On Adaptively Secure Protocols

Muthuramakrishnan Venkitasubramaniam

University of Rochester, Rochester, NY 14611, USA

Abstract. Adaptive security captures the capability of an adversary to adaptively affect a system during the course of its computation based on partial information gathered. In this work, we explore the theoretical complexity of achieving adaptive security in two settings:

1. **Adaptive UC-Secure Computation:** We provide a round-efficient compiler that transforms any stand-alone semi-honest adaptively secure multiparty computation to adaptive UC-security. Recently, Dana et. al (Asiacrypt 2013) showed how to acheive adaptive UC-security in any trusted setup under minimal assumptions. They achieve this by constructing an $O(n)$-round adaptively secure concurrent non-malleable commitment scheme. The main contribution of our work shows how to achieve the same in $O(1)$-rounds.

2. **Zero-Knowledge with Adaptive Inputs:** Lin and Pass in (TCC 2011) gave first constructions of concurrent non-malleable zero-know-ledge proofs secure w.r.t. *adaptively* chosen inputs in the plain model in a restricted setting, namely, where the adversary can only ask for proofs of *true* (adaptively-chosen) statements. We extend their definition to the fully-adaptive setting and show how to construct a protocol that satisfies this definition. As an independent contribution we provide a simple and direct compilation of any semi-honest secure protocol to a fully concurrently secure protocol under polynomial-time assumptions in the Angel-Based UC-Security.

1 Introduction

Adaptive security captures the capability of an adversary to adaptively affect a system during the course of its computation based on partial information gathered. In this work, we revisit the complexity of achieving two different notions of adaptive security: (1) Round complexity of achieving fully concurrent *adaptively* secure computation protocols, and, (2) Feasibility of concurrent non-malleable zero-knowledge protocols with *fully adaptively* chosen inputs.

Adaptive Secure Computation. The notion of *secure multi-party computation* allows m mutually distrustful parties to securely compute a functionality $f(\bar{x}) = (f_1(\bar{x}), ..., f_m(\bar{x}))$ of their corresponding private inputs $\bar{x} = x_1, ..., x_m$, such that party P_i receives the value $f_i(\bar{x})$. Loosely speaking, the security requirements are that the parties learn nothing more from the protocol than their

M. Abdalla and R. De Prisco (Eds.): SCN 2014, LNCS 8642, pp. 455–475, 2014.

prescribed output, and that the output of each party is distributed according to the prescribed functionality. This should hold even in the case that an arbitrary subset of the parties maliciously deviates from the protocol. The original setting in which secure multi-party protocols were investigated, however, only allowed the execution of a single instance of the protocol at a time; this is the so called *stand-alone setting*. A more realistic setting, is one which allows the concurrent execution of protocols. In the *concurrent setting*, many protocols are executed at the same time. This setting presents the new risk of a coordinated attack in which an adversary interleaves many different executions of a protocol and chooses its messages in each instance based on other partial executions of the protocol. The strongest (but also most realistic) setting for concurrent security—called *Universally Composable* (UC) security [Can01, PW01, DM00], or environmental-security—considers the execution of an unbounded number of protocols, running concurrently in an arbitrary, and adversarially controlled, network environment. Unfortunately, achieving UC security for most interesting functionalities is impossible unless some sort of trusted setup is assumed [CF01, CKL03, Lin03]. Previous works overcome this barrier either by using some trusted setup infrastructure or by relaxing the definition of security (we will see examples below).

When considering security in such settings, we refer to the adversary as *static* if it chooses the subset of the parties to corrupt at the beginning of the protocol and *adaptive* if it is allowed to corrupt on-the-fly. Adaptively secure multiparty computation protocols (in the non-erasure model) were first realized in Canetti, Feige, Goldreich and Naor [CFGN96] for the stand-alone case. Canetti, Lindell, Ostrovsky and Sahai [CLOS02] provided the first constructions of UC-secure protocols with adaptive security for "well-formed" functionalities in the *common reference string* model (CRS) where all parties have access to public reference string sampled from a pre-specified distribution. Subsequently, several results were obtained for both the static and adaptive case in other trusted-setup models and relaxed-security models. However, for a given functionality, realizing an adaptively secure protocol is significantly harder than realizing in the static case. In this work we focus on the round complexity of achieving adaptive security under minimal assumptions. In the static case Lin et. al [LPV08, LPV12] provide a round-efficient compiler from stand-alone semi-honest secure computation protocol to static UC-security under minimal assumptions. Since there exists $O(1)$-round protocols with static security in the semi-honest setting for most setup models, we can achieve the same with static UC-security as well. For adaptive-security, the best known semi-honest secure protocol requires $O(d_C)$-rounds for computing a circuit C, where d_C is the depth of the circuit.[1,2] When considering adaptive UC-security, for particular models such as the common-reference string (CRS) model, we know how to construct $O(d_C)$-round adaptive

[1] Informally, the depth of a circuit is the longest path from any input bit to the output bit.

[2] There are constant-round protocols if we assume erasures or restrict the adversary to corrupting a strict subset of the parties.

UC-secure protocols under minimal assumptions [CLOS02, DN02, CDSMW09]. For other models such as uniform reference string model [CLOS02], multi-string model [GO07], similar constructions exist with stronger assumptions (e.g., dense cryptosystems). However, when considering minimal assumptions, the work of [DSMRV13] shows how to achieve adaptive UC-security in any trusted setup in $O(nd_C)$-rounds where n is the length of the identifier of the parties. Thus the state-of-the-art shows a huge gap in the round-complexity required for achieving adaptive security under minimal assumptions in the semi-honest setting and UC-setting in most models. One of the main questions addressed in this work is

> *Is there a round-efficient transformation from stand-alone semi-honest adaptive security to adaptive UC-security under minimal assumptions?*

In this work, we answer this question in the affirmative and show how to obtain round-efficient compilation to obtain adaptively UC-secure protocols. Concretely, our work improves the round-complexity of constructing adaptively UC-secure protocols in most setup models under "minimal assumptions" and closes the gap of round-complexity between achieving such constructions in the semi-honest and fully-malicious UC setting in any setup.

Zero-Knowledge with Adaptively Chosen Inputs. Zero-knowledge interactive proofs [GMR89] are paradoxical constructs that allow one player (called the Prover) to convince another player (called the Verifier) of the validity of a mathematical statement $x \in L$, while providing *zero additional knowledge* to the Verifier. The notion of concurrent ZK, first introduced and achieved, by Dwork, Naor and Sahai [DNS04] considers the execution of zero-knowledge proofs in an asynchronous setting and concurrent setting. Non-malleable ZK was first introduced by Dolev, Dwork and Naor [DDN00] where they model an adversary as a man-in-the-middle participating in two executions, acting as a verifier in one execution (in the left) and as a prover in another execution (in the right). The notion of concurrent non-malleable ZK considers adversaries that participates in an unbounded number of executions as the prover and verifier. Barak, Prabharakan and Sahai [BPS06] gave the first ZK argument for **NP** that is concurrent non-malleable in the plain model (i.e. without any trusted set-up assumptions). Since then several works improve efficiency and/or construct concurrent non-malleable ZK proofs [OPV08, LPTV10, LP11b]. In all these works, the input statements for all executions on the right can be *adaptively* chosen by the adversary whereas on the left need to be selected *apriori*, i.e. at the beginning of the execution. In this work, we consider strengthening the security, by allowing the adversary to adaptively chose the statements it wishes to prove and verify. Lin and Pass [LP11b] show how to construct concurrent non-malleable ZK proofs in a restricted setting where the adversary can adaptively choose *only* true statements to receive proofs of. Furthermore, they argue that allowing the adversary to choose arbitrary statements will reveal more information and thus not be zero-knowledge (for languages in **NP**). That main question we address in this work is:

Is it possible to achieve fully adaptively chosen input concurrent non-malleable zero-knowledge protocols ($\mathcal{CNMZK}$) with meaningful security?

We again answer this question in the affirmative and show that we can define ZK with meaningful security in such a setting and show how to construct such protocols in the plain model. Previous works that achieve $\mathcal{CNMZK}$ with fully adaptive input selection have relied on some sort of trusted set-up models such as Bare-Public Key [OPV08], Common Reference String [SCO+01, CF01, DN02], etc. For the weaker case of witness indistinguishability, Ostrovksy et. al [OPV06] present a restricted version[3] of input-adaptive concurrent non-malleable witness indistinguishable argument in the plain model. As an independent contribution we show the power of our definition of ZK by showing how to achieve UC-security using super-polynomial helpers analogous to [CLP10] under polynomial-time assumptions in the angel-based security model of Prabhakaran and Sahai [PS04].

1.1 Our Results

We provide a round-efficient compilation from stand-alone adaptively secure semi-honest protocols to adaptive UC-security. We achieve this by constructing a round-efficient concurrent non-malleable equivocal commitment, a primitive defined in [DSMRV13]. From previous works [CLOS02, DN00] we know that every functionality can be compiled into a protocol in the ideal-commitment functionality hybrid model.[4] In [DSMRV13] they show (assuming the existence of simulatable public-key encryption schemes) that any t-round protocol in this hybrid model can be securely realized using an $O(t_p)$-round puzzle and $O(t_c)$ round concurrent non-malleable equivocal commitment scheme in $O(t(t_p + t_c))$ rounds. We prove the following theorem

Theorem 1 (Informal). *For any setup model $\mathcal{T}$, assume the existence of a $O(t_1)$-round adaptive puzzle and one-way functions, then there exists a $O(t_1)$ round concurrent non-malleable equivocal commitment scheme.*

Since in most models, there exists $O(1)$-round puzzles, our result combined with the work of [DSMRV13] shows that in most models any t-round protocol can be realized securely with adaptive UC-security in $O(t)$-rounds, thus yielding a round-preserving transformation. Furthermore, we obtain improvements in round-complexity in several setup models. Concretely, we obtain the following corollary:

Corollary 1 (Informal). *Assuming the existence of simulatable public-key encryption, any well-formed circuit C can be realized with adaptive UC-security in $O(d_C)$-rounds in Common Reference String, Uniform Reference String, Imperfect Reference String (Sunspots), Multi-String and Bounded Concurrent model, where d_C is the depth of circuit C.*

[3] This definition implicitly has a similar restriction to [LP11b] where the statements in the left need to be true statements.

[4] In this hybrid model, all parties have access to the ideal commitment functionality.

We also obtain corresponding improvements for relaxed-models of security, such as quasi-polynomial simulation and non-uniform simulation, where we need to make additional assumptions on the simulatable public-key encryption. We remark, this matches previous constructions in the Common-Reference String model. For the rest of the models, to achieve an $O(d_C)$-round construction we either did not have a construction (eg, bounded-concurrent) or needed additional assumptions (uniform reference string, sunspots). Our work closes the gap in the round-complexity required to achieve adaptively-secure protocols in the semi-honest setting and the fully malicious UC setting in most setup models.

Concurrent Non-Malleable Zero-Knowledge with Adaptive Inputs: As our second contribution, we define fully adaptive concurrent non-malleable zero-knowledge (a definition inspired by [LP11b] and angel-based security model of Prabhakaran and Sahai [PS04]) and show how to construct a protocol satisfying the definition. Our definition will allow for a man-in-the-middle adversary to adaptively request proofs of arbitrary **NP** statements, may it be true or false. For true statements, it will be provided with a proof using an efficient prover P using the witness to the statement and for false statements, it will be provided with a "fake" proof by a prover that potentially runs in exponential time. To construct a protocol, we rely on the recently introduced CCA-secure commitment schemes of [CLP10]. More formally, we prove the following theorem.

Theorem 2 (Informal). *Assuming the existence of one-way functions, there exists a $\tilde{O}(\log n)$-round fully-adaptive $\mathcal{CNMZK}$argument.*

As an independent contribution, we show that our $\mathcal{ACNMZK}$ protocol can be used to securely realize any functionality in the angel-based model introduced by Prabhakaran and Sahai [PS04]. Canetti, Lin and Pass [CLP10] were the first to show how to achieve UC-security using the CCA-secure commitments in the angel model under polynomial-time assumptions.

The high-level idea of our protocol follows previous approaches where the functionality is first compiled to a protocol in the IdealZK-hybrid model.[5] Then, in a second step, it is complied to a protocol in a slightly weaker ideal functionality called MemberZK which is the ideal zero-knowledge proof of membership protocol. Finally, a protocol realizing the MemberZK functionality is constructed assuming some setup. We show that any protocol satisfying our definition of $\mathcal{ACNMZK}$ realizes MemberZK in the angel-model. We additionally point out a subtlety that arises when compiling IdealZK into MemberZK which was not previously addressed and show how our protocol handles this. We also believe that a direct compilation, as in [GMW87], of any semi-honest protocol to full UC-security in the angel model is possible by simply requiring honest parties to provide a zero-knowledge proof using our $\mathcal{ACNMZK}$ protocol after every step.

[5] In the IdealZK-hybrid all parties have access to an ideal zero-knowledge proof-of-knowledge protocol.

1.2 Our Techniques

Proving UC-security essentially reduces to proving concurrent non-malleability and concurrent simulation. In [LPV09], Lin et. al inroduced the notion of a UC-puzzle that captures the concurrent simulation requirement of UC-security and showed how to achieve it using any setup under minimal assumptions. In [DSMRV13], Dana et. al extended the notion of UC-puzzle to adaptive security and showed how to achieve adaptive UC-security using an adaptive UC-puzzle and a special kind of commitment scheme known as an *equivocal non-malleable commitments*. Roughly speaking, such commitments require that no man-in-the-middle adversary participating as a sender and receiver in multiple concurrent commitments and decommitments, be able to break the binding property of the commitment scheme. Similar notions have been studied previously in the works of [CIO98, CKOS01] for the limited case of bounded concurrency and non-interactive commitments and in [OPV09, CVZ10] where they restrict the commitments and decommitments to be performed in two distinct phases that do not overlap in time.

In more detail, a tag-based commitment scheme (i.e., commitment scheme that takes an identifier—called the tag—as an additional input) is *concurrent non-malleable w.r.t opening* if for every man-in-the-middle adversary A that participates in several interactions with honest committers as a receiver (called *left* interactions) as well as several interactions with honest receivers as a committer (called *right* interactions), there exists a simulator S that can simulate the left interactions, while extracting the commitments made by the adversary in the right interactions (whose identifiers are different from all the left identifiers) before the adversary decommits.

Equivocal commitments can be constructed easily using trusted set-up. The basic idea is to provide the simulator with a trapdoor with which it can equivocate as well as extract the commitments on the right. (by e.g., relying on encryption). However, to ensure non-malleability, most constructions in literature additionally impose CCA-security or provide independent trapdoors for every interaction. In [DSMRV13], they show how to construct a concurrent non-malleable commitment scheme in any trusted set-up. More precisely, they construct a commitment scheme in any setup that admits a UC-puzzle, a formulation introduced by Lin et al [LPV08] for the case of static security that captures precisely the simulation requirement. However, the round complexity of their protocol is linear in the length of the identities. In this work we show how to construct constant-round protocols that are concurrently non-malleable. Moreover, our construction seemingly only relies on the stand-alone non-malleable commitment schemes in the static setting. In particular, using the $O(1)$-round scheme of [LP11b] we provide a $O(1)$-round concurrent non-malleable commitment w.r.t opening in any setup with a UC-puzzle.

Comparison with [DSMRV13]. The work of [DSMRV13] focused on constructing protocols with adaptive UC-security in any model under minimal assumptions. In particular their work showed how to minimize complexity assumptions and

the "trust" required in the setup analogous to [LPV09] for the static setting. In this work, we focus on obtaining round-efficient protocols. The novelty in our reduction over the work in [DSMRV13] is that we are able to directly reduce the security proof to the non-malleability of the underlying non-malleable commitment scheme while [DSMRV13] rely on a particular non-malleable commitment scheme (that of [DDN00, LPV08]) and provide a security reduction that is tailored to this scheme. Our proof is modular and in our opinion simpler than that of [DSMRV13]. Although the main application of equivocal non-malleable commitments is in achieving adaptive UC-security, these commitments may also be useful for other applications such as concurrent non-malleable zero knowledge secure under adaptive corruptions or obtaining composable protocols. We believe that this notion in some sense extends the notion of security w.r.t selective opening attacks to the non-malleable setting and our protocols might be useful in such contexts as well.

$\mathcal{CNMZK}$ with Adaptively Chosen Inputs. Lin and Pass [LP11b] consider the scenario where a man-in-the-middle adversary participates in unbounded zero-knowledge interactions as a verifier on the left and prover on the right and can adaptively choose statements for left and right interactions with the restriction that only true statements are chosen for left interactions. They call a zero-knowledge protocol in such a setting $\mathcal{CNMZK}$ with *Adaptive Input Selection* if for every adversary A, there exists a computationally efficient simulator-extractor that can simulate both the left and right interactions for A, while outputting a witness for every statement proved by the adversary in the right interactions. As pointed out in [LP11b, OPV06], having this restriction for protocols in the plain model seems inherent in light of the impossibility result of Lindell [Lin03]. To circumvent the impossibility result one can rely on some sort of trusted setup such as the Bare-Public Key (BPK) or Common Reference String (CRS) model. Indeed there are fully input-adaptive $\mathcal{CNMZK}$ arguments in these models [SCO$^+$01, CF01, DN02, OPV08].

In this work, we want a definition in the plain model that allows for concurrent composition. Our definition is inspired by the definitions of [LP11b] and the angel-based security model of Prabhakaran and Sahai [PS04, CLP10]. We allow the adversary to adaptively chose any statement in the left and require the existence of a simulator extractor that can achieve the same indistinguishability guarantees of [LP11b]. If we allow the adversary to choose false statements, a proof cannot be provided and hence will fall into the impossibility result of Lindell [Lin03]. Instead we will require that there be an exponential-time strategy that will provide proofs of false statements. Now for such a setting, our definition of $\mathcal{CNMZK}$ with fully adaptive input-selection requires for every adversary there be a simulator-extractor that can simulate both the left and right interactions for A, while outputting a witness for every statement proved by the adversary in the right interactions. Relying on CCA-secure commitments [CLP10, GLP$^+$12], we show how to construct an $\tilde{O}(\log n)$-round $\mathcal{CNMZK}$ protocol with fully adaptive-inputs selection. CCA-secure commitments were introduced by Canetti, et. al [CLP10] to construct UC-secure protocol with

angel-based security. We also prove that CCA-secure commitments are equivalent to $\mathcal{CNMZK}$ with fully adaptive-input selection.

2 Preliminaries

We assume familiarity with interactive protocols, commitment schemes. Some of the definitions are presented almost verbatim from [DSMRV13].

We adopt a variant of language-based commitment schemes introduced by Lindell et. al [LZ09]. Roughly speaking, in such commitments the sender and receiver share a common input, a statement x from an **NP** language L. The main difference from standard commitments is that the binding property of the commitment scheme. Informally, the binding property of the scheme asserts that any adversary violating the binding can be used to extract an **NP**-witness for the statement. We require a variant We present below the definition of a language-based equivocal commitment scheme which is a slight variant of such commitment schemes (See [DSMRV13] for the formal definition).

Definition 1 (Language-Based Equivocal Commitments). *Let L be an* **NP***-Language and $\mathcal{R}$, the associated* **NP***-relation. A language-based commitment scheme $\langle S, R\rangle$ for L is said to be equivocal, if there exists a tuple of algorithms $(\tilde{S}, \mathsf{Adap})$ such that the following holds:*

Special-Hiding: *For every (expected) PPT machine R^*, it holds that, the following ensembles are computationally indistinguishable over $n \in N$.*

- $\{\mathsf{sta}^{R^*}_{\langle S,R\rangle}(x, v_1, z)\}_{n\in N, x\in L\cap\{0,1\}^n, w\in\mathcal{R}(x), v_1\in\{0,1\}^n, z\in\{0,1\}^*}$
- $\{\mathsf{sta}^{R^*}_{\langle \tilde{S},R\rangle}(x, w, z)\}_{n\in N, x\in L\cap\{0,1\}^n, w\in\mathcal{R}(x), v_1\in\{0,1\}^n, z\in\{0,1\}^*}$

  where $\mathsf{sta}^{R^*}_{\langle \tilde{S},R\rangle}(x, w, z)$ denotes the random variable describing the output of $R^*(x, z)$ after receiving a commitment using $\langle \tilde{S}, R\rangle$.

Equivocability: *Let τ be the transcript of the interaction between R and $\tilde{S}$ on common input $x \in L \cap \{0,1\}^n$ and private input $w \in \mathcal{R}(x)$ and random tape $r \in \{0,1\}^*$ for $\tilde{S}$. Then for any $v \in \{0,1\}^n$, $\mathsf{Adap}(x, w, r, \tau, v)$ produces a random tape r' such that (r', v) serves as a valid decommitment for C on transcript τ.*

In [DSMRV13] (relying on the work of [LZ09]), show how to construct such schemes using one-way functions.

2.1 Definition of Equivocal Non-Malleable Commitments

Let $\langle C, R\rangle$ be a commitment scheme, and let $n \in N$ be a security parameter. Consider man-in-the-middle adversaries that are participating in left and right interactions in which $m = \text{poly}(n)$ commitments take place. We compare between a *man-in-the-middle* and a *simulated* execution. In the man-in-the-middle execution, the adversary A is simultaneously participating in m left and right interactions. In the left interactions the man-in-the-middle adversary A interacts

with C receiving commitments to values $v_1, \ldots, v_m$, using identities $\mathsf{id}_1, \ldots, \mathsf{id}_m$ of its choice. It must be noted here that values $v_1, \ldots, v_m$ are provided to committer on the left prior to the interaction. In the right interaction A interacts with R attempting to commit to a sequence of related values again using identities of its choice $\tilde{\mathsf{id}}_1, \ldots, \tilde{\mathsf{id}}_m$; $\tilde{v}_i$ is set to the value decommitted by A in the j^{th} right interaction. If any of the right commitments are invalid its committed value is set to $\perp$. For any i such that $\tilde{\mathsf{id}}_i = \mathsf{id}_j$ for some j, set $\tilde{v}_i = \perp$—i.e., any commitment where the adversary uses the same identity as one of the honest committers is considered invalid. Additionally, the adversary can adaptively choose a session in the left executions that has completed the commitment phase to be *decommitted*. Let $\mathsf{mim}^A_{\langle C,R\rangle}(v_1, \ldots, v_m, z)$ denote a random variable that describes the values $\tilde{v}_1, \ldots, \tilde{v}_m$ and the view of A, in the above experiment.

In the simulated execution, a simulator S interacts only with receivers on the right as follows:

1. Whenever the commitment phase of j^{th} interaction with a receiver on the right is completed, S outputs a value $\tilde{v}_j$ as the alleged committed value in a special-output tape.
2. During the interaction, S may output a partial view for a man-in-the-middle adversary whose right-interactions are identical to S's interaction so far. If the view contains a left interaction where the i^{th} commitment phase is completed and the decommitment is requested, then a value v_i is provided as the decommitment.
3. Finally, S outputs a view and values $\tilde{v}_1, \ldots, \tilde{v}_m$. Let $\mathsf{sim}^S_{\langle C,R\rangle}(1^n, v_1, \ldots, v_m, z)$ denote the random variable describing the view output by the simulation and values $\tilde{v}_1, \ldots, \tilde{v}_m$.

Definition 2. *A commitment scheme $\langle C, R\rangle$ is said to be* concurrent non-malleable w.r.t. opening *if for every polynomial $p(\cdot)$, and every probabilistic polynomial-time man-in-the-middle adversary A that participates in at most $m = p(n)$ concurrent executions, there exists a probabilistic polynomial time simulator S such that the following ensembles are computationally indistinguishable over $n \in N$:*

- $\left\{\mathsf{mim}^A_{\langle C,R\rangle}(v_1, \ldots, v_m, z)\right\}_{n\in N, v_1,\ldots,v_m\in\{0,1\}^n, z\in\{0,1\}^*}$, *and*
- $\left\{\mathsf{sim}^S_{\langle C,R\rangle}(1^n, v_1, \ldots, v_m, z)\right\}_{n\in N, v_1,\ldots,v_m\in\{0,1\}^n, z\in\{0,1\}^*}$

We will use the slightly relaxed definition where all the values committed to the adversary in the left interaction are sampled independently from an arbitrary distribution $\mathcal{D}$ fixed *apriori*. We call a commitment scheme an *equivocal non-malleable commitment scheme* if it is both a language-based equivocal commitment scheme and is concurrent non-malleable w.r.t. opening.

2.2 Adaptive UC-Puzzles

Informally, an adaptive UC-puzzle is a protocol $\langle \mathsf{S}, \mathrm{R}\rangle$ between two players–a *sender* S and a *receiver* R – and a PPT computable relation $\mathcal{R}$, such that the following two properties hold:

Soundness. No efficient receiver R^* can successfully complete an interaction with S and also obtain a "trapdoor" y, such that $\mathcal{R}(\text{TRANS}, y) = 1$, where TRANS is the transcript of the interaction.

Statistical UC-simulation with adaptive corruptions. For every efficient adversary $\mathcal{A}$ participating in a polynomial number of concurrent executions with receivers R (i.e., $\mathcal{A}$ is acting as a puzzle sender in all these executions) and at the same time communicating with an environment $\mathcal{Z}$, there exists a simulator $\mathcal{S}$ that is able to statistically simulate the view of $\mathcal{A}$ for $\mathcal{Z}$, while at the same time outputting trapdoors to all successfully completed puzzles. Moreover, $\mathcal{S}$ successfully simulates the view even when $\mathcal{A}$ may adaptively corrupt the receivers.

2.3 Fully Input-Adaptive Concurrent Non-Malleable Zero-Knowledge

Our definition of fully input-adaptive concurrent non-malleable zero-knowledge is closely based on the definition of adaptive concurrent non-malleable zero-knowledge from [LP11a] (which in turn are based on [BPS06, PR05]). The main difference is that we consider adversaries that is allowed to adaptively select true and false statements to receive proofs of.

Let $\langle P, V \rangle$ be an interactive argument for a language $L \in \mathbf{NP}$ with witness relation R_L and exponential time cheating prover $\hat{P}$, and let n be the security parameter. Consider a man-in-the-middle adversary A that participates in many left and right interactions in which $m = m(n)$ proofs take place. In the left interactions, the adversary A verifies the validity of the statements $x_1, \ldots, x_m$ by interacting with a special prover $\hat{P}$, using identities $\mathsf{id}_1, \ldots, \mathsf{id}_m$. In the right-interactions, A proves the validity of statements $\tilde{x}_1, \ldots, \tilde{x}_m$ to an honest verifier V, using identities $\tilde{\mathsf{id}}_1, \ldots, \tilde{\mathsf{id}}_m$. Prior to the interactions, all parties receive as common input the security parameter in unary 1^n, and A receives as auxiliary input $z \in \{0,1\}^*$. Furthermore, at the beginning of each left (respectively right) interaction, the adversary adaptively selects the statement x_i (respectively $\tilde{x}_i$) and the identity id_i (respectively $\tilde{\mathsf{id}}_i$). For each left-interaction, the special-prover $\hat{P}$ behaves as follows:

1. If the statement x_i chosen by A is false, then $\hat{P}$ runs the exponential-time cheating strategy,
2. If the statement x_i chosen by A is true, then $\hat{P}$ picks a witness $w_i \in R_L(x_i)$ and runs the honest prover strategy P with private input w_i.

Let $\text{VIEW}_{A,\hat{P}}$ denote a random variable that describes the view of A in the above experiment. Informally, an interactive argument (with exponential time cheater strategy) is fully-adaptive concurrent non-malleable zero-knowledge ($\mathcal{ACNMZK}$) if for all man-in-the-middle adversary A, there exists a probabilistic polynomial time machine (called the simulator-extractor) that can simulate both the left and right interactions of A, while outputting a witness for every statement proved by the adversary in the right interactions.

Definition 3. *An interactive argument (P, V) for a language L with witness-relation R_L and exponential-time prover $\hat{P}$ is said to be fully input-adaptive concurrent non-malleable zero-knowledge if for every polynomial m, and every probabilistic polynomial-time man-in-the-middle adversary A that participates in at most $m = m(n)$ concurrent executions, there exists a probabilistic polynomial time machine S, such that,*

1. *The ensembles $\left\{\text{VIEW}_{A,\hat{P}}(n,z)\right\}_{n\in\mathbf{N},z\in\{0,1\}^*}$ and $\{S_1(1^n,z)\}_{n\in\mathbf{N},z\in\{0,1\}^*}$ are computationally indistinguishable over $n \in \mathbf{N}$*
2. *Let $z \in \{0,1\}^*$ and (VIEW, $\boldsymbol{w}$) denote the output of $S(1^n, z)$. Let $\tilde{x}_1, \ldots, \tilde{x}_m$ be the statements of the right interactions in* VIEW*, and let $\mathsf{id}_1, \ldots, \mathsf{id}_m$ and $\tilde{\mathsf{id}}_1, \ldots, \tilde{\mathsf{id}}_m$ be the identities of the left-interactions and right-interactions in* VIEW*. Then for every $i \in [m]$, if the i^{th} right-interaction is accepting and $\tilde{\mathsf{id}}_i \neq \mathsf{id}_j$ for any $j \in [m]$, $\boldsymbol{w}$ contains w_i such that $R_L(\tilde{x}_i, w_i) = 1$.*

Remark 1. We remark that it would be impossible to achieve proofs according to our definition, since we allow for an exponential-time prover to convince verifiers of false statements.

3 EQNMCom Based on [LP11b]

Our starting point is any stand-alone non-malleable commitment scheme that follows that Feige-Shamir paradigm for the hiding property and the "simulation-soundness" paradigm of Sahai[Sah99] for non-malleability. More precisely, in a scheme following the Feige-Shamir paradigm, there is a trapdoor phase where, possibly through rewinding the receiver, a trapdoor can be obtained and a proof phase where the committer proves either knowledge of the value committed or knowledge of the trapdoor. To prove non-malleability, or simulation-soundness, these schemes provide a mechanism to rewind the left interaction of a man-in-the middle adversary to obtain a trapdoor while ensuring the right interactions remains "sound".

While our techniques are more generally applicable, in this work, we present a protocol based on the constant-round non-malleable commitment protocol of Lin and Pass [LP11b]. Their scheme relies on fixed-length signature scheme $\Pi = (Gen, Sign, Ver)$, zero-knowledge argument of knowledge, witness-indistinguishable special-sound proofs and we assume the readers familiarity with these primitives. Our protocol is obtained from the protocol from [LP11b] by replacing the non-interactive commitment com with the language-based equivocal commitment scheme EQCom_x (see Definition 1) from [DSMRV13] and the Stage 3 protocol with the adaptively-secure witness-indistinguishable proof of knowledge (WIPOK).

In more detail, to achieve equivocability, as in [DSMRV13], we rely on a variant of Feige-Shamir's trapdoor commitment scheme. Let x be an **NP**-statement. The sender commits to bit b by running the honest-verifier simulator for Blum's Hamiltonian Circuit protocol [Blu86] on input the statement x and the verifier message b, generating the transcript (a, b, z), and finally outputting a as

its commitment. In the decommitment phase, the sender reveals the bit b by providing both b, z. To achieve adaptively secure WIPOK protocol (we refer [DSMRV13] for a formal definition and construction) we rely on the protocol of Lindell-Zarosim [LZ09].

3.1 Equivocal Non-Malleable Commitment Scheme (EQNMCom) in any Setup

Given any setup $\mathcal{T}$ with an adaptive UC-Puzzle, we prove that $\Pi = \langle S, R \rangle$ described below is an equivocal non-malleable commitment scheme when combined with the adaptive UC-puzzle. In more detail, consider the following protocol: Let $(\langle \mathsf{S}_{\mathsf{puz}}, \mathrm{R}_{\mathsf{puz}} \rangle, \mathcal{R})$ be an adaptive UC-puzzle in setup $\mathcal{T}$. The protocol $\overline{\Pi}$ proceeds in two phases on common input the identity $\mathsf{id} \in \{0,1\}^{\ell}$ of the committer, and private-input string r for the committer and security parameter n.

Preamble Phase: An adaptive UC-Puzzle interaction $\langle \mathsf{S}_{\mathsf{puz}}, \mathrm{R}_{\mathsf{puz}} \rangle$ on input 1^n where $\mathsf{S}_{\mathsf{com}}$ is the receiver and $\mathsf{R}_{\mathsf{com}}$ is the sender. Let $x = \mathrm{TRANS}$ be the transcript of the messages exchanged.

Commitment Phase: The parties run protocol $\langle S, R \rangle$ with common input x and identifier id. S plays the part of sender with input r.

Our construction relies on the equivocal commitment scheme $\langle S, R \rangle$ constructed based on the non-malleability commitment scheme of Lin and Pass [LP11b]. For the purpose of describing the simulator we will only rely on the fact that the protocol has a round where the honest committer sends a commitment to the value in its input using EQCom_x and a proof phase where there are one or more adaptively secure WIPOK instantiations based on statement x from where the committer proves knowledge of the value committed using EQCom_x in that interaction. We now show that the protocol $\overline{\Pi}$ is concurrent non-malleable w.r.t opening w.r.t i.i.d commitments.

Theorem 1. *Commitment scheme $\overline{\Pi}$ described above is concurrent non-malleable w.r.t. opening with independent and identically distributed (i.i.d) commitments.*

First we describe the simulator and then prove correctness. Let A be a concurrent man-in-the-middle adversary that on input 1^n participates in at most $m(n)$ left-interactions as a receiver, receiving commitments from an honest committer whose values are chosen uniformly from distribution D and at most $m(n)$ right-interactions as a committer. On a high-level, S internally incorporates A and emulates an execution with A as follows: For all puzzle interactions where A^* controls the sender, S follows the puzzle simulator's strategy to simulate the puzzle and obtains a witness which it stores. In the right puzzle interactions, Sim simply forwards the messages to an external receiver. For every left interaction, Sim internally generates the messages using the code of special committer (guaranteed by the scheme), i.e. equivocate in the commitment phase with the witness w obtained from the puzzle interactions. When a decommitment is requested by

A, Sim outputs the current partial view of the transcript of messages exchanged by A in a special-output tape. Then, it receives a value v from outside to be decommitted to in the left interaction. Internally, it runs the Adap algorithm guaranteed by the equivocal commitment scheme to generate a decommitment to v and feeds it to A.

Whenever the commitment phase is completed with a receiver on the right, Sim temporary stalls the main-execution and tries to extract the value committed to by A in this interaction. For this, Sim selects a random adaptively secure WIPOK from that interaction and *rewinds* A to the point just before which A receives the challenge-message in the WIPOK. Sim supplies a new challenge message and continues simulation. In the rewinding, the right interactions are simulated as before (i.e. honestly) while the left interactions are simulated differently. Instead of equivocating the session, Sim follows the honest committer's strategy with value v, where v is the actual value decommitted to in left interaction if one exists from the main-execution or chosen uniformly at random from $\mathcal{D}$.[6] If in the rewinding, A provides a valid response for the selected WIPOK of the right interaction, then using the special-sound property of the WIPOK, Sim extracts the witness used in the WIPOK, which contains the committed value. If the adversary fails to provide a valid response for the particular WIPOK in the right interaction, Sim cancels the current rewinding and starts a new rewinding by supplying a new challenge.

The proof of correctness of the simulator is expressed in the following lemma.

Lemma 1. *The following ensembles are computationally indistinguishable*

$$\left\{(v_1 \ldots, v_m) \leftarrow D^n : \mathsf{mim}^A_{\langle S,R\rangle}(v_1, \ldots, v_m, z)\right\}_{n\in N, z\in\{0,1\}^*}$$

$$\left\{(v_1 \ldots, v_m) \leftarrow D^n : \mathsf{sim}^S_{\langle S,R\rangle}(1^n, v_1, \ldots, v_m, z)\right\}_{n\in N, z\in\{0,1\}^*}$$

Proof. We consider a sequence of intermediate hybrid experiments $H_0, \ldots, H_m$. In experiment H_i, we consider a simulator Sim^i that knows all the values $(v_1, \ldots, v_n)$ being committed to in the left interactions. On input z, Sim^i proceeds as follows: It follows Sim's strategy in the first i left interactions both in the main phase and rewinding phase. For the other left interactions, i.e. $j = i+1, \ldots, n$, Sim simulates the main phase by equivocating as before, but in the rewinding phase follows the honest committers strategy using v_j. Let $\mathsf{hyb}^i(1^n, v_1, \ldots, v_m, z)$ denote the output of Sim^i in H_i. It follows from description that $\mathsf{hyb}^m(1^n, v_1, \ldots, v_m, z)$ and $\mathsf{sim}^S_{\langle S,R\rangle}(1^n, v_1, \ldots, v_m, z)$ are indistinguishable.

In hybrid experiment H_0, all puzzles are simulated and the simulator with input $(v_1, \ldots, v_m)$ proceeds exactly as the real simulator in the main execution phase, whereas in the rewindings, it uses the code of the honest committer to commit to v_i. We consider an intermediate hybrid $\overline{H_0}$ which proceeds exactly like H_0 with the exception that the simulator uses the code of the honest committer in the main execution phase as well with value v_i for left session i. Let the output

[6] Sim can generate such messages for any value v, since by adaptive security, Sim can obtain random coins for an honest committer and any value v that is consistent with any partial transcript generated by the equivocator.

of this experiment be $\overline{\mathsf{hyb}^0}$. The proof of the Lemma follows from the next three claims using a standard hybrid argument.

Claim. $\mathsf{mim}^A_{\langle S,R\rangle}(v_1,\ldots,v_m,z) \approx \overline{\mathsf{hyb}^0}(1^n,v_1,\ldots,v_m,z)$

Proof Sketch. The proof of this claim essentially follows from the statistical simulation of the puzzle. In fact, this step is identical to the one presented as part of the proof in [DSMRV13] (see Claim 1 in [DSMRV11]) and will be presented in the full version.

Claim. $\overline{\mathsf{hyb}^0}(1^n,v_1,\ldots,v_m,z) \approx \mathsf{hyb}^0(1^n,v_1,\ldots,v_m,z)$

Proof. Recall that in hybrid $\overline{H_0}$, the simulator follows the honest committer strategy in the main and rewinding phases for left interactions while in H_0 it equivocates in the main phase and then follows honest committer strategy in the rewindings using the random coins generated for a honest committer by the Adap algorithm of the commitment scheme. Recall that the simulator in all hybrids have all the values to be committed in the left interactions at the beginning of the experiment.

Now, assume for contradiction there exists a distinguisher that can distinguish the outputs in both the experiments with probability $\frac{1}{p(n)}$ for some polynomial $p(\cdot)$. As in the previous claim, we can conclude that in H_0 the value extracted and the value decommitted to in some right interaction by the adversary is different with probability $\frac{1}{p(n)}$. Now, we consider a truncated version of both the hybrids where both the experiments are cut-off after the simulator runs T steps. Using the same argument as in the previous claim, we can conclude that the simulator runs in expected polynomial time in both the experiments. Let $t(n)$ be an upper bound for the expected running time in both experiments. We set T to be $2t(n)p(n)$. Using Markov's inequality, it holds that the distinguisher distinguishes the truncated experiments with probability at least $\frac{1}{2p(n)}$. We will show that the adversary A can be used to violate the pseudo-randomness of the commitments under Com and thus arrive at a contradiction.

Next, following [DSMRV13], we rely on the existence of a committer strategy C^* and distributions D_0 and D_1 such that on input the witness to the puzzle-statement and a sequence of strings from distribution D_b can commit to a value v such that messages are distributed identically to the honest committer's strategy if $b = 0$ and according to the equivocating strategy if $b = 1$. In fact one of these distributions is simply commitments to the bit 0 while the other is a the uniform random string.

Now consider a simulator S^* that additionally on input T samples $s_1,\ldots,s_T$ from D_b, proceeds exactly as in $\overline{H_0}$ with the exception that for all left interactions the simulator uses the committer strategy C^* with samples $s_1,\ldots,s_T$ in the main and rewinding phase. By construction, it follows that the output of S^* with samples from D_b is identical to truncated version of $\overline{H_0}$ when $b = 0$ and H_0 when $b = 1$. Therefore, by running D on the output of S^* we can distinguish D_0 from D_1 which is a contradiction.

Claim. $\mathsf{hyb}^0(1^n, v_1, \ldots, v_m, z) \approx \mathsf{hyb}^m(1^n, v_1, \ldots, v_m, z)$

Proof. Assume for contradiction, there exists an adversary A, distinguisher D, polynomial $p(\cdot)$ such that, for infinitely many n, D distinguishes the ensembles in the claim with probability at least $\frac{1}{p(n)}$. Recall that in hyb^0, the value successfully decommitted by the adversary in every right interaction is the value extracted by the simulator. Furthermore, the view output by the simulator in both hyb^0 and hyb^m are identical since the main execution is simulated in an identical manner. Hence, if D can distinguish, it must hold that the value extracted in some right interaction in hyb^m must not be the value decommitted to by the adversary with probability at least $\frac{1}{p(n)}$. Whenever this happens in an execution in the k^{th} right interaction, we will say that the adversary *cheats* in the k^{th} right interaction. Therefore, given our hypothesis, there must exist an i such that the difference in probability that A cheats in the k^{th} interaction is at least $\frac{1}{p_1(n)}$ between hyb^{i-1} and hyb^i for some polynomial $p_1(\cdot)$. Consider a random transcript truncated at the end the of k^{th} interaction. Then, the simulation strategies of the i^{th} left interaction according to hyb^{i-1} and hyb^i in the rewindings must yield different values extracted in k^{th} right interaction with probability at least $\frac{1}{p_1(n)}$. We can further impose the condition that the adversary has not corrupted or requested the decommitment of the left i^{th} interaction before the k^{th} right interaction has completed.[7]

We now consider truncated experiments $\overline{\mathsf{hyb}}_A^{i-1}$ and $\overline{\mathsf{hyb}}_A^{i}$ where the execution is stopped after the commitment phase of the k^{th} right interaction. We define the output of $\overline{\mathsf{hyb}}_A^{i-1}$ and $\overline{\mathsf{hyb}}_A^{i}$ as the partial view in the main execution and the value extracted in the k^{th} right interaction.

Recall that the only difference between $\overline{\mathsf{hyb}}_A^{i-1}$ and $\overline{\mathsf{hyb}}_A^{i}$ is the simulation strategies of the i^{th} left interaction in the rewindings. More precisely, the first $i-1$ left interactions are simulated by choosing either a random sample from D or the actual value (if it has already been decommitted to in the partial view of the main execution) in both $\overline{\mathsf{hyb}}_A^{i-1}$ and $\overline{\mathsf{hyb}}_A^{i}$, but the i^{th} left interaction is simulated using a fixed commitment chosen ahead of time in $\overline{\mathsf{hyb}}_A^{i-1}$ and using a random commitment from D in $\overline{\mathsf{hyb}}_A^{i}$. Observe that, since the adversary cheats with small probability in $\overline{\mathsf{hyb}}_A^{i-1}$, it holds that for random samples chosen for the first $i-1$ left interactions and a fixed commitment for the i^{th} interaction, the value extracted in the k^{th} right interaction is the same with high probability. Hence, there must exist particular values $\boldsymbol{v}_{-i}$, v_i and $\overline{v_i}$ such that if the simulator uses the values $\boldsymbol{v}_{-i}$ and v_i for the left interactions (call this experiment E_1) and values $\boldsymbol{v}_{-i}$ and $\overline{v_i}$ in another experiment call it E_2, the probability that the values extracted in the k^{th} right interaction is different in both the experiments with probability at least $\frac{1}{p_2(n)}$ for some polynomial $p_2(\cdot)$.

[7] Conditioned on the the adversary corrupting the i^{th} left party the outputs of $\overline{\mathsf{hyb}}_A^{i-1}$ and $\overline{\mathsf{hyb}}_A^{i}$ are identical.

We now proceed as in the previous claim where we consider a simulator S_1^* that with T samples $s_1, \ldots, s_T$ from D_b and values $\boldsymbol{v}$, runs the adversary until the k^{th} right interaction and rewinds the k^{th} right interaction with values $\boldsymbol{v}$ on the left. From the indistinguishability of D_0 and D_1, we can conclude that the value extracted by the adversary in the k^{th} right interaction must be the same when the samples come from D_0 and D_1. By construction, when the samples $s_1, \ldots, s_T$ come from D_0 and values for left interactions are $\boldsymbol{v} = \boldsymbol{v}_{-i} \cup \{v_i\}$ the simulation proceeds identical to E_1 and when the samples come from D_0 and values for left interactions $\boldsymbol{v} = \boldsymbol{v}_{-i} \cup \{\overline{v_i}\}$ the simulation is identical to E_2. Let the corresponding experiments when the samples $s_1, \ldots, s_T$ come from D_1 and values from $\boldsymbol{v} = \boldsymbol{v}_{-i} \cup \{v_i\}$ and $\boldsymbol{v} = \boldsymbol{v}_{-i} \cup \{\overline{v_i}\}$ be E_1^* and E_2^*. In E_1^* and E_2^* all the commitments in the left are honestly generated. So we can consider corresponding experiments E_1^{**} and E_2^{**} where the puzzles are not simulated. By indistinguishability of D_0 and D_1, the values extracted from E_1^* and E_2^* will be different with non-negligible probability. By statistical-closeness of the puzzle simulation we can now conclude that values extracted in E_1^{**} and E_2^{**} are statistically close to E_1^* and E_2^* respectively and thus different with non-negligible probability. We are now ready to violate the stand-alone non-malleability of $\langle S, R \rangle$. Observe that in experiments E_1^{**} and E_2^{**} all values used in left interactions except the value in the i^{th} left interaction are the same. We now construct a man-in-the-middle adversary that forwards the i^{th} left interaction and k^{th} right interaction and this violates the non-malleability of $\langle S, R \rangle$.

This concludes the proof of Lemma 1 and Theorem 1.

3.2 Round-Efficient Adaptively Secure UC-Protocols

On a high-level constructing UC-secure protocol proceeds in two steps: (1) First, every functionality is compiled into a protocol in the $\mathcal{F}_{\mathsf{mcom}}$-hybrid model.[8] This step follows as corollary from previous works [CLOS02, DN00] and the round-complexity is preserved upto constant factors. (2) In the second step, assuming the existence of a UC-*puzzle* and a EQNMCom protocol, any protocol in the $\mathcal{F}_{\mathsf{mcom}}$ functionality can be securely realized in the real-model. This step was formalized and proved in [DSMRV13] and captured in the following Lemma.

Lemma 2 (Lemma 5[DSMRV13]). *Let Π' be an t_p-round protocol in the $\mathcal{F}_{\mathsf{mcom}}$-hybrid model. Assume the existence of a t_{puz}-round adaptive puzzle in a $\mathcal{G}$-hybrid model, a t_c-round* EQNMCom *protocol $\langle S, R \rangle$ and a simulatable PKE scheme. Then, there exists an $O(t_p(t_{puz} + t_c))$-round protocol Π in the $\mathcal{G}$-hybrid such that, for every uniform $\mathcal{PPT}$ adversary $\mathcal{A}$, there exists a simulator $\mathcal{A}'$, such that, no environment $\mathcal{Z} \in \mathcal{C}_{\mathsf{env}}$ can distinguish the real execution with $\mathcal{A}$ in $\mathcal{G}$-hybrid and the simulator $\mathcal{A}'$ in the $\mathcal{F}_{\mathsf{mcom}}$-hybrid.*

In [DSMRV13], it was shown how to construct $O(1)$-round adaptive puzzles for various models. In the previous section we showed how to construct a

[8] In the $\mathcal{F}_{\mathsf{mcom}}$-hybrid, all parties have access to the ideal commitment functionality called $\mathcal{F}_{\mathsf{mcom}}$ functionality.

$O(1)$-round EQNMCOM protocol based on any $O(1)$-round adaptive puzzle and one-way functions. Hence for these models, our work combined with previous Lemma yields an adaptive UC-secure protocol whose round complexity is $O(t_p)$. From previous works [DN00, CLOS02], we know that every well-formed functionality represented as a circuit C, can be realized in the $\mathcal{F}_{\mathsf{mcom}}$-hybrid in $O(d_C)$-rounds where d_C is the depth of the circuit C. Thus, we obtain the following corollary.

Corollary 31 *Assuming the existence of simulatable public-key encryption, any well-formed circuit C with depth d_C can be realized with adaptive UC-security in $O(d_C)$-rounds*

1. *in Common Reference String, Uniform Reference String, Imperfect Reference String (Sunspots) model, and Bounded-Concurrent model, or*
2. *with Quasi-polynomial Simulation and Non-uniform Simulation*

For more details on the puzzle, we refer the reader to [DSMRV13]. We remark that in [DSMRV13], they provide puzzles for the tamper-proof model and the timing model as well. The simulators for these puzzles are not straight-line and have been excluded from this work for simplicity. However, analogous to [DSMRV13], we believe it is possible to extend our result to these models.

4 An $\mathcal{ACNMZK}$ Argument

In this section we construct a fully adaptive concurrent non-malleable zero-knowledge proof based on one-way functions. The construction is inspired by the CCA-secure commitment scheme in [GLP+12]. Let $\langle C, R\rangle$ be a tag-based commitment scheme that is $O(1)$-robust $\mathsf{CCA-secure}$ w.r.t decommitment oracle $\mathcal{O}$. Let NMCom be a tag-based non-malleable commitment scheme (there exists $O(1)$-round protocols for such commitments [LP11b, Goy11]). Let $\langle P_{\mathsf{swi}}, V_{\mathsf{swi}}\rangle$ be a public-coin strong-$\mathcal{WI}$ argument of knowledge.

We now describe $\langle P, V\rangle$, our fully input-adaptive concurrent non-malleable zero-knowledge protocol. Protocol $\mathcal{ACNMZK}$ for a language $L \in \mathbf{NP}$ proceeds in four stages given a security parameter n, a common input statement $x \in \{0,1\}^n$, an identity id, and a private input $w \in R_L(x)$ to the Prover.

Stage 1. The Prover commits to w using $\langle C, R\rangle$. Let τ_1 be the commitment-transcript.

Stage 2. The Verifier chooses a random string $r \in \{0,1\}^n$ and commits to r using $\langle C, R\rangle$.

Stage 3. The Prover commits to 0^n using the NMCom scheme. Let τ_2 be the commitment-transcript.

Stage 4. The Prover proves using $\langle P_{\mathsf{swi}}, V_{\mathsf{swi}}\rangle$ that either
- $\exists y$ s.t. $y \in R_L(x)$ and τ_1 is a valid commitment to y as per $\langle C, R\rangle$, or
- τ_2 is a valid commitment to r as per NMCom

Completeness and Soundness follows using standard techniques. On a high-level, our simulator-extractor S proceeds in two stages. First we construct an oracle hybrid-simulator $\tilde{S}$ that incorporate A internally and emulates a man-in-the-middle interaction with A. $\tilde{S}$ with oracle access to $\mathcal{O}$, will forward all the commitments made by the adversary A in Stage 1 of right-interactions and Stage 2 of left interactions to $\mathcal{O}$. To simulate left interactions $\tilde{S}$ commits to 0^n in Stage 1. Upon receiving the decommitment r from $\mathcal{O}$ of the commitment made by the adversary in Stage 2, $\tilde{S}$ will commit to r in Stage 3 using the honest-committer strategy and use r as the fake witness in Stage 4 $\mathcal{WI}$-proof. Finally $\tilde{S}$ will output the view of A from its internal emulation along with all the decommitments obtained from $\mathcal{O}$ for the Stage 1 commitments made by A in the right interactions. These will serve as the witnesses corresponding to the statements proved A in the right-interactions. Using $\tilde{S}$ we construct the actual simulator-extractor S. This essentially follows from robust CCA-security of $\langle C, R\rangle$ w.r.t $\mathcal{O}$. Recall that 0-robust CCA-security implies that for every adversary with oracle access to $\mathcal{O}$ there exists a stand-alone simulator that outputs the same. Applying this to $\tilde{S}$ we obtain S. Proving correctness of the simulator relies on standard techniques and is presented in the full version. We additionally show in the full version how to construct a CCA-secure commitment scheme using an $\mathcal{ACNMZK}$ protocol.

4.1 Achieving UC-Security with Super-Polynomial Helpers

It follows from the works of [Lin03, Pas04] that constructing UC-secure protocols for any functionality boils down to realizing the zero-knowledge proof-of-membership functionality, often referred to as the MemberZK functionality. Consider an oracle-helper $\mathcal{H}$ that will provide proof of any statement (adaptively chosen) using the $\mathcal{ACNMZK}$ protocol described above to the adversary but not allow the adversary to use it to prove false statements to honest verifiers.[9] We show how to realize MemberZK-functionality in the angel-model where all parties have access to $\mathcal{H}$. The protocol is simply requiring the prover to use the $\mathcal{ACNMZK}$ protocol to prove the statement.

A subtle issue arises when compiling a protocol in the IdealZK-hybrid to the MemberZK-hybrid. The security reduction proves that for every adversary that only sends true statements to the MemberZK-functionality (also known as non-abusing adversaries) in the MemberZK-hybrid there is a simulator in the IdealZK-hybrid. This proof inherently assumes that an adversary remains non-abusing while receiving false proofs. Previous works that follow this paradigm do not prove this additional requirement when realizing the MemberZK-functionality.[10] We remark that our compilation of the protocol in the MemberZK-hybrid using a $\mathcal{ACNMZK}$-protocol by definition achieves the additional requirement.

[9] This can be achieved analogous to [CLP10], by providing the helper with the identity of corrupt parties and checking the session-id of the interaction before providing the proof.

[10] The result and the constructions presented in these works are nevertheless secure, since a direct and more cumbersome proof can prove their correctness.

References

[Blu86] Blum, M.: How to prove a theorem so no one else can claim it. In: Proceedings of the International Congress of Mathematicians, pp. 1444–1451 (1986)

[BPS06] Barak, B., Prabhakaran, M., Sahai, A.: Concurrent non-malleable zero knowledge. In: FOCS, pp. 345–354 (2006)

[Can01] Canetti, R.: Universally composable security: A new paradigm for cryptographic protocols. In: FOCS, pp. 136–145 (2001)

[CDSMW09] Choi, S.G., Dachman-Soled, D., Malkin, T., Wee, H.: Simple, black-box constructions of adaptively secure protocols. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 387–402. Springer, Heidelberg (2009)

[CF01] Canetti, R., Fischlin, M.: Universally composable commitments. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 19–40. Springer, Heidelberg (2001)

[CFGN96] Canetti, R., Feige, U., Goldreich, O., Naor, M.: Adaptively secure multi-party computation. In: STOC, pp. 639–648 (1996)

[CIO98] Di Crescenzo, G., Ishai, Y., Ostrovsky, R.: Non-interactive and non-malleable commitment. In: STOC, pp. 141–150 (1998)

[CKL03] Canetti, R., Kushilevitz, E., Lindell, Y.: On the limitations of universally composable two-party computation without set-up assumptions. In: Biham, E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, pp. 68–86. Springer, Heidelberg (2003)

[CKOS01] Di Crescenzo, G., Katz, J., Ostrovsky, R., Smith, A.: Efficient and non-interactive non-malleable commitment. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, pp. 40–59. Springer, Heidelberg (2001)

[CLOS02] Canetti, R., Lindell, Y., Ostrovsky, R., Sahai, A.: Universally composable two-party and multi-party secure computation. In: STOC, pp. 494–503 (2002)

[CLP10] Canetti, R., Lin, H., Pass, R.: Adaptive hardness and composable security in the plain model from standard assumptions. In: FOCS, pp. 541–550 (2010)

[CVZ10] Cao, Z., Visconti, I., Zhang, Z.: Constant-round concurrent non-malleable statistically binding commitments and decommitments. In: Nguyen, P.Q., Pointcheval, D. (eds.) PKC 2010. LNCS, vol. 6056, pp. 193–208. Springer, Heidelberg (2010)

[DDN00] Dolev, D., Dwork, C., Naor, M.: Nonmalleable cryptography. SIAM J. Comput. 30(2), 391–437 (2000)

[DM00] Dodis, Y., Micali, S.: Parallel reducibility for information-theoretically secure computation. In: Bellare, M. (ed.) CRYPTO 2000. LNCS, vol. 1880, pp. 74–92. Springer, Heidelberg (2000)

[DN00] Damgård, I.B., Nielsen, J.B.: Improved non-committing encryption schemes based on a general complexity assumption. In: Bellare, M. (ed.) CRYPTO 2000. LNCS, vol. 1880, pp. 432–450. Springer, Heidelberg (2000)

[DN02] Damgård, I.B., Nielsen, J.B.: Perfect hiding and perfect binding universally composable commitment schemes with constant expansion factor. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 581–596. Springer, Heidelberg (2002)

[DNS04] Dwork, C., Naor, M., Sahai, A.: Concurrent zero-knowledge. J. ACM 51(6), 851–898 (2004)

[DSMRV11] Dachman-Soled, D., Malkin, T., Raykova, M., Venkitasubramaniam, M.: Adaptive and concurrent secure computation from new notions of non-malleability. IACR Cryptology ePrint Archive, 2011:611 (2011)

[DSMRV13] Dachman-Soled, D., Malkin, T., Raykova, M., Venkitasubramaniam, M.: Adaptive and concurrent secure computation from new adaptive, non-malleable commitments. In: Sako, K., Sarkar, P. (eds.) ASIACRYPT 2013, Part I. LNCS, vol. 8269, pp. 316–336. Springer, Heidelberg (2013)

[GLP+12] Goyal, V., Lin, H., Pandey, O., Pass, R., Sahai, A.: Round-efficient concurrently composable secure computation via a robust extraction lemma. Cryptology ePrint Archive, Report 2012/652 (2012)

[GMR89] Goldwasser, S., Micali, S., Rackoff, C.: The knowledge complexity of interactive proof systems. SIAM J. Comput. 18(1), 186–208 (1989)

[GMW87] Goldreich, O., Micali, S., Wigderson, A.: How to play any mental game or a completeness theorem for protocols with honest majority. In: STOC, pp. 218–229 (1987)

[GO07] Groth, J., Ostrovsky, R.: Cryptography in the multi-string model. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 323–341. Springer, Heidelberg (2007)

[Goy11] Goyal, V.: Constant round non-malleable protocols using one way functions. In: STOC, pp. 695–704 (2011)

[Lin03] Lindell, Y.: Bounded-concurrent secure two-party computation without setup assumptions. In: STOC, pp. 683–692 (2003)

[LP11a] Lin, H., Pass, R.: Concurrent non-malleable zero knowledge with adaptive inputs. In: Ishai, Y. (ed.) TCC 2011. LNCS, vol. 6597, pp. 274–292. Springer, Heidelberg (2011)

[LP11b] Lin, H., Pass, R.: Constant-round non-malleable commitments from any one-way function. In: STOC, pp. 705–714 (2011)

[LPTV10] Lin, H., Pass, R., Tseng, W.-L.D., Venkitasubramaniam, M.: Concurrent non-malleable zero knowledge proofs. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 429–446. Springer, Heidelberg (2010)

[LPV08] Lin, H., Pass, R., Venkitasubramaniam, M.: Concurrent non-malleable commitments from any one-way function. In: Canetti, R. (ed.) TCC 2008. LNCS, vol. 4948, pp. 571–588. Springer, Heidelberg (2008)

[LPV09] Lin, H., Pass, R., Venkitasubramaniam, M.: A unified framework for concurrent security: universal composability from stand-alone non-malleability. In: STOC, pp. 179–188 (2009)

[LPV12] Pass, R., Lin, H., Venkitasubramaniam, M.: A unified framework for UC from only OT. In: Wang, X., Sako, K. (eds.) ASIACRYPT 2012. LNCS, vol. 7658, pp. 699–717. Springer, Heidelberg (2012)

[LZ09] Lindell, Y., Zarosim, H.: Adaptive zero-knowledge proofs and adaptively secure oblivious transfer. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 183–201. Springer, Heidelberg (2009)

[OPV06] Ostrovsky, R., Persiano, G., Visconti, I.: Concurrent non-malleable witness indistinguishability and its applications. Electronic Colloquium on Computational Complexity (ECCC) 13(095) (2006)

[OPV08] Ostrovsky, R., Persiano, G., Visconti, I.: Constant-round concurrent non-malleable zero knowledge in the bare public-key model. In: Aceto, L., Damgård, I., Goldberg, L.A., Halldórsson, M.M., Ingólfsdóttir, A., Walukiewicz, I. (eds.) ICALP 2008, Part II. LNCS, vol. 5126, pp. 548–559. Springer, Heidelberg (2008)

[OPV09] Ostrovsky, R., Persiano, G., Visconti, I.: Simulation-based concurrent non-malleable commitments and decommitments. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 91–108. Springer, Heidelberg (2009)

[Pas04] Pass, R.: Bounded-concurrent secure multi-party computation with a dishonest majority. In: STOC, pp. 232–241 (2004)

[PR05] Pass, R., Rosen, A.: Concurrent non-malleable commitments. In: FOCS, pp. 563–572 (2005)

[PS04] Prabhakaran, M., Sahai, A.: New notions of security: achieving universal composability without trusted setup. In: STOC, pp. 242–251 (2004)

[PW01] Pfitzmann, B., Waidner, M.: A model for asynchronous reactive systems and its application to secure message transmission. In: IEEE Symposium on Security and Privacy, p. 184 (2001)

[Sah99] Sahai, A.: Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In: FOCS, pp. 543–553 (1999)

[SCO+01] De Santis, A., Di Crescenzo, G., Ostrovsky, R., Persiano, G., Sahai, A.: Robust non-interactive zero knowledge. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 566–598. Springer, Heidelberg (2001)