Concurrent Non-Malleable Commitments*

Rafael Pass † Alon Rosen ‡

Abstract

We present a non-malleable commitment scheme that retains its security properties even when concurrently executed a polynomial number of times. That is, a man-in-the-middle adversary who is simultaneously participating in multiple concurrent *commitment phases* of our scheme, both as a sender and as a receiver, cannot make the values he commits to depend on the values he receives commitments to. Our result is achieved without assuming an a-priori bound on the number of executions and without relying on any set-up assumptions.

Our construction relies on the existence of *standard* claw-free permutations and only requires a constant number of communication rounds.

1 Introduction

The notion of commitment is central in cryptographic protocol design. Often described as the "digital" analogue of sealed envelopes, commitment schemes enable a party, known as the *sender*, to commit itself to a value while keeping it secret from the *receiver*. This property is called *hiding*. Furthermore, the commitment is *binding*, and thus in a later stage when the commitment is opened, it is guaranteed that the "opening" can yield only a single value determined in the committing stage.

For some applications, the above security guarantees are not sufficient and additional properties are required. For instance, the definition of commitments does not rule out the possibility that an adversary, upon seeing a commitment to a specific value v, is able to commit to a related value (say, $v - 1$), even though it does not know the actual value of v. This kind of attack might have devastating consequences if the underlying application relies on the *independence* of committed values (e.g., consider a case in which the commitment scheme is used for securely implementing a contract bidding mechanism). The state of affairs is even worsened by the fact that many of the known commitment schemes are actually susceptible to this kind of attack.

1.1 Non-Malleable Commitments

In order to address the above concerns, Dolev, Dwork and Naor (DDN) introduced the concept of *non-malleable commitments* [16]. Loosely speaking, a commitment scheme is said to be non-malleable if no adversary can succeed in the attack described above. That is, it is infeasible for the adversary to maul a commitment to a value v into a commitment to a "related" value $\tilde{v}$.

The first non-malleable commitment protocol was constructed by Dolev, Dwork and Naor [16]. The security of their protocol relies on the existence of one-way functions, and requires $O(\log n)$ rounds of interaction, where $n \in N$ is a security parameter. A more recent result by Barak presents

*Preliminary version appeared in FOCS 2005, pages 563-572, 2005.

†Cornell University. E-mail: `pass@csail.mit.edu`. Part of work done while at CSAIL, MIT.

‡Harvard DEAS, Cambridge, MA. E-mail: `alon@eecs.harvard.edu`. Part of work done while at CSAIL, MIT.

a constant-round protocol for non-malleable commitment, whose security relies on the existence of trapdoor permutations and hash functions that are collision-resistant against sub-exponential sized circuits [2]. Even more recently, Pass and Rosen present a constant-round protocol for the same task, assuming only collision resistant hash function secure against polynomial sized circuits [42].

1.2 Concurrent Non-Malleable Commitments

The basic definition of non-malleable commitments only considers a scenario in which two executions take place at the same time. A natural extension of this scenario (already suggested in [16]) is one in which more than two invocations of the commitment protocol take place concurrently. In the concurrent scenario, the adversary is receiving commitments to multiple values $v_1, \ldots, v_m$, while attempting to commit to related values $\tilde{v}_1, \ldots, \tilde{v}_m$. As argued in [16], non-malleability with respect to two executions can be shown to guarantee *individual* independence of any $\tilde{v}_i$ from any v_j. However, it does not rule out the possibility of an adversary to create *joint* dependencies between more than a single individual pair (see [16], Section 3.4.1 for an example in the context of non-malleable encryption). Resolving this issue has been stated as a major open problem in [16].

Partially addressing this issue, Pass demonstrates the existence of commitment schemes that remain non-malleable under *bounded concurrent* composition [40]. That is, for any (predetermined) polynomial $p(\cdot)$, there exists a non-malleable commitment that remains secure as long as it is not executed more than $p(n)$ times, where $n \in N$ is a security parameter.

One evident disadvantage of the above solution is that it requires that the number of executions is fixed *before* the protocol is specified, or otherwise no security guarantee is provided. Less evidently, the length of the messages in the protocols has to grow linearly with the number of executions. Thus, from both a theoretical and a practical point of view, the solution is still not satisfactory. What we would like to have is a *single* protocol that preserves its non-malleability even when it is executed concurrently for *any* (not predetermined) polynomial number of times.

1.3 Our Results

We present a new protocol for *concurrent non-malleable* commitments. Our protocol remains non-malleable even when concurrently executed an (unbounded) polynomial number of times. We do not rely on any kind of set-up assumption (such as the existence of a common reference string).

The resulting commitment is *statistically binding*, and satisfies non-malleability *with respect to commitment*. The former condition implies that, except with negligible probability, a transcript of a commitment corresponds to a unique value, whereas the latter implies that, upon concurrently participating in polynomially many commitments, both as a receiver and as a sender, the adversary is not able to *commit* to a sequence of related values.[1] Here we assume that the adversary does not get to see the de-commitment to any of the values he is receiving a commitment to until he is done with committing to all of his values.

Theorem 1 (Concurrent non-malleable commitment) *Suppose that there exists a family of pairs of claw-free permutations.*[2] *Then, there exists a constant-round statistically-binding commit-*

[1] In a different variant, called non-malleable commitment *with respect to opening* [19], the adversary is considered to have succeeded only if it manages to *de-commit* to a related value. This paper only considers the notion of non-malleability with respect to commitments.

[2] The existence of claw-free permutations follows from the assumption that factoring Blum integers is hard (or from the hardness of finding discrete-logarithms modulo a prime). They are required for obtaining *perfectly* hiding-commitments, as well as collision resistant hashing.

ment scheme that is concurrently non malleable with respect to commitment.

To the best of our knowledge, this result yields the first instance of a non-trivial protocol that simultaneously satisfies non-malleability and unbounded concurrency without relying on set-up assumptions.

Additional contributions. Our proof also yields the first commitment scheme that satisfies non-malleability using a *strict* polynomial-time simulator (a.k.a. strict non-malleability) with respect to commitment.[3] By this we mean that the simulation used to prove non-malleability runs in strict (as opposed to expected) polynomial time. This was the security notion originally defined (but not achieved in) the DDN paper [16].

Our definitions of non-malleable commitments are somewhat different (stronger) than the ones appearing in the DDN paper [16]. Specifically, we formalize the notion of two values being unrelated through the concept of computational indistinguishability (rather than using polynomial time computable relations). The main reason for strengthening the definition is that it yields a notion that is more intuitive and easier to work with (especially in the concurrent setting). We stress that any protocol satisfying our definition also satisfies the original one.

Techniques and ideas. Our construction follows the paradigm introduced by Pass and Rosen, of using a protocol for non-malleable zero-knowledge in order to obtain (single execution) non-malleable commitments [42], and relies on the "message-length" technique of Pass [40]. While our construction relies on the same high-level structure, the analysis of the protocol is significantly different. The central observation that enables the analysis is that concurrent simulation of the underlying (non-malleable) zero-knowledge protocol is not actually necessary for proving concurrent non-malleability of our commitments. Indeed, for our analysis to go through, it will be sufficient to simulate only a *single* execution of the underlying zero-knowledge protocol. This will be performed while concurrently extracting *multiple* witnesses for the statements proved by the adversary. We call the above property *one-many simulation extractability.* We prove that this property is indeed satisfied by a variant of the non-malleable zero-knowledge protocols of [40, 42]. To show this, we rely on a *non-black box* simulation argument, which is delicately combined with a *black-box* extraction technique. (Here we use the fact that concurrent extraction is significantly easier than concurrent simulation (cf. [33]).)

1.4 Related Work

A large body of previous work deals with the construction of non-malleable protocols assuming various kinds of trusted set-up. Known constructions include non-malleable commitment schemes assuming the existence of a common reference string [19, 11], as well as non-malleable commitment schemes and non-interactive non-malleable $\mathcal{ZK}$ protocols assuming the existence of a common random string [15, 14, 13].

Several of the above works explicitly address the issue of multiple executions of non-malleable schemes [13, 11, 9] (also called *reusability* in the terminology of [11]). Perhaps most notable amongst the works addressing concurrency, is the one on Universally composable commitments [9]. Universal composability implies concurrent non-malleability. However, it is impossible to construct universally composable commitments without making set-up assumptions [9].

[3] This should not be confused with a previous result showing the existence of commitment schemes that are strictly non-malleable with respect to *opening* [42].

Other related works involve the task of session-key generation in a setting where the honest parties share a password that is taken from a relatively small dictionary [22, 39, 3]. These protocols are designed having a man-in-the-middle adversary in mind, and only require the usage of a "mild" set-up assumption (namely the existence of a "short" password). Some of these works explicitly address the issue of multiple protocol execution (cf. [22]), but their treatment is limited to the case of sequential composition. A treatment of the full concurrent case appears in [31] (see also [10, 3]), but it relies on the existence of a common reference string.

2 Preliminaries

2.1 Basic notation

We let N denote the set of all integers. For any integer $m \in N$, denote by $[m]$ the set $\{1, 2, \ldots, m\}$. For any $x \in \{0,1\}^*$, we let $|x|$ denote the size of x (i.e., the number of bits used in order to write it). For two machines M, A, we let $M^A(x)$ denote the output of machine M on input x and given oracle access to A. The term negligible is used for denoting functions that are (asymptotically) smaller than one over any polynomial. More precisely, a function $\nu(\cdot)$ from non-negative integers to reals is called negligible if for every constant $c > 0$ and all sufficiently large n, it holds that $\nu(n) < n^{-c}$.

2.2 Witness Relations

We recall the definition of a witness relation for an $\mathcal{NP}$ language [20].

Definition 2.1 (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 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 Computational indistinguishability and statistical closeness

The following definition of (computational) indistinguishability originates in the seminal paper of Goldwasser and Micali [26].

Let X be a countable set of strings. A probability ensemble indexed by X is a sequence of random variables indexed by X. Namely, any $A = \{A_x\}_{x \in X}$ is a random variable indexed by X.

Definition 2.2 ((Computational) Indistinguishability) *Let X and Y be countable sets. Two ensembles $\{A_{x,y}\}_{x \in X, y \in Y}$ and $\{B_{x,y}\}_{x \in X, y \in Y}$ are said to be computationally indistinguishable over X, if for every probabilistic "distinguishing" machine D whose running time is polynomial in its first input, there exists a negligible function $\nu(\cdot)$ so that for every $x \in X, y \in Y$:*

$$|\Pr[D(x, y, A_{x,y}) = 1] - \Pr[D(x, y, B_{x,y}) = 1]| < \nu(|x|)$$

$\{A_{x,y}\}_{x \in X, y \in Y}$ and $\{B_{x,y}\}_{x \in X, y \in Y}$ are said to be statistically close over X if the above condition holds for all (possibly unbounded) machines D.

2.4 Interactive Proofs, Zero-Knowledge and Witness-Indistinguishability

We use the standard definitions of interactive proofs (and interactive Turing machines) [27, 20] and arguments [4]. Given a pair of interactive Turing machines, P and V, we denote by $\langle P, V\rangle(x)$ the random variable representing the (local) output of V when interacting with machine P on common input x, when the random input to each machine is uniformly and independently chosen.

Definition 2.3 (Interactive Proof System) *A pair of interactive machines $\langle P, V\rangle$ is called an* interactive proof system *for a language L if machine V is polynomial-time and the following two conditions hold with respect to some negligible function $\nu(\cdot)$:*

- Completeness: *For every $x \in L$,*
$$\Pr[\langle P, V\rangle(x) = 1] \geq 1 - \nu(|x|)$$
- 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.*

Zero-knowledge. An interactive proof is said to be *zero-knowledge* ($\mathcal{ZK}$) if it yields nothing beyond the validity of the assertion being proved. This is formalized by requiring that the view of every probabilistic polynomial-time adversary V^* interacting with the honest prover P can be simulated by a probabilistic polynomial-time machine S (a.k.a. the *simulator*). The idea behind this definition is that whatever V^* might have learned from interacting with P, he could have actually learned by himself (by running the simulator S).

The notion of $\mathcal{ZK}$ was introduced by Goldwasser, Micali and Rackoff [27]. To make $\mathcal{ZK}$ robust in the context of protocol composition, Goldreich and Oren [25] suggested to augment the definition so that the above requirement holds also with respect to all $z \in \{0,1\}^*$, where both V^* and S are allowed to obtain z as auxiliary input. The verifier's view of an interaction consists of the common input x, followed by its random tape and the sequence of prover messages the verifier receives during the interaction. We denote by $\text{view}^P_{V^*}(x, z)$ a random variable describing $V^*(z)$'s view of the interaction with P on common input x.

Definition 2.4 (Zero-knowledge) *Let $\langle P, V\rangle$ be an interactive proof system. We say that $\langle P, V\rangle$ is* zero-knowledge*, if for every probabilistic polynomial-time interactive machine V^* there exists a probabilistic polynomial-time algorithm S such that the ensembles $\{\text{view}^P_{V^*}(x, z)\}_{x\in L, z\in\{0,1\}^*}$ and $\{S(x, z)\}_{x\in L, z\in\{0,1\}^*}$ are computationally indistinguishable over L.*

A stronger variant of zero-knowledge is one in which the output of the simulator is statistically close to the verifier's view of real interactions. We focus on *argument* systems, in which the soundness property is only guaranteed to hold with respect to polynomial time provers.

Definition 2.5 (Statistical zero-knowledge) *Let $\langle P, V\rangle$ be an interactive argument system. We say that $\langle P, V\rangle$ is* statistical zero-knowledge*, if for every probabilistic polynomial-time V^* there exists a probabilistic polynomial-time S such that the ensembles $\{\text{view}^P_{V^*}(x, z)\}_{x\in L, z\in\{0,1\}^*}$ and $\{S(x, z)\}_{x\in L, z\in\{0,1\}^*}$ are statistically close over L.*

In case that the ensembles $\{\text{view}^P_{V^*}(x, z)\}_{x\in L, z\in\{0,1\}^*}$ and $\{S(x, z)\}_{x\in L, z\in\{0,1\}^*}$ are identically distributed, the protocol $\langle P, V\rangle$ is said to be *perfect* zero-knowledge.

Witness Indistinguishability. An interactive proof is said to be *witness indistinguishable* ($\mathcal{WI}$) if the verifier's view 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 view 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 views are computationally indistinguishable.

Let V^* be a probabilistic polynomial time adversary interacting with the prover, and let $\text{view}^P_{V^*}(x, w, z)$ denote V^*'s view of an interaction in which the witness used by the prover is w (where the common input is x and V^*'s auxiliary input is z).

Definition 2.6 (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 probabilistic polynomial-time interactive machine V^* and for every two sequences $\{w^1_x\}_{x\in L}$ and $\{w^2_x\}_{x\in L}$, such that $w^1_x, w^2_x \in R_L(x)$ for every $x \in L$, the probability ensembles $\{\text{view}^P_{V^*}(x, w^1_x, z)\}_{x\in L, z\in\{0,1\}^*}$ and $\{\text{view}^P_{V^*}(x, w^2_x, z)\}_{x\in L, z\in\{0,1\}^*}$ are computationally indistinguishable over L.*

In case that the ensembles $\{\text{view}^P_{V^*}(x, w^1_x, z)\}_{x\in L, z\in\{0,1\}^*}$ and $\{\text{view}^P_{V^*}(x, w^2_x, z)\}_{x\in L, z\in\{0,1\}^*}$ are identically distributed, the proof system $\langle P, V\rangle$ is said to be witness *independent*.

2.5 Universal Arguments

Universal arguments (introduced in [5] and closely related to the notion of CS-proofs [34]) are used in order to provide "efficient" proofs to statements of the form $y = (M, x, t)$, where y is considered to be a true statement if M is a non-deterministic machine that accepts x within t steps. The corresponding language and witness relation are denoted $L_{\mathcal{U}}$ and $R_{\mathcal{U}}$ respectively, where the pair $((M, x, t), w)$ is in $R_{\mathcal{U}}$ if M (viewed here as a two-input deterministic machine) accepts the pair (x, w) within t steps. Notice that every language in $\mathcal{NP}$ is linear time reducible to $L_{\mathcal{U}}$. Thus, a proof system for $L_{\mathcal{U}}$ allows us to handle all $\mathcal{NP}$-statements. In fact, a proof system for $L_{\mathcal{U}}$ enables us to handle languages that are presumably "beyond" $\mathcal{NP}$, as the language $L_{\mathcal{U}}$ is $\mathcal{NE}$-complete (hence the name universal arguments).[4]

Definition 2.7 (Universal argument) *A pair of interactive Turing machines (P, V) is called a* universal argument *system if it satisfies the following properties:*

- Efficient verification: *There exists a polynomial p such that for any $y = (M, x, t)$, the total time spent by the (probabilistic) verifier strategy V, on common input y, is at most $p(|y|)$. In particular, all messages exchanged in the protocol have length smaller than $p(|y|)$.*

- Completeness by a relatively efficient prover: *For every $((M, x, t); w)$ in $R_{\mathcal{U}}$,*

  $$\Pr[(P(w), V)(M, x, t) = 1] = 1$$

  Furthermore, there exists a polynomial q such that the total time spent by $P(w)$, on common input (M, x, t), is at most $q(T_M(x, w)) \le q(t)$, where $T_M(x, w)$ denotes the running time of M on input (x, w).

[4]Furthermore, every language in $\mathcal{NEXP}$ is polynomial-time (but not linear-time) reducible to $L_{\mathcal{U}}$

- Computational Soundness: *For every polynomial size circuit family* $\{P_n^*\}_{n\in N}$*, and every triplet* $(M,x,t) \in \{0,1\}^n \setminus L_{\mathcal{U}}$*,*

$$\Pr[(P_n^*, V)(M;x;t) = 1] < \nu(n)$$

where $\nu(\cdot)$ *is a negligible function.*

- Weak proof of knowledge: *For every positive polynomial* p *there exists a positive polynomial* p' *and a probabilistic polynomial-time oracle machine* E *such that the following holds: for every polynomial-size circuit family* $\{P_n^*\}_{n\in N}$*, and every sufficiently long* $y = (M;x;t) \in \{0,1\}^*$ *if* $\Pr[(P_n^*; V)(y) = 1] > 1/p(|y|)$ *then*

$$\Pr[\exists w = w_1, \ldots w_t \in R_{\mathcal{U}}(y) \;\; s.t. \;\; \forall i \in [t], \; E_r^{P_n^*}(y;i) = w_i] > \frac{1}{p'(|y|)}$$

where $R_{\mathcal{U}}(y) \stackrel{\text{def}}{=} \{w : (y,w) \in R_{\mathcal{U}}\}$ *and* $E_r^{P_n^*}(\cdot,\cdot)$ *denotes the function defined by fixing the random-tape of* E *to equal* r*, and providing the resulting* E_r *with oracle access to* P_n^**.*

2.6 Commitment Schemes

Commitment schemes are used to enable a party, known as the *sender*, to commit itself to a value while keeping it secret from the *receiver* (this property is called hiding). Furthermore, the commitment is binding, and thus in a later stage when the commitment is opened, it is guaranteed that the "opening" can yield only a single value determined in the committing phase. Commitment schemes come in two different flavors, statistically-binding and statistically-hiding. We sketch the properties of each one of these flavors. Full definitions can be found in [20].

Statistically-binding: In statistically binding commitments, the binding property holds against unbounded adversaries, while the hiding property only holds against computationally bounded (non-uniform) adversaries. Loosely speaking, the statistical-binding property asserts that, with overwhelming probability over the coin-tosses of the receiver, the transcript of the interaction fully determines the value committed to by the sender. The computational-hiding property guarantees that the commitments to any two different values are computationally indistinguishable.

Statistically-hiding: In statistically-hiding commitments, the hiding property holds against unbounded adversaries, while the binding property only holds against computationally bounded (non-uniform) adversaries. Loosely speaking, the statistical-hiding property asserts that commitments to any two different values are statistically close (i.e., have negligible statistical distance). In case the statistical distance is 0, the commitments are said to be *perfectly-hiding*. The computational-binding property guarantees that no polynomial time machine is able to open a given commitment in two different ways.

Non-interactive perfectly-binding commitment schemes can be constructed using any 1–1 one-way function (see Section 4.4.1 of [20]). Allowing some minimal interaction (in which the receiver first sends a single random initialization message), statistically-binding commitment schemes can be obtained from any one-way function [35, 30]. We will think of such commitments as a *family* of non-interactive commitments, where the description of members in the family will be the initialization message. Statistically-hiding commitment schemes can be constructed from any one-way function [36, 28], but constant-round schemes are only know to exists under stronger assumptions; specifically, assuming collision-resistant hashfunctions [12, 29]. Perfectly-hiding commitment schemes, on the other hand, can be constructed from one-way permutations [37], but constant-round schemes are only known under stronger assumptions, such as the existence of a collection of certified claw-free permutations [21].

2.7 Proofs of Knowledge

Informally an interactive proof is a proof of knowledge if the prover convinces the verifier not only of the validity of a statement, but also that it possesses a witness for the statement. This notion is formalized by the introduction of an machine E, called a knowledge extractor. As the name suggests, the extractor E is supposed to extract a witness from any malicious prover P^* that succeeds in convincing an honest verifier. More formally,

Definition 2.8 (Proof of Knowledge) *Let (P, V) be an interactive proof system for the language L with witness relation R_L. We say that (P, V) is a proof of knowledge if there exists a polynomial q and a probabilistic oracle machine E, such that for every probabilistic polynomial-time interactive machine P^*, there exists some negligible function $\mu(\cdot)$ such that for every $x \in L$ and every $y, r \in \{0,1\}^*$ such that $\Pr[\langle P^*_{x,y,r}, V(x)\rangle = 1] > 0$, where $P^*_{x,y,r}$ denotes the machine P^* with common input fixed to x, auxiliary input fixed to y and random tape fixed to r, the following holds*

1. *The expected number of steps taken by $E^{P^*_{x,y,r}}$ is bounded by*
$$\frac{q(|x|)}{\Pr[\langle P^*_{x,y,r}, V(x)\rangle = 1]}$$
where $E^{P^*_{x,y,r}}$ denotes the machine E with oracle access to $P^*_{x,y,r}$.

2. *Furthermore,*
$$\Pr[\langle P^*_{x,y,r}, V(x)\rangle = 1 \wedge E^{P^*_{x,y,r}} \notin R_L(x)] \leq \mu(|x|)$$

The machine E is called a (knowledge) extractor.

We remark that as our definition only considers computationally bounded provers, we only get a "computationally convincing" notion of a proof of knowledge (a.k.a *arguments of knowledge*) [4]. We additionally point out that our definition is slightly different from the definition of [6] in that we require that the expected running-time of the extractor is always bounded by $\mathrm{poly}(|x|)/p$, where p denotes the success probability of P^*, whereas [6] allows for some additional slackness in the running-time. On the other hand, whereas [6] requires the extractor to always output a valid witness, we instead allow the extractor to fail with some negligible probability. We will rely on the following theorem:

Theorem 2.9 ([7, 4]) *Assume the existence of claw-free permutations. Then there exists a constant-round public-coin witness independent argument of knowledge for $\mathcal{NP}$.*

Indeed, standard techniques can be used to show that the parallelized version of the protocol of [7], using perfectly-hiding commitments, is an argument of knowledge (as defined above). As usual: the knowledge extractor E proceeds by feeding new "challenges" to the prover P^* until it gets two accepting transcripts. If the two accepting challenges contain the same challenge, or if the prover manages to open up a commitment in two different ways, the extractor outputs `fail`; otherwise it can extract a witness.

3 Concurrent Non-Malleable Commitments

Non-malleable commitments were introduced by Dolev, Dwork and Naor (DDN) [16]. Our definitions of non-malleability are somewhat stronger that the ones proposed by DDN [16]. Specifically, we formalize the notion of two values being unrelated through the concept of computational indistinguishability (rather than using polynomial time computable relations).

3.1 The General Setting

Let $\langle C, R\rangle$ be a commitment scheme and consider a man-in-the-middle adversary A that is simultaneously participating in multiple concurrent executions of $\langle C, R\rangle$. Executions in which A is playing the role of the receiver are said to belong to the left interaction, whereas executions in which A is playing the role of the sender are said to belong to the right interaction. We assume for simplicity, and without loss of generality, that the number of commitment schemes taking place in the left and right interactions is identical. The total number of the interactions in which the adversary is involved (either as a sender or as a receiver) is not a-priori bounded by any polynomial (though it is assumed to be polynomial in the security parameter). We assume that the adversary does not get to see the de-commitment to any of the values he is receiving a commitment to until he is done with committing to all of his values.

Besides controlling the messages that it sends in the left and right interactions, A has control over their scheduling. In particular, it may delay the transmission of a message in one interaction until it receives a message (or even multiple messages) in the other interaction. It can also arbitrarily interleave messages that belong to different executions within an interaction.

The adversary A is trying to take advantage of his participation in the commitments taking place in the left interaction in order to commit to a related value in the right interaction. The honest sender and receiver are not necessarily aware to the existence of the adversary, and might be under the impression that they are interacting one with the other. We let $v_1, \ldots, v_m$ denote the values committed to in the left interaction and $\tilde{v}_1, \ldots, \tilde{v}_m$ denote the values committed to in the right interaction.[5] The above scenario is depicted in Figure 1 (with no explicit demonstration of possible interleavings of messages between different executions).

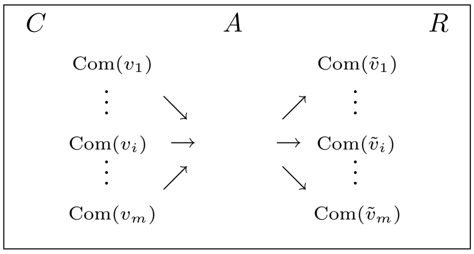

Figure 1: A concurrent man-in-the-middle adversary.

The traditional definition of non-malleable commitments [16] considers the case when $m = 1$. Loosely speaking, it requires that the left interaction does not "help" the adversary A in committing to a value $\tilde{v}_1$ that is somehow correlated with the value v_1. In this work we focus on non-malleability with respect to commitment [16], where the adversary is said to succeed if it manages to *commit*

[5]The adversary does not necessarily "know" the values $\tilde{v}_1, \ldots, \tilde{v}_m$. However, since the commitment is statistically binding, this value is (almost always) well defined.

to a related value (even without being able to later de-commit to this value). Note that this notion makes sense only in the case of statistically-binding commitments.

3.2 Non-Malleability via Indistinguishability

Following the simulation paradigm [26, 27, 23, 24], the notion of non-malleability is formalized by comparing between a *man-in-the-middle* and a *simulated* execution. In the man-in-the-middle execution the adversary is simultaneously acting as a receiver in one interaction and as a committer in another interaction. In the simulated execution the adversary is engaged in a single interaction where it is acting as a committer.

The original definition of non malleability required that for any polynomial time computable (non-reflexive) relation $\mathcal{R}$, the value $\tilde{v}$ committed to by the adversary in the simulated execution is no (significantly) less likely to satisfy $\mathcal{R}(v, \tilde{v}) = 1$ than the value committed to by the adversary in the man-in-the-middle execution [16].

To facilitate the formalization for $m > 1$, we choose to adopt a slightly different definitional approach and will actually require an even stronger condition (which we are still able to satisfy with our protocol). Specifically, we require that for any adversary in a man-in-the-middle execution, there exists an adversary that commits to essentially the same value in the simulated execution. By essentially the same value, we mean that the value committed to by the simulator is computationally indistinguishable from the value committed to by the adversary in the man-in-the middle execution.

Since copying cannot be ruled out, we will only be interested in the case where copying is not considered success. We therefore impose the condition that whenever the adversary has *fully* copied a transcript of an interaction in which it acts as a receiver, the value $\tilde{v}$ that he has committed to in the corresponding execution is set to be a special "failure" symbol, denoted $\perp$.

3.3 The Actual Definition

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.

The man-in-the-middle 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 $v_1, \ldots, v_m$. In the right interaction A interacts with R attempting to commit to a sequence of related values $\tilde{v}_1, \ldots, \tilde{v}_m$. Prior to the interaction, the values $v_1, \ldots, v_m$ are given to C as local input. A receives an auxiliary input z, which may contain a-priori information about $v_1, \ldots, v_m$. Let $\mathsf{mim}^A_{\mathsf{com}}(v_1, \ldots, v_m, z)$ denote a random variable that describes the values $\tilde{v}_1, \ldots, \tilde{v}_m$ to which the adversary has committed in the right interaction. (Since we are dealing with statistically binding commitments, $\tilde{v}_1, \ldots, \tilde{v}_m$ are (almost always) well defined. Whenever the value of the commitment is not uniquely defined (which can happen with some negligible probability in case of statistically binding commitments), the value of the commitment is defined to be $\perp$.) If the transcript of the i^{th} right commitment is identical to the transcript of *any* of the left interactions (which means that adversary has fully copied a specific commitment that has taken place on the left), the value $\tilde{v}_i$ is set to be $\perp$.[6]

[6] This approach allows $\tilde{v}_i = v_j$ for some $i, j \in \{1, \ldots, m\}$, as long as the man-in-the-middle does not fully copy the messages from one of the left executions. This is in contrast to the original definition which does not handle the case of $\tilde{v}_i = v_j$ (as $\mathcal{R}$ is non-reflexive). This means that the new approach takes into consideration a potentially larger class of attacks.

The simulated execution. In the simulated execution a simulator S directly interacts with R. As in the man-in-the-middle execution, the values $v_1, \ldots, v_m$ are chosen prior to the interaction and S receives some a-priori information about $v_1, \ldots, v_m$ as part of its auxiliary input z. We let $\mathsf{sim}^S_{\mathsf{com}}(z)$ denote a random variable that describes the values committed to in the output of S (which consists of a sequence of values $\tilde{v}_1, \ldots, \tilde{v}_m$).

Definition 3.1 *A commitment scheme $\langle C, R\rangle$ is said to be* concurrent non-malleable with respect to commitment *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 $\{0,1\}^*$:*

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

It can be seen that for $m = 1$ any protocol that satisfies Definition 3.1 also satisfies the original (relation based) definition of non-malleability of [16]. Loosely speaking, this is because the existence of a polynomial time computable relation $\mathcal{R}$ that violates the original definition of non-malleability could be used to distinguish between the values of $\mathsf{mim}^A_{\mathsf{com}}(v, z)$ and $\mathsf{sim}^S_{\mathsf{com}}(z)$. We additionally point out that even when only considering the case when $m = 1$, Definition 3.1 is stronger than relation-based definition in that it prevents an adversary from producing a different commitment to the same value that it receives a commitment to; the original definition did not consider this as a successful attack.[7]

3.4 One-Many Concurrent Non-Malleable Commitments

A seemingly more relaxed (and thus potentially easier to satisfy) notion of concurrent non-malleable commitments is one in which the man-in-the-middle adversary A engages in only a *single* commitment protocol in the left interaction (but still polynomially many in the right interaction). Such a notion is a special case of Definition 3.1 in which the adversary A participates in only one commitment session on the left hand side (instead of m sessions).

A commitment protocol that satisfies the relaxed definition is said to be *one-many concurrent non-malleable*. As we argue below, the relaxed notion "in essence" implies full-fledged non-malleability. In particular, in order to construct concurrent non-malleable commitments, it will be sufficient to come up with a protocol that is one-many concurrent non-malleable. To formalize this composability property we need to restrict our attention to certain "natural" commitment schemes.[8] We say that a commitment scheme is *natural* if any commitment where the committer aborts (sending $\perp$) before the end of the protocols makes the commitment invalid (i.e., a commitment to $\perp$).

Proposition 3.2 *Let $\langle C, R\rangle$ be a natural one-many concurrent non-malleable commitment. Then, $\langle C, R\rangle$ is also a (full fledged) concurrent non-malleable commitment.*

[7] In contrast, the definition of non-malleable *encryption* of [16] indeed also prevents this type of attack. In a sense, our definition is thus more in line with their definition of non-malleability for encryptions schemes.

[8] Similar technical restrictions are needed to show composability of non-malleable encryption [43].

Proof: Let A be a man-in-the-middle adversary that participates in at most $m = p(n)$ concurrent executions. We show the existence of a simulator S such that the following ensembles are computationally indistinguishable over $\{0,1\}^*$:

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

The simulator S proceeds as follows on input z. S incorporates $A(z)$ and internally emulates all the left interactions for A by simply *honestly* committing to the string 0^n (i.e., in order to emulate the i^{th} left interaction, S executes the algorithm C on input 0^n). Messages from the right interactions are instead forwarded externally, with the following exception: whenever A wishes to send the last message q_i in the i'th right session, S "holds-on" to it without (yet) forwarding it externally. Finally, when S has completed the emulation of all left interactions for A, it checks whether A fully copied any of the left executions. For each execution i where A fully copied one of the left executions, S externally sends $\perp$ as its last message in the i'th right execution (to invalidate that commitment); for all other executions j, A instead sends the final message q_j.

We show that the values that S commits to are indistinguishable from the values that A commits to. Suppose, for contradiction, that this is not the case. That is, there exists a polynomial-time distinguisher D and a polynomial $p(n)$ such that for infinitely many n, there exist strings $v_1, \ldots, v_m \in \{0,1\}^n, z \in \{0,1\}^*$ such that

$$\Pr\left[D(\mathsf{mim}^A_{\mathsf{com}}(v_1, \ldots, v_m, z) = 1\right] - \Pr\left[D(\mathsf{sim}^S_{\mathsf{com}}(z)) = 1\right] \geq \frac{1}{p(n)}$$

Fix a generic n for which this happens. We provide a hybrid argument that will contradict the one-many non-malleability of $\langle C, R\rangle$. The "hybrid" random variable $\mathsf{hyb}_k(v_1, \ldots, v_m, z)$ involves an execution where $A(z)$ is participating in m left and m right interactions, and is defined in the following way:

- For $j \leq k$, the j^{th} session in the left interaction consists of a commitment to 0^n.
- For $j > k$, the j^{th} session in the left interaction consists of a commitment to v_j.
- Output the values $\tilde{v}_1, \ldots, \tilde{v}_m$ committed to by A in the right interactions with an honest R. As in the definition of $\mathsf{mim}_{\mathsf{com}}$ (see Definition 3.1), the value $\tilde{v}_i$ of a commitment c_i is set to $\perp$ if A fully copies one of the left executions.

Note that the values $\tilde{v}_1, \ldots, \tilde{v}_m$ are not efficiently computable, but are well defined nevertheless. Just as in Definition 3.1, if a commitment can be opened to two (or more) different values, we set its value to $\perp$. By construction, it directly follows that:

$$\mathsf{hyb}_0(v_1, \ldots, v_m, z) = \mathsf{mim}^A_{\mathsf{com}}(v_1, \ldots, v_m, z)$$

By additionally relying on the naturality property of $\langle C, R\rangle$, it holds that:[9]

$$\mathsf{hyb}_m(v_1, \ldots, v_m, z) = \mathsf{sim}^S_{\mathsf{com}}(z)$$

[9]The naturality property together with the construction of S is used to make sure that S produces a commitment to $\perp$ whenever A copies the left interaction.

It follows by a standard hybrid argument that there exists an $i \in [m]$ such that

$$\Pr\left[D(\mathsf{hyb}_{i-1}(v_1,\ldots,v_m,z)=1\right] - \Pr\left[D(\mathsf{hyb}_i(v_1,\ldots,v_m,z))=1\right] \geq \frac{1}{p(n)m}$$

Note that the only difference between the experiments $\mathsf{hyb}_{i-1}(v_1,\ldots,v_m,z)$ and $\mathsf{hyb}_i(v_1,\ldots,v_m,z)$, is that in the former A receives a commitment to v_i in session i, whereas in the latter it receives a commitment to 0^n. Now, consider the one-many adversary $\tilde{A}$ that when receiving $\tilde{z} = (i, v_1, \ldots, v_m, z)$ as auxiliary input proceeds as follows. $\tilde{A}$ internally incorporates $A(z)$ and emulates the left and right interactions for A.

1. $\tilde{A}$ forwards messages in its j^{th} right execution directly to and from A (as part of its j^{th} right execution) with the following exception: whenever A wishes to send the last message q_i in the i'th right execution, $\tilde{A}$ holds on to it without forwarding it externally.

2. $\tilde{A}$ forwards messages from its left session directly to and from A (as part of its i^{th} session).

3. $\tilde{A}$ emulates all left sessions $j \neq i$, by committing to v_j if $j > i$, and committing to 0^n otherwise.

4. Whenever $\tilde{A}$ has completed the emulation of all left executions, it checks whether A fully copied any the left executions; for each such copied execution i, $\tilde{A}$ sends $\perp$ as its last message in the i'th right execution, and otherwise sends q_i.

It follows directly from the construction and the natural property of $\langle C, R\rangle$ that

$$\mathsf{mim}^{\tilde{A}}_{\mathsf{com}}(v_i, \tilde{z}) = \mathsf{hyb}_{i-1}(v_1,\ldots,v_m,z)$$

$$\mathsf{mim}^{\tilde{A}}_{\mathsf{com}}(0, \tilde{z}) = \mathsf{hyb}_{i}(v_1,\ldots,v_m,z)$$

This contradicts the fact that there exists a simulator $\tilde{S}$ for $\tilde{A}$ such that both:

1. $\mathsf{mim}^{\tilde{A}}_{\mathsf{com}}(v_i, \tilde{z})$ and $\mathsf{sim}^{\tilde{S}}_{\mathsf{com}}(\tilde{z})$ are indistinguishable, and

2. $\mathsf{mim}^{\tilde{A}}_{\mathsf{com}}(0, \tilde{z})$ and $\mathsf{sim}^{\tilde{S}}_{\mathsf{com}}(\tilde{z})$ are indistinguishable.

We conclude that $\langle C, R\rangle$ is not one-many concurrent non-malleable. ■

4 The Protocol

Our construction of concurrent non-malleable commitments follows the paradigm introduced by Pass and Rosen for obtaining (single execution) non-malleable commitments [42]. The commit phase of the Pass–Rosen protocol consists of having the sender engage in a (standard) statistically binding commitment with the receiver and thereafter also provide a *non-malleable* $\mathcal{ZK}$ proof of knowledge of the value committed to. The reveal phase consists of sending the de-commitment information of the statistically binding commitment used in the commit phase.

The basic scenario in which non-malleable $\mathcal{ZK}$ protocols take place involves a man-in-the-middle adversary A that is simultaneously participating in two executions of the protocol. These executions are called the left and the right interaction.

The left interaction is tagged by an identity string $\text{TAG} \in \{0,1\}^n$, and the right interaction is tagged by an identity $\widetilde{\text{TAG}} \in \{0,1\}^n$. The instructions of the protocol executed in each of the interactions depend on the corresponding identities TAG and $\widetilde{\text{TAG}}$. (The way in which the identity

$$P_{\text{TAG}} \overset{x\in L}{\Longrightarrow} A \overset{\tilde{x}\in L}{\Longrightarrow} V_{\widetilde{\text{TAG}}}$$

Figure 2: The man-in-the-middle adversary.

strings are determined and used will become clear at a later stage.) We let $\langle P_{\text{TAG}}, V_{\text{TAG}}\rangle$ denote a protocol execution in which the identity TAG is used as a tag.

In the left interaction, the adversary A is verifying the validity of a statement x by interacting with an honest prover P_{TAG} using a protocol $\langle P_{\text{TAG}}, V_{\text{TAG}}\rangle$. In the right interaction A proves the validity of a statement $\tilde{x}$ to the honest verifier $V_{\widetilde{\text{TAG}}}$. The statement $\tilde{x}$ proved in the right interaction is chosen by A, possibly depending on the messages it receives in the left interaction. As in the case of concurrent non-malleable commitments, A has control over the scheduling of the messages.

Loosely speaking, a protocol is non-malleable if for any such man-in-the-middle adversary, there exists a "stand-alone" prover that convinces the verifier in the right interaction with essentially the same success probability as the adversary does in a man-in-the-middle execution. See [42] and Section 4.1 for further details.

4.1 Non-Malleability and Simulation-Extractability

In [42] it is shown that a commitment is non-malleable provided that the underlying $\mathcal{ZK}$ protocols satisfy a *simulation extractability* property (a strengthening of non-malleability). Loosely speaking, simulation extractability requires that for any man-in-the-middle adversary A, there exists a simulator-extractor that can simulate both the left and the right interactions for A, while outputting a witness for the statement proved by the adversary in the right interaction.

For the purpose of the current work we will need to show that the $\mathcal{ZK}$ protocols used in the compilation satisfy an even stronger property, which we call *one-many simulation-extractability.* This is a strengthening of the simulation extractability property in that it guarantees simulation and extraction (of all witnesses on the right) even if there is an *unbounded* number of concurrent *right* interactions (but still with only one left interaction).

As we will show later, a (non-interactive) commitment scheme that is compiled with one-many simulation-extractable $\mathcal{ZK}$ will result in a one-many concurrent non-malleable commitment protocol $\langle C, R\rangle$. By Proposition 3.2, this implies that $\langle C, R\rangle$ is also concurrent non-malleable.

Let A be a man-in-the middle adversary that is simultaneously participating in one left interaction of $\langle P_{\text{TAG}}, V_{\text{TAG}}\rangle$ while acting as verifier, and an (unbounded) polynomial number of right-interactions of $\langle P_{\widetilde{\text{TAG}}_i}, V_{\widetilde{\text{TAG}}_i}\rangle_{i=1}^{m}$ while acting as prover.

Let $\mathsf{view}_A(x, z, \text{TAG})$ denote the *joint* view of $A(x, z)$ and the honest verifier $V_{\widetilde{\text{TAG}}}$ when A is verifying a left-proof of the statement x, using identity TAG, and and proving in the right interactions statements of its choice using tags of its choice. (The view consists of the messages sent and received by A in both left and right interactions, and the random coins of A, and $V_{\widetilde{\text{TAG}}}$).[10] Given a function $t = t(n)$ we use the notation $\{\cdot\}_{x,z,\text{TAG}}$ as shorthand for $\{\cdot\}_{x\in L, z\in\{0,1\}^*, \text{TAG}\in\{0,1\}^{t(|x|)}}$.

Definition 4.1 (One-many Simulation-extractability) *A family $\{\langle P_{\text{TAG}}, V_{\text{TAG}}\rangle\}_{\text{TAG}\in\{0,1\}^*}$ of interactive proofs for L is said to be* simulation extractable *with tags of length $t = t(n)$ if for any polynomial $p(\cdot)$ and any man-in-the-middle adversary A that participates in one left interaction*

[10] Since the messages sent by A are fully determined given the code of A and the messages it receives, including them as part of the view is somewhat redundant. The reason we have chosen to do so is for convenience of presentation.

and at most $m = p(n)$ right interactions, there exists a probabilistic expected poly-time machine S such that:

1. *The probability ensembles $\{S_1(x, z, \text{TAG})\}_{x,z,\text{TAG}}$ and $\{\mathsf{view}_A(x, z, \text{TAG})\}_{x,z,\text{TAG}}$ are statistically close over L, where $S_1(x, z, \text{TAG})$ denotes the first output of $S(x, z, \text{TAG})$.*

2. *Let $x \in L, z \in \{0,1\}^*, \text{TAG} \in \{0,1\}^{t(|x|)}$, and let $(\mathsf{view}, \bar{w})$ denote the output of $S(x, z, \text{TAG})$ (on input some random tape). Let $\tilde{x}_1 \dots, \tilde{x}_m$ be the right-execution statements appearing in view and let $\tilde{\text{TAG}}_1 \dots \tilde{\text{TAG}}_m$ denote the correspoding right-execution tags. Then, for any $i \in [m]$ such that the i^{th} right-execution in view is accepting AND $\text{TAG} \neq \tilde{\text{TAG}}_i$, $\bar{w}$ contains a witness w_i so that $R_L(\tilde{x}_i, w_i) = 1$.*

4.2 A Simulation Extractable Protocol

We now describe our construction of simulation extractable protocols. At a high level, the construction proceeds in two steps:

1. For any $n \in N$, construct a family $\{\langle P_{\mathsf{tag}}, V_{\mathsf{tag}}\rangle\}_{\mathsf{tag}\in[2n]}$ of simulation-extractable arguments with tags of length $t(n) = \log n + 1$.

2. For any $n \in N$, use the family $\{\langle P_{\mathsf{tag}}, V_{\mathsf{tag}}\rangle\}_{\mathsf{tag}\in[2n]}$ to construct a family $\{\langle P_{\text{TAG}}, V_{\text{TAG}}\rangle\}_{\text{TAG}\in\{0,1\}^n}$ of simulation extractable arguments with tags of length $t(n) = n$.

The construction of the family $\{\langle P_{\mathsf{tag}}, V_{\mathsf{tag}}\rangle\}_{\mathsf{tag}\in[2n]}$ relies on Barak's non black-box techniques for obtaining constant-round public-coin $\mathcal{ZK}$ for $\mathcal{NP}$ [1], and are very similar in structure to the $\mathcal{ZK}$ protocols used by Pass in [40]. Overall, the construction of $\langle P_{\text{TAG}}, V_{\text{TAG}}\rangle$ is essentially identical to the construction of the simulation-extractable protocols in [42]; the only difference is that in the current paper we will replace statistically hiding commitments with *perfectly* hiding commitments, and statistically witness indistinguishable arguments with witness *independent* arguments.

Let $n \in N$, and let $T : N \to N$ be a function that satisfies $T(n) = n^{\omega(1)}$. Our construction relies on a "special" $\mathbf{NTIME}(T(n))$ relation, which we denote by $\mathbf{R}_{\mathsf{sim}}$. It also makes use of a "special-purpose" universal argument ($UARG$) [18, 17, 32, 34, 5, 42]. Let $\{\mathcal{H}_n\}_n$ be a family of hash functions where a function $h \in \mathcal{H}_n$ maps $\{0,1\}^*$ to $\{0,1\}^n$, and let **Com** be a perfectly hiding commitment scheme for strings of length n, where for any $\alpha \in \{0,1\}^n$, the length of $\mathbf{Com}(\alpha)$ is upper bounded by $2n$. The relation $\mathbf{R}_{\mathsf{sim}}$ is described in Figure 3.

Instance: A triplet $\langle h, c, r\rangle \in \mathcal{H}_n \times \{0,1\}^n \times \{0,1\}^{\text{poly}(n)}$.

Witness: A program $\Pi \in \{0,1\}^*$, a string $y \in \{0,1\}^*$ and a string $s \in \{0,1\}^{\text{poly}(n)}$.

Relation: $\mathbf{R}_{\mathsf{sim}}(\langle h, c, r\rangle, \langle \Pi, y, s\rangle) = 1$ if and only if:

1. $|y| \leq |r| - n$.
2. $c = \mathbf{Com}(h(\Pi); s)$.
3. $\Pi(y) = r$ within $T(n)$ steps.

Figure 3: The relation $\mathbf{R}_{\mathsf{sim}}$.

Remark 1 (Simplifying assumption) *For simplicity of exposition, we view* **Com** *as a one-message perfectly hiding commitment scheme (even though such commitments cannot exist). In reality,* **Com** *would be taken to be a 2-message commitment scheme (which may be based on collections of claw-free permutations* [21]*).*

The construction of our protocol employs a universal argument that is specially tailored for our purposes (a variant of which has already appeared in [42]). The main distinguishing features of this universal argument, which we call the *special purpose* argument, are: (1) it is witness *independent*; and (2) it will enable us to prove that our protocols satisfy the proof of knowledge property of Definition 2.8.[11] Let $\langle P_{\mathsf{pWI}}, V_{\mathsf{pWI}}\rangle$ be a witness independent argument of knowledge, and let $\langle P_{\mathsf{UA}}, V_{\mathsf{UA}}\rangle$ be a 4-message, public-coin universal argument where the length of the messages is upper bounded by n.[12] The special purpose $UARG$, which we denote by $\langle P_{\mathsf{sUA}}, V_{\mathsf{sUA}}\rangle$, handles statements of the form $(x, \langle h, c_1, c_2, r_1, r_2\rangle)$, where the triplets $\langle h, c_1, r_1\rangle$ and $\langle h, c_2, r_2\rangle$ correspond to instances for $\mathbf{R}_{\mathsf{sim}}$. The protocol $\langle P_{\mathsf{sUA}}, V_{\mathsf{sUA}}\rangle$ is described in Figure 4.

Parameters: Security parameter 1^n.

Common Input: $x \in \{0,1\}^n$, $\langle h, c_1, c_2, r_1, r_2\rangle$ where for $i \in \{1,2\}$, $\langle h, c_i, r_i\rangle$ is an instance for $\mathbf{R}_{\mathsf{sim}}$.

Stage 1 (Encrypted UARG):

$V \to P$: Send $\alpha \stackrel{\text{R}}{\leftarrow} \{0,1\}^n$.

$P \to V$: Send $\widehat{\beta} = \mathbf{Com}(0^n)$.

$V \to P$: Send $\gamma \stackrel{\text{R}}{\leftarrow} \{0,1\}^n$.

$P \to V$: Send $\widehat{\delta} = \mathbf{Com}(0^n)$.

Stage 2 (Body of the proof):

$P \leftrightarrow V$: A witness independent argument of knowledge $\langle P_{\mathsf{pWI}}, V_{\mathsf{pWI}}\rangle$ proving the OR of the following statements:

1. $\exists\, w \in \{0,1\}^{\mathrm{poly}(|x|)}$ so that $R_L(x, w) = 1$.
2. $\exists\, \langle \beta, \delta, s_1, s_2\rangle$ so that:
   - $\widehat{\beta} = \mathbf{Com}(\beta; s_1)$.
   - $\widehat{\delta} = \mathbf{Com}(\delta; s_2)$.
   - $(\alpha, \beta, \gamma, \delta)$ is an accepting transcript for $\langle P_{\mathsf{UA}}, V_{\mathsf{UA}}\rangle$ proving the statement:
     - $\exists\, \langle i, \Pi, y, s\rangle$ so that $\mathbf{R}_{\mathsf{sim}}(\langle h, c_i, r_i\rangle, \langle \Pi, y, s\rangle) = 1$

Figure 4: A special-purpose universal argument $\langle P_{\mathsf{sUA}}, V_{\mathsf{sUA}}\rangle$.

4.2.1 A family of $2n$ protocols

We next present a family of protocols $\{\langle P_{\mathsf{tag}}, V_{\mathsf{tag}}\rangle\}_{\mathsf{tag}\in[2n]}$ (with tags of length $t(n) = \log n + 1$).[13] These protocols are a "two-slot" version of Barak's $\mathcal{ZK}$ arguments [1]. (The idea of using a multiple

[11] The "weak" proof of knowledge property of a universal argument (as defined in [5]) is not sufficient for our purposes. Specifically, while in a weak proof of knowledge it is required that the extractor succeeds with probability that is polynomially related to the success probability of the prover, in our proof of security we will make use of an extractor that succeeds with probability negligibly close to the success probability of the prover.

[12] Both witness independent arguments of knowledge, and 4-message, public-coin, universal arguments can be constructed assuming a family $\mathcal{H}_n$ of claw-free permutations (cf. [20] and [32, 34, 5]).

[13] A closer look at the construction will reveal that it will in fact work for any $t(n) = O(\log n)$. The choice of $t(n) = \log n + 1$ is simply made for the sake of concreteness (as in our constructions it is the case that $\mathsf{tag} \in [2n]$).

slot version of Barak's protocol already appeared in [41, 40], and the "message-length" technique appeared in [40].) There are two aspects in which the protocols presented here differ from the protocol of [40]: the new protocols satisfy a (1) *perfect* secrecy property (note that this is also different from [42], where secrecy is statistical), and (2) a *proof of knowledge* property.

Let **Com** be a perfectly-hiding commitment scheme for strings of length n, where for any $\alpha \in \{0,1\}^n$, the length of $\mathbf{Com}(\alpha)$ is upper bounded by $2n$. Let $\mathbf{R}_{\mathsf{sim}}$ be the perfect variant of the relation R_{sim}, and let $\langle P_{\mathsf{sUA}}, V_{\mathsf{sUA}} \rangle$ be the special purpose universal argument. Protocol $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}} \rangle$ is described in Figure 5.

Common Input: An instance $x \in \{0,1\}^n$

Parameters: Security parameter 1^n, length parameter $\ell(n)$.

Tag String: $\mathsf{tag} \in [2n]$.

Stage 0 (Set-up):

$V \to P$: Send $h \xleftarrow{\text{R}} \mathcal{H}_n$.

Stage 1 (Slot 1):

$P \to V$: Send $c_1 = \mathbf{Com}(0^n)$.

$V \to P$: Send $r_1 \xleftarrow{\text{R}} \{0,1\}^{\mathsf{tag}\cdot\ell(n)}$.

Stage 1 (Slot 2):

$P \to V$: Send $c_2 = \mathbf{Com}(0^n)$.

$V \to P$: Send $r_2 \xleftarrow{\text{R}} \{0,1\}^{(2n+1-\mathsf{tag})\cdot\ell(n)}$.

Stage 2 (Body of the proof):

$P \Leftrightarrow V$: A special-purpose *UARG* $\langle P_{\mathsf{sUA}}, V_{\mathsf{sUA}} \rangle$ proving the OR of the following statements:

1. $\exists\, w \in \{0,1\}^{\mathrm{poly}(|x|)}$ s.t. $R_L(x,w) = 1$.
2. $\exists\, \langle \Pi, y, s \rangle$ s.t. $\mathbf{R}_{\mathsf{Sim}}(\langle h, c_1, r_1 \rangle, \langle \Pi, y, s \rangle) = 1$.
3. $\exists\, \langle \Pi, y, s \rangle$ s.t. $\mathbf{R}_{\mathsf{Sim}}(\langle h, c_2, r_2 \rangle, \langle \Pi, y, s \rangle) = 1$.

Figure 5: Protocol $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}} \rangle$.

Note that the only difference between two protocols $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}} \rangle$ and $\langle P_{\tilde{\mathsf{tag}}}, V_{\tilde{\mathsf{tag}}} \rangle$ is the length of the verifier's "next messages": in fact, the length of those messages in $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}} \rangle$ is a parameter that depends on tag (as well as on the length parameter $\ell(n)$). This property will be crucial for the analysis of these protocols in the man in the middle setting.

4.2.2 A family of 2^n protocols

Relying on the protocol family $\{\langle P_{\mathsf{tag}}, V_{\mathsf{tag}} \rangle\}_{\mathsf{tag} \in [2n]}$, we now show how to construct a family $\{\langle P_{\text{TAG}}, V_{\text{TAG}} \rangle\}_{\text{TAG} \in \{0,1\}^n}$ with tags of length $t(n) = n$. The protocols are *constant-round* and involve n parallel executions of $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}} \rangle$, with appropriately chosen tags. This new family of protocols is denoted $\{\langle P_{\text{TAG}}, V_{\text{TAG}} \rangle\}_{\text{TAG} \in \{0,1\}^n}$ and is described in Figure 6.

Notice that $\langle P_{\text{TAG}}, V_{\text{TAG}} \rangle$ has a constant number of rounds (since each $\langle P_{\mathsf{tag}_i}, V_{\mathsf{tag}_i} \rangle$ is constant-round). Also notice that for $i \in [n]$, the length of $\mathsf{tag}_i = (i, \text{TAG}_i)$ is

$$|i| + |\text{TAG}_i| = \log n + 1 = \log 2n.$$

Common Input: An instance $x \in \{0,1\}^n$

Parameters: Security parameter 1^n, length parameter $\ell(n)$

Tag String: $\text{TAG} \in \{0,1\}^n$. Let $\text{TAG} = \text{TAG}_1, \dots, \text{TAG}_n$.

The protocol:

$P \leftrightarrow V$: For all $i \in \{1, \dots, n\}$ (in parallel):

1. Set $\mathsf{tag}_i = (i, \text{TAG}_i)$.
2. Run $\langle P_{\mathsf{tag}_i}, V_{\mathsf{tag}_i} \rangle$ with common input x and length parameter $\ell(n)$.

V: Accept if and only if all runs are accepting.

Figure 6: Protocol $\langle P_{\text{TAG}}, V_{\text{TAG}} \rangle$.

Viewing (i, TAG_i) as elements in $[2n]$ we infer that the length of verifier messages in $\langle P_{\mathsf{tag}_i}, V_{\mathsf{tag}_i} \rangle$ is upper bounded by $2n\ell(n)$. Hence, as long as $\ell(n) = \text{poly}(n)$ the length of verifier messages in $\langle P_{\text{TAG}}, V_{\text{TAG}} \rangle$ is $2n^2\ell(n) = \text{poly}(n)$.

As shown in [42], for any $\text{TAG} \in 2^n$, the protocol $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}} \rangle$ is an interactive argument. In fact, what they show is a stronger statement. Namely, that the protocols $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}} \rangle$ are arguments of knowledge (as in Definition 2.8). What [42] actually show is how to prove the above assuming a family of hash functions that is collision resistant against $T(n) = n^{\omega(1)}$-sized circuits. As mentioned in [42], by slightly modifying $\mathbf{R}_{\mathsf{sim}}$, one can prove the same statement under the more standard assumption of collision resistance against polynomial-sized circuits (c.f. [5]).

Proposition 4.2 (Argument of knowledge [42]) *Let $\langle P_{\mathsf{pWI}}, V_{\mathsf{pWI}} \rangle$ and $\langle P_{\mathsf{UA}}, V_{\mathsf{UA}} \rangle$ be the protocols used in the construction of $\langle P_{\mathsf{sUA}}, V_{\mathsf{sUA}} \rangle$. Suppose that $\{\mathcal{H}_n\}_n$ is collision resistant for $T(n)$-sized circuits, that* **Com** *is perfectly hiding, that $\langle P_{\mathsf{pWI}}, V_{\mathsf{pWI}} \rangle$ is a witness independent argument of knowledge, and that $\langle P_{\mathsf{UA}}, V_{\mathsf{UA}} \rangle$ is a universal argument. Then, for any $\text{TAG} \in \{0,1\}^n$, $\langle P_{\text{TAG}}, V_{\text{TAG}} \rangle$ is a perfect zero-knowledge argument of knowledge.*

The main technical contribution of the current paper consists of proving that the protocol $\{\langle P_{\text{TAG}}, V_{\text{TAG}} \rangle\}_{\text{TAG} \in \{0,1\}^n}$ is one-many simulation extractable.

Lemma 4.3 (Main technical lemma) *Let $\langle P_{\mathsf{pWI}}, V_{\mathsf{pWI}} \rangle$ and $\langle P_{\mathsf{UA}}, V_{\mathsf{UA}} \rangle$ be the protocols used in the construction of $\langle P_{\mathsf{sUA}}, V_{\mathsf{sUA}} \rangle$. Suppose that $\{\mathcal{H}_n\}_n$ is collision resistant for $T(n)$-sized circuits, that* **Com** *is perfectly hiding, that $\langle P_{\mathsf{pWI}}, V_{\mathsf{pWI}} \rangle$ is a witness independent argument of knowledge, that $\langle P_{\mathsf{UA}}, V_{\mathsf{UA}} \rangle$ is a universal argument and that $\ell(n) \geq 2n^3 + n$. Then, $\{\langle P_{\text{TAG}}, V_{\text{TAG}} \rangle\}_{\text{TAG} \in \{0,1\}^n}$ is one-many simulation extractable with tags of length $t(n) = n$.*

Before we go on and prove Lemma 4.3, we turn to describe our commitment protocol and to show how one many simulation extractability is used in order to establish its one-many concurrent non-malleability. The full proof of Lemma 4.3 can be found in Section 5.

4.3 The Commitment Protocol

Using protocols from $\{\langle P_{\text{TAG}}, V_{\text{TAG}} \rangle\}_{\text{TAG} \in \{0,1\}^n}$ as a subroutine, we present the construction of concurrent non-malleable commitments. Let $\{\text{Com}_r\}_{r \in \{0,1\}^*}$ be a family of *non-interactive* statistically binding commitment schemes (e.g., Naor's commitment [35]). Let $(\mathsf{Gen}, \mathsf{Sign}, \mathsf{Verify})$ be a one-time signature scheme secure against a chosen-message attack. Consider the following protocol (which is a variant of the non-malleable commitment of Pass and Rosen [42]).[14]

> **Security Parameter:** 1^n.
>
> **String to be committed to**: $v \in \{0,1\}^n$.
>
> **Commit Phase:**
>
> $R \to C$: Send a uniformly chosen $r \in \{0,1\}^n$.
>
> $C \to R$: Let $vk, sk \leftarrow \mathsf{Gen}(1^n)$. Pick uniformly $s \in \{0,1\}^n$.
> Set $\text{TAG} = vk$ and send $c = \text{Com}_r(v; s), \text{TAG}$.
>
> $C \leftrightarrow R$: Prove using $\langle P_{\text{TAG}}, V_{\text{TAG}} \rangle$ that there exist $v, s \in \{0,1\}^n$ so that $c = \text{Com}_r(v; s)$.
>
> $C \to R$: Let T denote the transcript of the above interaction.
> Compute $\sigma = \mathsf{Sign}(sk, T)$ and send σ.
>
> R: Verify that $\langle P_{\text{TAG}}, V_{\text{TAG}} \rangle$ is accepting and that $\mathsf{Verify}(vk, T, \sigma) = 1$.
>
> **Reveal Phase:**
>
> $C \to R$: Send v and s.
>
> R: Verify that $c = \text{Com}_r(v; s)$.

Figure 7: Concurrent non-malleable commitment - $\langle C, R \rangle$.

As argued in [42], the statistical binding property of $\langle C, R \rangle$ follows directly from the statistical binding of Com. The computational hiding property follows from the computational hiding of Com, as well as from the (stand alone) $\mathcal{ZK}$ property of $\langle P_{\text{TAG}}, V_{\text{TAG}} \rangle$. Hence, we have:

Proposition 4.4 ([42]) *Suppose that* $\{\text{Com}_r\}_{r \in \{0,1\}^*}$ *is a family of* non-interactive *statistically binding commitment schemes, and that all members in the family* $\{\langle P_{\text{TAG}}, V_{\text{TAG}} \rangle\}_{\text{TAG} \in \{0,1\}^n}$ *are (stand-alone) zero-knowledge. Then,* $\langle C, R \rangle$ *is a statistically-binding commitment protocol.*

4.4 Concurrent Non-Malleablity

Relying on the one-many simulation extractability of $\langle P_{\text{TAG}}, V_{\text{TAG}} \rangle$, we next argue that $\langle C, R \rangle$ is one-many non-malleable.

Theorem 4.5 *Suppose that* $\{\text{Com}_r\}_{r \in \{0,1\}^*}$ *is a family of non-interactive statistically binding commitment schemes, and that* $\{\langle P_{\text{TAG}}, V_{\text{TAG}} \rangle\}_{\text{TAG} \in \{0,1\}^n}$ *is one-many simulation extractable and natural. Then,* $\langle C, R \rangle$ *is a natural one-many concurrent non-malleable commitment.*

Proof: We start by noting that it follows directly from the fact that the last message of $\langle C, R \rangle$ is supposed to be a signature on the transcript (in case the commitment is valid), that any commitment where the committer aborts before the last round is invalid—in other words, $\langle C, R \rangle$ is natural.

Next, consider a man-in-the-middle adversary A that participates in one left execution and $m = m(n)$ right executions. We assume without loss of generality that A is deterministic (this is

[14] The difference between this protocol and the protocol of [42] is that here we also employ a signature scheme. We note that the important difference, nevertheless, lies in the analysis of the protocol.

w.l.o.g since A can obtain its "best" random tape as auxiliary input). Consider the simulator S that proceeds as follows on input z. S incorporates $A(z)$ and internally emulates the left interactions for A by simply *honestly* committing to the string 0^n (i.e., S executes the algorithm C on input 0^n). Messages from the right interactions are instead forwarded externally, with the following exception: whenever A wishes to send the last message q_i in the i'th right session, S "holds-on" to it without (yet) forwarding it externally. Finally, when S has completed the emulation of the left interactions for A, it checks whether A fully copied any of the left executions. For each execution i where A fully copied any of the left executions, S externally sends $\perp$ as its last message in the i'th right execution (to invalidate that commitment); for all other executions j, $\tilde{A}$ instead sends the final message q_j. We show that the following distributions are indistinguishable over $\{0,1\}^*$.

- $\left\{\mathsf{mim}^A_{\mathsf{com}}(v,z)\right\}_{v\in\{0,1\}^n, n\in N, z\in\{0,1\}^*}$
- $\left\{\mathsf{sim}^S_{\mathsf{com}}(z)\right\}_{v\in\{0,1\}^n, n\in N, z\in\{0,1\}^*}$

Suppose, for contradiction, that this is not the case. That is, there exists a polynomial-time distinguisher D and a polynomial $p(n)$ such that for infinitely many n, there exists strings $v \in \{0,1\}^n, z \in \{0,1\}^*$ such that

$$\Pr\left[D(\mathsf{mim}^A_{\mathsf{com}}(v,z) = 1\right] - \Pr\left[D(\mathsf{sim}^S_{\mathsf{com}}(z)) = 1\right] \geq \frac{1}{p(n)}$$

Fix a generic n for which this happens. We show how this contradicts the simulation-extractability property of $\langle P_{\text{TAG}}, V_{\text{TAG}}\rangle$. We start by providing an (oversimplified) sketch. On a high-level, the proof consists of the following steps:

1. We first note that since the commit phase of (C,R) "essentially" only consists of a statement (c,r) (i.e., the commitment) and a proof of the validity of (c,r), A can be interpreted as a one-many simulation-extractability adversary A' for $\langle P_{\text{TAG}}, V_{\text{TAG}}\rangle$.

2. It follows from the simulation-extractability property of $\langle P_{\text{TAG}}, V_{\text{TAG}}\rangle$ that there exist a combined simulator-extractor S' for A' that outputs a view that is statistically close to that of A', while at the same time outputting witnesses to all accepting right proofs.

3. Since the view output by the simulator-extractor S' is *statistically* close to the view of A' in the real interaction, it follows that also the *values* committed to in that view are are statistically close to the values committed to by A'. (Note that computational indistinguishability would not have been enough to argue the indistinguishability of these values, since they are not efficiently computable from the view.)

4. It also follows that except with negligible probability, the simulator-extractor S' will output also the witnesses to all accepting right executions.[15] We conclude that S' additionally outputs the values *committed to* in the right executions.

5. We finally note that if D can distinguish between the values committed to by A and by S, then D can also distinguish the second output (which consists of the committed values) of S' when run on input a commitment (using Com) to v, and the second output of S' when run on input a commitment to 0. This contradicts the hiding property of Com.

[15] More precisely, the simulator-extractor only outputs witnesses to all right-executions that use a different tag than the left interaction. We rely on the use of the digital signature to handle the case when A copies the tag of the left interaction.

We proceed to a formal proof. One particular complication that arises with the above proof sketch is that in the construction of $\langle C, R\rangle$ we are relying on the use of a family of commitment schemes $\{\mathrm{Com}_r\}_{r\in\{0,1\}^*}$ and not a single non-interactive commitment scheme. To address this issue we make use of non-uniformity to show the existence of particular "prefix" of right-interactions such that A always chooses the instance Com_r in its left interaction, yet A commits to different values when receiving a commitments to v (as in mim) and 0 (as in sim).

More precisely, since in both experiments mim and sim are identical up until the point where A sends its first message in the left interaction, there must exists some *fixed* prefix transcript τ of A's right interactions such that

1. A sends its first message r_τ in its *left* interaction directly after receiving the messages in τ (as part of its right executions).

2. D distinguishes between $\mathsf{mim}^A_{\mathsf{com}}(v, z)$ and $\mathsf{sim}^S_{\mathsf{com}}(z)$ with probability $1/p(n)$, conditioned on the event that the right executions are consistent with τ.

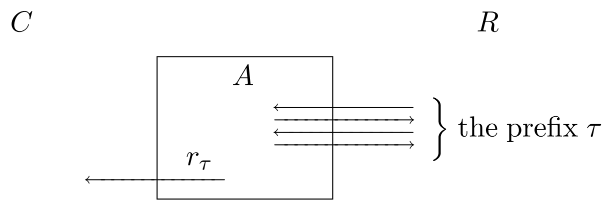

Figure 8: The prefix τ before r_τ is sent.

Note that since the receiver R sends the first message in the protocol, τ contains all first messages r sent by the honest R in all m right executions. Let $r_\tau^{(1)}, .., r_\tau^{(m)}$ denote these first messages. Let $\mathtt{committed}_\tau \subseteq [m]$ denote the set of executions i such that $A(z)$ has sent its first message in the i'th execution in τ. (Recall that the first message sent by A, playing the "role" of C in execution i, consists of a commitment using $\mathrm{Com}_{r_\tau^{(i)}}$.) For each $i \in \mathtt{committed}_\tau$, let $\mathtt{value}_\tau(i)$ denote the value committed to in the first message sent by $A(z)$ in execution i in τ. (If this value is not uniquely defined, set $\mathtt{value}_\tau(i) = \perp$). Additionally, for each $i \in \mathtt{committed}_\tau$, let $\mathtt{coins}_\tau(i)$ be a random tape for R such that D distinguishes between $\mathsf{mim}^A_{\mathsf{com}}(v, z)$ and $\mathsf{sim}^S_{\mathsf{com}}(z)$ with probability $\frac{1}{p(n)}$, conditioned on the event that the right executions are consistent with τ and that in each right execution $i \in \mathtt{committed}_\tau$, R uses the random tape $\mathtt{coins}_\tau(i)$. (It follows by an averaging argument that such random tapes must exist).

Given the partial transcript τ we next proceed in the following three steps:

1. We first define a simulation-extractability adversary A'.

2. We next show that A' can be used to violate the non-malleability property of Com.

3. In the final step, we show how to use the simulator-extractor S' for A' to violate the hiding property of Com.

Step 1: Defining a simulation-extractability adversary A'. We define a one-many simulation extractability adversary A' for $\langle P_{\text{TAG}}, V_{\text{TAG}}\rangle$. On a high-level A' internally incorporates A on input the transcript τ, externally forwards all messages that are part of $\langle P_{\text{TAG}}, V_{\text{TAG}}\rangle$ while internally handling all other messages. Additionally, A' internally emulates (without externally forwarding) all messages that are part of executions $i \in \texttt{committed}_\tau$ (note that messages in these executions cannot be externally forwarded as the execution of protocol $\langle P_{\text{TAG}}, V_{\text{TAG}}\rangle$ could potentially have already begun).

We proceed to a formal description of A'. To simplify notation we assume that the machine A' that has the values $\tau, \texttt{committed}_\tau, \texttt{coins}_\tau$ hard-coded in its description (as it can always receive them as part of its auxiliary input z'). On input x', TAG', z' (i.e., $A'(z')$ expects to receive a proof of the statement x' using tag TAG'), where $x' = (c, r_\tau)$, and $z' = (z, sk)$ such that sk is the corresponding secret key for the signature verification key TAG', A' then internally incorporates $A(z)$ and emulates the left and right interactions for A in the following manner.

1. It starts by feeding A all messages in τ as part of its right executions.
2. For each $i \in \texttt{committed}_\tau$, A' internally emulates the (rest of the) i'th right execution for A by honestly following the strategy of R using the random tape $\texttt{coins}_\tau(i)$.
3. For each $i \in [m]$ such that $i \notin \texttt{committed}_\tau$, A' externally forwards messages in the i'th right interaction, as follows. Whenever A sends its first message c_i, vk_i, in the i^{th} execution, A externally forwards $(c_i, r_\tau^{(i)})$ as its statement and vk_i as its tag. Thereafter, it externally forwards all the messages from $\langle P_{\text{TAG}}, V_{\text{TAG}}\rangle$. (Note that since $i \notin \texttt{committed}_\tau$, A has not yet sent any messages in execution i; thus, A is expected to produce a proof of the statement $(c', r_\tau^{(i)})$ using TAG' as tag).
4. Messages in A's left interaction are instead forwarded externally as part of A''s left interaction. Once A has sent its first message r_τ, A' start by feeding it c, TAG', and next externally forwards all remaining messages (i.e., all the messages that are part of $\langle P_{\text{TAG}}, V_{\text{TAG}}\rangle$). Once the execution of $\langle P_{\text{TAG}}, V_{\text{TAG}}\rangle$ has concluded A' signs the transcript of the left interaction using sk (received as auxiliary input) and feeds the signature to A. (Recall that whereas A' receives only a proof as part of its left interaction, A expects to see a commitment using $\langle C, R\rangle$. It is therefore essential that A' adds a signature to the proof.)

Step 2: Show that A' violates non-malleability of Com**.** Define the following experiment $\mathsf{hyb}_1(v')$:

1. Pick $vk, sk \leftarrow \mathsf{Gen}(1^n)$, and pick uniformly $s \in \{0,1\}^n$.
2. Let $c = \text{Com}_{r_\tau}(v', s)$.
3. Let $x' = (c, r_\tau), \text{TAG}' = vk, z' = (z, sk)$.
4. Emulate an execution for $A'(x', \text{TAG}', z')$ by honestly providing a proof of x' (using tag TAG' and the witness (v', s)) as part of its left interaction, and honestly verifying all right interactions.
5. Finally, given the view of A' in the above emulation, reconstruct the view of A in the emulation by A'. Output the pair $(\mathsf{view}, \bar{v})$ where view denotes the reconstructed view of A and $\bar{v}$ denotes the values committed to in the view of A. (As in Definition 3.1, if a commitment is undefined, invalid, or if the transcript of the commitment is identical to the transcript of the commitment received by A on the left, its value is set to $\perp$). Note that although the values committed to are not necessarily efficiently computable from the view of A, they are determined.

Note that $\mathsf{hyb}_1(\cdot)$ is not efficiently samplable, since the last step in the description of hyb_1 is not efficient. However, except for that last step, every other operation in hyb_1 is indeed efficient. (This will be useful to us at a later stage).

Claim 4.6
$$\Pr\Big[D(\mathsf{hyb}_1(v)) = 1\Big] - \Pr\Big[D(\mathsf{hyb}_1(0^n)) = 1\Big] \geq \frac{1}{p(n)}$$

Proof: By construction of A' and hyb_1, it directly follows that: the view of A in $\mathsf{hyb}_1(v)$ is identically distributed to the view of A in $\mathsf{mim}^A_{\mathsf{com}}(v, z)$, conditioned on the event that the right executions (in the view of A) are consistent with τ and $\mathtt{coins}_\tau$. Additionally, by relying on the natural property of $\langle C, R\rangle$ (which guarantees that S outputs a commitment to $\perp$ whenever A fully copies any of the left executions), it holds that: the view of A in $\mathsf{hyb}_1(0^n)$ is identically distributed to the view of A in $\mathsf{sim}^S_{\mathsf{com}}(z)$, conditioned on the event that the right executions (in the view of A) are consistent with τ and $\mathtt{coins}_\tau$. Since the outputs of hyb_1, mim and sim are computed by applying the same function to the view of A, in the corresponding experiments, it follows that

1. The output of $\mathsf{hyb}_1(v)$ is identically distributed to the output of $\mathsf{mim}^A_{\mathsf{com}}(v, z)$, conditioned on the event that that the right executions (in the view of A) are consistent with τ and $\mathtt{coins}_\tau$.

2. The output of $\mathsf{hyb}_1(0^n)$ is identically distributed to the output of $\mathsf{sim}^S_{\mathsf{com}}(v, z)$, conditioned on the event that that the right executions (in the view of A) are consistent with τ and $\mathtt{coins}_\tau$.

The claim now follows from the fact that D distinguishes $\mathsf{mim}^A_{\mathsf{com}}(v, z)$ and $\mathsf{sim}^S_{\mathsf{com}}(z)$ with probability $\frac{1}{p(n)}$, conditioned on the right executions being consistent with τ and $\mathtt{coins}_\tau$. ■

Step 3: Show that the simulator for A' violates the hiding property of Com. We next use the simulator-extractor S' for A' to construct a (non-uniformly) *efficiently computable* experiment that is statistically close to hyb_1. Towards this goal, we first define an additional hybrid experiment $\mathsf{hyb}_2(\cdot)$. $\mathsf{hyb}_2(\cdot)$ proceeds just as hyb_1, except that instead of emulating the left and right interactions for A', hyb_2 runs the combined simulator-extractor S' for A' to generate the view of A'. (The second output of hyb_2 is, however, still computed as in hyb_1—i.e., hyb_2 ignores the second output of S').

Claim 4.7 *The ensembles* $\{\mathsf{hyb}_1(v')\}_{v' \in \{0,1\}^*}$ *and* $\{\mathsf{hyb}_2(v')\}_{v' \in \{0,1\}^*}$ *are stat. close over* $\{0,1\}^*$.

Proof: It directly follows from the statistical indistinguishability property of S', that the first output of $\mathsf{hyb}_1(\cdot)$ is statistically close to the first output of $\mathsf{hyb}_2(\cdot)$. The claim is concluded by observing that the second output in both experiments are determined in the same way (as a function of the first output). ■

Remark 2 *Note that the proof of Claim 4.7 inherently relies on the* statistical *indistinguishability property of S'. Indeed, if the simulation had only been computationally indistinguishable, we would not have been able to argue indistinguishability of the second output of* $\mathsf{hyb}_1(\cdot)$ *and* $\mathsf{hyb}_2(\cdot)$. *This follows from the fact that the second output (which consists of the* actual *committed values) is not efficiently computable from the first output (i.e., the view).*

We next define the final experiment $\mathsf{hyb}_3(\cdot)$. $\mathsf{hyb}_3(\cdot)$ proceeds just as hyb_2 with the exception that instead of setting its second output, $\bar{v}$, to the actual values committed to in the view of A, hyb_3 efficiently computes values $\bar{v} = v_1, .., v_m$ as follows. Recall that the combined-simulator extractor S' outputs both a view and witnesses to all accepting right interactions. For each accepting right interaction i in the reconstructed view of A, hyb_3 lets $v_i = \mathtt{value}_\tau(i)$ if $i \in \mathtt{committed}_\tau$ and otherwise set v_i to be consistent with the witness output for execution i by the combined simulator-extractor S'. For all right executions j for which the reconstructed view of A is rejecting, instead set $v_j = \perp$.[16] Note that in contrast to $\mathsf{hyb}_2(\cdot)$, $\mathsf{hyb}_3(\cdot)$ is efficiently computable.

Claim 4.8 *The ensembles* $\{\mathsf{hyb}_2(v')\}_{v' \in \{0,1\}^*}$ *and* $\{\mathsf{hyb}_3(v')\}_{v' \in \{0,1\}^*}$ *are stat. close over* $\{0,1\}^*$*.*

Proof: Recall that the only difference between the experiments $\mathsf{hyb}_2(\cdot)$ and $\mathsf{hyb}_3(\cdot)$ is the way the second output is computed; in the former it is computed as the actual values committed to in the view of A, whereas in the latter it is computed by relying on $\mathtt{value}_\tau$, and the witnesses output by the simulator-extractor S'. We show that except with negligible probability these values are identical, which concludes the claim.

First, note that in any given view of A, it "trivially" holds that hyb_3 outputs the correct value for all rejecting right-executions, and all accepting right-executions i such that $i \in \mathtt{committed}_\tau$. It only remains to consider accepting right-executions i such that $i \notin \mathtt{committed}_\tau$. For these executions the value v_i computed by hyb_3 is obtained from the witnesses output by the simulator-extractor S'.

Assume, first, that A *never* is able to violates the security of $(\mathsf{Gen}, \mathsf{Sign}, \mathsf{Verify})$, or is able to construct a commitment using Com that can be opened up in two different ways (i.e., violate the statistical binding property of Com). Under these assumptions, we show that the values computed by hyb_3 are identical to the actual values committed to in the view of A.

By the definition of the simulator-extractor, it holds that the witnesses output by S', for all accepting right interactions which use a different tag than the one used in the left interaction, are valid. In other words, values for all (non-rejected) right-commitments that use a verification key vk for the signature scheme that is different from the one used in the left-commitment, are extracted. Furthermore, by our assumption that A is not able to break the statistical binding property of Com it follows that the extracted values are identical to the actual values committed to in the view of A. Additionally, note that values for right-commitments that have *exactly* the same transcript as the left-commitment are "trivially" extracted (as they are just $\perp$). It only remains to analyze what happens to (non-rejecting) right-commitments that use the same verification key as the left-commitment, but a different transcripts. In this case, A must have been able to produce a signature on a new message, given only a randomly generated verification key, and single signed message of its choice (namely the transcript on the left). This contradicts our assumption that A does not forge signatures.

We conclude that, conditioned on the event that A is not able to forge a signature, or is able to break the statistical binding property of Com, $\mathsf{hyb}_2(\cdot)$ and $\mathsf{hyb}_3(\cdot)$ are identical. The claim now follows by the security of $(\mathsf{Gen}, \mathsf{Sign}, \mathsf{Verify})$ and Com. ■

Combining the above two claims we thus get:

Claim 4.9 *The ensembles* $\{\mathsf{hyb}_1(v')\}_{v' \in \{0,1\}*}$ *and* $\{\mathsf{hyb}_3(v')\}_{v' \in \{0,1\}*}$ *are stat. close over* $\{0,1\}^*$*. Furthermore,* $\mathsf{hyb}_3(\cdot)$ *is (non-uniformly) efficiently computable.*

Finally, by combining Claim 4.9 with Claim 4.6 we get that D distinguishes the second outputs of $\mathsf{hyb}_3(v)$ and $\mathsf{hyb}_3(0^n)$ with inverse polynomial probability. However, since hyb_3 is (non-uniformly) efficiently samplable, this contradicts the (non-uniform) hiding property of $\mathrm{Com}_{\mathrm{r}_\tau}$. More formally, define the distinguisher D' that proceeds as follows on input a commitment c' using Com_{r_τ}:

[16]Note that an interaction that is accepting in the view of A' can still be rejecting in the view of A, since in the latter we additionally require a valid signature.

1. D' performs the same operations as hyb_3, except that instead of generating the commitment c, it simply sets $c = c'$. Let $(\mathsf{view}, \bar{v})$ denote the output when executing $\tilde{H}$ in this manner.

2. Finally, D' outputs $D(\mathsf{view}, \bar{v})$.

It directly follows from the construction that $D'(c')$ is identically distributed to $D(\mathsf{hyb}_3(0^n))$ when c' is a random commitment to 0^n, and identically distributed to $D(\mathsf{hyb}_3(v))$ when c' is a random commitment to v. We conclude that D'—which is efficient—distinguishes commitments (using Com_{r_τ}) to 0^n and v. ■

5 Simulation-Extractability

We now turn to prove Lemma 4.3. We start by proving an analogous lemma for the "small" family $\{\langle P_{\mathsf{tag}}, V_{\mathsf{tag}}\rangle\}_{\mathsf{tag}\in[2n]}$ of $2n$ protocols. Then we show how to extend the analysis to the family $\{\langle P_{\mathrm{TAG}}, V_{\mathrm{TAG}}\rangle\}_{\mathrm{TAG}\in\{0,1\}^n}$.

Lemma 5.1 *Suppose that $\{\mathcal{H}_n\}_n$ is collision resistant for $T(n)$-sized circuits, that* **Com** *is perfectly hiding, that $\langle P_{\mathsf{pWI}}, V_{\mathsf{pWI}}\rangle$ is a witness independent argument of knowledge, that $\langle P_{\mathsf{UA}}, V_{\mathsf{UA}}\rangle$ is a universal argument and that $\ell(n) \geq 2n^2 + n$. Then, $\{\langle P_{\mathsf{tag}}, V_{\mathsf{tag}}\rangle\}_{\mathsf{tag}\in[2n]}$ is one-many simulation extractable with tags of length $t(n) = \log n + 1$.*

Proof: Recall that one-many simulation-extractability (Definition 4.1) means that there exists a combined simulator-extractor $S = (\mathsf{SIM}, \mathsf{EXT})$ that is able to simulate both the left and the right interactions for a man-in-the-middle adversary A, while simultaneously extracting witnesses to the m statements proved in the right interaction. The construction of S is fairly complex. To keep things simple, we decompose the description of the simulator into three simulation procedures, where each procedure relies on the previous (simpler) ones:

Basic simulator. This consists of the simulator that is used to establish the traditional (stand-alone) zero-knowledge property of $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}}\rangle$. The simulator is similar to the one used in Barak's original protocol [1].

Alternative simulator. This consists of the simulator that is used for establishing the "simulation soundness" (cf. [44]) of $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}}\rangle$. The simulator is designed to work in the presence of a man-in-the-middle adversary that is conducting a single left interaction of $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}}\rangle$ concurrently with a single right interaction of $\langle P_{\tilde{\mathsf{tag}}}, V_{\tilde{\mathsf{tag}}}\rangle$. It guarantees that an adversary whose left view consists of a simulated execution of $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}}\rangle$ cannot break the soundness of $\langle P_{\tilde{\mathsf{tag}}}, V_{\tilde{\mathsf{tag}}}\rangle$. The simulator is essentially identical to the one used by Pass [40].

Simulator-extractor. The description of the simulation extraction procedure $S = (\mathsf{SIM}, \mathsf{EXT})$ relies on the previous two simulators. The simulator SIM relies on the basic simulator, whereas the extractor EXT (which also employs a simulation of the left interaction) makes use of the alternative simulator.

We turn to provide a description of the above simulation procedures. (We only provide a brief sketch of the basic and alternative simulators, and assume familiarity with the protocols of [1] and [40]. For completeness, the description of the simulator-extractor is nevertheless self-contained.)

5.1 Basic Simulator

Given the program, V^*_{tag}, of an adversary verifier, the basic simulator acts as follows. In Stage 1 of the protocol (i.e., in Slots 1 and 2), the simulator proceeds by committing to the program $\Pi \stackrel{\text{def}}{=} V^*_{\mathsf{tag}}$. Let s_1, s_2 the randomness used for the commitments.

In Stage 2 of the protocol, the simulator proves that it committed to the program of the verifier in Slot 1. More concretely, the simulator uses the tuple $\langle \Pi, c_1, s_1 \rangle$ as a witness for $\langle h, c_1, r_1 \rangle \in L_{\mathsf{sim}}$ (where L_{sim} is the language that corresponds to R_{sim}). This is a valid witness, since: (1) by the definition of Π it holds that $\Pi(c_1) = r_1$, and (2) as long as $\ell(n) \geq 3n$, for every $\mathsf{tag} \in [n]$, $|r_i| - |c_i| = \ell(n) - |c_i| \geq n$.

5.2 Alternative Simulator

The alternative simulator is constructed having a man-in-the middle adversary A in mind. Consider an A that manages to violate the soundness of protocol $\langle P_{\tilde{\mathsf{tag}}}, V_{\tilde{\mathsf{tag}}} \rangle$, while verifying a simulated proof of $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}} \rangle$. We show how to construct a cheating prover P^* for a single instance of $\langle P_{\tilde{\mathsf{tag}}}, V_{\tilde{\mathsf{tag}}} \rangle$ by forwarding A's messages in $\langle P_{\tilde{\mathsf{tag}}}, V_{\tilde{\mathsf{tag}}} \rangle$ to an external honest verifier V and internally simulating the messages of $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}} \rangle$ for A. The problem that arises in the attempt to simulate is that the code of the external verifier V is not available to the simulator. This means that a stand-alone simulation of the protocol $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}} \rangle$ cannot be completed as it is, since it explicitly requires possession of a "short" program Π that would have generated the corresponding verifier messages.

On a high-level, the prover P^* simulates the left interaction in the following way.[17] In Slot 1 of the protocol, the simulator proceeds by committing to the program $\Pi_1 \stackrel{\text{def}}{=} A$. So far its instructions are just like the basic simulator. In Slot 2, however, the simulator commits to a program Π_2 which consists of both the code of A and all messages A has received from $V_{\tilde{\mathsf{tag}}}$ in the right interaction. In Stage 2 of the protocol, the simulator attempts to prove that it committed to the program of the verifier in either Slot 1 or Slot 2. The simulator will succeed in this task provided that there exists a "short" message y (the actual required length of y is determined by the tag tag and the slot number) such that $\Pi_1(y) = r_1$ or $\Pi_2(y) = r_2$, where r_1, r_2 denote the challenges receives in Slot 1 and 2 respectively (of $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}} \rangle$).

Note that except for the "long" challenges $\tilde{r_1}, \tilde{r_2}$ sent by the verifier of $\langle P_{\tilde{\mathsf{tag}}}, V_{\tilde{\mathsf{tag}}} \rangle$ we do have a description of all messages sent to the adversary A that is shorter than $\ell(n) - n$ (since $\ell(n) = \ell'(n) + n$, where $\ell'(n)$ upper bounds the total length of both prover and verifier messages, except for the challenges r_1, r_2). In order to show that we can still perform a simulation, even in the presence of these messages (for which we do not have a short description), we use the fact that it is sufficient to have a short description of the messages sent in *one* of the slots of $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}} \rangle$. As in [40], we separate between two different schedulings:

There exists one "free" slot j in $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}} \rangle$ in which neither of $\tilde{r_1}$, $\tilde{r_2}$ are contained. In this case the "free" slot j in $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}} \rangle$ can be used to perform a basic simulation (since in this case the simulator did indeed produce a commitment c_j to the code of a machine that on input c_j outputs the challenge r_j in slot j).

The messages r_1, r_2 in $\langle P_{\tilde{\mathsf{tag}}}, V_{\tilde{\mathsf{tag}}} \rangle$ occur in slot 1,2 respectively in $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}} \rangle$. By the construction of the protocols it follows that the length of either the first or the second challenge in $\langle P_{\tilde{\mathsf{tag}}}, V_{\tilde{\mathsf{tag}}} \rangle$ is at least $\ell(n)$ bits longer than the corresponding challenge in $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}} \rangle$. Thus

[17] We provide a detailed description of the actual simulation procedure when we later apply it in the construction of the simulator-extractor.

there exist a slot j in $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}} \rangle$ such that even if we include the verifier's challenge $\tilde{r}_j$ from the protocol $\langle P_{\tilde{\mathsf{tag}}}, V_{\tilde{\mathsf{tag}}} \rangle$ in the description y, we still have $\ell(n) - n$ bits to describe all other messages.

5.3 Simulator-Extractor

Consider a man-in-the-middle adversary A. We assume without loss of generality that A is deterministic and has the auxiliary input z hardwired in. Let k denote the number of rounds in $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}} \rangle$, and let m be an upper-bound on the number of right interactions that A participates in. We describe a combined simulator-extractor $S = (\mathsf{SIM}, \mathsf{EXT})$, that proceeds as follows on input x, z, and TAG.

5.3.1 Simulation of view

We start by describing a machine SIM that simulates the view of A. This requires simulating all the left and the right interactions for A. In the right interactions SIM acts as a verifier. Thus, simulation is straightforward, and is performed by simply playing the role of an honest verifier in all the executions of the protocol. In the left interaction, on the other hand, SIM is supposed to act as a prover, and thus the simulation task is more involved. Towards its goals, SIM acts as follows.

1. For all $i \in [m]$, pick random $\bar{r}_i = (r_{i,1}, \ldots, r_{i,k})$ honest verifier messages for the right interactions. Messages in the right interactions are then emulated by playing the role of the honest verifiers with the fixed random messages $\bar{r}_1, \ldots, \bar{r}_m$. That is, in order to emulate the j^{th} message in the i^{th} right interaction, SIM forwards the message $r_{i,j}$ to A.

2. The left interaction is simulated as follows. SIM views the execution of A and the emulation of the right interactions (with the fixed messages $\bar{r}_1, \ldots, \bar{r}_m$) as a *stand-alone* verifier for the left interaction and applies a close variant of the basic simulator to this interaction. Let $\Pi(\cdot)$ denote the joint code of A and the emulation of the right interactions (including the coins $\bar{r}_1, \ldots, \bar{r}_m$). Whereas the basic simulator would have committed to $\Pi(\cdot)$, we instead let SIM commit to a program $\Pi'(b, \cdot)$ that is defined as follows:

   (a) if $b = 0$, execute $\Pi(\cdot)$;

   (b) if $b = 1$, execute $\Pi(\cdot)$ with the exception that messages $\bar{r}_i = (r_{i,1}, \ldots, r_{i,k})$ (i.e., messages of the i^{th} right interaction) are not emulated, but rather received externally as input.

   Thereafter, SIM proceeds exactly as the basic simulator, by additionally using both $b = 0$ and Π' as a witness in stage 2 of the protocol. More concretely, SIM starts by computing $h = \Pi(\cdot)$. It then generates prover commitments $c_1 = \mathrm{Com}(h(\Pi'); s_1)$ and $c_2 = \mathrm{Com}(h(\Pi'); s_2)$, where $s_1, s_2 \xleftarrow{\mathrm{R}} \{0,1\}^{\mathrm{poly}(n)}$. Using c_1 and c_2, it computes $r_1 = \Pi(c_1)$, and $r_2 = \Pi(c_1, c_2)$. Combining the messages together, this results in a Stage 1 transcript $\tau_1 = \langle h, c_1, r_1, c_2, r_2 \rangle$.

   By definition of $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}} \rangle$, the transcript τ_1 induces a Stage 2 *WIUARG* with $\Pi(h, c_1, c_2)$ as verifier and $(x, \langle h, c_1, r_1 \rangle, \langle h, c_2, r_2 \rangle)$ as common input. Using $\langle \Pi', (0, h, c_1), s_1 \rangle$ as witness for the statement $\langle h, c_1, r_1 \rangle \in L_{\mathsf{sim}}$, the SIM follows the prescribed prover strategy of the *WIUARG* and produces a convincing stage 2 transcript τ_2. Since $r_1 = \Pi(c_1) = \Pi'(0, c_1)$ and since $|0| + |c_1| \leq \ell(n)$ it follows that SIM can always succeed in this task.

Figure 9 demonstrates the definition of SIM, as well as of the program $\Pi'(b, \cdot)$ (for simplicity the various sessions are depicted as if they were executed sequentially).

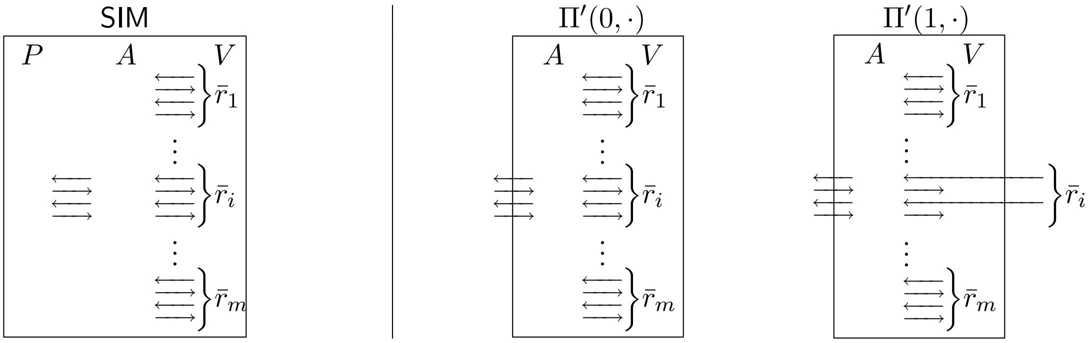

Figure 9: The simulator SIM and the program $\Pi'(b,\cdot)$.

5.3.2 Extraction of witnesses

Once the view of A has been simulated, we turn to the extraction of witnesses to the statements proved by A. Note that we need to extract witnesses to all concurrent right interactions. Towards this goal we rely on a variant of Lindell's concurrent extraction technique [33], combined with the alternative simulator technique described in Section 5.2. In a sense, this can be seen as a (non-trivial) extension of the method of Pass and Rosen [42] (which was used to show a similar property for the simpler case of only one right interaction).

The machine EXT fixes the random coins of the simulator SIM and iteratively extracts witnesses for each of the right interactions. More specifically, EXT starts by sampling a random execution of SIM, using random coins $\bar{s}, \bar{r}$. Let $x_1, \ldots, x_m$ be the inputs corresponding to the m sessions that have taken place in the right interaction.

For each $i \in [m]$ such that the i^{th} right session was *not* accepting, EXT will assume that no witness exists for the corresponding statement x_i, and will refrain from extraction. For all $i \in [m]$ so that the i^{th} right session *is* accepting in this execution of SIM, and for which the tag of the i^{th} session is different from the tag of the left session, EXT will attempt to extract a witness for the statement x_i being proved in the corresponding session.

To do so EXT constructs a stand-alone prover P_i for the i^{th} right interaction $\langle P_{\tilde{\mathsf{tag}}_i}, V_{\tilde{\mathsf{tag}}_i} \rangle$, and from which it will later attempt to extract the witness. In principle, the prover P_i will follow SIM's actions using the same random coins $\bar{s}, \bar{r}$ used for initially sampling the execution of SIM. However, P_i's execution will differ from SIM's execution in the following important ways:

1. Messages in the i^{th} right session are no longer emulated internally, but forwarded externally.

2. In the simulation of the left protocol, use the alternative simulator from Section 5.2 in order to complete stage 2 of the protocol.

Figure 10 demonstrates the definition of P_i (as before, the sessions are depicted sequentially).

The reason for using the alternative simulation instead of the basic one is that the latter might not be able to commit to the external messages of the i^{th} right interaction (as it might not know these messages at the time it commits). Note that the way the simulation within P_i is defined, the program committed to in Stage 1 is Π'. To enable the alternative simulation with a commitment to Π' in Stage 1, we let the simulator additionally provide the input $b = 1$ to Π' as part of the witness in Stage 2 (this enables Π' to depend on the external messages in the i^{th} right session). The alternative simulation technique, combined with the fact that there is only *one* external interaction on the right hand side, are what eventually enables the simulation to go through.

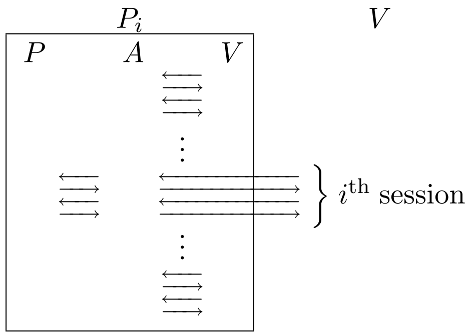

Figure 10: The prover P_i.

The actual witness used by the simulator in Stage 2 depends on the scheduling of the messages. We distinguish between the following cases, depending on where the i^{th} session has started with relation to the messages c_1, c_2 in the left hand side protocol (see Figure 11).

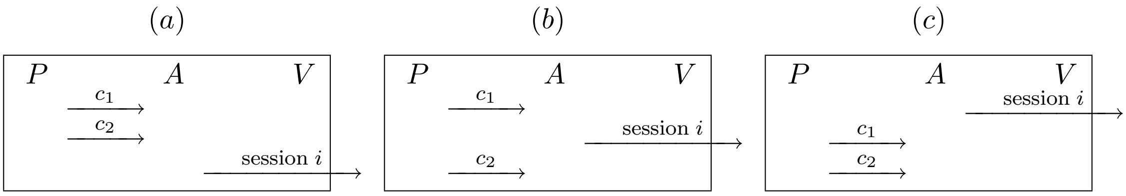

Figure 11: Three possible "starting points" for session i.

In each corresponding case, EXT acts as follows:

Both c_1 and c_2 have been sent before session i begins (Figure 11.a). In this event Slot 1 has no external messages and the basic simulation can be performed, i.e., EXT can use Π' as a witness for $\langle h, c_1, r_1\rangle \in L_{\mathsf{sim}}$ in Stage 2 (just as in SIM).

c_1 has been sent but not c_2 (Figure 11.b). Let M_1, M_2 denote the "external" messages A receives on the right hand side in Slot 1 and Slot 2 of the left interaction, respectively (see Figure 12 for two "representative" schedulings). In this case, we define $\Pi'_2(b, \cdot) = \Pi'(b, M_1, \cdot)$ and let EXT send $c_2 = \mathrm{Com}(\Pi'_2; s)$ (whereas c_1 is defined just as in SIM).

Consider a Stage 1 transcript $\tau_1 = \langle h, c_1, r_1, c_2, r_2\rangle$ of the left interaction. By the construction of $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}}\rangle$, and from the fact that $\tilde{\mathsf{tag}}_i$ is different from tag, it must be the case that either $|M_1| \leq |r_1| - \ell(n)$ or $|M_2| \leq |r_2| - \ell(n)$. This is implied by the following simple fact.

Fact 5.2 *If* $\mathsf{tag} \neq \tilde{\mathsf{tag}}$ *then there exists* $i \in \{1, 2\}$ *so that* $|\tilde{r}_i| \leq |r_i| - \ell(n)$.

In particular, either $|M_1| + |c_1| + n \leq |r_1| - n$ or $|M_2| + |c_1| + n \leq |r_2| - n$. Furthermore, c_1 is a commitment to Π' and c_2 is a commitment to Π'_2, and $r_1 = \Pi'(1, (c_1, M_1))$ and $r_2 = \Pi'_2(1, (c_1, M_2))$. Thus, either $w_1 = \Pi', 1, (c_1, M_1), s_1$ is a valid witness for $\langle h, c_1, r_1\rangle \in L_{\mathsf{sim}}$ or $w_2 = \Pi'_2, 1, (c_2, M_2), s_2$ is a valid witness for $\langle h, c_2, r_2\rangle \in L_{\mathsf{sim}}$. If the former is true, EXT follows the prescribed prover using w_1 as witness, and otherwise uses w_2 as witness.

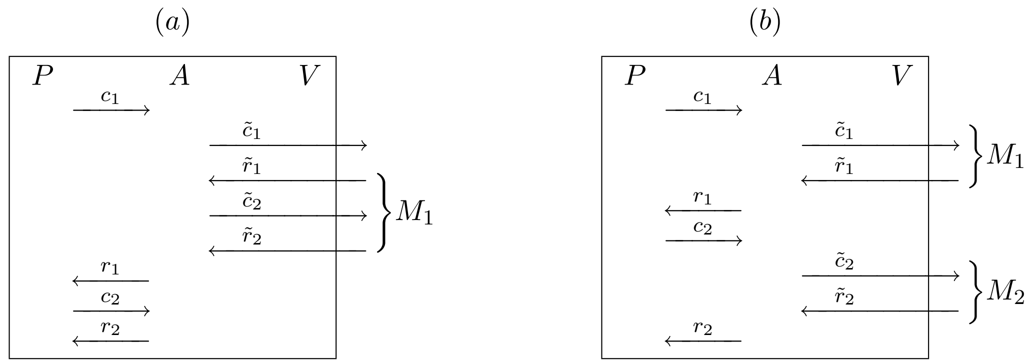

Figure 12: Two "representative" schedulings.

Neither of c_1 or c_2 have been sent (Figure 11.c). In this case EXT first generates a commitment c_1 just as SIM would, i.e., lets c_1 be a commitment to Π', and then performs the same operations as in the previous case.

It follows from the description above that the simulation employed by P_i on the left interaction is always able to convince A of the validity of the statement proved on the left interaction.

Once P_i is constructed, EXT can apply the (stand-alone) extractor, guaranteed by the proof of knowledge property of $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}}\rangle$, to P_i and extract a witness to the statement x_i. In the unlikely event that the extraction failed in any of the m executions, EXT outputs `fail`, and otherwise it outputs all the extracted witnesses.

Remark 3 *It is important to have a P_i for the* entire *protocol $\langle P_{\tilde{\mathsf{tag}}_i}, V_{\tilde{\mathsf{tag}}_i}\rangle$ (and not just for $\langle P_{UA}, V_{UA}\rangle$). This is in order to argue that the witness extracted is a witness for x_i and not a witness to $\langle h, c_1, r_1\rangle \in L_{\mathsf{sim}}$ or to $\langle h, c_2, r_2\rangle \in L_{\mathsf{sim}}$ (which could indeed be the case if we fixed the messages $\langle h, c_1, r_1, c_2, r_2\rangle$ in advance).*

The output of S. Finally the combined simulator-extractor S outputs `fail` whenever EXT does. Otherwise, S outputs whatever SIM outputs, followed by whatever EXT outputs.

Remark 4 *It is important to have both SIM and EXT use the same simulator program S (with same random coins) in their respective executions. Otherwise, we are not guaranteed that the statement $\tilde{x}$ appearing in the output of SIM is the same one EXT extracts a witness from.*[18]

5.4 Correctness of the simulation-extraction

We proceed to show the correctness of the combined simulator-extractor $S = (\mathsf{SIM}, \mathsf{EXT})$. We start by showing that the view of A in the simulation by SIM is identical to its view in an actual interaction. Let the random variable $\mathsf{SIM}(x, z, \textsc{tag})$ denote the view of A in the simulation by SIM performed in the execution of $S(x, z, \textsc{tag})$.

Claim 5.3 *$\{\mathsf{SIM}(x, z, \textsc{tag})\}_{x\in L, z\in\{0,1\}^*, \mathsf{tag}\in[2|x|]}$ and $\{\mathsf{view}_A(x, z, \textsc{tag})\}_{x\in L, z\in\{0,1\}^*, \mathsf{tag}\in[2|x|]}$ are identically distributed.*

[18] The statement $\tilde{x}_i$ will remain unchanged because $\tilde{x}_i$ occurs prior to any message in $\langle P_{\tilde{\mathsf{tag}}_i}, V_{\tilde{\mathsf{tag}}_i}\rangle$ (and hence does not depend on the external messages received by P_i).

Proof: The claim follows from (1) the perfect zero-knowledge property of $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}} \rangle$, and (2) the fact that the emulation of the right interactions by SIM is perfect. Specifically, consider the following hybrid experiments.

1. Let H_0 denote the view of A in the simulated execution.

2. Let H_1 denote the view of A in a simulated execution when letting the simulator use the true witness w for x in the special-purpose $UARG$, $\langle P_{\mathsf{sUA}}, V_{\mathsf{sUA}} \rangle$, in Stage 2 (instead of using the "fake" witness"). Thus the only difference between H_0 and H_1 is the choice of the witness used in $\langle P_{\mathsf{sUA}}, V_{\mathsf{sUA}} \rangle$.

3. Let H_2 denote the real execution. Note that the only difference between H_1 and H_2 is that in H_1 A receives commitments c_1, c_2 to a program Π', whereas in H_2 it receives a commitments to the string 0^k.

Sub Claim 5.4 *H_0 is identically distributed to H_1*

Proof: The claim follows from the witness independent property of $\langle P_{\mathsf{pWI}}, V_{\mathsf{pWI}} \rangle$ used in Stage 2. More precisely, assume for contradiction that H_0 is not identically distributed to H_1. Then there must exists some Stage 1 transcript, such that the proofs generated in Stage 2 in H_0 and H_1 are not identically distributed, in contradiction to the witness independence property (here we use the fact that there exist two possible witnesses for stage 2 – one is the witness used by the simulator and the other is w). ■

Sub Claim 5.5 *H_1 is identically distributed to H_2*

Proof: The claim directly follows from the perfect hiding property of the commitments used to generate c_1 and c_2. ■

This completes the proof of Claim 5.3. ■

We proceed to show that EXT outputs `fail` with negligible probability. Let the random variable $\mathsf{EXT}(x, z, \mathsf{tag})$ denote the output of EXT in the execution of $S(x, z, \mathsf{tag})$.

Claim 5.6 *There exists a negligible function $\mu(\cdot)$ such that for every $x \in L, z \in \{0,1\}^*, \mathsf{tag} \in [2|x|]$*

$$\Pr\left[\mathsf{EXT}(x, z, \mathsf{tag}) = \mathtt{fail}\right] \leq \mu(|x|)$$

Proof: Recall that EXT proceeds by constructing stand alone provers P_i, and then applying the (stand-alone) extractor, guaranteed by the proof of knowledge property of $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}} \rangle$, to P_i, in order to extract witnesses w_i to the statements x_i. Note that EXT outputs `fail` only in the event that extraction from one of the right interactions $i \in [m]$ fails. By the proof of knowledge property of $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}} \rangle$ it holds that for each execution i extraction for execution i fails only with negligible probability (recall that the extractor is only invoked when A provides an accepting proof in execution i). Since the extraction procedure is repeated at most m times (at most once per right-interaction), we conclude (by the Union Bound) that the probability that extraction fails for any of the right interactions is negligible. ■

By combining Claim 5.3 and 5.6, we conclude that, the first output of S is statistically close to A's view in an actual interaction. As in Definition 4.1, let the random variable $S_1(x, z, \textsc{tag})$ denote the first output of $S(x, z, \textsc{tag})$.

Claim 5.7 *$\{S_1(x, z, \text{TAG})\}_{x \in L, z \in \{0,1\}^*, \mathsf{tag} \in [2|x|]}$ and $\{\mathsf{view}_A(x, z, \text{TAG})\}_{x \in L, z \in \{0,1\}^*, \mathsf{tag} \in [2|x|]}$ are statistically close over L.*

Proof: Recall that the first output of S consists of the view of A as generated by SIM. By claim 5.3 if follows that the first output of S is identically distributed to a "real" interaction, conditioned on the event that S does not output `fail`. However, since this event only happens when EXT outputs `fail`, which by claim 5.6 only happens with negligible probability, the claim follows. ■

We proceed to show the correctness of the extraction.

Claim 5.8 *Let $x \in L, z \in \{0,1\}^*, \text{TAG} \in \{0,1\}^{2|x|}$, and let $(\mathsf{view}, \bar{w})$ denote the output of $S(x, z, \text{TAG})$ (on input some random tape). Let $\tilde{x}_1 \ldots, \tilde{x}_m$ be the right-execution statements appearing in view and let $\tilde{\mathsf{tag}}_1 \ldots \tilde{\mathsf{tag}}_m$ denote the correspoding right-execution tags. Then, for any $i \in [m]$ such that the i^{th} right-execution in view is accepting AND $\mathsf{tag} \neq \tilde{\mathsf{tag}}_i$, $\bar{w}$ contains a witness w_i so that $R_L(\tilde{x}_i, w_i) = 1$.*

Proof: First, note that since S always outputs `fail` whenever the extraction by EXT fails, the claim trivially holds in the event that the extraction by EXT fails. Consider, next, the case when extraction by EXT does not fail. Recall that EXT performs extraction for all right-executions which satisfy the properties described in the hypothesis (i.e., accepting proofs and different tags). Furthermore, for each such interaction i, the stand-alone prover P_i—constructed by EXT—uses the same random coins as SIM in order to emulate all the interactions before session i begins. In addition, the prescribed actions for the simulation by EXT are identical to the prescribed actions for the simulation by SIM. This means that the statement proved by P_i will be identical to the statement proved in the view output by SIM. Finally, by our assumption that the extraction by EXT does not fail, we conclude that a valid witness for the statement proved by P_i is extracted. This concludes the claim. ■

We conclude the proof by bounding the running time of the combined simulator-extractor S.

Claim 5.9 *$S(x, z, \mathsf{tag})$ runs in expected polynomial time (in $|x|$).*

Proof: We start by noting that the running time of SIM is polynomial. Recall that the program Π' committed to by SIM is of size $\text{poly}(n)$. It thus directly follows that simulation of Stage 1 messages can be done in polynomial time. Furthermore, it follows that the verification time of R_{sim} on the instance $\langle h, c_1, r_1 \rangle$ is polynomial in n. Finally, by the relative prover efficiency of $\langle P_{\mathsf{UA}}, V_{\mathsf{UA}} \rangle$ it holds that the simulator can generate also Stage 2 message in polynomial time.

It now only remains to show that the *expected* running time of EXT is also polynomial. Recall that EXT proceeds by first sampling a view using SIM and then proceeds to extract witnesses in all *accepting* right executions. We show that for every right execution i, the expected running-time needed to extract a witness from that execution is polynomially bounded. Since the number of right interactions is polynomially bounded, we conclude by linearity of expectations that the total expected running time of the combined simulator-extractor $S = (\mathsf{SIM}, \mathsf{EXT})$ is polynomially bounded.

Let $\mathtt{view}_i$ denote the partial view for A in an emulation by SIM up until A is about to start its i^{th} right execution. Let $p_i(\mathtt{view}_i)$ denote the probability that A produces an accepting proof in the i^{th} right execution in the simulation by SIM, given that SIM has fed to A the view $\mathtt{view}_i$. Let $p'_i(\mathtt{view}_i)$ denote the probability that A produces an accepting proof in the i^{th} right execution in the simulation by P_i (constructed in EXT), given that EXT has fed A the view $\mathtt{view}_i$.

Sub Claim 5.10 *Let* $\mathtt{view}_i$ *denote the partial view for* A *in a emulation by* SIM *up until* A *is about to start its* i^{th} *right execution. Then,* $p_i(\mathtt{view}_i) = p'_i(\mathtt{view}_i)$*.*

Proof: The claim follows from the *perfect* indistinguishability of the basic simulator used by SIM, and the alternative simulator used by EXT (this is proved similarly to Claim 5.3). ■

Note that if we only assume that $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}}\rangle$ is statistical zero-knowledge, we could only conclude that $p'_i(\mathtt{view_i})$ is *negligibly close* to $p_i(\mathtt{view_i})$. This would not be sufficient to bound the running-time of the simulator (as this would have introduced difficulties similar to the ones discussed in [21]).

By the proof of knowledge property of $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}}\rangle$ it holds that for any partial view $\mathtt{view}_i$ up until A is about to start its i^{th} right execution, the expected running-time of the extractor is bounded by

$$\frac{\text{poly}(n)}{p'_i(\mathtt{view}_i)}$$

Since the probability of invoking the extraction procedure given this partial view is $p_i(\mathtt{view}_i)$, the expected number of steps used to extract a witness is[19]

$$p_i(\mathtt{view_i})\frac{\text{poly}(n)}{p'_i(\mathtt{view_i})} = p_i(\mathtt{view_i})\frac{\text{poly}(n)}{p_i(\mathtt{view_i})} = \text{poly}(n)$$

We conclude that the expected time needed to extract the witness in the i^{th} right execution is polynomially bounded. The claim follows. ■

This completes the proof of Lemma 5.1. ■

5.5 Analyzing the Family of 2^n Protocols

Relying on the proof from Section 5.3 we now argue that the family $\{\langle P_{\text{TAG}}, V_{\text{TAG}}\rangle\}_{\text{TAG}\in\{0,1\}^n}$ is also one-many simulation extractable. The key for demonstrating this is to show that the protocols $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}}\rangle$ are simulation-extractable as long as the number of left interactions is *a-priori bounded* (in contrast to the single left interaction considered in Def. 4.1), and even if the number of right interactions is unbounded.

Specifically, consider a man-in-the middle adversary A that is simultaneously participating in $k = k(n)$ left interactions of $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}}\rangle$, acting as verifier, and an (unbounded) polynomial number of right-interactions of $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}}\rangle$, acting as prover. Let $\mathsf{view}_A(x, z, \mathsf{tag})$ denote the view of $A(x, z)$ when receiving left-proofs of statements $\bar{x} = x_1, \ldots, x_k$, using identity strings $\bar{\mathsf{tag}} = \mathsf{tag}_1, \ldots, \mathsf{tag}_k$, and proving statements of its choice in the right interaction (using tags of its choice). Given a function $t = t(n)$ and some $k \in N$, we use the notation $\{\cdot\}_{\bar{x},z,\bar{\mathsf{tag}}}$ as shorthand for $\{\cdot\}_{x_1,\ldots,x_k\in L, z\in\{0,1\}^*, \mathsf{tag}_1,\ldots,\mathsf{tag}_k\in\{0,1\}^{t(|x|)}}$.

Definition 5.11 (Bounded-many Simulation-extractability) *A family* $\{\langle P_{\mathsf{tag}}, V_{\mathsf{tag}}\rangle\}_{\mathsf{tag}\in\{0,1\}^*}$ *of interactive proofs for the language* L *, is said to be* k-bounded-many simulation extractable *with tags of length* $t = t(n)$ *if for* any *polynomial* $p(\cdot)$ *and any man-in-the-middle adversary* A *that participates in* $k = k(n)$ *left interactions and at most* $m = p(n)$ *right interactions, there exists a probabilistic expected poly-time machine* S *such that:*

1. *The probability ensembles* $\{S_1(\bar{x}, z, \bar{\mathsf{tag}})\}_{\bar{x},z,\bar{\mathsf{tag}}}$ *and* $\{\mathsf{view}_A(\bar{x}, z, \bar{\mathsf{tag}})\}_{\bar{x},z,\bar{\mathsf{tag}}}$ *are statistically close over* L*, where* $S_1(\bar{x}, z, \bar{\mathsf{tag}})$ *denotes the first output of* $S(\bar{x}, z, \bar{\mathsf{tag}})$*.*

[19] It is here that complications arise in the case when $p'_i \neq p_i$. Note that the expected number of steps is no longer guaranteed to be polynomial in this case, *even* if p'_i is negligibly close to p_i.

2. *Let* $x_1, \ldots, x_k \in L, z \in \{0,1\}^*, \mathsf{tag}_1 \ldots \mathsf{tag}_k \in \{0,1\}^{t(|x|)}$*, and let* $(\mathsf{view}, \bar{w})$ *denote the output of* $S(\bar{x}, z, \bar{\mathsf{tag}})$ *(on input some random tape). Let* $\tilde{x}_1 \ldots, \tilde{x}_m$ *be the right-execution statements appearing in* view *and let* $\tilde{\mathsf{tag}}_1, \ldots, \tilde{\mathsf{tag}}_m$ *denote the correspoding right-execution tags. Then, for any* $i \in [m]$ *such that the* i^{th} *right-execution in* view *is accepting AND for all* $j \in [k]$ $\tilde{\mathsf{tag}}_i \neq \mathsf{tag}_j$*,* $\bar{w}$ *contains a witness* w_i *so that* $R_L(\tilde{x}_i, w_i) = 1$*.*

Lemma 5.12 *Suppose that* $\{\mathcal{H}_n\}_n$ *is collision resistant for* $T(n)$*-sized circuits, that* **Com** *is perfectly hiding, that* $\langle P_{\mathsf{pWI}}, V_{\mathsf{pWI}} \rangle$ *is a witness independent argument of knowledge, that* $\langle P_{\mathsf{UA}}, V_{\mathsf{UA}} \rangle$ *is a universal argument and that* $\ell(n) \geq 2n^3 + n$*. Then,* $\{\langle P_{\mathsf{tag}}, V_{\mathsf{tag}} \rangle\}_{\mathsf{tag} \in [2n]}$ *is* n*-bounded-many simulation extractable with tags of length* $t(n) = \log n + 1$*.*

Proof: The proof is essentially identical to the proof of one-many simulation-extractability of $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}} \rangle$ (Lemma 5.1). The only difference is that in the simulation by SIM (and EXT), the message r_1 in the i^{th} left execution can no longer be computed as $\Pi(c_1)$, but is in fact defined as $\Pi(M)$ where M denotes all left-hand side prover messages that have occurred before r_1 (the same holds analogously for r_2). This creates a potential problem when simulating the Stage 2 messages in the j^{th} left protocol.

The key observation is that the *total* length of all prover messages on the left interaction does not exceed $2n^3$ (here we assume w.l.o.g. that the length of all prover messages in a session is upper bounded by n^2). Thus SIM (as well as EXT) can include *all* left-hand side prover messages sent to A before the message r_1 (or r_2 depending on what "slot" the simulator uses) as part of the witness for either $\langle h, c_1, r_1 \rangle \in L_{\mathsf{sim}}$ or $\langle h, c_2, r_2 \rangle \in L_{\mathsf{sim}}$. ■

Our main technical lemma (Lemma 4.3) is finally obtained by combining Lemmata 5.1 and 5.12.

Lemma 4.3 (Main technical lemma) *Suppose that* $\{\mathcal{H}_n\}_n$ *is collision resistant for* $T(n)$*-sized circuits, that* **Com** *is perfectly hiding, that* $\langle P_{\mathsf{pWI}}, V_{\mathsf{pWI}} \rangle$ *is a witness independent argument of knowledge, that* $\langle P_{\mathsf{UA}}, V_{\mathsf{UA}} \rangle$ *is a universal argument and that* $\ell(n) \geq 2n^3+n$*. Then,* $\{\langle P_{\text{TAG}}, V_{\text{TAG}} \rangle\}_{\text{TAG} \in \{0,1\}^n}$ *is one-many simulation extractable with tags of length* $t(n) = n$*.*

Proof: Consider a man-in-the-middle adversary A that is verifying a statement x with identity string $\text{TAG} = \text{TAG}_1, \ldots, \text{TAG}_n$ in the left interaction while proving m statements $\tilde{x}_1, \ldots, \tilde{x}_m$ in the right interaction, where for $i \in [m]$ the i^{th} right session has identity string $\tilde{\text{TAG}}^i = \tilde{\text{TAG}}^i_1, \ldots, \tilde{\text{TAG}}^i_n$. We show how to construct a simulator-extractor $S = (\mathsf{SIM}, \mathsf{EXT})$ that simulates the view of A while extracting all the witnesses for statements $\tilde{x}_i$ for which $\tilde{\text{TAG}}^i \neq \text{TAG}$.

First, observe that for any $i \in [m]$ so that $\tilde{\text{TAG}}^i \neq \text{TAG}$, there exist $i_0 \in [n]$ for which $(i_0, \tilde{\text{TAG}}^i_{i_0}) \neq (j, \text{TAG}_j)$ for all $j \in [n]$ (just take the i_0 for which $\tilde{\text{TAG}}^i_{i_0} \neq \text{TAG}_{i_0}$). Let $\tilde{\mathsf{tag}}_i = (i_0, \tilde{\text{TAG}}^i_{i_0})$.

Given a one-many adversary A for $\langle P_{\text{TAG}}, V_{\text{TAG}} \rangle$, we, next, construct an n-many adversary A' for $\langle P_{\mathsf{tag}}, V_{\mathsf{tag}} \rangle$ that runs n parallel sessions in the left interaction and $m' = mn$ concurrent sessions in the right interaction. The inputs and identity strings for the various sessions are defined as follows:

Left sessions. For $j \in [n]$ the common input of the j^{th} left session is $x_j = x$ and the identity string is $\bar{\mathsf{tag}} = (j, \text{TAG}_j)$.

Right sessions. For $(i, j) \in [m] \times [n]$, the input to the $(i, j)^{\text{th}}$ right session is $\tilde{x}_j$ and the identity string is $\tilde{\mathsf{tag}}^i_j = (j, \tilde{\text{TAG}}^i_j)$.

By Lemma 5.12 there exists a simulator S' that produces a view that is statistically close to the real view of A', and outputs witnesses to all right executions for which the tag is different from all of $(1, \text{TAG}_1), \ldots, (n, \text{TAG}_n)$. Relying on S', we construct the simulator-extractor S. $S(x, z, \text{TAG})$

proceeds as follows. It parses TAG as $\text{TAG} = \text{TAG}_1, \ldots, \text{TAG}_n$, where $\text{TAG}_i \in \{0,1\}$. For $i \in [n]$, let $x_i = x$, $\mathsf{tag}_i = (i, \text{TAG}_i)$. Let $(\mathsf{view}, \bar{w})$ denote the output of $S'(x_1, \ldots x_n, z, \mathsf{tag}_1, \ldots \mathsf{tag}_n)$. Additionally, let $\tilde{x}_1 \ldots, \tilde{x}_{m'}$ be the right-execution statements appearing in view and $\tilde{\text{TAG}}^1 \ldots \tilde{\text{TAG}}^{m'}$ the correspoding right-execution tags. As observed above, for any $i \in [m]$ so that $\tilde{\text{TAG}}^i \neq \text{TAG}$, there exists some identity $\tilde{\mathsf{tag}}_i$ that A' uses in the proof of the i^{th} right interaction which is different than *all* identities $(\mathsf{tag}_1, \ldots, \mathsf{tag}_n)$ used in the n left interactions. Thus, for every $i \in [m]$ so that $\tilde{\text{TAG}}_i \neq \text{TAG}$, S can successfully find a witness for the statement $\tilde{x}_i$ in $\bar{w}$. S finally outputs view and the witnesses obtained above. The correctness of the simulator-extractor S follows directly from the construction. ■

6 Acknowledgments

We are grateful to Silvio Micali and Moni Naor for many helpful discussions and encouragement. Thanks also to the anonymous reviewers for their careful reading and thoughtful comments.

References

[1] B. Barak. How to go Beyond the Black-Box Simulation Barrier. In *42nd FOCS*, pages 106–115, 2001.

[2] B. Barak. Constant-Round Coin-Tossing or Realizing the Shared Random String Model. In *43rd FOCS*, pages 345-355, 2002.

[3] B. Barak, R. Canetti, Y. Lindell, R. Pass and T. Rabin. Secure Computation Without Authentication. In *CRYPTO 2005*.

[4] G. Brassard, D. Chaum and C. Crépeau. Minimum Disclosure Proofs of Knowledge. *JCSS*, Vol. 37, No. 2, pages 156–189, 1988. in *27th FOCS*, 1986.

[5] B. Barak and O. Goldreich. Universal Arguments and their Applications. *17th CCC*, pages 194–203, 2002.

[6] M. Bellare and O. Goldreich. On Defining Proofs of Knowledge. In *CRYPTO'92*, Springer (LNCS 740), pages 390–420, 1993.

[7] M. Blum. Coin Flipping by Telephone. In *CRYPTO 1981*, pages 11-15, 1981.

[8] M. Bellare, R. Impagliazzo and M. Naor. Does Parallel Repetition Lower the Error in Computationally Sound Protocols? In *38th FOCS*, pages 374–383, 1997.

[9] R. Canetti and M. Fischlin. Universally Composable Commitments. In *CRYPTO 2001*, pages 19–40, 2001.

[10] R. Canetti, S. Halevi, J. Katz, Y. Lindell, P. D. MacKenzie. Universally Composable Password-Based Key Exchange. In *EUROCRYPT 2005*, pages 404-421, 2005.

[11] I. Damgård and J. Groth. Non-interactive and Reusable Non-Malleable Commitment Schemes. In *35th STOC*, pages 426-437, 2003.

[12] I. Damgård, T. Pedersen and B. Pfitzmann. On the Existence of Statistically Hiding Bit Commitment Schemes and Fail-Stop Signatures. In *CRYPTO 1993*, pages 250–265, 1993.

[13] A. De Santis, G. Di Crescenzo, R. Ostrovsky, G. Persiano and A. Sahai. Robust Non-interactive Zero Knowledge. In *CRYPTO 2001*, pages 566-598, 2001.

[14] G. Di Crescenzo, J. Katz, R. Ostrovsky and A. Smith. Efficient and Non-interactive Non-malleable Commitment. In *EUROCRYPT 2001*, pages 40-59, 2001.

[15] G. Di Crescenzo, Y. Ishai and R. Ostrovsky. Non-Interactive and Non-Malleable Commitment. In *30th STOC*, pages 141-150, 1998

[16] D. Dolev, C. Dwork and M. Naor. Non-Malleable Cryptography. *SIAM Jour. on Computing*, Vol. 30(2), pages 391–437, 2000. Preliminary version in *23rd STOC*, pages 542-552, 1991

[17] U. Feige, D. Lapidot and A. Shamir. Multiple Noninteractive Zero Knowledge Proofs under General Assumptions. *Siam Jour. on Computing 1999*, Vol. 29(1), pages 1-28.

[18] U. Feige and A. Shamir. Witness Indistinguishability and Witness Hiding Protocols. In *22nd STOC*, p. 416–426, 1990.

[19] M. Fischlin and R. Fischlin. Efficient Non-malleable Commitment Schemes. *In CRYPTO 2000*, pages 413-431, 2000.

[20] O. Goldreich. *Foundation of Cryptography – Basic Tools*. Cambridge University Press, 2001.

[21] O. Goldreich and A. Kahan. How to Construct Constant-Round Zero-Knowledge Proof Systems for NP. *Jour. of Cryptology*, Vol. 9, No. 2, pages 167–189, 1996.

[22] O. Goldreich and Y. Lindell. Session-Key Generation Using Human Passwords Only. In *CRYPTO 2001*, p. 408-432, 2001.

[23] O. Goldreich, S. Micali and A. Wigderson. Proofs that Yield Nothing But Their Validity or All Languages in NP Have Zero-Knowledge Proof Systems. *JACM*, Vol. 38(1), pages 691–729, 1991.

[24] O. Goldreich, S. Micali and A. Wigderson. How to Play any Mental Game - A Completeness Theorem for Protocols with Honest Majority. In *19th STOC*, p. 218–229, 1987.

[25] O. Goldreich and Y. Oren. Definitions and Properties of Zero-Knowledge Proof Systems. *Jour. of Cryptology*, Vol. 7, No. 1, pages 1–32, 1994.

[26] S. Goldwasser and S. Micali. Probabilistic Encryption. *JCSS*, Vol. 28(2), pages 270-299, 1984.

[27] S. Goldwasser, S. Micali and C. Rackoff. The Knowledge Complexity of Interactive Proof Systems. *SIAM Jour. on Computing*, Vol. 18(1), pages 186–208, 1989.

[28] I. Haitner and O. Reingold. Statistically-hiding commitment from any one-way function. In *38th STOC*, pages 1–10, 2007.

[29] S. Halevi and S. Micali. Practical and Provably-Secure Commitment Schemes from Collision-Free Hashing. In *Crypto96*, Springer LNCS 1109, pages 201–215, 1996.

[30] J. Håstad, R. Impagliazzo, L.A. Levin and M. Luby. Construction of Pseudorandom Generator from any One-Way Function. *SIAM Jour. on Computing*, Vol. 28 (4), pages 1364–1396, 1999.

[31] J. Katz, R. Ostrovsky, M. Yung. Efficient Password-Authenticated Key Exchange Using Human-Memorable Passwords. In *EUROCRYPT 2001*, pages 475-494, 2001.

[32] J. Kilian. A Note on Efficient Zero-Knowledge Proofs and Arguments. In *24th STOC*, pages 723–732, 1992.

[33] Y. Lindell. Bounded-Concurrent Secure Two-Party Computation Without Setup Assumptions. In *34th STOC*, pages 683–692, 2003.

[34] S. Micali. CS Proofs. *SIAM Jour. on Computing*, Vol. 30 (4), pages 1253–1298, 2000.

[35] M. Naor. Bit Commitment using Pseudorandomness. *Jour. of Cryptology*, Vol. 4, pages 151–158, 1991.

[36] M. Nguyen, S. Ong and S. Vadhan. Statistical Zero-Knowledge Arguments for NP from Any One-Way Function. In *37th FOCS*, pages 3–14, 2006.

[37] M. Naor, R. Ostrovsky, R. Venkatesan and M. Yung. Perfect Zero-Knowledge Arguments for NP Using any One-Way Permutation. *Jour. of Cryptology*, Vol. 11, pages 87–108, 1998.

[38] M. Naor and M. Yung. Universal One-Way Hash Functions and their Cryptographic Applications. In *21st STOC*, pages 33–43, 1989.

[39] M. Nguyen and S. Vadhan. Simpler Session-Key Generation from Short Random Passwords. In *1st TCC*, p. 428–445, 2004.

[40] R. Pass. Bounded-Concurrent Secure Multi-Party Computation with a Dishonest Majority. In *36th STOC*, 2004, pages 232-241, 2004.

[41] R. Pass and A. Rosen. Bounded-Concurrent Secure Two-Party Computation in a Constant Number of Rounds. In *34th FOCS*, pages 404-413, 2003.

[42] R. Pass and A. Rosen. New and Improved Constructions of Non-Malleable Cryptographic Protocols. In *37th STOC*, 2004, pages 533-542, 2005.

[43] R. Pass, A. Shelat, V. Vaikuntanathan. Relations Among Notions of Non-malleability for Encryption Schemes. To appear in *AsiaCrypt'07*, 2007.

[44] A. Sahai. Non-Malleable Non-Interactive Zero Knowledge and Adaptive Chosen-Ciphertext Security. *In 40th FOCS*, pages 543-553, 1999.