the pre-sampled part F of the random oracle RO_1, the query is either answered randomly as expected (by tossing internal random coins or returning the random string returned by $\tilde{\mathcal{A}}$), or answered using the unique committed value (found in W or obtained from $\tilde{\mathcal{A}}$). Therefore S emulates the execution of R perfectly, and we have that, for every $n \in N$, S breaks $\mathsf{Chal}_{m(n)}$ with advantage $1/2a(n)$.

Experiment $\mathcal{E}_3$: $\mathcal{E}_3$ proceeds identically to $\mathcal{E}_2$ except that instead of sampling RO_1 respecting the pre-sampled part F, it simply samples it randomly. That is, the experiment proceeds as follows:

(1) $RO \xleftarrow{\$} \mathbf{RO}_n$, $z = \mathsf{Z}(\mathcal{A}_n^{RO})$, $F = \mathsf{Samp}(z, \xi(n))$, $W = \mathsf{W}(F)$, $y = (z, F, W)$.

(2) $RO_2 \xleftarrow{\$} \mathbf{RO}_n$.

(3) Run $S^{\tilde{\mathcal{A}}_n^{RO_2}}(1^{m(n)}, y)$ with $\mathsf{Chal}_{m(n)}$.

We argue that the execution of S in $\mathcal{E}_3$ proceeds identically to that in $\mathcal{E}_2$: Note that by construction S never asks any query q to $\tilde{\mathcal{A}}$ (and thus in turn to its internal random oracle) that is answered by F; therefore queries from S are answered randomly by both RO_1 and RO_2. Thus we have that the views of S are the same in $\mathcal{E}_2$ and $\mathcal{E}_3$, and for every $n \in N$, S breaks $\mathsf{Chal}_{m(n)}$ with advantage $1/2a(n)$.

[12]The receiver's messages in $\mathcal{T}$ are not contained in q, but generated before by S as in step (a).

Experiment $\mathcal{E}_4$: $\mathcal{E}_4$ proceeds identically to $\mathcal{E}_3$ except that, instead of sampling the non-uniform advice y for the reduction S at random in Step 1, fix the best advice y^* that maximizes the winning probability of S in Step 3 of $\mathcal{E}_3$; furthermore, modify S so that whenever it needs to toss some random coins, it instead uses a (separate) random oracle to generate them deterministically. More precisely, Let $\hat{S}$ be a non-uniform deterministic machine that proceeds identically to S, except that, $\hat{S}$ always receives a fixed advice y^*, and it additionally has access to a random oracle RO_3 so that whenever it needs some random coins, it evaluates the random oracle RO_3 on the partial transcript of messages it has obtained from the external challenger $\mathsf{Chal}_{m(n)}$, the ideal adversary $\tilde{\mathcal{A}}^{RO_2}$, as well as the reduction R that it emulates internally, and use the output as the random coins. (As we shall see later in the proof of Claim 1 that it is important that $\hat{S}$ derives its random coins depending on the messages from R it emulates internally). Then the experiment proceeds as follow:

(1) Fix the best y^*.

(2) $RO_2 \xleftarrow{\$} \mathbf{RO}_n$, $RO_3 \xleftarrow{\$} \mathbf{RO}_n$.

(3) Run $\hat{S}^{RO_3,\tilde{\mathcal{A}}_n^{RO_2}}(1^{m(n)}, y^*)$ with $\mathsf{Chal}_{m(n)}$.

It follow from the fact that y^* is fixed to maximize the winning probability of S and the outputs of the random oracle RO_3 are uniformly random, we have that $\hat{S}$ achieves at least the same advantage as S; thus for every $n \in N$, $\hat{S}$ breaks $\mathsf{Chal}_{m(n)}$ with advantage $1/2a(n)$.

Experiment $\mathcal{E}_5$: $\mathcal{E}_5$ proceeds identically to $\mathcal{E}_4$ except that, instead of running $\hat{S}^{RO_3}$ with access to the ideal adversary $\tilde{\mathcal{A}}$, it runs a meta-reduction $\overline{B}$ with oracle access to $\hat{S}^{RO_3}$, which is constructed using the meta-reduction B provided by Theorem 11. Then, the experiment $\mathcal{E}_5$ proceeds as follow:

(1) Fix the best y^*.

(2) $RO_3 \xleftarrow{\$} \mathbf{RO}_n$.

(3) Run $\overline{B}^{\hat{S}^{RO_3}(1^{m(n)},y^*)}(1^{m(n)}, 1^{4a(n)})$ with $\mathsf{Chal}_{m(n)}$.

Before describing the construction of $\overline{B}$, we first establish some basic facts about the winning probability of B: For every $n \in N$ and $RO_3 \in \mathbf{RO}_n$, let $\mathsf{Ad}(n, RO_3)$ denote the advantage of $B(1^{m(n)}, 1^{4a(n)})$ with oracle access to $\hat{S}^{RO_3}(1^{m(n)}, y^*)$ when interacting with $\mathsf{Chal}_{m(n)}$. We will show later that the following claim holds.

Claim 1. *For every infinitely many $n \in N$, with probability at least $1/8a(n)$ over the choice of the random oracle $RO_3 \xleftarrow{\$} \mathbf{RO}_n$, the advantage $\mathsf{Ad}(n, RO_3)$ of B is at least $1/p(m(n) \cdot 4a(n))$.*

Given this claim, we construct the meta-reduction $\overline{B}$ as follows. Let $v(n) = p(m(n) \cdot 4a(n))$. On input $(1^{m(n)}, 1^{4a(n)})$ and with oracle access to $\hat{S}^{RO_3}(1^{m(n)}, y^*)$, $\overline{B}$ first estimate the advantage $\mathsf{Ad}(n, RO_3)$ of B with error at most $1/100v(n)$, by emulating a sufficiently large number of interactions between B (with input $(1^{m(n)}, 1^{4a(n)})$ and oracle access to $\hat{S}^{RO_3}(1^{m(n)}, y^*)$) and $\mathsf{Chal}_{m(n)}$ (also emulated internally) and taking average of the advantage of B in these executions; this can be done in time polynomial in the running time of $\mathsf{Chal}_{m(n)}$ and B and $v(n)$. Then, if the estimated advantage is larger than $3/4v(n)$, $\overline{B}$ runs B to attack the external challenger $\mathsf{Chal}_{m(n)}$; otherwise, $\overline{B}$ runs the trivial strategy that achieves winning probability τ against $\mathsf{Chal}_{m(n)}$. We claim that:

Claim 2. *There is a polynomial v', such that, for every infinitely many $n \in N$, the meta-reduction $\overline{B}$ in $\mathcal{E}_5$ breaks $\mathsf{Chal}_{m(n)}$ with advantage at least $1/v'(n)$.*

Proof. The proof of the claim follows similarly to the analysis of the advantage of $\overline{R}$ in Lemma 1. Consider the following two cases:

- First, if the estimated advantage of B is larger than $3/4v(n)$, it holds that except with negligible probability the actual advantage $\mathsf{Ad}(n, RO_3)$ of B is at least $1/2v(n)$; in this case, $\overline{R}$ emulates the execution of B interacting with $\mathsf{Chal}_{m(n)}$ perfectly, and achieves advantage at least $1/2v(n)$; thus,

$$\Pr[RO_3 \leftarrow \mathbf{RO}_n \ : \ \overline{B}^{\hat{S}^{RO_3}(1^{m(n)},y^*)}(1^{m(n)}, 1^{4a(n))}) \text{ breaks } \mathsf{Chal}_{m(n)} \mid \text{Case 1 occurs }] > \tau + 1/2v(n) - \text{negl}(n)$$

- Second, if the estimated advantage is smaller than $3/4v(n)$, $\overline{R}$ plays the trivial strategy and thus breaks the challenger with probability at least τ.

As shown in Claim 1, for infinitely many $n \in N$, with probability at least $1/8a(n)$ over the choice of RO_3, the advantage $\mathsf{Ad}(n, RO_3)$ of B conditioned on this RO_3 is sampled is at least $1/v(n)$; whenever such a RO_3 is sampled, with overwhelming probability the advantage estimated by $\overline{B}$ is greater than $3/4v(n)$; therefore, the first case occurs with probability at least $1/8a(n) - \text{negl}(n)$. Thus, the overall advantage of $\overline{B}$ is at least $1/8a(n) \cdot 1/2v(n) - \text{negl}(n) > 1/17a(n)v(n) = 1/v'(n)$. □

Experiment $\mathcal{E}_6$: $\mathcal{E}_6$ proceeds identically to $\mathcal{E}_5$ except that, instead of sampling the oracle RO_3 at random, it runs a machine E that on input $(1^{m(n)}, y^*)$, emulates the execution of $\overline{B}$ with oracle access to $\hat{S}^{RO_3}$ perfectly, by running $\overline{B}(1^{m(n)}, 1^{4a(n)})$ and $\hat{S}(1^{m(n)}, y^*)$ internally, and emulating RO_3 through lazy evaluation. That is, the experiment proceeds as follows:

(1) Fix the best y^*.

(2) Run $E(1^{m(n)}, y^*)$ with $\mathsf{Chal}_{m(n)}$.

It is easy to see that E is efficient and satisfies that for every sufficiently large $n \in N$, E breaks $\mathsf{Chal}_{m(n)}$ with advantage $1/v'(n)$.

□

Proof of Claim 1. Assume for contradiction that for sufficiently large $n \in N$, with probability greater than $1-1/8a(n)$ over the choice of the random oracle $RO_3 \xleftarrow{\$} \mathbf{RO}_n$, the advantage $\mathsf{Ad}(n, RO_3)$ of B is less than $1/p(m(n) \cdot 4a(n))$. In this case, we show that we can construct a machine violating the soundness of the trapdoor setting sub-protocol $\langle \mathsf{Sen}, \mathsf{Rec} \rangle$.

We first observe that since in $\mathcal{E}_4$ it holds that (for every $n \in N$) $\hat{S}$ breaks $\mathsf{Chal}_{m(n)}$ with advantage $1/2a(n)$; therefore, with probability at least $1/4a(n)$ over the choice of RO_3, $\hat{S}$ breaks $\mathsf{Chal}_{m(n)}$ with advantage $1/4a(n)$. Then by our hypothesis, we have that for every sufficiently large $n \in N$, with probability at least $1/8a(n)$ over the choice of RO_3, the following two condition holds:

- $\hat{S}^{RO_3, \tilde{A}^{RO_2}}(1^{m(n)}, y^*)$ breaks $\mathsf{Chal}_{m(n)}$ with advantage at least $1/4a(n)$, but
- $B^{\hat{S}^{RO_3}(1^{m(n)},y^*)}(1^{m(n)}, 1^{4a(n)})$ breaks $\mathsf{Chal}_{m(n)}$ with advantage smaller than $1/p(m(n) \cdot 4a(n))$.

Then, for every RO_3 for which the above holds, the uniform security reduction M provided by Theorem 11 guarantees the following:

- $M^{\hat{S}^{RO_3}(1^{m(n)},y^*)}(1^{m(n)}, 1^{4a(n)})$ violates the special soundness of $\langle P_s, V_s\rangle$ w.r.t. R_L with probability $1/q(m(n) \cdot 4a(n))$, that is, it outputs a statement $x = (c', \mathcal{T}) \in L \cap \{0,1\}^n$ and two accepting transcripts (α, β, γ) and $(\alpha, \beta', \gamma')$ of $\langle P_s, V_s\rangle$ proof of x, satisfying that $\beta \neq \beta'$, but, the deterministic extractor of $\langle P_s, V_s\rangle$ fails to extract the unique witness of $w = R_L(x)$, that is, the unique decommitment of c', with probability $1/q(m \cdot 4a(n))$. Moreover, the statement $x = (c', \mathcal{T})$ that M outputs is contained in one of the query answers returned from its oracle $\hat{S}^{RO_3}(1^{m(n)}, y^*)$.

Therefore, with probability $1/u(n) = 1/8a(n) \cdot q(m \cdot 4a(n))$, it holds that after an execution of M with a randomly sampled RO_3, producing $x = (c', \mathcal{T})$ and two transcripts, the value extracted from the two transcripts is not the unique decommitment of c'.

Recall that by construction, the reduction $\hat{S}^{RO_3}(1^{m(n)}, y^*)$ (which proceeds similarly to the reduction S in $\mathcal{E}_3$) sends queries q' of the form $(x, m'_1, \cdot, m'_i)$ containing statements $x = (c', \mathcal{T})$ as a query to the ideal adversary $\tilde{\mathcal{A}}$ (now emulated by M); in the execution of $\hat{S}$, these statements $(c', \mathcal{T})$ are generated as follows (see the description of S in experiment $\mathcal{E}_3$ for more details): All committer's messages c' and $m_1, \cdots, m_k$ in $\mathcal{T}$ are generated by the reduction R emulated internally by $\hat{S}$ *in sequence*, and all receiver's messages $a_1, \cdots, a_k$ in $\mathcal{T}$ are generated by evaluating the random oracle RO_3 on the partial transcript of messages that $\hat{S}$ has received from the external challenger $\mathsf{Chal}_{m(n)}$, its ideal adversary $\tilde{\mathcal{A}}$, and the reduction R that it emulates internally; therefore, in particular, a_i is the evaluation of RO_3 on a string t containing $m_1, \cdots, m_{i-1}$. Notice that $(a_1, m_1, \cdots, a_k, m_k)$ is a transcript of the trapdoor-setting sub-protocol $\langle \mathsf{Sen}, \mathsf{Rec}\rangle$. We now construct a machine D that violates the soundness of $\langle \mathsf{Sen}, \mathsf{Rec}\rangle$.

The machine D on input 1^n, interacts with the honest receiver Rec externally, and internally emulates an execution of $M^{\hat{S}^{RO_3}(1^{m(n)},y^*)}(1^{m(n)}, 1^{4a(n)})$ by emulating RO_3 for M using lazy evaluation, except the following: During the execution of M, D tries to guess which sender's messages $m_1, \cdots, m_k$ generated by R will be output as a part of the statement x that M outputs at the end, and forwards them externally to Rec; for every reply a'_i to m_i from Rec, D sets $a_i = a'_i$ as the answer to the corresponding query from $\hat{S}$ for generating the reply a_i to R (instead of generating a_i using the random oracle emulated using lazy evaluation). We note that the fact that $\hat{S}$ generates a_i by querying a string t containing all previous sender's messages $m_1, \cdots, m_{i-1}$ guarantees that $\hat{S}$ would never ask for a_i before all $m_1, \cdots, m_{i-1}$ are generated, and thus D always succeeds in obtaining a'_i from Rec before a_i is requested since it can always forward $m_1, \cdots, m_{i-1}$ externally first. Finally, M outputs $x = (c', \mathcal{T})$ and the two transcripts (α, β, γ), $(\alpha, \beta', \gamma')$. D checks whether it has guessed correctly all the sender's message in x, if not it aborts; otherwise, it extracts a value o from the two transcripts using the extractor of $\langle P_s, V_s\rangle$, and outputs o if it is a "trapdoor" for the transcript $\mathcal{T}$, and aborts otherwise.

Now we argue that D after an interaction with Rec outputs a valid trapdoor with non-negligible probability. First, we note that D emulates the execution of M perfectly since all answers from RO_3 are emulated perfectly through either lazy evaluation or forwarding from the external Rec which sends public coins according to the protocol $\langle \mathsf{Sen}, \mathsf{Rec}\rangle$; therefore with probability at least $1/u(n)$, the value o extracted at the end is not a decommitment of c'. Then it follows from the actual special-soundness property of $\langle P_s, V_s\rangle$ w.r.t. $R_{L'}$ that except with negligible probability, o must be a "trapdoor". Furthermore, since the protocol $\langle \mathsf{Sen}, \mathsf{Rec}\rangle$ is constant-round, the probability that D guesses all the sender's messages $m_1, \cdots, m_k$ contained in x with some non-negligible probability $1/v(n)$. Therefore with probability at least $1/2u(n)v(n)$, the value o extracted is a valid "trapdoor"

for the interaction with Rec. Therefore D violates the soundness of $\langle \mathsf{Sen}, \mathsf{Rec} \rangle$ and this gives a contradiction. □

References

[Bar01] Boaz Barak. How to go beyond the black-box simulation barrier. In *FOCS*, volume 0, pages 106–115, 2001.

[BCC88] Gilles Brassard, David Chaum, and Claude Crépeau. Minimum disclosure proofs of knowledge. *J. Comput. Syst. Sci.*, 37(2):156–189, 1988.

[BCL+05] Boaz Barak, Ran Canetti, Yehuda Lindell, Rafael Pass, and Tal Rabin. Secure computation without authentication. In *CRYPTO*, pages 361–377, 2005.

[BCNP04] Boaz Barak, Ran Canetti, Jesper Buus Nielsen, and Rafael Pass. Universally composable protocols with relaxed set-up assumptions. In *FOCS*, pages 186–195, 2004.

[Bea91] Donald Beaver. Foundations of secure interactive computing. In *CRYPTO*, pages 377–391, 1991.

[BG02] Boaz Barak and Oded Goldreich. Universal arguments and their applications. In *Computational Complexity*, pages 162–171, 2002.

[BG08] Boaz Barak and Oded Goldreich. Universal arguments and their applications. *SIAM J. Comput.*, 38(5):1661–1694, 2008.

[BGGL01] Boaz Barak, Oded Goldreich, Shafi Goldwasser, and Yehuda Lindell. Resettably-sound zero-knowledge and its applications. In *FOCS*, pages 116–125, 2001.

[BS05] Boaz Barak and Amit Sahai. How to play almost any mental game over the net - concurrent composition via super-polynomial simulation. In *FOCS*, pages 543–552, 2005.

[Can00] Ran Canetti. Security and composition of multiparty cryptographic protocols. *Journal of Cryptology*, pages 143–202, 2000.

[Can01] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *FOCS*, pages 136–145, 2001.

[Can04] Ran Canetti. Universally composable signature, certification, and authentication. In *CSFW*, pages 219–, 2004.

[CDPW07] Ran Canetti, Yevgeniy Dodis, Rafael Pass, and Shabsi Walfish. Universally composable security with global setup. In *TCC*, pages 61–85, 2007.

[CF01] Ran Canetti and Marc Fischlin. Universally composable commitments. In *CRYPTO*, pages 19–40, 2001.

[CKL03] Ran Canetti, Eyal Kushilevitz, and Yehuda Lindell. On the limitations of universally composable two-party computation without set-up assumptions. In *EUROCRYPT*, pages 68–86, 2003.

[CLMP12] Kai-Min Chung, Huijia Lin, Mohammad Mahmoody, and Rafael Pass. On the power of nonuniformity in proofs of security. To Appear ITCS 2013, 2012.

[CLOS02] Ran Canetti, Yehuda Lindell, Rafail Ostrovsky, and Amit Sahai. Universally composable two-party and multi-party secure computation. In *STOC*, pages 494–503, 2002.

[CLP10] Ran Canetti, Huijia Lin, and Rafael Pass. Adaptive hardness and composable security in the plain model from standard assumptions. In *FOCS*, pages 541–550, 2010.

[CLP12] Ran Canetti, Huijia Lin, and Omer Paneth. Public coin concurrent zero-knowledge in the global hash model. Manuscript, 2012.

[DDN00] Danny Dolev, Cynthia Dwork, and Moni Naor. Nonmalleable cryptography. *SIAM Journal on Computing*, 30(2):391–437, 2000.

[DGS09] Yi Deng, Vipul Goyal, and Amit Sahai. Resolving the simultaneous resettability conjecture and a new non-black-box simulation strategy. In *FOCS*, pages 251–260, 2009.

[DM00] Yevgeniy Dodis and Silvio Micali. Parallel reducibility for information-theoretically secure computation. In *CRYPTO*, pages 74–92, 2000.

[DNS04] Cynthia Dwork, Moni Naor, and Amit Sahai. Concurrent zero-knowledge. *J. ACM*, 51(6):851–898, 2004.

[DOP05] Yevgeniy Dodis, Roberto Oliveira, and Krzysztof Pietrzak. On the generic insecurity of the full domain hash. In *CRYPTO*, pages 449–466, 2005.

[FS87] Amos Fiat and Adi Shamir. How to prove yourself: practical solutions to identification and signature problems. In *CRYPTO*, pages 186–194, 1987.

[FS90] Uriel Feige and Adi Shamir. Witness indistinguishable and witness hiding protocols. In *STOC*, pages 416–426, 1990.

[GJ10] Vipul Goyal and Abhishek Jain. On the round complexity of covert computation. In *STOC*, pages 191–200, 2010.

[GL90] Shafi Goldwasser and Leonid A. Levin. Fair computation of general functions in presence of immoral majority. In *CRYPTO*, pages 77–93, 1990.

[GMR89] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, 1989.

[GMW91] Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. *J. ACM*, 38(3):690–728, 1991.

[Gol01] Oded Goldreich. *Foundations of Cryptography — Basic Tools*. Cambridge University Press, 2001.

[Gol04] Oded Goldreich. *Foundations of Cryptography — Basic Applications*. Cambridge University Press, 2004.

[GW11] Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In *STOC*, pages 99–108, 2011.

[HH09] Iftach Haitner and Thomas Holenstein. On the (im)possibility of key dependent encryption. In *TCC*, pages 202–219, 2009.

[HILL99] Johan Håstad, Russell Impagliazzo, Leonid Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28:12–24, 1999.

[IR88] Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In *CRYPTO*, pages 8–26, 1988.

[Kat07] Jonathan Katz. Universally composable multi-party computation using tamper-proof hardware. In *EUROCRYPT*, pages 115–128, 2007.

[KLP07] Yael Tauman Kalai, Yehuda Lindell, and Manoj Prabhakaran. Concurrent composition of secure protocols in the timing model. *J. Cryptology*, 20(4):431–492, 2007.

[LP12] Huijia Lin and Rafael Pass. Black-box constructions of composable protocols without set-up. In *CRYPTO*, pages 461–478, 2012.

[LPV08] Huijia Lin, Rafael Pass, and Muthuramakrishnan Venkitasubramaniam. Concurrent non-malleable commitments from any one-way function. In *TCC*, pages 571–588, 2008.

[LPV09] Huijia Lin, Rafael Pass, and Muthuramakrishnan Venkitasubramaniam. A unified framework for concurrent security: universal composability from stand-alone non-malleability. In *STOC*, pages 179–188, 2009.

[MMY06] Tal Malkin, Ryan Moriarty, and Nikolai Yakovenko. Generalized environmental security from number theoretic assumptions. In *TCC*, pages 343–359, 2006.

[MPR06] Silvio Micali, Rafael Pass, and Alon Rosen. Input-indistinguishable computation. In *FOCS*, pages 367–378, 2006.

[MR91] Silvio Micali and Phillip Rogaway. Secure computation (abstract). In *CRYPTO*, pages 392–404, 1991.

[Nao91] Moni Naor. Bit commitment using pseudorandomness. *Journal of Cryptology*, 4:151–158, 1991.

[Nao03] Moni Naor. On cryptographic assumptions and challenges. In *CRYPTO*, pages 96–109, 2003.

[Pas03a] Rafael Pass. Simulation in quasi-polynomial time, and its application to protocol composition. In *EUROCRYPT*, pages 160–176, 2003.

[Pas03b] Rafael Pass. Simulation in quasi-polynomial time, and its application to protocol composition. In *EUROCRYPT*, pages 160–176, 2003.

[Pas11] Rafael Pass. Limits of provable security from standard assumptions. In Lance Fortnow and Salil P. Vadhan, editors, *STOC*, pages 109–118. ACM, 2011.

[PR05] Rafael Pass and Alon Rosen. Concurrent non-malleable commitments. In *FOCS*, pages 563–572, 2005.

[PR08] Manoj Prabhakaran and Mike Rosulek. Cryptographic complexity of multi-party computation problems: Classifications and separations. In *CRYPTO*, pages 262–279, 2008.

[PRT11] Rafael Pass, Alon Rosen, and Wei-Lung Dustin Tseng. Public-coin parallel zero-knowledge for np. *J. Cryptology*, 2011.

[PS04] Manoj Prabhakaran and Amit Sahai. New notions of security: achieving universal composability without trusted setup. In *STOC*, pages 242–251, 2004.

[PW00] Birgit Pfitzmann and Michael Waidner. Composition and integrity preservation of secure reactive systems. In *ACM Conference on Computer and Communications Security*, pages 245–254, 2000.

[RK99] Ransom Richardson and Joe Kilian. On the concurrent composition of zero-knowledge proofs. In *Eurocrypt*, pages 415–432, 1999.

[Rom90] John Rompel. One-way functions are necessary and sufficient for secure signatures. In *STOC*, pages 387–394, 1990.

[RS91] Charles Rackoff and Daniel R. Simon. Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack. In *CRYPTO*, pages 433–444, 1991.

[RTV04] Omer Reingold, Luca Trevisan, and Salil P. Vadhan. Notions of reducibility between cryptographic primitives. In *TCC*, pages 1–20, 2004.

[RV10] Guy N. Rothblum and Salil P. Vadhan. Are pcps inherent in efficient arguments? *Computational Complexity*, 19(2):265–304, 2010.

[Unr07] Dominique Unruh. Random oracles and auxiliary input. In Alfred Menezes, editor, *CRYPTO*, volume 4622 of *Lecture Notes in Computer Science*, pages 205–223. Springer, 2007.