Concurrent Non-Malleable Commitments (and More) in 3 Rounds

Michele Ciampi
DIEM
Università di Salerno
ITALY
mciampi@unisa.it

Rafail Ostrovsky
UCLA
Los Angeles
rafail@cs.ucla.edu

Luisa Siniscalchi
DIEM
Università di Salerno
ITALY
lsiniscalchi@unisa.it

Ivan Visconti
DIEM
Università di Salerno
ITALY
visconti@unisa.it

Abstract

The round complexity of commitment schemes secure against man-in-the-middle attacks has been the focus of extensive research for about 25 years. The recent breakthrough of Goyal, Pandey and Richelson [STOC 2016] showed that 3 rounds are sufficient for (one-left, one-right) non-malleable commitments. This result matches a lower bound of [Pas13]. The state of affairs leaves still open the intriguing problem of constructing 3-round concurrent non-malleable commitment schemes.

In this paper we solve the above open problem by showing how to transform any 3-round (one-left one-right) non-malleable commitment scheme (with some extractability property) in a 3-round concurrent non-malleable commitment scheme. Our transform makes use of complexity leveraging and when instantiated with the construction of [GPR16] gives a 3-round concurrent non-malleable commitment scheme from one-way permutations secure w.r.t. subexponential-time adversaries.

We also show how our 3-round concurrent non-malleable commitment scheme can be used for 3-round arguments of knowledge and in turn for 3-round identification schemes secure against concurrent man-in-the-middle attacks.

Keywords: non-malleability, commitments, PoKs, identification schemes.

Contents

1 Introduction

Commitment schemes are fundamental in Cryptography. They require a sender to fix a message that can not be changed anymore, but that will remain hidden to a receiver until the sender decides to reveal it.

In order to model modern real-world adversaries, commitment schemes have been proposed with additional security properties. Here we consider the intriguing question of constructing a scheme that remains secure against man-in-the-middle (MiM) attacks: a non-malleable (NM) commitment scheme [DDN91].

This fascinating setting is much harder to deal with than the classic stand-alone setting. Indeed while we know 1-round and 2-round regular commitment schemes under various assumptions ([GL89, NY89, Ped91, Nao91, HM96]), Pass proved that NM commitments[1] require at least 3 rounds [Pas13] when security is proved through a black-box reduction to a falsifiable (polynomial or subexponential time) hardness assumption. Instead in different and more controversial models (e.g., assuming the existence of a trusted random string, by modeling hash functions as random oracles, by weakening the security definition admitting an inefficient challenger) we know constructions of non-interactive NM commitments [DG03, PPV08].

The round complexity of NM commitment schemes in the standard model has puzzled researchers for long time. Starting from the construction of [DDN91] that required a logarithmic number of rounds, various constant-round schemes were proposed [Bar02, PR05b, PR05a, PR08, PW10, LP11, Goy11, GLOV12, GRRV14, BGR$^+$15, LP15, COSV16]. Interestingly Ciampi et al. in [COSV16] show a 4-round commitment scheme that is secure also when the adversary mounts a concurrent MiM attack, a setting that corresponds to what can actually happen when sender and receiver are connected through a communication network like the Internet. In such a much more interesting setting a MiM adversary receives multiple commitments from senders and sends his commitments to multiple receivers.

1.1 Towards 3-Round (Concurrent) NM Commitments

The existence of 3-round NM commitment schemes is an important question first because 3 is the best possible constant (in light of the lower bound of [Pas13]), and second because 3 is the smallest number of rounds for a primitive that often makes use of commitment schemes: proofs of knowledge.

The importance of obtaining 3-round (and not just any constant-round) NM commitments motivated the very recent and innovative work of [GPR16] that, by just relying on any non-interactive commitment scheme and exploiting the power of non-malleable codes in the split-state model, shows a 3-round NM commitment scheme. Interestingly, such construction is not claimed to be secure against concurrent man-in-the-middle attacks. Therefore the following natural and important question remains open.

Main Open Question: *Can we construct a 3-round concurrent non-malleable commitment scheme matching the lower bound of [Pas13]?*

[1] We consider the notion of NM commitment w.r.t. commitment.

1.2 Other 3-Round Challenges

We list here 3 other interesting settings where no 3-round construction is known against concurrent MiM adversaries.

- Proofs[2] of knowledge are very useful in Cryptography. They have been studied in particular when there are only 3 rounds and the verifier just sends random bits (e.g., Σ-protocols [CDS94, Dam10], Blum's protocol for Hamiltonicity [Blu86]). Despite their importance, there is no construction for 3-round proofs of knowledge (PoK) that is sufficiently secure under concurrent MiM attacks. This is due to the fact that such attacks are in general extremely difficult to deal with. Even though there exist constructions with a constant number of rounds, the case of just 3 rounds so far has remained unsolved.

- In [LS90][3] Lapidot and Shamir proposed a 3-round public-coin witness indistinguishable PoK for NP (the LS protocol) where the input (except its size) is needed only when playing the 3rd round. This special completeness property named "delayed input" in [CPS$^+$16a, CPS$^+$16b] has been critically used in many applications (e.g., [KO04, DPV04, HV16, YZ07, YYZ10, Wee10, GMPP16, HV16, MV16, COSV16]), and in [CPS$^+$16a, CPS$^+$16b] it was considered for the OR composition of Σ-protocols instead of relying on LS. When a PoK is used as sub-protocol the delayed-input feature is instrumental to give a better round complexity to the external protocol. An additional features of delayed-input protocols is that they allow to shift large part of the computation to an off-line phase. Unfortunately the LS protocol and the PoKs of [CPS$^+$16a, CPS$^+$16a] are not secure against concurrent MiM attacks and this penalizes those applications where both round complexity and security against concurrent MiM attacks are important.

- We notice that identification schemes have been often proposed (e.g., [FFS87]) through the paradigm of proving "knowledge" of a secret[4]. Under this formulation there are constant-round constructions that are proven secure against concurrent MiM attacks [BFGM01]. However no 3-round scheme known in literature is proven secure in presence of a concurrent MiM adversary.

1.3 Results of This Work

In this work we study 3-round commitment scheme in presence of concurrent MiM attacks and solve in the positive the above open problems.

3-Round concurrent NM commitment schemes. In the main result of this submission, we show a transform that on input any 3-round NM commitment scheme[5] gives a 3-round concurrent

[2] For simplicity in the informal part of the paper we will not make a strict distinction between proofs and arguments. In the formal part we will use appropriate terms.

[3] See [OV12] for a detailed description of [LS90].

[4] Other notions based on signature or decryption capabilities are considered weaker since in some applications the verifier wants to make sure that the prover is the actual entity matching the announced identity. Indeed without a PoK a prover could give some partial information about his secret to others that can still succeed in convincing the verifier, even though they do not know the full secret.

[5] We also require the scheme to be extractable. Extractability often comes for free since it is commonly used in the non-malleability proof.

NM commitment scheme. The construction of [GPR16] can be used to instantiate our transform, therefore obtaining a 3-round concurrent NM commitment scheme based on any one-way permutation secure against subexponential-time adversaries. This result solves the main open question. Moreover our scheme (still when instantiated with the one of [GPR16] and using a proper one-way permutation) is public coin and (if desired[6]) has the delayed-input property.

Our transform extends the security of the underlying commitment scheme to multiple receivers. It is known that this implies security also with multiple senders [LPV08]. The crucial idea of our transform is to combine the underlying NM commitment scheme along with a one-time pad, to produce a commitment of a message that by itself, in case of a malleability attack, will have sufficient structure to be recognized by a distinguisher in the session in which it appears. Therefore a successful concurrent MiM even playing multiple commitments with multiple receivers will have to maul the underlying commitment scheme in at least one session. Since the message has sufficient structure with respect to that single session, we are able to translate the concurrent MiM attack into a non-concurrent MiM that violates the security of the underlying (non-concurrent) NM commitment scheme. We will implement the idea of committing to a message with structure by forcing a successful concurrent MiM to commit to the solution of a puzzle in at least one session. We will use complexity leveraging to show that the attack of the concurrent MiM is indistinguishable from the attack of a polynomial-time simulator that plays with receivers only.

Just for completeness, we also show an explicit concurrent MiM adversary $\mathcal{A}$ for the scheme of [GPR15]. The crucial point here, following a technique of [FMNV14] is that the scheme of [GPR15] allows $\mathcal{A}$ to spread the message committed by the honest sender over several commitments that the adversary sends to multiple receivers. The scheme presented in [GPR16] is slightly different and became available after our work was already submitted, therefore when describing $\mathcal{A}$ we stick with [GPR15].

3-round arguments of knowledge and ID schemes against concurrent MiM attacks. We notice that our 3-round concurrent NM commitment scheme is a commit-and-prove argument of knowledge (AoK). This means that one can see our scheme as a commitment followed by an AoK about the committed value. By applying a simple change to the statement of the underlying AoK we obtain a 3-round concurrent NM witness-indistinguishable AoK (concurrent NMWIAoK) a notion introduced in [OPV08] and later on extended in [LPV09]. We stress that the delayed-input and public-coin properties of our commitment scheme are preserved by our concurrent NMWIAoK.

Notice that AoKs under standard assumptions require at least 3 rounds. The simulation-based notion for concurrent non-malleable arguments of knowledge (i.e., concurrent NM zero knowledge) requires at least a polylogarithmic number of rounds [CKPR01, BPS06] when the simulator is black box. In [OPV08] it is shown how to get concurrent NM zero knowledge (NMZK) in the bare public-key (BPK) model [CGGM00, MR01] with just two executions of a concurrent NMWIAoK. Therefore here we directly obtain a round-efficient concurrent NMZKAoK in the BPK model. Notice also that by making use of the delayed-input feature the simulator can extend a main thread avoiding issues due to aborting adversaries as discussed in [SV12, ORSV13].

Finally, we notice that one can get an identification scheme secure in the PoK sense in the concurrent[7] setting of [BFGM01] as well as under the stronger definition based on matching conversations

[6] Our transform can be instantiated in two ways. In the former the message to commit is required already when playing the first round, while in the latter the message to commit is required when playing the third round only.

[7] In [BFGM01] a notion called CR2 is proposed to deal with concurrent MiM attacks and reset attacks. Reset

of [BR93, Kat02] naturally extended to multiple concurrent sessions. Following [OPV08, COSV12], the key idea consists in using an identity that has two possible secrets such that knowledge of one witness does not allow to compute the other one in polynomial time. Then, by using our implementation of a concurrent NMWIAoK for proving knowledge of a secret associated to such identity we obtain a 3-round identification scheme secure against concurrent MiM attacks.

Challenges for future work. The existence of OWPs is a standard falsifiable hardness assumption. Our scheme relies on a strengthening of this standard assumption w.r.t. subexponential-time adversaries. Notice that the lower bound of [Pas13] still applies in case of subexponential-time hardness, therefore our 3-round concurrent non-malleable scheme is round optimal. Various natural and fascinating questions on commitments and proofs of knowledge remain open after our work and as such we think our results will motivate further research. Examples of open questions about concurrent NM commitments are the following: 1) the existence of 3-round schemes based on standard falsifiable hardness assumptions w.r.t. polynomial-time adversaries only; 2) the existence of 3-round schemes with black-box use of primitives; 3) the existence of practical schemes.

2 Notation, Definitions and Tools

We denote the security parameter by λ and use "$|$" as concatenation operator (i.e., if a and b are two strings then by $a|b$ we denote the concatenation of a and b). For a finite set Q, $x \leftarrow Q$ denotes the algorithm that chooses x from Q with uniform distribution. Usually we use the abbreviation PPT that stays for probabilistic polynomial-time. We use $\mathsf{poly}(\cdot)$ to indicate a generic polynomial function of the input.

A *polynomial-time relation* Rel (or *polynomial relation*, in short) is a subset of $\{0,1\}^* \times \{0,1\}^*$ such that membership of (x, w) in Rel can be decided in time polynomial in $|x|$. For $(x, w) \in \mathsf{Rel}$, we call x the *instance* and w a *witness* for x. For a polynomial-time relation Rel, we define the NP-language L_{Rel} as $L_{\mathsf{Rel}} = \{x | \exists w : (x, w) \in \mathsf{Rel}\}$. Analogously, unless otherwise specified, for an NP-language L we denote by $\mathsf{Rel_L}$ the corresponding polynomial-time relation (that is, $\mathsf{Rel_L}$ is such that $L = L_{\mathsf{Rel_L}}$).

Let A and B be two interactive probabilistic algorithms A and B. We denote by $\langle A(\alpha), B(\beta)\rangle(\gamma)$ the distribution of B's output after running on private input β with A using private input α, both running on common input γ. Typically, one of the two algorithms receives 1^λ as input. A *transcript* of $\langle A(\alpha), B(\beta)\rangle(\gamma)$ consists of the messages exchanged during an execution where A receives a private input α, B receives a private input β and both A and B receive a common input γ. Moreover, we will refer to the *view* of A as the messages it received during the execution of $\langle A(\alpha), B(\beta)\rangle(\gamma)$, along with its randomness and its input. We denote by A_r an algorithm A that receives as randomness r. We say that a protocol (A, B) is public coin if B sends to A random bits only.

A function $\nu(\cdot)$ from non-negative integers to reals is called negligible, if for every constant $c > 0$ and all sufficiently large $\lambda \in \mathbb{N}$ we have $\nu(\lambda) < \lambda^{-c}$.

Definition 1 (One-way function (OWF)). *A function $f : \{0,1\}^\star \rightarrow \{0,1\}^\star$ is called one way if the following two conditions hold:*

attack were also considered in the notion CR1+ introduced in [BPSV08]. Since reset attacks are out of the scope of this work, we will focus on concurrent MiM attacks only.

- *there exist a deterministic polynomial-time algorithm that on input y in the domain of f outputs $f(y)$;*
- *for every* PPT *algorithm $\mathcal{A}$ there exists a negligible function ν, such that for every auxiliary input $z \in \{0,1\}^{\mathsf{poly}(\lambda)}$:*

$$\text{Prob}\left[\, y \leftarrow \{0,1\}^\star : \mathcal{A}(f(y), z) \in f^{-1}(f(y)) \,\right] < \nu(\lambda).$$

We say that a OWF f is a one-way permutation (OWP) *if f is a permutation.*

We will require that an algorithm that runs in time $\tilde{T} = 2^{\lambda^\alpha}$ for some positive constant $\alpha < 1$, can invert a OWP f. In this case we say that f is $\tilde{T}$-breakable.

Definition 2 (Computational indistinguishability)**.** *Let $X = \{X_\lambda\}_{\lambda \in \mathbb{N}}$ and $Y = \{Y_\lambda\}_{\lambda \in \mathbb{N}}$ be ensembles, where X_λ's and Y_λ's are probability distribution over $\{0,1\}^l$, for same $l = \mathsf{poly}(\lambda)$. We say that $X = \{X_\lambda\}_{\lambda \in \mathbb{N}}$ and $Y = \{Y_\lambda\}_{\lambda \in \mathbb{N}}$ are* computationally indistinguishable, *denoted $X \approx Y$, if for every* PPT *distinguisher $\mathcal{D}$ there exist a negligible function ν such that for sufficiently large $\lambda \in \mathbb{N}$,*

$$\left|\text{Prob}\left[\, t \leftarrow X_\lambda : \mathcal{D}(1^\lambda, t) = 1 \,\right] - \text{Prob}\left[\, t \leftarrow Y_\lambda : \mathcal{D}(1^\lambda, t) = 1 \,\right]\right| < \nu(\lambda).$$

We note that in the usual case where $|X_\lambda| = \Omega(\lambda)$ and λ can be derived from a sample of X_λ, it is possible to omit the auxiliary input 1^λ. In this paper we also use the definition of *Statistical Indistinguishability*. This definition is the same as Definition 2 with the only difference that the distinguisher $\mathcal{D}$ is unbounded. In this case use $X \equiv_s Y$ to denote that two ensembles are statistically indistinguishable.

We note that in the usual case where $|X_\lambda| = \Omega(\lambda)$ and the length λ can be derived from a sample of X_λ, it is possible to omit the auxiliary input 1^λ.

Definition 3 (Delayed-input proof/argument system)**.** *A pair of* PPT *interactive algorithms $\Pi = (\mathcal{P}, \mathcal{V})$ constitutes a* proof system *(resp., an* argument system*) for an NP-language L, if the following conditions hold:*

Completeness: *For every $x \in L$ and w such that $(x, w) \in \mathsf{Rel_L}$, it holds that:*

$$\text{Prob}\left[\, \langle \mathcal{P}(w), \mathcal{V} \rangle(x) = 1 \,\right] = 1.$$

Soundness: *For every interactive (resp.,* PPT *interactive) algorithm $\mathcal{P}^\star$, there exists a negligible function ν such that for every $x \notin L$ and every z:*

$$\text{Prob}\left[\, \langle \mathcal{P}^\star(z), \mathcal{V} \rangle(x) = 1 \,\right] < \nu(|x|).$$

A proof/argument system $\Pi = (\mathcal{P}, \mathcal{V})$ for an NP-language L, enjoys delayed-input *completeness if $\mathcal{P}$ needs x and w only to compute the last round and $\mathcal{V}$ needs x only to compute the output. Before that, $\mathcal{P}$ and $\mathcal{V}$ run having as input only the size of x.*

The notion of delayed-input completeness was defined in [CPS+16b]. We say that the transcript τ of an execution of $(\mathcal{P}, \mathcal{V})$ is *accepting* if $\mathcal{V}$ outputs 1. An interactive protocol $\Pi = (\mathcal{P}, \mathcal{V})$ is *public coin* if, at every round, $\mathcal{V}$ at each round simply tosses a predetermined number of coins (random challenge) and sends them to $\mathcal{P}$.

Definition 4 (Witness Indistinguishable (WI)). *An argument/proof system* $\Pi = (\mathcal{P}, \mathcal{V})$*, is* Witness Indistinguishable (WI) *for a relation* Rel *if, for every malicious* PPT *verifier* $\mathcal{V}^\star$*, there exists a negligible function* ν *such that for all* x, w, w' *such that* $(x, w) \in \mathsf{Rel}$ *and* $(x, w') \in \mathsf{Rel}$ *it holds that:*

$$\left|\text{Prob}\left[\ \langle \mathcal{P}(w), \mathcal{V}^\star\rangle(x) = 1\ \right] - \text{Prob}\left[\ \langle \mathcal{P}(w'), \mathcal{V}^\star\rangle(x) = 1\ \right]\right| < \nu(|x|).$$

The notion of a *perfect* WI proof system is obtained by requiring $\nu(|x|) = 0$.

Obviously one can generalize the above definitions of WI to their natural adaptive-input variants, where the adversarial verifier can select the statement and the witnesses adaptively, before the prover plays the last round.

Definition 5 (Proof of Knowledge [LP11]). *A proof system* $\Pi = (\mathcal{P}, \mathcal{V})$ *is a* proof of knowledge *(PoK) for the relation* Rel_L *if there exist a probabilistic expected polynomial-time machine* E*, called the extractor, such that for every algorithm* $\mathcal{P}^\star$*, there exists a negligible function* $\nu(\lambda)$*, every statement* $x \in \{0,1\}^\lambda$*, every randomness* $r \in \{0,1\}^\star$ *and every auxiliary input* $z \in \{0,1\}^\star$*,*

$$\text{Prob}\left[\ \langle \mathcal{P}^\star_r(z), \mathcal{V}\rangle(x) = 1\ \right] \leq \text{Prob}\left[\ w \leftarrow \mathsf{E}^{\mathcal{P}^\star_r(z)}(x) : (x, w) \in \mathsf{Rel}_\mathsf{L}\ \right] + \nu(\lambda).$$

We also say that an argument system Π *is a* argument of knowledge *(AoK) if the above condition holds w.r.t. any* PPT $\mathcal{P}^\star$*.*

In this paper we also consider the *adaptive-PoK* property. The adaptive-PoK property ensures that the PoK property still holds when a malicious prover can choose the statement adaptively at the last round. In this case, to be consistent with Definition 5 where the extractor algorithm E takes as input the statement proved by $\mathcal{P}^\star$, we have to consider a different extractor algorithm. This extractor algorithm takes as input the randomness r of $\mathcal{P}$, the randomness r' of $\mathcal{V}$ and outputs the witness for $x \in L$, where x is selected by $\mathcal{P}^\star_r$ when interacting with $\mathcal{V}_{r'}$.

In this paper we use the 3-round public-coin WI Proof of Knowledge (WIPoK) proposed by Lapidot and Shamir [LS90], that we denote by LS. LS enjoys delayed-input completeness since the inputs for both $\mathcal{P}$ and $\mathcal{V}$ are needed only to play the last round, and only the length of the instance is needed earlier. The LS protocol is also sound when a malicious prover can choose the statement adaptively at the third round. We refer to this property as adaptive soundness. LS also enjoys the property of adaptive PoK and adaptive WI.

2.1 Commitment Schemes

Definition 6 (Commitment Scheme). *Given a security parameter* 1^λ*, a commitment scheme* $(\mathsf{Sen}, \mathsf{Rec})$ *is a two-phase protocol between two* PPT *interactive algorithms, a sender* Sen *and a receiver* Rec*. In the commitment phase* Sen *on input a message* m *interacts with* Rec *to produce a commitment* $\mathtt{com}$*. In the decommitment phase,* Sen *sends to* Rec *a decommitment information* $\mathtt{d}$ *such that* Rec *accepts* m *as the commitment of* $\mathtt{com}$*.*

Formally, we say that $\mathsf{CS} = (\mathsf{Sen}, \mathsf{Rec})$ *is a perfectly binding commitment scheme if the following properties hold:*

Correctness:

- *Commitment phase. Let* $\mathtt{com}$ *be the commitment of the message* m *(i.e.,* $\mathtt{com}$ *is the transcript of an execution of* $\mathsf{CS} = (\mathsf{Sen}, \mathsf{Rec})$ *where* Sen *runs on input a message* m*). Let* $\mathtt{d}$ *be the private output of* Sen *in this phase.*

- *Decommitment phase*[8]*.* Rec *on input* m *and* $\mathtt{d}$ *accepts* m *as decommitment of* $\mathtt{com}$*.*

Hiding([Lin10]): *for a* PPT *adversary* $\mathcal{A}$ *and a randomly chosen bit* $b \in \{0,1\}$*, consider the following hiding experiment* $\mathsf{ExpHiding}^b_{\mathcal{A},\mathsf{CS}}(\lambda)$*:*

- *Upon input* 1^λ*, the adversary* $\mathcal{A}$ *outputs a pair of messages* m_0, m_1 *that are of the same length.*
- Sen *on input the message* m_b *interacts with* $\mathcal{A}$ *to produce a commitment of* m_b*.*
- $\mathcal{A}$ *outputs a bit* b' *and this is the output of the experiment.*

For any PPT *adversary* $\mathcal{A}$*, there exist a negligible function* ν*, such that:*

$$\left|\text{Prob}\left[\ \mathsf{ExpHiding}^0_{\mathcal{A},\mathsf{CS}}(\lambda) = 1\ \right] - \text{Prob}\left[\ \mathsf{ExpHiding}^1_{\mathcal{A},\mathsf{CS}}(\lambda) = 1\ \right]\right| < \nu(\lambda).$$

Binding: *for every commitment* $\mathtt{com}$ *generated during the commitment phase by a possibly malicious unbounded sender* $\mathsf{Sen}^\star$ *interacting with an honest receiver* Rec*, there exists at most one message* m *that* Rec *accepts as decommitment of* $\mathtt{com}$*.*

We also consider the definition of a commitment scheme where the hiding property still holds against an adversary $\mathcal{A}$ running in time bounded by $T = 2^{\lambda^\alpha}$ for some positive constant $\alpha < 1$. In this case we will say that a commitment scheme is T-hiding. We will also say that a commitment scheme is $\tilde{T}$-breakable to specify that an algorithm running in time $\tilde{T} = 2^{\lambda^\beta}$, for some positive constant $\beta < 1$, recovers the (if any) only message that can be successfully decommitment.

In the rest of the paper we also use a non-interactive commitment schemes, with secure parameter λ. In this case we consider a commitment scheme as a pair of PPT algorithms $(\mathsf{NISen}, \mathsf{NIRec})$ where:

- NISen takes as input $(m;\sigma)$, where $m \in \{0,1\}^{\mathsf{poly}(\lambda)}$ is the message to be committed and $\sigma \leftarrow \{0,1\}^\lambda$ is randomness, and outputs the commitment $\mathtt{com}$ and the decommitment $\mathtt{dec}$;
- NIRec takes as input $(\mathtt{dec}, \mathtt{com}, m)$ and outputs 1 if it accepts m as a decommitment of $\mathtt{com}$ and 0 otherwise.

3-Round extractable commitment schemes. Informally, a 3-round commitment scheme is extractable if there exists an efficient extractor that having black-box access to any efficient malicious sender ExCom* that successfully performs the commitment phase, outputs the only committed string that can be successfully decommitted.

Definition 7 (3-Round Extractable Commitment Scheme [PW09, GLOV12])**.** *A 3-round perfectly binding commitment scheme* $\mathtt{ExCS}$ = (ExCom, ExRec) *is an* extractable commitment scheme *if given oracle access to any malicious sender* ExCom**, there exists an expected* PPT *extractor* $\mathtt{Ext}$ *that outputs a pair* $(\tau, \sigma^\star)$ *such that the following properties hold:*

- **Simulatability:** *the simulated view* τ *is identically distributed to the view of* ExCom* *(when interacting with an honest* ExRec*) in the commitment phase.*
- **Extractability:** *there exists no decommitment of* τ *to* σ*, where* $\sigma \neq \sigma^\star$*.*

[8] In this paper we consider a non-interactive decommitment phase only.

2.2 Non-Malleable Commitment Schemes

Here we follow [LPV08][9]. Let $\mathsf{\Pi} = (\mathsf{Sen}, \mathsf{Rec})$ be a statistically binding commitment scheme. Consider MiM adversaries that are participating in left and right sessions in which $\mathsf{poly}(\lambda)$ commitments take place. We compare between a MiM and a simulated execution. In the MiM execution the adversary $\mathcal{A}$, with auxiliary information z, is simultaneously participating in $\mathsf{poly}(\lambda)$ left and right sessions. In the left sessions the MiM adversary $\mathcal{A}$ interacts with Sen receiving commitments to values $m_1, \ldots, m_{\mathsf{poly}(\lambda)}$ using identities $\mathtt{id}_1, \ldots, \mathtt{id}_{\mathsf{poly}(\lambda)}$ of its choice. In the right session $\mathcal{A}$ interacts with Rec attempting to commit to a sequence of related values $\tilde{m}_1, \ldots, \tilde{m}_{\mathsf{poly}(\lambda)}$ again using identities of its choice $\tilde{\mathtt{id}}_1, \ldots, \tilde{\mathtt{id}}_{\mathsf{poly}}(\lambda)$. If any of the right commitments is invalid, or undefined, its value is set to $\perp$. For any i such that $\tilde{\mathtt{id}}_i = \mathtt{id}_j$ for some j, set $\tilde{m}_i = \perp$ (i.e., any commitment where the adversary uses the same identity of one of the honest senders is considered invalid). Let $\mathsf{mim}_{\mathsf{\Pi}}^{\mathcal{A}, m_1, \ldots, m_{\mathsf{poly}(\lambda)}}(z)$ denote a random variable that describes the values $\tilde{m}_1, \ldots, \tilde{m}_{\mathsf{poly}(\lambda)}$ and the view of $\mathcal{A}$, in the above experiment. In the simulated execution, an efficient simulator S directly interacts with Rec. Let $\mathsf{sim}_{\mathsf{\Pi}}^{S}(1^\lambda, z)$ denote the random variable describing the values $\tilde{m}_1, \ldots, \tilde{m}_{\mathsf{poly}(\lambda)}$ committed by S, and the output view of S; whenever the view contains in the i-th right session the same identity of any of the identities of the left session, then $\tilde{m}_i$ is set to $\perp$.

In all the paper we denote by $\tilde{\delta}$ a value associated with the right session (where the adversary $\mathcal{A}$ plays with a receiver MMRec) where δ is the corresponding value in the left session. For example, the sender commits to v in the left session while $\mathcal{A}$ commits to $\tilde{v}$ in the right session.

Definition 8 (Concurrent NM commitment scheme [LPV08])**.** *A commitment scheme is* concurrent NM with respect to commitment *(or a many-many NM commitment scheme) if, for every* PPT *concurrent MiM adversary $\mathcal{A}$, there exists a* PPT *simulator S such that for all $m_i \in \{0,1\}^{\mathsf{poly}(\lambda)}$ for $i = \{1, \ldots, \mathsf{poly}(\lambda)\}$ the following ensembles are computationally indistinguishable:*

$$\{\mathsf{mim}_{\mathsf{\Pi}}^{\mathcal{A}, m_1, \ldots, m_{\mathsf{poly}(\lambda)}}(z)\}_{z \in \{0,1\}^\star} \approx \{\mathsf{sim}_{\mathsf{\Pi}}^{S}(1^\lambda, z)\}_{z \in \{0,1\}^\star}.$$

As in [LPV08] we also consider relaxed notions of concurrent non-malleability: one-many and one-one NM commitment schemes. In a one-many NM commitment scheme, $\mathcal{A}$ participates in one left and polynomially many right sessions. In a one-one (i.e., a stand-alone secure) NM commitment scheme, we consider only adversaries $\mathcal{A}$ that participate in one left and one right session. We will make use of the following proposition of [LPV08].

Proposition 1. *Let* $(\mathsf{Sen}, \mathsf{Rec})$ *be a one-many NM commitment scheme. Then,* $(\mathsf{Sen}, \mathsf{Rec})$ *is also a concurrent (i.e., many-many) NM commitment scheme.*

We also consider the definition of a NM commitment scheme secure against a MIM $\mathcal{A}$ running in time bounded by $T = 2^{\lambda^\alpha}$ for some positive constant $\alpha < 1$. In this case we will say that a commitment scheme is T-non-malleable.

When the identity is selected by the sender then the above id-based definitions guarantee non-malleability without ids as long as the MiM does not behave like a proxy (an unavoidable attack). Indeed the sender can pick as $\mathtt{id}$ the public key of a strong signature scheme signing the transcript. The MiM will have to use a different $\mathtt{id}$ or to break the signature scheme.

[9] In this paper we will consider only NM commitments w.r.t. commitments. Difficulties on achieving also the notion of NM w.r.t. decommitments were discussed in [OPV09, CVZ10].

2.3 3-Round One-One NM Commitment Scheme

As main tool we need a 3-round one-one NM commitment scheme (NMCS) that enjoys the extractability property. In the rest of the paper we will refer to such a commitment scheme as $\Pi_{\mathsf{NM}} = (\mathsf{Sen}_{\mathsf{NM}}, \mathsf{Rec}_{\mathsf{NM}})$.

In [GPR16] the authors provide the first 3-round one-one NM commitment scheme. Their scheme enjoys also the extractability property[10] and public coin. We will call a 3-round one-one NM commitment scheme as $\Pi_{\mathsf{NM}} = ((\mathrm{Sen}^1_{\mathsf{NM}}, \mathrm{Sen}^2_{\mathsf{NM}}), \mathsf{Rec}_{\mathsf{NM}})$ where:

- the algorithm $\mathrm{Sen}^1_{\mathsf{NM}}$ takes as input $(\mathtt{id}, m; \rho)$, where $\mathtt{id} \in \{0,1\}^\lambda$ is the identity, m is the message to be committed and $\rho \leftarrow \{0,1\}^\lambda$ is a randomness, and outputs a that is the first round of the commitment scheme to be sent to the receiver;
- the algorithm $\mathrm{Sen}^2_{\mathsf{NM}}$ takes as input $(\mathtt{id}, \mathsf{c}, m; \rho)$, where c is the second round received by Rec, m is the message to be committed, $\mathtt{id}$ is the same identity received as input by $\mathrm{Sen}^1_{\mathsf{NM}}$, ρ is the randomness, and outputs $(\mathsf{z}, \mathtt{dec})$ where z is the last round of the commitment, and $\mathtt{dec}$ is the decommitment value.

The reveal phase consists in sending $\mathtt{dec}$ and m to the receiver. The receiver $\mathsf{Rec}_{\mathsf{NM}}$, on input the randomness it used during the commitment phase, the transcript $\mathtt{com} = (\mathsf{a}, \mathsf{c}, \mathsf{z}, \mathtt{id})$, m and $\mathtt{dec}$ outputs 1 if $\mathtt{dec}$ is valid w.r.t. $\mathtt{com}$ and m and outputs 0 otherwise.

2.4 The LS Proof of Knowledge and NMWI Argument Systems

In this paper we use the 3-round public-coin WI adaptive proof of knowledge (see Sec. 5 for a formal definition) proposed by Lapidot and Shamir [LS90], that we denote by LS. LS is delayed-input since the inputs for the prover and the verifier are needed only to play the last round, while only the size of the common input is needed earlier. For this reason we will refer to a prover $\mathcal{P}$ as a pair $(\mathsf{P}^1, \mathsf{P}^2)$. More formally, LS for a relation Rel is a pair $\Pi = (\mathcal{P} = (\mathsf{P}^1, \mathsf{P}^2), \mathcal{V})$, with security parameter λ, where $\mathcal{P}$ executes the algorithms P^1 and P^2 defined as follows. The algorithm P^1, takes as input $(\ell; \alpha)$, ℓ is the instance length and $\alpha \leftarrow \{0,1\}^\lambda$ is the randomness, and outputs the 1st round of the LS protocol. The algorithm P^2 takes as input $(x, w, c; \alpha)$, where x, w are such that $(x, w) \in \mathsf{Rel}$, c is the challenge sent by $\mathcal{V}$ and α is the randomness[11] and outputs the 3rd round of the LS protocol.

In this paper we also consider a definition where the WI property of LS still holds against a distinguisher with running time bounded by $T = 2^{\lambda^\alpha}$ for some constant positive constant $\alpha < 1$. In this case we say that the instantiation of LS is T-witness indistinguishable (T-WI).

Witness indistinguishability and MiM attacks. The definition of non-malleable witness indistinguishability (NMWI) given in [OPV08] requires that the witness *encoded in the proof* given by the MiM $\mathcal{A}$ be independent of the witness used by the honest prover in his proof. The concept of witness encoded in a proof becomes clear when considering *commit-and-prove* argument systems. In such arguments, the transcript includes a commitment that encodes in an unambiguous way the witness used by the prover. In a NMWI commit-and-prove argument system, the witness encoded in the proof produced by the $\mathcal{A}$ must be independent of the witness used (and thus encoded by the honest prover) in the proof in which $\mathcal{A}$ acts as a verifier. Similarly to the case of non-malleable

[10] Extractability is informally stated in Claim 12 of [GPR15].

[11] The same α is passed to P^1 and P^2 so that P^2 can reconstruct the state of P^1.

commitments, one can give a definition with or without sessions ids. Here we use the one without sessions ids since it is more useful in the applications.

Let $\mathcal{A}$ be a MiM interacting in the left proof with $\mathcal{P}$ that is running on input x and witness w. In the right proof $\mathcal{A}$ interacts with $\mathcal{V}$ on common input $\tilde{x}$ chosen by $\mathcal{A}$. We denote by z the auxiliary information available to $\mathcal{A}$. NMWI is defined in terms of the random variable $\mathsf{wmim}^{\mathcal{A}}(x, w, z)$ that is the view of $\mathcal{A}$ that we denote by $\mathsf{View}^{\mathcal{P}}_{\mathcal{A}}(x, w, z)$ (i.e., the view of $\mathcal{A}$ when running with z as auxiliary input and playing with $\mathcal{P}$ that runs on input (x, w)) and the witness encoded in the right proof given by $\mathcal{A}$. If the right proof is not accepting or the transcript is identical to the one of the left proof then the witness encoded is $\perp$; otherwise the string w committed by $\mathcal{A}$ in the right proof is returned. In other words, $\mathsf{wmim}^{\mathcal{A}}(x, w, z)$ consists of the view of $\mathcal{A}$ and the witness encoded in the right proof unless the proof is not accepting.

Definition 9 (NMWI argument system [OPV08])**.** *A commit-and-prove argument system* $\Pi = (\mathcal{P}, \mathcal{V})$ *for an NP-language L and corresponding relation* $\mathsf{Rel_L}$ *is a* non-malleable witness indistinguishable *argument if, for all* PPT *man-in-the-middle adversaries $\mathcal{A}$, for all $x \in L$ and all w, w' such that* $\mathsf{Rel_L}(x, w)$, $\mathsf{Rel_L}(x, w')$ *for all auxiliary information z it holds that* $\{\mathsf{wmim}^{\mathcal{A}}(x, w, z)\} \approx \{\mathsf{wmim}^{\mathcal{A}}(x, w', z)\}$.

The above notion extends in a straight-forward way to the case of a concurrent MiM adversary trivially. Formally, the concurrent MiM $\mathcal{A}$ opens $\mathsf{poly}(\lambda)$ left and right proofs each with a common input of length $\mathsf{poly}(\lambda)$. $\mathcal{A}$ interacts in the i-th left proof with an instance of the honest prover $\mathcal{P}$ on common input "$x_i \in L$" and private input w_i such that $\mathsf{Rel_L}(x_i, w_i)$. In the j-th right proof $\mathcal{A}$ is interacting with the honest verifier V on common input $\tilde{x}_j$ of its choice.

Let X, W be respectively the sequence of instances and witnesses in input to $\mathcal{P}$ in the left proofs. Now the distribution $\mathsf{wmim}^{\mathcal{A}}(X, W, z)$ consists of the view $\mathsf{View}^{\mathcal{P}}_{\mathcal{A}}(X, W, z)$ of $\mathcal{A}$ along with a sequence $(\tilde{w}_1, \ldots, \tilde{w}_{\mathsf{poly}(\lambda)})$ and it holds that: if the j-th right proof is not accepting or the transcript is identical to the one of a left proof then $\tilde{w}_j = \perp$; otherwise, $\tilde{w}_j$ is the witness encoded in the j-th right proof.

Definition 10 (cNMWI argument [OPV08])**.** *A commit-and-prove argument system* $\Pi = (\mathcal{P}, \mathcal{V})$ *for an NP-language L and corresponding relation* $\mathsf{Rel_L}$ *is* concurrent non-malleable witness indistinguishable *if, for all* PPT *concurrent man-in-the-middle adversaries $\mathcal{A}$, for all sequences X of* $\mathsf{poly}(\lambda)$ *elements of L of length* $\mathsf{poly}(\lambda)$*, for all sequences W and W' of witnesses for X, and for all auxiliary information z it holds that* $\{\mathsf{wmim}^{\mathcal{A}}(X, W, z)\} \approx \{\mathsf{wmim}^{\mathcal{A}}(X, W', z)\}$.

3 3-Round Concurrent Non-Malleable Commitments

In this section we show the main result of this work, a transform that starting from a 3-round extractable one-one non-malleable commitment scheme outputs a 3-round concurrent non-malleable commitment scheme.

3.1 Informal Description

Our transform takes as input a 3-round extractable one-one NM commitment scheme Π_{NM}, a OWP f, a non-interactive perfectly binding commitment scheme NI, the 3-round delayed-input adaptive WI/PoK LS and outputs a 3-round fully concurrent (i.e., many-many) NM commitment scheme $\Pi_{\mathsf{MMCom}} = (\mathsf{MMSen}, \mathsf{MMRec})$.

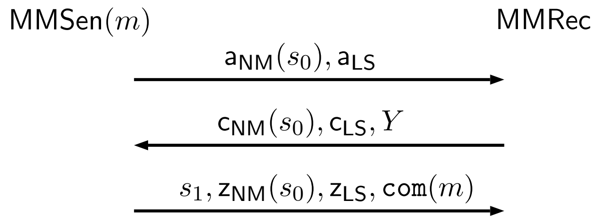
MMSen(m)
MMRec
$\mathsf{a}_{\mathsf{NM}}(s_0), \mathsf{a}_{\mathsf{LS}}$
$\mathsf{c}_{\mathsf{NM}}(s_0), \mathsf{c}_{\mathsf{LS}}, Y$
$s_1, \mathsf{z}_{\mathsf{NM}}(s_0), \mathsf{z}_{\mathsf{LS}}, \mathtt{com}(m)$

- Y is an element taken from the range of the OWP f.
- $\mathtt{com}(m)$ is the perfectly binding commitment of m computed using NI.
- $(\mathsf{a}_{\mathsf{NM}}(s_0), \mathsf{c}_{\mathsf{NM}}(s_0), \mathsf{z}_{\mathsf{NM}}(s_0)) = \tau$ is the transcript of the execution of the NM commitment scheme Π_{NM} when the sender commits to s_0.
- $(\mathsf{a}_{\mathsf{LS}}, \mathsf{c}_{\mathsf{LS}}, \mathsf{z}_{\mathsf{LS}}) = \pi$ is the transcript of LS proving knowledge of either m and the randomness used to compute $\mathtt{com}$, or of $(s_0, \mathtt{dec})$, s.t. $f(s_0 \oplus s_1) = Y$ and $\mathtt{dec}$ is a valid decommitment of s_0 w.r.t. τ.

Figure 1: Informal description of our 3-round concurrent NM commitment scheme.

Let m be the message that MMSen wants to commit. The high-level idea of our compiler is depicted in Fig. 1. The sender MMSen, on input the session-id $\mathtt{id}$ and the message m, computes the 1st round of the protocol by running LS and sending the 1st round of NM to commit to a random message s_0 using $\mathtt{id}$ as session-id. In the 2nd round the receiver MMRec sends the challenges of NM and LS, also sends a random value Y in the range of the OWP f[12]. In the last round MMSen commits to message m using NI, therefore obtaining $\mathtt{com}$, then computes the last round of NM, completes the transcript of LS, and finally sends a random string s_1. The protocol LS is used by MMSen to prove to MMRec that either she knows message m and the randomness used to compute $\mathtt{com}$, or she knows the values $(s_0, \mathtt{dec})$, such that $f(s_0 \oplus s_1) = Y$ and $\mathtt{dec}$ is a valid decommitment to s_0 w.r.t. the commitment computed using Π_{NM}. We observe that MMSen needs m only when computing the 3rd round, therefore our construction enjoys delayed-input correctness.

3.2 Our Compiler

Our compiler needs the following tools:

1. a OWP f that is secure against PPT adversaries and $\tilde{T}_f$-breakable;
2. a non interactive perfectly binding commitment scheme $\mathsf{NI} = (\mathsf{NISen}, \mathsf{NIRec})$ that is T_{NI}-hiding and $\tilde{T}_{\mathsf{NI}}$-breakable;
3. a 3-round extractable *one-one* NM commitment scheme $\Pi_{\mathsf{NM}} = (\mathsf{Sen}_{\mathsf{NM}} = (\mathrm{Sen}^1_{\mathsf{NM}}, \mathrm{Sen}^2_{\mathsf{NM}}), \mathsf{Rec}_{\mathsf{NM}})$ that is T_{NM}-hiding/non-malleable, and $\tilde{T}_{\mathsf{NM}}$-breakable;

[12] When sampling from the range of f corresponds to picking a random string, we have that our commitment scheme is public coin.

4. the LS proof system $\mathsf{LS} = (\mathcal{P} = (\mathsf{P}^1, \mathsf{P}^2), \mathcal{V})$ for the language

$$L = \big\{\big((a, c, z), Y, s_1, \mathtt{com}, \mathtt{id}\big) : \exists\ (m, \sigma) \text{ s.t. } \mathtt{com} = \mathsf{NISen}(m; \sigma) \texttt{ OR}\big(\exists(\rho, s_0)$$
$$\text{s.t. } a = \text{Sen}^1_{\mathsf{NM}}(\mathtt{id}, s_0; \rho) \texttt{ AND } z = \text{Sen}^2_{\mathsf{NM}}(\mathtt{id}, c, s_0; \rho) \texttt{ AND } Y = f(s_0 \oplus s_1)\big)\big\}$$

that is T_{LS}-WI for the corresponding relation $\mathsf{Rel}_{\mathsf{L}}$.

Let λ be the security parameter of our scheme. We will use wlog λ also as security parameter for the hardness to invert f with respect to polynomial time adversaries. Then we consider the following hierarchy of security levels for the above tools: $T_f << T_{\mathsf{NI}} << \sqrt{T_{\mathsf{NM}}} << T_{\mathsf{NM}} << \sqrt{T_{\mathsf{LS}}} << T_{\mathsf{LS}}$ where by "$T << T'$" we mean that "$T \cdot \mathsf{poly}(\lambda) < T'$". We also require that:

- NI is T_{NI}-hiding, but is also $\tilde{T}_{\mathsf{NI}} = \sqrt{T_{\mathsf{NM}}}$-breakable;
- Π_{NM} is T_{NM} hiding/non-malleable, but the hiding is also $\tilde{T}_{\mathsf{NM}} = \sqrt{T_{\mathsf{LS}}}$-breakable.

Now we need to define different security parameters, one for each tool involved in the security proof to be consistent with the hierarchy of security levels defined above (a similar use of security parameters has been proposed in [PW10]). Given the security parameter λ of our scheme, we will make use of the following security parameters (all polynomially related to λ and such that the above hierarchy of security levels holds): λ for f, λ_{NI} for NI, λ_{NM} for Π_{NM}, λ_{LS} for LS.

We denote by Params the function that on input λ outputs $(\lambda_{\mathsf{NI}}, \lambda_{\mathsf{NM}}, \lambda_{\mathsf{LS}}, \ell)$ where ℓ is the size of the theorem to be proved using LS[13]. Our concurrent NM commitment scheme $\Pi_{\mathsf{MMCom}} = (\mathsf{MMSen}, \mathsf{MMRec})$ is fully described in Fig. 2.

Theorem 1. *Suppose there exist OWPs secure against subexponential-time adversaries, then* Π_{MMCom} *is a perfectly binding delayed-input commitment scheme.*

Proof. **Correctness.** The delayed-input correctness of Π_{MMCom} follows by inspection considering the delayed-input completeness of LS, and the correctness of Π_{NM} and NI.

Binding. To prove the binding property we only observe that the message given in output in the decommitment phase of Π_{MMCom} is the message committed using NI. Moreover the decommitment phase of Π_{MMCom} coincides with the decommitment of NI. Since NI is perfectly binding we have that therefore Π_{MMCom} is perfectly binding too.

Hiding. The hiding property follows directly from the non-malleability property proved in Theorem 2. Indeed the proof of Theorem 2 does not rely on the hiding of Π_{MMCom}. □

3.3 Proof of Non-Malleability

In this section we prove our main theorem.

Theorem 2. *Suppose there exist OWPs secure against subexponential-time adversaries, then* Π_{MMCom} *is concurrent (i.e., many-many) non-malleable.*

Proof. Since we can use Proposition 1, we only need to prove that our commitment enjoys one-many non-malleability. More formally with respect to a one-many adversary $\mathcal{A}$, we need to show that for all $m \in \{0,1\}^{\mathsf{poly}(\lambda)}$ it holds that:

$$\{\mathsf{mim}^{\mathcal{A},m}_{\Pi_{\mathsf{MMCom}}}(z)\}_{z\in\{0,1\}^\star} \approx \{\mathsf{sim}^{S}_{\Pi_{\mathsf{MMCom}}}(1^\lambda, z)\}_{z\in\{0,1\}^\star}$$

[13] To compute 1st and 2nd round of LS only the length ℓ of the instance is required.

Common input: Security parameters: λ, $(\lambda_{\mathsf{NI}}, \lambda_{\mathsf{NM}}, \lambda_{\mathsf{LS}}, \ell) = \mathsf{Params}(\lambda)$.
MMSen's identity: $\mathtt{id} \in \{0,1\}^{\lambda}$.
Input to MMSen**:** $m \in \{0,1\}^{\mathsf{poly}\{\lambda\}}$.

Commitment Phase:

1. $\mathsf{MMSen} \rightarrow \mathsf{MMRec}$
   - 1.1. Pick $s_0 \in \{0,1\}^{\lambda}$.
   - 1.2. Pick a randomness $\rho \in \{0,1\}^{\lambda_{\mathsf{NM}}}$ and compute $\mathsf{a}_{\mathsf{NM}} = \mathrm{Sen}^1_{\mathsf{NM}}(\mathtt{id}, s_0; \rho)$.
   - 1.3. Pick a randomness $\alpha \in \{0,1\}^{\lambda_{\mathsf{LS}}}$ and compute $\mathsf{a}_{\mathsf{LS}} = \mathsf{P}^1(\ell; \alpha)$.
   - 1.4. Send $(\mathsf{a}_{\mathsf{NM}}, \mathsf{a}_{\mathsf{LS}})$ to MMRec.
2. $\mathsf{MMRec} \rightarrow \mathsf{MMSen}$
   - 2.1. Pick a randomness γ and run $\mathsf{Rec}_{\mathsf{NM}}$ on input $(\mathtt{id}, \mathsf{a}_{\mathsf{NM}}; \gamma)$ to obtain c_{NM}.
   - 2.2. Pick a randomness β and run $\mathcal{V}$ to obtain c_{LS}.
   - 2.3. Pick a random $y \in \{0,1\}^{\lambda}$ and compute $Y = f(y)$.
   - 2.4. Send $(\mathsf{c}_{\mathsf{NM}}, \mathsf{c}_{\mathsf{LS}}, Y)$ to MMSen.
3. $\mathsf{MMSen} \rightarrow \mathsf{MMRec}$
   - 3.1. Pick a randomness $\sigma \in \{0,1\}^{\lambda_{\mathsf{NI}}}$ and compute $(\mathtt{com}, \mathtt{dec}) = \mathsf{NISen}(m; \sigma)$.
   - 3.2. Pick $s_1 \leftarrow \{0,1\}^{\lambda}$.
   - 3.3. Compute $(\mathsf{z}_{\mathsf{NM}}, \mathtt{dec}_{\mathsf{NM}}) = \mathrm{Sen}^2_{\mathsf{NM}}(\mathtt{id}, \mathsf{c}_{\mathsf{NM}}, s_0; \rho)$.
   - 3.4. Set $x = \big((\mathsf{a}_{\mathsf{NM}}, \mathsf{c}_{\mathsf{NM}}, \mathsf{z}_{\mathsf{NM}}), Y, s_1, \mathtt{com}, \mathtt{id}\big)$ and $w = (m, \sigma, \perp, \perp)$ with $(|x| = \ell)$. Run $\mathsf{z}_{\mathsf{LS}} = \mathsf{P}^2(x, w, \mathsf{c}_{\mathsf{LS}}; \alpha)$ where x is the theorem to be proven and w is the witness.
   - 3.5. Send $(\mathsf{z}_{\mathsf{NM}}, \mathtt{com}, \mathsf{z}_{\mathsf{LS}}, s_1)$ to MMRec.
4. MMRec: Set $x = \big((\mathsf{a}_{\mathsf{NM}}, \mathsf{c}_{\mathsf{NM}}, \mathsf{z}_{\mathsf{NM}}), Y, s_1, \mathtt{com}, \mathtt{id}\big)$ and abort iff $(\mathsf{a}_{\mathsf{LS}}, \mathsf{c}_{\mathsf{LS}}, \mathsf{z}_{\mathsf{LS}})$ is not accepting for $\mathcal{V}$ with respect to x.

Decommitment Phase:

1. $\mathsf{MMSen} \rightarrow \mathsf{MMRec}$: Send $(\mathtt{dec}, m, \mathtt{dec}_{\mathsf{NM}}, s_0)$ to MMRec.
2. MMRec: Accept m as the committed message iff
   - 2.1. $\mathsf{NIRec}(\mathtt{dec}, \mathtt{com}, m) = 1$ and
   - 2.2. $\mathsf{Rec}_{\mathsf{NM}}$ on input γ, $(\mathsf{a}_{\mathsf{NM}}, \mathsf{c}_{\mathsf{NM}}, \mathsf{z}_{\mathsf{NM}}, \mathtt{id})$, s_0 and $\mathtt{dec}_{\mathsf{NM}}$ outputs 1.

Figure 2: Our 3-round concurrent NM commitment scheme.

where S is the simulator depicted in Fig. 3.

This means that the real execution in which the sender runs MMSen to commit to a message m must be indistinguishable with respect to an execution in which a simulator S runs internally the MiM adversarial $\mathcal{A}$ sending a commitment of 0^λ, and then forwards the messages that $\mathcal{A}$ sends in the right sessions to receivers $\mathsf{MMRec}_1, \dots, \mathsf{MMRec}_{\mathsf{poly}(\lambda)}$.

We remark that in the security proof we denote by $\tilde{\delta_i}$ a value associated with the i-th right session (where the adversary $\mathcal{A}$ plays with a receiver MMRec_i with $i \in \{1, \dots, \mathsf{poly}(\lambda)\}$) where δ is the corresponding value in the left session. For example, the sender commits to v in the left session while $\mathcal{A}$ commits to $\tilde{v}_i$ in the i-th right session.

To prove the indistinguishability of the above two experiments we proceed by showing 3 hybrid experiments[14] $\mathcal{H}_i^m(z)$ with $i = 1, 2, 3$, where m is the message committed in the left session. Following [LP11] we denote by $\{\mathsf{mim}^{\mathcal{A}}_{\mathcal{H}_i^m}(z)\}_{z \in \{0,1\}^\star}$ the random variable describing the view of the MiM $\mathcal{A}$ combined with the value it commits in the right interaction in hybrid $\mathcal{H}_i^m(z)$ (as usual, the committed value is replaced by $\perp$ if the right interaction does not correspond to a commitment that can be successfully opened or if $\mathcal{A}$ has copied the identity of the left interaction).

The first hybrid is the experiment in which in the left session MMSen commits to m, while in the right session we run $\mathsf{MMRec}_1, \dots, \mathsf{MMRec}_{\mathsf{poly}(\lambda)}$ for the rights sessions played by $\mathcal{A}$. We refer to this hybrid experiment as $\mathcal{H}_1^m(z)$, details follow below.

$\boldsymbol{\mathcal{H}_1^m(z)}$.

Left session:

1. First round.
   - 1.1. Pick $s_0 \leftarrow \{0,1\}^\lambda$.
   - 1.2. Compute $\mathsf{a}_{\mathsf{NM}} = \mathrm{Sen}^1_{\mathsf{NM}}(\mathtt{id}, s_0; \rho)$.
   - 1.3. Compute $\mathsf{a}_{\mathsf{LS}} = \mathsf{P}^1(1^{\lambda_{\mathsf{LS}}}, \ell; \alpha)$.
   - 1.4. Send $(\mathsf{a}_{\mathsf{NM}}, \mathsf{a}_{\mathsf{LS}})$ to $\mathcal{A}$.
2. Third round, upon receiving $(\mathsf{c}_{\mathsf{NM}}, \mathsf{c}_{\mathsf{LS}}, Y)$ from $\mathcal{A}$.
   - 2.1. Compute $(\mathtt{com}, \mathtt{dec}) = \mathsf{NISen}(m; \sigma)$.
   - 2.2. Pick $s_1 \leftarrow \{0,1\}^\lambda$.
   - 2.3. Compute $(\mathsf{z}_{\mathsf{NM}}, \mathtt{dec}_{\mathsf{NM}}) = \mathrm{Sen}^2_{\mathsf{NM}}(\mathtt{id}, \mathsf{c}_{\mathsf{NM}}, s_0; \rho)$.
   - 2.4. Set $x = \big((\mathsf{a}_{\mathsf{NM}}, \mathsf{c}_{\mathsf{NM}}, \mathsf{z}_{\mathsf{NM}}), Y, s_1, \mathtt{com}, \mathtt{id}\big)$ and $w = (m, \sigma, \perp, \perp)$ with $(|x| = \ell)$. Run $\mathsf{z}_{\mathsf{LS}} = \mathsf{P}^2(x, w, \mathsf{c}_{\mathsf{LS}}; \alpha)$.
   - 2.5. Send $(\mathsf{z}_{\mathsf{NM}}, \mathtt{com}, \mathsf{z}_{\mathsf{LS}}, s_1)$ to $\mathcal{A}$.

Right sessions: act as a proxy between $\mathcal{A}$ and $\mathsf{MMRec}_1, \dots, \mathsf{MMRec}_{\mathsf{poly}(\lambda)}$.

We have that for all $m \in \{0,1\}^{\mathsf{poly}(\lambda)}$ $\{\mathsf{mim}^{\mathcal{A}}_{\mathcal{H}_1^m}(z)\}_{z \in \{0,1\}^\star}$ clearly corresponds to $\{\mathsf{mim}^{\mathcal{A},m}_{\Pi_{\mathsf{MMCom}}}(z)\}_{z \in \{0,1\}^\star}$. Before we move on with the sequence of hybrid experiments we need to prove that, for all $i \in \{1, \dots, \mathsf{poly}(\lambda)\}$ $\mathcal{A}$ does not manage to invert any values $\tilde{Y}_i$ in the right sessions by sending a value $\tilde{s_1}_i$ such that $f(\tilde{s_0}_i \oplus \tilde{s_1}_i) = \tilde{Y}_i$ where $\tilde{s_0}_i$ is the message committed in the i-th right session through NM.

Lemma 1. *Let p_i be the probability that in the i-th right session, for $i \in \{1, \dots, \mathsf{poly}(\lambda)\}$, $\mathcal{A}$ sends a value $\tilde{s_1}_i$ such that $f(\tilde{s_1}_i \oplus \tilde{s_0}_i) = \tilde{Y}_i$ where $\tilde{s_0}_i$ is the value committed using NM. Then $p_i < \nu(\lambda)$ for some negligible function ν.*

[14] We will describe the hybrid experiments in a succinct way focusing on the key steps (e.g., omitting sampling of randomness, generation of parameters $\lambda_{\mathsf{NI}}, \lambda_{\mathsf{NM}}, \lambda_{\mathsf{LS}}, \ell$).

Proof. Suppose by contradiction that for a right session i the claim does not hold. We can construct an adversary $\mathcal{A}_f$ that inverts the OWP f in polynomial time. Formally we consider a challenger $\mathcal{C}_f$ of f that chooses a random Y in the range of f and sends it to $\mathcal{A}_f$. $\mathcal{A}_f$ wins if it gives as output y such that $Y = f(y)$. Before describing the adversary we need to consider the augmented machine $\mathcal{S}_{\mathsf{n}\rightarrow\mathsf{1}}$ that will be used by $\mathcal{A}_f$. $\mathcal{S}_{\mathsf{n}\rightarrow\mathsf{1}}$ internally executes $\mathcal{A}$, and interacts with an external receiver $\mathsf{Rec}_{\mathsf{ext}}$ of the protocol Π_{NM} acting as the sender. Formally $\mathcal{S}_{\mathsf{n}\rightarrow\mathsf{1}}$ acts as follows.

$\boldsymbol{\mathcal{S}_{\mathsf{n}\rightarrow\mathsf{1}}}(Y, \varphi, z)$

1. Act in the left session with $\mathcal{A}$ (that runs using randomness φ) as in $\mathcal{H}_1^m(z)$.
2. For all $j \neq i \in \{1, \dots \mathsf{poly}(\lambda)\}$ run MMRec_j as in $\mathcal{H}_1^m(z)$. Instead run MMRec_i as described in steps 3, 4 and 5.
3. Upon receiving the 1st round of the i-th right session $(\tilde{\mathsf{a}}_{\mathsf{NM}_i}, \tilde{\mathsf{a}}_{\mathsf{LS}_i})$ from $\mathcal{A}$, send $\tilde{\mathsf{a}}_{\mathsf{NM}_i}$ to $\mathsf{Rec}_{\mathsf{ext}}$.
4. Upon receiving $\tilde{\mathsf{c}}_{\mathsf{NM}_\mathsf{i}}$ from $\mathsf{Rec}_{\mathsf{ext}}$, run as follows:

    4.1. Run $\mathcal{V}$ to obtain $\tilde{\mathsf{c}}_{\mathsf{LS}_\mathsf{i}}$.

    4.2. Set $\tilde{Y}_i = Y$.

    4.3. Send $(\tilde{\mathsf{c}}_{\mathsf{NM}_\mathsf{i}}, \tilde{\mathsf{c}}_{\mathsf{LS}_\mathsf{i}}, \tilde{Y}_i)$ to $\mathcal{A}$.

5. Upon receiving the 3rd round of the i-th right session $(\tilde{\mathsf{z}}_{\mathsf{NM}_i}, \tilde{\mathtt{com}}_i, \tilde{\mathsf{z}}_{\mathsf{LS}_i}, \tilde{s_1}_i)$,

    set $\tilde{x} = \big((\tilde{\mathsf{a}}_{\mathsf{NM}_i}, \tilde{\mathsf{c}}_{\mathsf{NM}_i}, \tilde{\mathsf{z}}_{\mathsf{NM}_i}), \tilde{Y}, \tilde{s_1}_i, \tilde{\mathtt{com}}_i, \tilde{\mathtt{id}}\big)$ and abort iff $(\tilde{\mathsf{a}}_{\mathsf{LS}_i}, \tilde{\mathsf{c}}_{\mathsf{LS}_i}, \tilde{\mathsf{z}}_{\mathsf{LS}_i})$ is not accepting for $\mathcal{V}$ with respect to $\tilde{x}$.
6. Send $\tilde{\mathsf{z}}_{\mathsf{NM}_i}$ to $\mathsf{Rec}_{\mathsf{ext}}$.

Notice that the above execution of $\mathcal{S}_{\mathsf{n}\rightarrow\mathsf{1}}$ is distributed identically to $\mathcal{H}_1^m(z)$ when $\mathsf{Rec}_{\mathsf{ext}}$ plays identically as honest receiver. Now we can conclude the proof of this lemma by describing how $\mathcal{A}_f$ works. $\mathcal{A}_f$ runs the extractor of Π_{NM} using $\mathcal{S}_{\mathsf{n}\rightarrow\mathsf{1}}$ as sender (recall that an extractor of Π_{NM} plays only having access to a sender of Π_{NM}). We have that the extractor with non-negligible probability outputs the committed message of an execution that inverts f. By using the randomness φ, $\mathcal{A}_f$ can reconstruct the view of $\mathcal{A}$ and retrive the value $\tilde{s_1}_i$. Therefore $\mathcal{A}$ running in polynomial time[15] outputs with non-negligible probability the value $y = \tilde{s_0}_i \oplus \tilde{s_1}_i$ such that $f(y) = Y$. □

We now consider the second hybrid experiment $\mathcal{H}_2^m(z)$ where in the left session, after receiving Y from $\mathcal{A}$, the sender in time T_f finds a value y such that $Y = f(y)$. Then the sender sets and sends $s_1 = y \oplus s_0$, where s_0 is the value committed using Π_{NM}. The only difference between this hybrid experiment and $\mathcal{H}_1^m(z)$ is that $\mathcal{H}_2^m(z)$ runs in time sub-exponential in λ, and the value s_1 is equal to $y \oplus s_0$ where $Y = f(y)$. Formally $\mathcal{H}_2^m(z)$ is the following experiment.

$\boldsymbol{\mathcal{H}_2^m(z)}$.

Left session:

[15] The extractor is an expected polynomial-time algorithm while $\mathcal{A}_f$ must be a strict polynomial-time algorithm. Therefore $\mathcal{A}_f$ will run the extractor up to a given upperbounded number of steps that is higher than the expected running time of the extractor. Obviously with non-negligible probability the *truncated* extraction procedure will be completed successfully and this is sufficient for $\mathcal{A}_f$ to invert f. The same standard argument about truncating the execution of an expected polynomial-time algorithm will be needed later but for simplicity we will not repeat this discussion.

1. First round.
   - 1.1. Pick $s_0 \leftarrow \{0,1\}^\lambda$.
   - 1.2. Compute $\mathsf{a}_{\mathsf{NM}} = \mathrm{Sen}^1_{\mathsf{NM}}(\mathtt{id}, s_0; \rho)$.
   - 1.3. Compute $\mathsf{a}_{\mathsf{LS}} = \mathsf{P}^1(1^{\lambda_{\mathsf{LS}}}, \ell; \alpha)$.
   - 1.4. Send $(\mathsf{a}_{\mathsf{NM}}, \mathsf{a}_{\mathsf{LS}})$ to $\mathcal{A}$.
2. Third round. Upon receiving $(\mathsf{c}_{\mathsf{NM}}, \mathsf{c}_{\mathsf{LS}}, Y)$ from $\mathcal{A}$.
   - 2.1. Compute $(\mathtt{com}, \mathtt{dec}) = \mathsf{NISen}(m; \sigma)$.
   - 2.2. Run in time T_f to compute y such that $Y = f(y)$.
   - 2.3. Set $s_1 = y \oplus s_0$.
   - 2.4. Compute $(\mathsf{z}_{\mathsf{NM}}, \mathtt{dec}_{\mathsf{NM}}) = \mathrm{Sen}^2_{\mathsf{NM}}(\mathtt{id}, \mathsf{c}_{\mathsf{NM}}, s_0; \rho)$.
   - 2.5. Set $x = \big((\mathsf{a}_{\mathsf{NM}}, \mathsf{c}_{\mathsf{NM}}, \mathsf{z}_{\mathsf{NM}}), Y, s_1, \mathtt{com}, \mathtt{id}\big)$ and $w = (m, \sigma, \perp, \perp)$ with $(|x| = \ell)$. Run $\mathsf{z}_{\mathsf{LS}} = \mathsf{P}^2(x, w, \mathsf{c}_{\mathsf{LS}}; \alpha)$.
   - 2.6. Send $(\mathsf{z}_{\mathsf{NM}}, \mathtt{com}, \mathsf{z}_{\mathsf{LS}}, s_1)$ to MMRec.

Right sessions: Act as a proxy between $\mathcal{A}$ and $\mathsf{MMRec}_1, \ldots, \mathsf{MMRec}_{\mathsf{poly}(\lambda)}$.

When switching from $\mathcal{H}^m_1(z)$ to $\mathcal{H}^m_2(z)$ we will make sure that the following two properties hold.

1. For all message $m \in \{0,1\}^{\mathsf{poly}(\lambda)}$ it holds that $\mathsf{mim}^{\mathcal{A}}_{\mathcal{H}^m_1}(z) \approx \mathsf{mim}^{\mathcal{A}}_{\mathcal{H}^m_2}(z)$.[16]
2. Let p_i be the probability that in the i-th right session of $\mathcal{H}_2$, for $i \in \{1, \ldots, \mathsf{poly}(\lambda)\}$, $\mathcal{A}$ sends a value $\tilde{s_1}_i$ such that $f(\tilde{s_1}_i \oplus \tilde{s_0}_i) = \tilde{Y}_i$ where $\tilde{s_0}_i$ is the value committed using NM. Then $p_i < \nu(\lambda)$ for some negligible function ν.

We now prove that the above two properties hold.

Lemma 2. *For all message $m \in \{0,1\}^{\mathsf{poly}(\lambda)}$ it holds that $\mathsf{mim}^{\mathcal{A}}_{\mathcal{H}^m_1}(z) \approx \mathsf{mim}^{\mathcal{A}}_{\mathcal{H}^m_2}(z)$.*

Proof. Suppose by contradiction that the distribution of $\mathsf{mim}^{\mathcal{A}}_{\mathcal{H}^m_1}(z)$ is distinguishable from $\mathsf{mim}^{\mathcal{A}}_{\mathcal{H}^m_2}(z)$; this means that there exists a distinguisher $\mathcal{D}$ that can tell apart such two distributions. We now use $\mathcal{D}$ and $\mathcal{A}$ to construct an adversary $\mathcal{A}_{\mathsf{Hiding}}$ that breaks the hiding of Π_{NM} in time $\mathsf{poly}(\lambda) \cdot T_{\mathsf{NI}}$ therefore reaching a contradiction[17]. Let $\mathcal{C}_{\mathsf{Hiding}}$ be the challenger of the hiding game, we consider two randomly chosen challenge messages (m_0, m_1) sent to $\mathcal{C}_{\mathsf{Hiding}}$. We now provide a formal description of the adversary $\mathcal{A}_{\mathsf{Hiding}}$.

$\boldsymbol{\mathcal{A}}_{\mathsf{Hiding}}(m_0, m_1, z)$

1. Upon receiving the 1st round a_{NM} from $\mathcal{C}_{\mathsf{Hiding}}$, run as follows:
   - 1.1. Compute $\mathsf{a}_{\mathsf{LS}} = \mathsf{P}^1(1^{\lambda_{\mathsf{LS}}}, \ell; \alpha)$.
   - 1.2. Send $(\mathsf{a}_{\mathsf{NM}}, \mathsf{a}_{\mathsf{LS}})$ to $\mathcal{A}$.
2. Upon receiving $(\mathsf{c}_{\mathsf{NM}}, \mathsf{c}_{\mathsf{LS}}, Y)$ from $\mathcal{A}$, send c_{NM} to $\mathcal{C}_{\mathsf{NM}}$.
3. Upon receiving the 3rd round z_{NM} from $\mathcal{C}_{\mathsf{Hiding}}$, run as follows:
   - 3.1. Compute y such that $Y = f(y)$, set $s_1 = m_0 \oplus y$.

[16] To simplify the notation here, and in the rest of the proof, we will omit that the indistinguishability between two distributions must hold for every auxiliary input z.

[17] Recall that Π_{NM} is secure against adversaries running in time $\mathsf{poly}(\lambda) \cdot T_{\mathsf{NI}} < T_{\mathsf{NM}}$.

3.2. Compute $(\mathtt{com}, \mathtt{dec}) = \mathsf{NISen}(m; \sigma)$.

3.3. Set $x = \big((\mathsf{a}_{\mathsf{NM}}, \mathsf{c}_{\mathsf{NM}}, \mathsf{z}_{\mathsf{NM}}), Y, s_1, \mathtt{com}, \mathtt{id}\big)$ and $w = (m, \sigma, \perp, \perp)$ with $(|x| = \ell)$. Run $\mathsf{z}_{\mathsf{LS}} = \mathsf{P}^2(x, w, \mathsf{c}_{\mathsf{LS}}; \alpha)$.

3.4. Send $(\mathsf{z}_{\mathsf{NM}}, \mathtt{com}, \mathsf{z}_{\mathsf{LS}}, s_1)$ to $\mathcal{A}$.

4. Simulate $\mathsf{MMRec}_1, \ldots, \mathsf{MMRec}_{\mathsf{poly}(\lambda)}$ with $\mathcal{A}$ when $\mathcal{A}$ plays as a sender.
5. Let M be an empty tuple. For all $i \in \{1, \ldots, \mathsf{poly}(\lambda)\}$, consider $\tilde{\mathtt{com}}_i$, the non-interactive commitment received by MMRec_i, run in time T_{NI} to compute $\tilde{m}_i$ such that $\exists\ \tilde{\mathtt{dec}} : 1 = \mathsf{NIRec}(\tilde{\mathtt{com}}_i, \tilde{\mathtt{dec}}, \tilde{m}_i)$ and add $\tilde{m}_i$ to M.
6. Give M and the view of $\mathcal{A}$ to the distinguisher $\mathcal{D}$ and output what $\mathcal{D}$ outputs.

The proof ends with the observation that if $\mathcal{C}_{\mathsf{Hiding}}$ has committed to m_0 then the xor of the committed value with s_1 is equal to y such that $f(y) = Y$, like in $\mathcal{H}_2^m(z)$. If instead $\mathcal{C}_{\mathsf{Hiding}}$ has committed to m_1 then the xor of the committed value and s_1 is equal to a random value, like in $\mathcal{H}_1^m(z)$. □

Lemma 3. *Let p_i be the probability that in the i-th right session of $\mathcal{H}_2$, for $i \in \{1, \ldots, \mathsf{poly}(\lambda)\}$, $\mathcal{A}$ sends a value $\tilde{s}_{1i}$ such that $f(\tilde{s}_{1i} \oplus \tilde{s}_{0i}) = \tilde{Y}_i$ where $\tilde{s}_{0i}$ is the value committed using* NM*. Then $p_i < \nu(\lambda)$ for some negligible function ν.*

Proof. Suppose by contradiction that for a right session i the claim does not hold. We can construct a distinguisher $\mathcal{D}_{\mathsf{NM}}$ and an adversary $\mathcal{A}_{\mathsf{NM}}$ that break the non-malleability of Π_{NM}. Let $\mathcal{C}_{\mathsf{NM}}$ be the challenger of the NM commitment and let (m_0, m_1) be two randomly chosen challenge messages given to $\mathcal{C}_{\mathsf{NM}}$.

$\boldsymbol{\mathcal{A}}_{\mathsf{NM}}(m_0, m_1, z)$

Left session:

1. Act as $\mathcal{A}_{\mathsf{Hiding}}$ acts in the left session.

Right sessions:

1. For all $j \neq i \in \{1, \ldots, \mathsf{poly}(\lambda)\}$ run MMRec_j as in $\mathcal{H}_2^m(z)$. Instead run MMRec_i as described in steps 1.1, 1.2 and 1.3.

1.1. Forward $\tilde{\mathsf{a}}_{\mathsf{NM}_i}$ to $\mathsf{Rec}_{\mathsf{NM}}$.

1.2. Upon receiving $\tilde{\mathsf{c}}_{\mathsf{NM}}$ from $\mathsf{Rec}_{\mathsf{NM}}$, pick a random $\tilde{\mathsf{c}}_{\mathsf{LS}_i}$, pick a random $\tilde{Y}_i$ and send $(\tilde{\mathsf{c}}_{\mathsf{NM}_i}, \tilde{\mathsf{c}}_{\mathsf{LS}_i}, \tilde{Y}_i)$ to $\mathcal{A}$.

1.3. Upon receiving $\tilde{\mathsf{z}}_{\mathsf{NM}_i}$ from $\mathcal{A}$, send it to $\mathsf{Rec}_{\mathsf{NM}}$.

Let $\mathsf{mim}^{\mathcal{A}_{\mathsf{NM}}}(z)$ be the view of $\mathsf{mim}^{\mathcal{A}_{\mathsf{NM}}}(z)$ and the tuple of committed messages in the right session. The distinguisher $\mathcal{D}_{\mathsf{NM}}$ takes as input $\mathsf{mim}^{\mathcal{A}_{\mathsf{NM}}}(z)$ and acts as follows.

$\boldsymbol{\mathcal{D}}_{\mathsf{NM}}(\mathsf{mim}^{\mathcal{A}_{\mathsf{NM}}}(\boldsymbol{z}))$ **:** Let $\tilde{s}_{0i}$ be the committed message sent in the i-right session by $\mathcal{A}_{\mathsf{NM}}$ to MMRec. Reconstruct the output messages of $\mathcal{A}$ (using the same randomness of $\mathsf{mim}^{\mathcal{A}_{\mathsf{NM}}}(z)$) to pick $\tilde{s}_{1i}$. If $f(\tilde{s}_{1i} \oplus \tilde{s}_{0i}) = \tilde{Y}_i$ output 1 and output 0 otherwise. The proof ends with the observation that if $\mathcal{C}_{\mathsf{NM}}$ has committed to m_0 then the xor of the committed value with s_{1i} is equal to y such that $f(y) = Y$ like in $\mathcal{H}_2^m$. If instead $\mathcal{C}_{\mathsf{Hiding}}$ has committed to m_1 then the xor of the committed value with s_{1i} is equal to a random string as in $\mathcal{H}_1^m$. □

The third hybrid experiment that we consider is equal to $\mathcal{H}_2^m(z)$ with the difference that the LS proof system is executed using s_0 and the randomness of the non-malleable commitment of s_0. Recall that $f(s_0 \oplus s_1) = Y$. We observe that in the left session of $\mathcal{H}_2^m(z)$ it already holds that $f(s_0 \oplus s_1) = Y$, therefore we can switch the witness used in LS and complete the execution of the proof system. Formally $\mathcal{H}_3^m(z)$ is the following experiment.

$\boldsymbol{\mathcal{H}_3^m(z)}$.

Left sessions:

1. First round.
   - 1.1. Pick $s_0 \leftarrow \{0,1\}^\lambda$.
   - 1.2. Compute $\mathsf{a_{NM}} = \mathrm{Sen}^1_{\mathsf{NM}}(\mathtt{id}, s_0; \rho)$.
   - 1.3. Compute $\mathsf{a_{LS}} = \mathsf{P}^1(1^{\lambda_{\mathsf{LS}}}, \ell; \alpha)$.
   - 1.4. Send $(\mathsf{a_{NM}}, \mathsf{a_{LS}})$ to $\mathcal{A}$.
2. Third round. Upon receiving $(\mathsf{c_{NM}}, \mathsf{c_{LS}}, Y)$ from $\mathcal{A}$.
   - 2.1. Compute $(\mathtt{com}, \mathtt{dec}) = \mathsf{NISen}(m; \sigma)$.
   - 2.2. Run in time T_f to compute y such that $Y = f(y)$.
   - 2.3. Set $s_1 = s_0 \oplus y$.
   - 2.4. Compute $(\mathsf{z_{NM}}, \mathsf{dec_{NM}}) = \mathrm{Sen}^2_{\mathsf{NM}}(\mathtt{id}, \mathsf{c_{NM}}, s_0; \rho)$.
   - 2.5. Compute $(\mathtt{com}, \mathtt{dec}) = \mathsf{NISen}(1^{\lambda_{\mathsf{NI}}}, m; \sigma)$.
   - 2.6. Set $x = \big((\mathsf{a_{NM}}, \mathsf{c_{NM}}, \mathsf{z_{NM}}), Y, s_1, \mathtt{com}, \mathtt{id}\big)$ and $\underline{w = (\bot, \bot, s_0, \rho)}$ with $(|x| = \ell)$. Run $\mathsf{z_{LS}} = \mathsf{P}^2(x, w, \mathsf{c_{LS}}; \alpha)$.
   - 2.7. Send $(\mathsf{z_{NM}}, \mathtt{com}, \mathsf{z_{LS}}, s_1)$ to $\mathcal{A}$.

Right sessions: Act as a proxy between $\mathcal{A}$ and $\mathsf{MMRec}_1, \ldots, \mathsf{MMRec}_{\mathsf{poly}(\lambda)}$.

Even in this case we need to prove the following two properties.

1. For all message $m \in \{0,1\}^{\mathsf{poly}(\lambda)}$ it holds that $\mathsf{mim}^{\mathcal{A}}_{\mathcal{H}_2^m}(z) \approx \mathsf{mim}^{\mathcal{A}}_{\mathcal{H}_3^m}(z)$.
2. Let p_i be the probability that in the i-th right session of $\mathcal{H}_3$, for any $i \in \{1, \ldots, \mathsf{poly}(\lambda)\}$, $\mathcal{A}$ sends a value $\tilde{s_1}_i$ such that $f(\tilde{s_1}_i \oplus \tilde{s_0}_i) = \tilde{Y}_i$ where $\tilde{s_0}_i$ is the value committed using NM. Then $p_i < \nu(\lambda)$ for some negligible function ν.

Lemma 4. *For any message $m \in \{0,1\}^{\mathsf{poly}(\lambda)}$ it holds that $\mathsf{mim}^{\mathcal{A}}_{\mathcal{H}_2^m}(z) \approx \mathsf{mim}^{\mathcal{A}}_{\mathcal{H}_3^m}(z)$.*

Proof. Suppose by contradiction that there exist a adversary $\mathcal{A}$ and a distinguisher $\mathcal{D}$ that can tell apart such two distributions. We can use this adversary and the associated distinguisher to construct an adversary $\mathcal{A}_{\mathsf{LS}}$ for the T_{LS}-witness-indistinguishable property of the LS protocol. Let $\mathcal{C}_{\mathsf{LS}}$ be the WI challenger, the adversary works as follows. $\boldsymbol{\mathcal{A}}_{\mathsf{LS}}(z)$

1. Pick $s_0 \leftarrow \{0,1\}^\lambda$.
2. Compute $\mathsf{a_{NM}} = \mathrm{Sen}^1_{\mathsf{NM}}(\mathtt{id}, s_0; \rho)$.
3. Upon receiving $\mathsf{a_{LS}}$ from $\mathcal{C}_{\mathsf{LS}}$, send $(\mathsf{a_{NM}}, \mathsf{a_{LS}})$ to $\mathcal{A}$.
4. Upon receiving $(\mathsf{c_{NM}}, \mathsf{c_{LS}}, Y)$ from $\mathcal{A}$ run as follows:
   - 4.1. Run in time T_f to compute y such that $Y = f(y)$.
   - 4.2. Set $s_1 = s_0 \oplus y$.

   4.3. Compute $(\mathsf{z}_{\mathsf{NM}}, \mathtt{dec}_{\mathsf{NM}}) = \mathrm{Sen}^2_{\mathsf{NM}}(\mathtt{id}, \mathsf{c}_{\mathsf{NM}}, s_0; \rho)$.

   4.4. Compute $(\mathtt{com}, \mathtt{dec}) = \mathsf{NISen}(1^{\lambda_{\mathsf{NI}}}, m; \sigma)$.

   4.5. Set $x = \big((\mathsf{a}_{\mathsf{NM}}, \mathsf{c}_{\mathsf{NM}}, \mathsf{z}_{\mathsf{NM}}), Y, s_1, \mathtt{com}, \mathtt{id}\big)$, $w_0 = (\perp, \perp, s_0, \rho)$,
   $w_1 = (m, \sigma, \perp, \perp)$ and send $(x, \mathsf{c}_{\mathsf{LS}}, w_0, w_1)$ to $\mathcal{C}_{\mathsf{LS}}$.

5. Upon receiving z_{LS} from $\mathcal{C}_{\mathsf{LS}}$, send $(\mathsf{z}_{\mathsf{NM}}, \mathtt{com}, \mathsf{z}_{\mathsf{LS}})$ to $\mathcal{A}$.
6. Simulate $\mathsf{MMRec}_1, \dots, \mathsf{MMRec}_{\mathsf{poly}(\lambda)}$ with $\mathcal{A}$, when $\mathcal{A}$ plays as a sender.
7. Let M be an empty tuple. For all $i \in \{1, \dots, \mathsf{poly}(\lambda)\}$, consider $\tilde{\mathtt{com}}_i$, the non-interactive commitment received by MMRec_i, and run in time $\tilde{T}_{\mathsf{NI}}$ to compute $\tilde{m}_i$ such that $\exists\ \tilde{\mathtt{dec}} : 1 = \mathsf{NIRec}(\tilde{\mathtt{com}}_i, \tilde{\mathtt{dec}}, \tilde{m}_i)$ and add $\tilde{m}_i$ to M.
8. Give M and the view of $\mathcal{A}$ to the distinguisher $\mathcal{D}$.
9. Output what $\mathcal{D}$ outputs.

□

The proof ends with the observation that if $\mathcal{C}_{\mathsf{LS}}$ has has used as witness the randomness of the non-malleable commitment of the value s_0 such that $f(s_0 \oplus s_1) = Y$ then we are in the hybrid experiment $\mathcal{H}^m_3(z)$. If instead $\mathcal{C}_{\mathsf{LS}}$ has used as a witness the randomness used to compute the non-interactive commitment NI then we are in the hybrid experiment $\mathcal{H}^m_2(z)$.

Lemma 5. *Let p_i be the probability that in the i-th right session of $\mathcal{H}^m_3$, for $i \in \{1, \dots, \mathsf{poly}(\lambda)\}$, $\mathcal{A}$ sends a value $\tilde{s}_{1i}$ such that $f(\tilde{s}_{1i} \oplus \tilde{s}_{0i}) = \tilde{Y}_i$ where $\tilde{s}_{0i}$ is the value committed using* NM*. Then $p_i < \nu(\lambda)$ for some negligible function ν.*

Proof. Suppose by contradiction that for a right session i the claim does not hold, then we can construct an adversary $\mathcal{A}'_{\mathsf{LS}}$ for the T_{LS} witness-indistinguishable property of the LS protocol. Let $\mathcal{C}_{\mathsf{LS}}$ be the WI challenger, the adversary works as follows.

$\boldsymbol{\mathcal{A}'_{\mathsf{LS}}}(z)$

1. Pick $s_0 \leftarrow \{0,1\}^\lambda$.
2. Compute $\mathsf{a}_{\mathsf{NM}} = \mathrm{Sen}^1_{\mathsf{NM}}(\mathtt{id}, s_0; \rho)$.
3. Upon receiving a_{LS} from $\mathcal{C}_{\mathsf{LS}}$, send $(\mathsf{a}_{\mathsf{NM}}, \mathsf{a}_{\mathsf{LS}})$ to $\mathcal{A}$.
4. Upon receiving $(\mathsf{c}_{\mathsf{NM}}, \mathsf{c}_{\mathsf{LS}}, Y)$ from $\mathcal{A}$, run as follow:

   4.1. Run in time T_f to compute y such that $Y = f(y)$.

   4.2. Set $s_1 = s_0 \oplus y$.

   4.3. Compute $(\mathsf{z}_{\mathsf{NM}}, \mathtt{dec}_{\mathsf{NM}}) = \mathrm{Sen}^2_{\mathsf{NM}}(\mathtt{id}, \mathsf{c}_{\mathsf{NM}}, s_0; \rho)$.

   4.4. Compute $(\mathtt{com}, \mathtt{dec}) = \mathsf{NISen}(1^{\lambda_{\mathsf{NI}}}, m; \sigma)$.

   4.5. Set $x = \big((\mathsf{a}_{\mathsf{NM}}, \mathsf{c}_{\mathsf{NM}}, \mathsf{z}_{\mathsf{NM}}), Y, s_1, \mathtt{com}, \mathtt{id}\big)$, $w_0 = (\perp, \perp, s_0, \rho)$,
   $w_1 = (m, \sigma, \perp, \perp)$ and send $(x, \mathsf{c}_{\mathsf{LS}}, w_0, w_1)$ to $\mathcal{C}_{\mathsf{LS}}$.

5. Upon receiving z_{LS} from $\mathcal{C}_{\mathsf{LS}}$, send $(\mathsf{z}_{\mathsf{NM}}, \mathtt{com}, \mathsf{z}_{\mathsf{LS}})$ to $\mathcal{A}$.
6. Simulate $\mathsf{MMRec}_1, \dots, \mathsf{MMRec}_{\mathsf{poly}(\lambda)}$ with $\mathcal{A}$, when $\mathcal{A}$ plays as a sender.

7. Run in time $\tilde{T}_{\mathsf{NM}}$ to extract the value $\tilde{s_0}_i$ from the non-malleable commitment sent by $\mathcal{A}$ in the i-th session. Output 1 if $f(\tilde{s_0}_i \oplus \tilde{s_1}_i) = \tilde{Y}_i$ and output 0 otherwise.

The proof ends with the observation that if $\mathcal{C}_{\mathsf{LS}}$ has used $w_0 = (\perp, \perp, s_0, \rho)$ as a witness then $\mathcal{A}$ acts as in $\mathcal{H}_3^m(z)$, sending with non-negligible probability two shares such that the xor of them gives a puzzle solution. If $\mathcal{C}_{\mathsf{LS}}$ has used $w_1 = (m, \sigma, \perp, \perp)$ then the xor of the two shares is with overwhelming probability different from a puzzle solution as in $\mathcal{H}_2^m(z)$. □

The next hybrid experiment that we consider is $\mathcal{H}_3^0(z)$. The only differences between this hybrid experiment and $\mathcal{H}_3^m(z)$ is that the sender, using NI, commits to a message 0^λ instead of m. Formally the hybrid experiment is the following.

$\boldsymbol{\mathcal{H}_3^0(z)}$.

Left session:

1. First round.
   - 1.1. Pick $s_0 \leftarrow \{0,1\}^\lambda$.
   - 1.2. Compute $\mathsf{a}_{\mathsf{NM}} = \mathrm{Sen}^1_{\mathsf{NM}}(\mathtt{id}, s_0; \rho)$.
   - 1.3. Compute $\mathsf{a}_{\mathsf{LS}} = \mathsf{P}^1(\ell; \alpha)$.
   - 1.4. Send $(\mathsf{a}_{\mathsf{NM}}, \mathsf{a}_{\mathsf{LS}})$ to $\mathcal{A}$.
2. Third round. Upon receiving $(\mathsf{c}_{\mathsf{NM}}, \mathsf{c}_{\mathsf{LS}}, Y)$ from $\mathcal{A}$, run as follows:
   - 2.1. Run in time T_f to compute y such that $Y = f(y)$.
   - 2.2. Set $s_1 = s_0 \oplus y$.
   - 2.3. Compute $(\mathsf{z}_{\mathsf{NM}}, \mathsf{dec}_{\mathsf{NM}}) = \mathrm{Sen}^2_{\mathsf{NM}}(\mathtt{id}, \mathsf{c}_{\mathsf{NM}}, s_0; \rho)$.
   - 2.4. Compute $(\mathtt{com}, \mathtt{dec}) = \underline{\mathsf{NISen}(0^\lambda; \sigma)}$.
   - 2.5. Set $x = \big((\mathsf{a}_{\mathsf{NM}}, \mathsf{c}_{\mathsf{NM}}, \mathsf{z}_{\mathsf{NM}}), Y, s_1, \mathtt{com}, \mathtt{id}\big)$ and $w = (\perp, \perp, s_0, \rho)$ with $(|x| = \ell)$. Run $\mathsf{z}_{\mathsf{LS}} = \mathsf{P}^2(x, w, \mathsf{c}_{\mathsf{LS}}; \alpha)$.
   - 2.6. Send $(\mathsf{z}_{\mathsf{NM}}, \mathtt{com}, \mathsf{z}_{\mathsf{LS}}, s_1)$ to $\mathcal{A}$.

Right sessions: Act as a proxy between $\mathcal{A}$ and $\mathsf{MMRec}_1, \ldots, \mathsf{MMRec}_{\mathsf{poly}(\lambda)}$.

We now prove the following properties.

1. Let p_i be the probability that in the i-th right session of $\mathcal{H}_3^0$, for any $i \in \{1, \ldots, \mathsf{poly}(\lambda)\}$, $\mathcal{A}$ sends a value $\tilde{s_1}_i$ such that $f(\tilde{s_1}_i \oplus \tilde{s_0}_i) = \tilde{Y}_i$ where $\tilde{s_0}_i$ is the value committed using NM. Then $p_i < \nu(\lambda)$ for some negligible function ν.
2. For any message $m \in \{0,1\}^{\mathsf{poly}(\lambda)}$ it holds that $\mathsf{mim}^{\mathcal{A}}_{\mathcal{H}_3^m}(z) \approx \mathsf{mim}^{\mathcal{A}}_{\mathcal{H}_3^0}(z)$.

Lemma 6. *Let p_i be the probability that in the i-th right session of $\mathcal{H}_3^0$, for $i \in \{1, \ldots, \mathsf{poly}(\lambda)\}$, $\mathcal{A}$ sends a value $\tilde{s_1}_i$ such that $f(\tilde{s_1}_i \oplus \tilde{s_0}_i) = \tilde{Y}_i$ where $\tilde{s_0}_i$ is the value committed using* NM. *Then $p_i < \nu(\lambda)$ for some negligible function ν.*

Proof. Suppose by contradiction that there exists a right session $i \in \{1, \ldots, \mathsf{poly}(\lambda)\}$ in which $\mathcal{A}$ commit to a string $\tilde{s_0}$ such that $f(\tilde{s_0}_i \oplus \tilde{s_1}_i) = \tilde{Y}_i$ using Π_{NM}. Then we can construct an adversary $\mathcal{A}_{\mathsf{NI}}$ that breaks the hiding property of the non interactive commitment scheme NI. Let $\mathcal{C}_{\mathsf{NI}}$ be the challenger that on input $m_0 = 0^\lambda$ and $m_1 = m$, picks a random bit b, computes $(\mathtt{com}, \mathtt{dec}) = \mathsf{NISen}(1^{\lambda_{\mathsf{NI}}}, m_b; \sigma)$ and sends $\mathtt{com}$ to $\mathcal{A}_{\mathsf{NI}}$.

Before describing $\mathcal{A}_{\mathsf{NI}}$ we need to consider, as in the proof of Lemma 1, a machine $\mathcal{S}_{\mathsf{n}\to 1}$ that internally executes $\mathcal{A}$, and interacts with a receiver $\mathsf{Rec}_{\mathsf{ext}}$ of the protocol Π_{NM} acting as the sender.

Formally $\mathcal{S}_{\mathsf{n}\to 1}$ acts as follows.

```
$\boldsymbol{\mathcal{S}_{\mathsf{n}\to 1}}(\mathtt{com}, \varphi, z)$
    Run $\mathcal{A}$ using randomness $\varphi$.

    1. Pick $s_0 \leftarrow \{0,1\}^{\lambda}$.
    2. Compute $\mathsf{a}_{\mathsf{NM}} = \mathrm{Sen}^1_{\mathsf{NM}}(\mathtt{id}, s_0; \rho)$.
    3. Compute $\mathsf{a}_{\mathsf{LS}} = \mathsf{P}^1(1^{\lambda_{\mathsf{LS}}}, \ell; \alpha)$.
    4. Send $(\mathsf{a}_{\mathsf{NM}}, \mathsf{a}_{\mathsf{LS}})$ to $\mathcal{A}$.
    5. Upon receiving $(\mathsf{c}_{\mathsf{NM}}, \mathsf{c}_{\mathsf{LS}}, Y)$ from $\mathcal{A}$, run as follows:
        5.1. Run in time $T_f$ to compute $y$ such that $Y = f(y)$.
        5.2. Set $s_1 = s_0 \oplus y$.
        5.3. Compute $(\mathsf{z}_{\mathsf{NM}}, \mathtt{dec}_{\mathsf{NM}}) = \mathrm{Sen}^2_{\mathsf{NM}}(\mathtt{id}, \mathsf{c}_{\mathsf{NM}}, s_0; \rho)$.
        5.4. Set $x = \big((\mathsf{a}_{\mathsf{NM}}, \mathsf{c}_{\mathsf{NM}}, \mathsf{z}_{\mathsf{NM}}), Y, s_1, \mathtt{com}, \mathtt{id}\big)$ and $w = (\perp, \perp, s_0, \rho)$ with $(|x| = \ell)$. Run
             $\mathsf{z}_{\mathsf{LS}} = \mathsf{P}^2(x, w, \mathsf{c}_{\mathsf{LS}}; \alpha)$.
        5.5. Send $(\mathsf{z}_{\mathsf{NM}}, \mathtt{com}, \mathsf{z}_{\mathsf{LS}}, s_1)$ to $\mathcal{A}$.
    6. Let $i \in \{1, \dots, \mathsf{poly}(\lambda)\}$ be the right session that contradicts the claim. For all $j \neq i \in$
       $\{1, \dots \mathsf{poly}(\lambda)\}$ run $\mathsf{MMRec}_j$ as in $\mathcal{H}_4(m, z)$. Run $\mathsf{MMRec}_i$ as follows.
        6.1. Upon receiving the 1rd round of the $i$-th right session $(\tilde{\mathsf{a}}_{\mathsf{NM}_i}, \tilde{\mathsf{a}}_{\mathsf{LS}_i})$ from $\mathcal{A}$, send $\tilde{\mathsf{a}}_{\mathsf{NM}_i}$ to
             the external receiver $\mathsf{Rec}_{\mathsf{ext}}$.
        6.2. Upon receiving $\tilde{\mathsf{c}}_{\mathsf{NM}_{\mathsf{i}}}$ from $\mathsf{Rec}_{\mathsf{ext}}$, run as follows:
               i. Run $\mathcal{V}$ to obtain $\tilde{\mathsf{c}}_{\mathsf{LS}_{\mathsf{i}}}$.
              ii. Pick a random $\tilde{Y}_i$.
             iii. Send $(\tilde{\mathsf{c}}_{\mathsf{NM}_{\mathsf{i}}}, \tilde{\mathsf{c}}_{\mathsf{LS}_{\mathsf{i}}}, \tilde{Y}_i)$ to $\mathcal{A}$.
        6.3. Upon receiving the 3rd round of the $i$-th right session $(\tilde{\mathsf{z}}_{\mathsf{NM}_i}, \tilde{\mathtt{com}}_i, \tilde{\mathsf{z}}_{\mathsf{LS}_i}, \tilde{s_1}_i)$,
             set $\tilde{x} = \big((\tilde{\mathsf{a}}_{\mathsf{NM}_i}, \tilde{\mathsf{c}}_{\mathsf{NM}_i}, \tilde{\mathsf{z}}_{\mathsf{NM}_i}), \tilde{Y}, \tilde{s_1}_i, \tilde{\mathtt{com}}_i, \tilde{\mathtt{id}}\big)$ and abort iff $(\tilde{\mathsf{a}}_{\mathsf{LS}_i}, \tilde{\mathsf{c}}_{\mathsf{LS}_i}, \tilde{\mathsf{z}}_{\mathsf{LS}_i})$ is not accepted
             by $\mathcal{V}$ with respect to $\tilde{x}$.
        6.4. Send $\tilde{\mathsf{z}}_{\mathsf{NM}_i}$ to $\mathsf{Rec}_{\mathsf{ext}}$.
```

Now we can conclude the proof of this lemma by describing how $\mathcal{A}_{\mathsf{NI}}$ works. $\mathcal{A}_{\mathsf{NI}}$ runs the extractor of the protocol Π_{NM} using $\mathcal{S}_{\mathsf{n}\to 1}$ as sender (recall that an extractor of Π_{NM} plays only having access to a sender of Π_{NM}). Since the extractor with non-negligible probability outputs the committed message we have that $\mathcal{A}_{\mathsf{NI}}$ retrives $\tilde{s_0}_i$. Moreover $\mathcal{A}_{\mathsf{NI}}$ gets $\tilde{s_1}_i$ by reconstructing the view of $\mathcal{A}$ using the randomness φ. Since by contradiction $\mathcal{A}$ contradicts the claim of this lemma, we have that $\mathcal{A}_{\mathsf{NI}}$ can break the hiding of NI because $f(\tilde{s_0}_i \oplus \tilde{s_1}_i) = \tilde{Y}$ with non-negligible probability in $\mathcal{H}^0_3(z)$ where $m_0 = 0^{\lambda}$ is committed in $\mathtt{com}$, while the same happens with negligible probability only in $\mathcal{H}^m_3(z)$ where $m_1 = m$. Therefore if this happens, $\mathcal{A}_{\mathsf{NI}}$ outputs 0, otherwise $\mathcal{A}_{\mathsf{NI}}$ outputs a random bit. □

Lemma 7. *For any message $m \in \{0,1\}^{\mathsf{poly}(\lambda)}$ it holds that $\mathsf{mim}^{\mathcal{A}}_{\mathcal{H}_3^m}(z) \approx \mathsf{mim}^{\mathcal{A}}_{\mathcal{H}_3^0}(z)$.*

Proof. Suppose by contradiction that there exists a distinguisher $\mathcal{D}$ and an adversary $\mathcal{A}$ such that $\mathsf{mim}^{\mathcal{A}}_{\mathcal{H}_3^m}(z)$ is distinguishable from $\mathsf{mim}^{\mathcal{A}}_{\mathcal{H}_3^0}(z)$ then we can construct an adversary $\mathcal{A}_{\mathsf{NI}}$ that breaks the hiding property of the non-interactive commitment scheme NI. Let $\mathcal{C}_{\mathsf{NI}}$ be the challenger that on input $m_0 = 0^\lambda$ and $m_1 = m$, picks a random bit b, computes $(\mathtt{com}, \mathtt{dec}) = \mathsf{NISen}(1^{\lambda_{\mathsf{NI}}}, m_b; \sigma)$ and sends $\mathtt{com}$ to $\mathcal{A}_{\mathsf{NI}}$. Before describing $\mathcal{A}_{\mathsf{NI}}$, we consider the following experiment $\mathcal{E}_{m_b}(\varphi, \mathtt{com}, z)$.

$\boldsymbol{\mathcal{E}_{m_b}(\varphi, \mathtt{com}, z)}$.
The randomness required from all next steps is take from φ.

Run $\mathcal{A}(z)$.

Left session:

1. First round.
   1.1. Pick $s_0 \leftarrow \{0,1\}^\lambda$.
   1.2. Compute $\mathsf{a}_{\mathsf{NM}} = \mathrm{Sen}^1_{\mathsf{NM}}(\mathtt{id}, s_0; \rho)$.
   1.3. Compute $\mathsf{a}_{\mathsf{LS}} = \mathsf{P}^1(\ell; \alpha)$.
   1.4. Send $(\mathsf{a}_{\mathsf{NM}}, \mathsf{a}_{\mathsf{LS}})$ to $\mathcal{A}$.
2. Third round. Upon receiving $(\mathsf{c}_{\mathsf{NM}}, \mathsf{c}_{\mathsf{LS}}, Y)$ from $\mathcal{A}$, run as follows:
   2.1. Run in time T_f to compute y such that $Y = f(y)$.
   2.2. Set $s_1 = s_0 \oplus y$.
   2.3. Compute $(\mathsf{z}_{\mathsf{NM}}, \mathsf{dec}_{\mathsf{NM}}) = \mathrm{Sen}^2_{\mathsf{NM}}(\mathtt{id}, \mathsf{c}_{\mathsf{NM}}, s_0; \rho)$.
   2.4. Set $x = \big((\mathsf{a}_{\mathsf{NM}}, \mathsf{c}_{\mathsf{NM}}, \mathsf{z}_{\mathsf{NM}}), Y, s_1, \mathtt{com}, \mathtt{id}\big)$ and $w = (\bot, \bot, s_0, \rho)$ with $(|x| = \ell)$. Run $\mathsf{z}_{\mathsf{LS}} = \mathsf{P}^2(x, w, \mathsf{c}_{\mathsf{LS}}; \alpha)$.
   2.5. Send $(\mathsf{z}_{\mathsf{NM}}, \mathtt{com}, \mathsf{z}_{\mathsf{LS}}, s_1)$ to $\mathcal{A}$.

Right sessions: Act as a proxy between $\mathcal{A}$ and $\mathsf{MMRec}_1, \ldots, \mathsf{MMRec}_{\mathsf{poly}(\lambda)}$.

Now we are ready to describe the adversary $\mathcal{A}_{\mathsf{NI}}$ for the hiding of NI. $\mathcal{A}_{\mathsf{NI}}$ executes the following steps.

1. Let M be an empty tuple. $\mathcal{A}_{\mathsf{NI}}$ runs $\boldsymbol{\mathcal{E}_{m_b}(\varphi, \mathtt{com}, z)}$.

2. For all $i \in \{1, \ldots, \mathsf{poly}(\lambda)\}$, $\mathcal{A}_{\mathsf{NI}}$ runs the extractor of LS on the i-th right session of the execution of $\boldsymbol{\mathcal{E}_{m_b}(\varphi, \mathtt{com}, z)}$ obtaining $\tilde{m}_i$ and adds it to M.

3. Using the randomness φ, $\mathcal{A}_{\mathsf{NI}}$ reconstructs the view of $\mathcal{A}$ in the execution of $\boldsymbol{\mathcal{E}_{m_b}(\varphi, \mathtt{com}, z)}$. Use such view and M as input to $\mathcal{D}$.

4. Output what $\mathcal{D}$ outputs.

The proof ends with the observation that if $\mathcal{C}_{\mathsf{NI}}$ has committed to 0^λ then the view of $\mathcal{A}$ and the distribution of the committed messages coincide with $\mathcal{H}_3^0$, otherwise they coincide with $\mathcal{H}_3^m$. □

Common input: Security parameters: λ, $(\lambda_{\mathsf{NI}}, \lambda_{\mathsf{NM}}, \lambda_{\mathsf{LS}}, \ell) = \mathsf{Params}(\lambda)$. Identity: $\mathtt{id} \in \{0,1\}^{\lambda}$.

Internal simulation of the left session:

1. Pick $s_0 \leftarrow \{0,1\}^{\lambda}$.
2. Pick a randomness ρ, and compute $(\mathtt{dec}_{\mathsf{NM}}, \mathsf{a}_{\mathsf{NM}}) = \mathrm{Sen}^{1}_{\mathsf{NM}}(\mathtt{id}, s_0; \rho)$.
3. Pick a randomness α and compute $\mathsf{a}_{\mathsf{LS}} = \mathsf{P}^{1}(\ell; \alpha)$.
4. Send $(\mathsf{a}_{\mathsf{NM}}, \mathsf{a}_{\mathsf{LS}})$ to $\mathcal{A}$.
5. Upon receiving $(\mathsf{c}_{\mathsf{NM}}, \mathsf{c}_{\mathsf{LS}}, Y)$ from $\mathcal{A}$.

    5.1. Pick a randomness σ and compute $(\mathtt{com}, \mathtt{dec}) = \mathsf{NISen}(1^{\lambda_{\mathsf{NI}}}, 0^{\lambda}; \sigma)$.

    5.2. Pick $s_1 \leftarrow \{0,1\}^{\lambda}$.

    5.3. Compute $\mathsf{z}_{\mathsf{NM}} = \mathrm{Sen}^{2}_{\mathsf{NM}}(\mathtt{id}, \mathsf{c}_{\mathsf{NM}}, s_0; \rho)$.

    5.4. Set $x = \big((\mathsf{a}_{\mathsf{NM}}, \mathsf{c}_{\mathsf{NM}}, \mathsf{z}_{\mathsf{NM}}), Y, s_1, \mathtt{com}, \mathtt{id}\big)$ and $w = (0^{\lambda}, \sigma, \bot, \bot)$ with $(|x| = \ell)$. Run $\mathsf{z}_{\mathsf{LS}} = \mathsf{P}^{2}(x, w, \mathsf{c}_{\mathsf{LS}}; \alpha)$ where x is the theorem to be proven and w is the witness.

    5.5. Send $(\mathsf{z}_{\mathsf{NM}}, \mathtt{com}, \mathsf{z}_{\mathsf{LS}}, s_1)$ to $\mathcal{A}$.

Stand-alone commitment:

1. S acts as a proxy between $\mathcal{A}$ and MMRec_i for $i = 1, \ldots, \mathsf{poly}(\lambda)$.

Figure 3: The simulator S.

The entire security proof now is almost over because we have proved that for all $m \in \{0,1\}^{\mathsf{poly}(\lambda)}$ the following relation holds:

$$\{\mathsf{mim}^{\mathcal{A},m}_{\Pi_{\mathsf{MMCom}}}(z)\}_{z\in\{0,1\}^\star} = \{\mathsf{mim}^{\mathcal{A}}_{\mathcal{H}^m_1}(z)\}_{z\in\{0,1\}^\star} \approx \{\mathsf{mim}^{\mathcal{A}}_{\mathcal{H}^m_2}(z)\}_{z\in\{0,1\}^\star} \approx \{\mathsf{mim}^{\mathcal{A}}_{\mathcal{H}^m_3}(z)\}_{z\in\{0,1\}^\star} \approx \{\mathsf{mim}^{\mathcal{A}}_{\mathcal{H}^0_3}(z)\}_{z\in\{0,1\}^\star} \approx \{\mathsf{mim}^{\mathcal{A}}_{\mathcal{H}^0_2}(z)\}_{z\in\{0,1\}^\star} \approx \{\mathsf{mim}^{\mathcal{A}}_{\mathcal{H}^0_1}(z)\}_{z\in\{0,1\}^\star} = \{\mathsf{sim}^{S}_{\Pi_{\mathsf{MMCom}}}(1^{\lambda}, z)\}_{z\in\{0,1\}^\star}.$$

We show in Figure 3 the simulator S.

We observe that in this proof we had to consider a delayed-input version of our commitment scheme. Indeed, the sender can choose the message m to be committed by sending the non-interactive commitment $\mathtt{com}$ of the message m in the third round. It is easy to see that the same security proof still works when the non-interactive commitment is sent in the 1st round, but then clearly the delayed-input property is lost. □

4 More 3-Round Protocols Against Concurrent MiM Attacks

In this section we show how to obtain some forms of 3-round arguments of knowledge and of 3-round identification schemes that are secure against concurrent MiM attacks.

4.1 Non-Malleable WI Arguments of Knowledge

Our concurrent NM commitment scheme when instantiated without sessions ids, can be used to obtain almost directly a *commit-and-prove* AoK. Recall that in our scheme there is a non-interactive commitment $\mathtt{com}$ of m and then rest of the protocol is an AoK. This AoK is used by the sender to claim that either he knows the message committed in $\mathtt{com}$, or he committed through Π_{NM} to a share s_0 that allows to compute the solution of the puzzle.

In order to be fully compliant with the notion of commit-and-prove AoK, we just need to make a trivial change to the statement of the LS subprotocol. Given an instance $x \in L$ and a witness w the prover of our commit-and-prove AoK uses the non-interactive commitment to commit to w, and uses the rest to prove that either he knows the committed message w that moreover is a witness for $x \in L$ or again, he committed through Π_{NM} to a share s_0 that allows to compute the solution of the puzzle.

More formally, we define a commit-and-prove AoK $\mathsf{\Pi_{CaP}} = (\mathcal{P}_{\mathsf{CaP}}, \mathcal{V}_{\mathsf{CaP}})$ that corresponds to our concurrent NM commitment scheme with some minimal changes. First, $\mathcal{P}_{\mathsf{CaP}}$ and $\mathcal{V}_{\mathsf{CaP}}$ have as a common input an instance $x \in L$, where L is an NP-language. Second, $\mathcal{P}_{\mathsf{CaP}}$ has as private input w such that $(x, w) \in \mathsf{Rel_L}$. Third, $\mathcal{P}_{\mathsf{CaP}}$ runs the sender MMSen having as input w, while $\mathcal{V}_{\mathsf{CaP}}$ runs the receiver MMRec with the exception of running the LS subprotocol LS for:

$$L_{\mathsf{CaP}} = \Big\{\big(x,(a,c,z),Y,s_1,\mathtt{com},\mathtt{id}\big) : (\exists\ (w,\sigma) \text{ s.t. } \mathtt{com} = \mathsf{NISen}(w;\sigma) \ \mathtt{AND}\ (x,w) \in \mathsf{Rel_L}) \\ \mathtt{OR}\ \big(\exists(\rho,s_0) \text{ s.t. } a = \mathrm{Sen}^1_{\mathsf{NM}}(\mathtt{id},s_0;\rho)\ \mathtt{AND}\ z = \mathrm{Sen}^2_{\mathsf{NM}}(\mathtt{id},c,s_0;\rho)\ \mathtt{AND}\ Y = f(s_0 \oplus s_1)\big)\Big\}$$

that is WI for the relation

$$\mathsf{Rel}_{\mathsf{L_{CaP}}} = \Big\{\Big(\big(x,(a,c,z),Y,s_1,\mathtt{com},\mathtt{id}\big),(w,\sigma,s_0,\rho)\Big) : (\mathtt{com} = \mathsf{NISen}(w;\sigma)\ \mathtt{AND} \\ (x,w) \in \mathsf{Rel_L})\ \mathtt{OR}\ \big(a = \mathrm{Sen}^1_{\mathsf{NM}}(\mathtt{id},s_0;\rho)\ \mathtt{AND}\ z = \mathrm{Sen}^2_{\mathsf{NM}}(\mathtt{id},c,y;\rho)\ \mathtt{AND}\ Y = f(s_0 \oplus s_1)\big)\Big\}$$

We can now claim the following theorem.

Theorem 3. *Suppose there exist OWPs secure against subexponential-time adversaries, then* $\mathsf{\Pi_{CaP}}$ *is a 3-round concurrent NMWI argument of knowledge.*

Proof. The proof of this theorem is pretty straightforward given the previous proof for the concurrent non-malleability of our commitment scheme, therefore here we just point out the main intuition.

First of all, $\mathsf{\Pi_{CaP}}$ is clearly a commit-and-prove AoK. Indeed, there exists a commitment of the witness and there is an AoK proving that the committed message is a witness. In order to see this, notice that for any PPT malicious prover succeeding with non-negligible probability in proving a statement $x \in L$, the extractor of LS (of course this needs to be run against an augmented machine) would return (in expected polynomial time and with overwhelming probability) the committed witness since otherwise it would return a share s_0 that combined with s_1 allows to invert the OWP in polynomial time.

We can now focus on the concurrent NMWI property, and we can assume (by contradiction) that the adversary succeeds in encoding in the right sessions witnesses that are related to the witnesses encoded in the left sessions. Notice that the proof is almost identical to the one of Theorem 2. We can indeed prove the case of one prover and multiple verifiers (i.e., one-many), and then we can apply the fact that any one-many NMWIAoK is also a concurrent NMWIAoK. Indeed this was used in [OPV08] and follows similar arguments given in [PR05a, LPV08]. For the one-many case we can therefore follow the proof of Theorem 2 with the following trivial change. Instead of running hybrid experiments starting with a message m and ending with a message 0, in the proof of one-many concurrent NMWI we start with a witness w_0 and end with a witness w_1. Everything else remains untouched and all the reductions work directly. □

We finally notice that Π_{CaP} can be instantiated to be public-coin and delayed-input, precisely as our concurrent non-malleable commitment scheme. While what we discussed above applies to arguments only, techniques to obtain proofs can be found in [CVZ11].

Instances with just one witness and non-transferability. Recall that the definition of NMWI considers two experiments that differ only on the witness used by the prover. Therefore it is unclear which security is given by a NMWIAoK when the instance has only one witness. In order to understand the security guaranteed by Π_{CaP} in such a case, consider the proof of concurrent NMWI, and thus, in turn, consider the proof of concurrent non-malleability of our commitment scheme. Notice that while the sequence of hybrids goes from an experiment where the committed message is m to an experiment where the committed message is 0, there is an experiment $\mathcal{H}_3(\cdot, z)$ in which the committed message is irrelevant. Indeed, the entire execution is based on inverting the OWP, in encrypting it through the shares s_0 and s_1 and in using this witness in the execution of LS. This experiment can be seen as the execution of a quasi-polynomial time simulator that breaks the puzzle[18] following the approach of [Pas03][19]. Therefore following the same observations of [Pas03, Pas04] on the security offered by quasi-polynomial time simulation, our concurrent NMWIAoK even for instances with just one witness would not help the adversary in proving a statement whose witness is much harder to compute than breaking the puzzle.

The above discussion explains also the non-transferability flavor of Π_{CaP}. Indeed, at first sight, a MiM attack of an adversary $\mathcal{A}$ to an AoK should be an attempt of $\mathcal{A}$ to transfer the proof that it gets from the prover to a verifier. As such, an AoK that is secure against concurrent MiM attacks should provide some non-transferability guarantee. Since the success of $\mathcal{A}$ during a MiM attack can be replicated without a MiM attack by a quasi-polynomial time simulator, we have that Π_{CaP} guarantees non-transferability whenever computing the witnesses for the considered instances is assumed to be harder than breaking the puzzle.

Using NMWI for NMZK in the Bare Public-Key (BPK) model. In [OPV08] it is shown that a concurrent NMWIAoK Π gives directly a concurrent NMZKAoK in the BPK model. The construction is straightforward as it just consists of running Π twice, first from the verifier to the prover (proving knowledge of one out of two secrets) and then from the prover to the verifier (proving knowledge of either a witness for $x \in L$ or of one out of the two secrets of the verifier).

Our construction from Theorem 3 when combined with the construction of [OPV08] gives a candidate round-efficient concurrent NMZKAoK in the BPK model.

4.2 Identification Schemes

Identification schemes represent one of the most successful real-world applications of cryptographic protocols. We show here a 3-round identification scheme secure against concurrent MiM attacks following the concept of proving knowledge of a secret.

[18] The puzzle can be implemented through a OWP that can be inverted in quasi-polynomial time.

[19] The work of Pass did not take into account MiM attacks.

Identification schemes based on proving knowledge of a secret. The importance of this setting was for instance discussed in [COSV12] mentioning the following example. Consider a verifier $\mathcal{V}$ that provides a service to restricted group of provers $\mathcal{P}$. A malicious prover $\mathcal{P}^\star$ could give to another party B that is not part of the group, some partial information about his secret that is sufficient for B to obtain the service from $\mathcal{V}$, while still B does not know $\mathcal{P}^\star$'s secret. The paradigm of proving knowledge of a secret in an identification scheme allows to prevent attacks like the one just described. When the identification scheme consists in proving knowledge of a secret the sole fact that B convinces $\mathcal{V}$ is sufficient to claim that one can extract the whole secret from B. This implies that B obtained $\mathcal{P}^\star$'s secret corresponding to his identity, and thus B is actually $\mathcal{P}^\star$[20].

We now introduce a security definition that takes into account concurrent MiM attacks similarly to the definition CR2 (concurrent-reset on-line) of [BFGM01]. The definition of [BFGM01] also includes possible reset attacks in addition to allowing $\mathcal{A}$ to invoke multiple concurrent executions of the prover in the left sessions while $\mathcal{A}$ is interacting with the verifier. In the remaining part of this section we will ignore reset attacks since they are out of the purpose of our work. As described in [Kat02] in most network-based settings reset attacks are not an issue. Following the notation of [Kat02] we now give a formal security definitions for an identification scheme.

Definition 11. *Let $\Pi = (\mathcal{K}, \mathcal{P}, \mathcal{V})$ be a tuple of* PPT *algorithms. We say Π is an identification scheme secure against man-in-the-middle attacks if the following conditions hold:*

Correctness. *For all $(\mathsf{pk}, \mathsf{sk})$ output by $\mathcal{K}(1^\lambda)$, we have*

$$\text{Prob}\left[\ \langle \mathcal{P}(\mathsf{sk}), \mathcal{V}\rangle(\mathsf{pk}) = 1\ \right] = 1.$$

Security. *For all* PPT *adversaries $\mathcal{A}$ there exists a negligible function ν such that*

$$\text{Prob}\left[\ (\mathsf{pk}, \mathsf{sk}) \leftarrow \mathcal{K}(1^\lambda) : \langle \mathcal{A}^{\mathcal{P}(\mathsf{sk})}, \mathcal{V}\rangle(\mathsf{pk}) = 1 \ \texttt{AND} \ \ \tau \ \notin T\ \right] < \nu(\lambda),$$

where $\mathcal{A}$ has oracle access to a stateful (i.e., non-resettable) $\mathcal{P}(\mathsf{sk})$, T is defined as the transcripts set of the interactions between $\mathcal{P}(\mathsf{sk})$ and $\mathcal{A}$, and τ is defined as the transcript of one of the interactions between $\mathcal{A}$ and $\mathcal{V}$. All interactions can be arbitrarily interleaved and $\mathcal{A}$ controls the scheduling of the messages.

Identification scheme from NMWI. Our construction $\Pi_{\mathsf{ID}} = (\mathcal{K}_{\mathsf{ID}}, \mathcal{P}_{\mathsf{ID}}, \mathcal{V}_{\mathsf{ID}})$ follows the approach of [OPV08, COSV12]. Let $f : \{0,1\}^\lambda \rightarrow \{0,1\}^\lambda$ be a one-way permutation, let λ be the security parameter. The public key of $\mathcal{P}_{\mathsf{ID}}$ is the pair $(\mathsf{pk}_0, \mathsf{pk}_1)$, the secret key is sk_b for a randomly chosen bit b, such that $\mathsf{pk}_\mathsf{b} = f(\mathsf{sk}_b)$. Therefore the algorithm $\mathcal{K}_{\mathsf{ID}}$ takes as input the security parameter and outputs $((\mathsf{pk}_0, \mathsf{pk}_1), \mathsf{sk}_b)$ as described above. The protocol simply consists in $\mathcal{P}_{\mathsf{ID}}$ running our 3-round concurrent NMWIAoK Π_{CaP} with $\mathcal{V}_{\mathsf{ID}}$ to prove that it *knows* the pre-image of either pk_0 or pk_1. Formally, let $L_{\mathtt{id}}$ be the following language $L_{\mathtt{id}} = \{(y_0, y_1) : \exists\ x \ \in \{0,1\}^\lambda$ such that $y_0 = f(x) \vee y_1 = f(x)\}$, then the identification scheme consists of $\mathcal{P}_{\mathsf{ID}}$ proving the statement $(\mathsf{pk}_0, \mathsf{pk}_1) \in L_{\mathtt{id}}$ using Π_{CaP}. Fig. 4 summarizes our identification scheme. Now we can claim the following theorem.

Theorem 4. *Assume the existence of OWPs secure against subexponential-time adversaries then Π_{ID} is an identification scheme secure against concurrent MiM attacks.*

[20] This is instead not likely to happen in scenarios where the same secret key is used for other critical tasks such as signatures of any type of document.

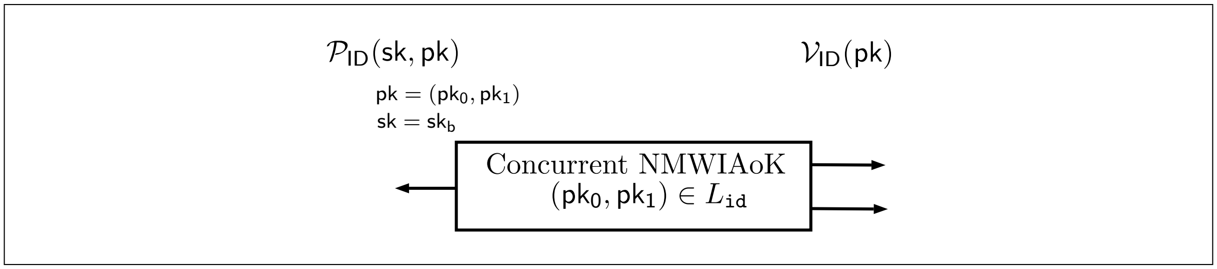
$\mathcal{P}_{\mathsf{ID}}(\mathsf{sk},\mathsf{pk})$
$\mathcal{V}_{\mathsf{ID}}(\mathsf{pk})$
$\mathsf{pk} = (\mathsf{pk}_0, \mathsf{pk}_1)$
$\mathsf{sk} = \mathsf{sk}_\mathsf{b}$
Concurrent NMWIAoK
$(\mathsf{pk}_0, \mathsf{pk}_1) \in L_{\mathtt{id}}$

Figure 4: Our 3-round identification scheme Π_{ID} from our 3-round concurrent NMWIAoK.

The proof is again straight-forward. If a PPT $\mathcal{A}$ succeeds then concurrent NMWI of Π_{CaP} guarantees that the witness that he encoded in the proof is independent of the one encoded in the proofs given by $\mathcal{P}$. Therefore by using the AoK property of Π_{CaP} we can invert f with non-negligible probability.

5 Concurrent Malleability of [GPR15]

Here, just for completeness, we briefly explain the intuition behind the fact that the 3-round NM commitment scheme $\Pi_{\mathsf{NM}} = (\mathsf{Sen}_{\mathsf{NM}}, \mathsf{Rec}_{\mathsf{NM}})$ of [GPR15] is malleable with respect to a concurrent MiM attack. In order to do this we follow the technique of [FMNV14].

We will describe a concurrent MiM adversary $\mathcal{A}$ and a distinguisher $\mathcal{D}$ that win in the non-malleability security game. We will refer to a NM commitment of the message m using the scheme Π_{NM} as $\mathsf{nmcom}(m)$. We stress that $\mathsf{nmcom}(m)$ is the result of a 3-round interaction between the sender $\mathsf{Sen}_{\mathsf{NM}}$ and the receiver $\mathsf{Rec}_{\mathsf{NM}}$. We start by describing the high-level idea of the protocol Π_{NM}. In the 1st round a left-state L is computed using a special split-state non-malleable code. Let $n = |\mathsf{L}|$. Then a non-interactive commitment $\mathtt{com}_{\mathsf{L}}$ of L is sent in the 1st round, while in the 3rd round the sender computes the right-state R corresponding to the message m and sends it in the clear. In parallel there is also a PoK of the message L committed in $\mathtt{com}_{\mathsf{L}}$. This PoK can be seen as a PoK of each bit of L. Therefore there are n PoKs where the j-th proof is used to prove knowledge of the bit L_j of L.

The actual scheme of [GPR15] is much more sophisticated than what we have just described, there are various other components but however they have no impact on the work done by our $\mathcal{A}$, so we will omit them from this short description. Essentially, we will show here that a simplified version of the scheme of [GPR15] is concurrently malleable. However all our arguments apply to their full scheme.

The proposed adversary $\mathcal{A}$ interacts with one sender $\mathsf{Sen}_{\mathsf{NM}}$ in the left session and with many receiver $\mathsf{Rec}_{\mathsf{NM}1}, \dots, \mathsf{Rec}_{\mathsf{NM}_{\mathsf{poly}(\lambda)}}$ in the right sessions. The behavior of $\mathcal{A}$ in the left and right session can be summarized as following.

Left session. $\mathsf{Sen}_{\mathsf{NM}}$ computes the 1st round of Π_{NM} as follows. First, he computes L, then he computes a perfectly binding commitment $\mathtt{com}_{\mathsf{L}}$ of L and computes n PoKs one for each bit of the message committed in $\mathtt{com}_{\mathsf{L}}$. In the last round of Π_{NM} $\mathsf{Sen}_{\mathsf{NM}}$ completes the n PoKs and sends R to $\mathcal{A}$ such that the pair (L, R) is a valid encoding of m according to the special non-malleable code. Hence in the left session $\mathcal{A}$ receives $\mathtt{com}_{\mathsf{L}}$, R and n PoKs one for each bit of the string committed in $\mathtt{com}_{\mathsf{L}}$, therefore a PoK for each bit L_j of L.

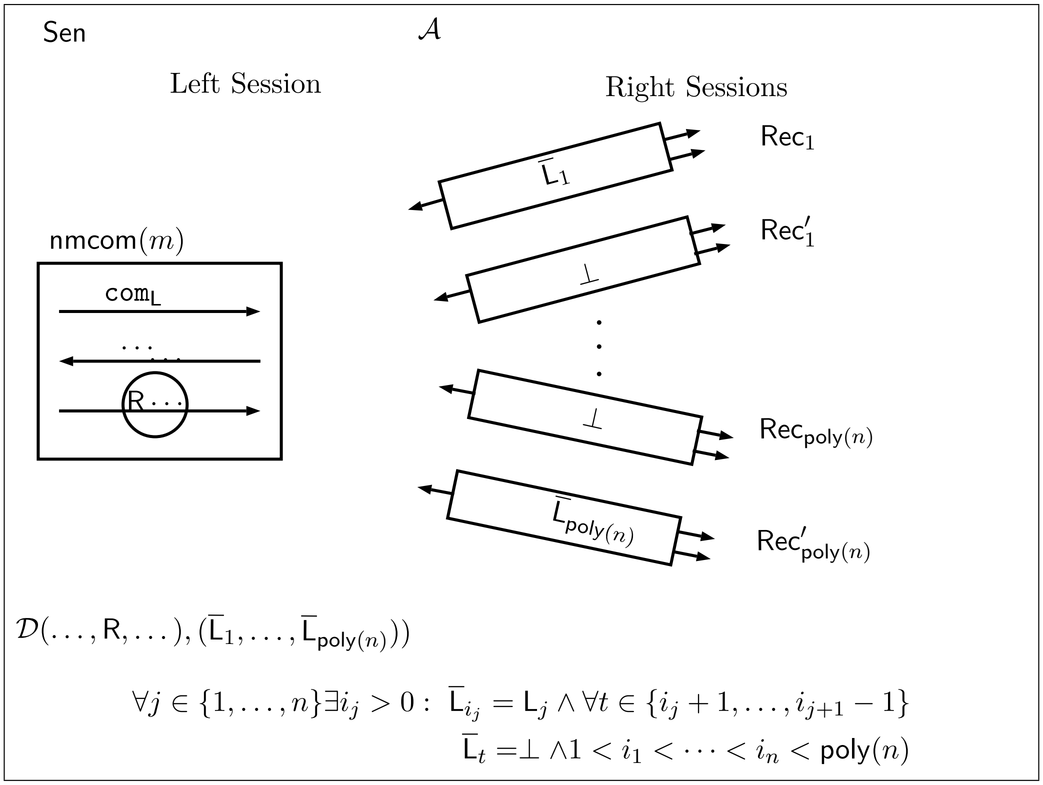
Sen
$\mathcal{A}$
Left Session
Right Sessions
$\mathsf{nmcom}(m)$
$\mathtt{com}_\mathsf{L}$
...
...
R...
$\overline{\mathsf{L}}_1$
Rec_1
$\perp$
Rec'_1
$\perp$
$\mathsf{Rec}_{\mathsf{poly}(n)}$
$\overline{\mathsf{L}}_{\mathsf{poly}(n)}$
$\mathsf{Rec}'_{\mathsf{poly}(n)}$
$\mathcal{D}(\ldots,\mathsf{R},\ldots),(\overline{\mathsf{L}}_1,\ldots,\overline{\mathsf{L}}_{\mathsf{poly}(n)}))$
$\forall j \in \{1,\ldots,n\} \exists i_j > 0 : \overline{\mathsf{L}}_{i_j} = \mathsf{L}_j \wedge \forall t \in \{i_j+1,\ldots,i_{j+1}-1\}$
$\overline{\mathsf{L}}_t = \perp \wedge 1 < i_1 < \cdots < i_n < \mathsf{poly}(n)$

Figure 5: The one-many MiM $\mathcal{A}$.

Right sessions. In the right sessions $\mathcal{A}$ interacts with $\mathsf{Rec}_{\mathsf{NM}1},\ldots,\mathsf{Rec}_{\mathsf{NM}\mathsf{poly}(\lambda)}$ mauling the commitments received on the left. More specifically, it starts $2n$ right sessions where n of them should correspond to $\mathsf{nmcom}(\mathsf{L}_1),\ldots,\mathsf{nmcom}(\mathsf{L}_n)$ such that $\mathsf{L} = \mathsf{L}_1 \ldots \mathsf{L}_n$, and the other n sessions should correspond to invalid commitments (we refer to such commitments as $\mathsf{nmcom}(\perp)$).

More precisely, our adversary computes, for each bit L_j of L, two NM commitments $\mathsf{nmcom}(1^\lambda)$, $\mathsf{nmcom}(0^\lambda)$ such that if $\mathsf{L}_j = 1$ then $\mathsf{nmcom}(0^\lambda)$ is invalid, otherwise $\mathsf{nmcom}(1^\lambda)$ is invalid. In order to poison one out of $\mathsf{nmcom}(0^\lambda)$ and $\mathsf{nmcom}(1^\lambda)$, $\mathcal{A}$ will rely on the PoK of L_j received on the left. The PoK of L_j will be plugged in the PoKs of $\mathsf{nmcom}(0^\lambda)$ and in the PoKs of $\mathsf{nmcom}(1^\lambda)$. More precisely one of the n PoKs of $\mathsf{nmcom}(0^\lambda)$ that correspond to a PoK of the bit 0 will be replaced with the PoK of L_j. The same approach is applied when $\mathcal{A}$ computes $\mathsf{nmcom}(1^\lambda)$ with the only difference that the PoK that $\mathcal{A}$ will replace corresponds to a PoK of a bit 1. In this way only one out of $\mathsf{nmcom}(0^\lambda)$ and $\mathsf{nmcom}(1^\lambda)$ still remain a valid commitment. In particular $\mathsf{nmcom}(\mathsf{L}_j)$ will remain a valid commitment while $\mathsf{nmcom}(1 - \mathsf{L}_j)$ will be poisoned and thus will correspond to an invalid commitment.

There is however a subtlety. Since the PoK played on the right is for one component copied from the PoK played on the left, it can be completed successfully with constant probability and the adversary has to abort the session if it can not complete the PoK. Therefore each of the above $2n$ right sessions could be repeated multiple times, but however the total amount of right sessions will still be polynomial in the security parameter.

Finally our distinguisher $\mathcal{D}$ given as input the committed bits $\mathsf{L}_1,\ldots,\mathsf{L}_n$ and R contained in the view of $\mathcal{A}$, can easily recover the message m committed in the left interaction.

6 Acknowledgments

We thank Vipul Goyal, and Silas Richelson for several remarkable discussions on [GPR16].

Research supported in part by "GNCS - INdAM", EU COST Action IC1306, NSF grants 1065276, 1118126 and 1136174, US-Israel BSF grant 2008411, OKAWA Foundation Research Award, IBM Faculty Research Award, Xerox Faculty Research Award, B. John Garrick Foundation Award, Teradata Research Award, and Lockheed-Martin Corporation Research Award. This material is based upon work supported in part by DARPA Safeware program. The views expressed are those of the authors and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

The work of the 1st, 3rd and 4th authors has been done in part while visiting UCLA.

References

[Bar02] Boaz Barak. Constant-round coin-tossing with a man in the middle or realizing the shared random string model. In *43rd Symposium on Foundations of Computer Science (FOCS 2002), 16-19 November 2002, Vancouver, BC, Canada, Proceedings*, pages 345–355, 2002. (Cited on page 3.)

[BFGM01] Mihir Bellare, Marc Fischlin, Shafi Goldwasser, and Silvio Micali. Identification protocols secure against reset attacks. In *Advances in Cryptology - EUROCRYPT 2001, International Conference on the Theory and Application of Cryptographic Techniques, Innsbruck, Austria, May 6-10, 2001, Proceeding*, pages 495–511, 2001. (Cited on pages 4, 5, and 28.)

[BGR$^+$15] Hai Brenner, Vipul Goyal, Silas Richelson, Alon Rosen, and Margarita Vald. Fast non-malleable commitments. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, Denver, CO, USA, October 12-6, 2015*, pages 1048–1057, 2015. (Cited on page 3.)

[Blu86] Manuel Blum. How to prove a theorem so no one else can claim it. In *In Proceedings of the International Congress of Mathematicians*, page 444451, 1986. (Cited on page 4.)

[BPS06] Boaz Barak, Manoj Prabhakaran, and Amit Sahai. Concurrent non-malleable zero knowledge. In *47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006), 21-24 October 2006, Berkeley, California, USA, Proceedings*, pages 345–354, 2006. (Cited on page 5.)

[BPSV08] Carlo Blundo, Giuseppe Persiano, Ahmad-Reza Sadeghi, and Ivan Visconti. Improved security notions and protocols for non-transferable identification. In *Computer Security - ESORICS 2008, 13th European Symposium on Research in Computer Security, Málaga, Spain, October 6-8, 2008. Proceedings*, volume 5283 of *Lecture Notes in Computer Science*, pages 364–378. Springer, 2008. (Cited on page 6.)

[BR93] Mihir Bellare and Phillip Rogaway. Entity authentication and key distribution. In *Advances in Cryptology - CRYPTO '93, 13th Annual International Cryptology Conference, Santa Barbara, California, USA, August 22-26, 1993, Proceedings*, pages 232–249, 1993. (Cited on page 6.)

[CDS94] Ronald Cramer, Ivan Damgård, and Berry Schoenmakers. Proofs of partial knowledge and simplified design of witness hiding protocols. In YvoG. Desmedt, editor, *Advances in Cryptology — CRYPTO '94*, volume 839 of *Lecture Notes in Computer Science*, pages 174–187. Springer Berlin Heidelberg, 1994. (Cited on page 4.)

[CGGM00] Ran Canetti, Oded Goldreich, Shafi Goldwasser, and Silvio Micali. Resettable zero-knowledge (extended abstract). In *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing, May 21-23, 2000, Portland, OR, USA*, pages 235–244, 2000. (Cited on page 5.)

[CKPR01] Ran Canetti, Joe Kilian, Erez Petrank, and Alon Rosen. Black-box concurrent zero-knowledge requires omega˜(log n) rounds. In *Proceedings on 33rd Annual ACM Symposium on Theory of Computing, July 6-8, 2001, Heraklion, Crete, Greece*, pages 570–579, 2001. (Cited on page 5.)

[COSV12] Chongwon Cho, Rafail Ostrovsky, Alessandra Scafuro, and Ivan Visconti. Simultaneously resettable arguments of knowledge. In *Theory of Cryptography - 9th Theory of Cryptography Conference, TCC 2012, Taormina, Sicily, Italy, March 19-21, 2012. Proceedings*, pages 530–547, 2012. (Cited on pages 6 and 28.)

[COSV16] Michele Ciampi, Rafail Ostrovsky, Luisa Siniscalchi, and Ivan Visconti. On round-efficient non-malleable protocols. Unpublished manuscript, 2016. (Cited on pages 3 and 4.)

[CPS$^+$16a] Michele Ciampi, Giuseppe Persiano, Alessandra Scafuro, Luisa Siniscalchi, and Ivan Visconti. Improved or-composition of sigma-protocols. In *Theory of Cryptography - 13th International Conference, TCC 2016-A, Tel Aviv, Israel, January 10-13, 2016, Proceedings, Part II*, pages 112–141, 2016. (Cited on page 4.)

[CPS$^+$16b] Michele Ciampi, Giuseppe Persiano, Alessandra Scafuro, Luisa Siniscalchi, and Ivan Visconti. Online/offline OR composition of sigma protocols. In *Advances in Cryptology - EUROCRYPT 2016 - 35th Annual Internati onal Conference on the Theory and Applications of Cryptographic Techni ques, Vienna, Austria, May 8-12, 2016, Proceedings, Part II*, pages 63–92, 2016. (Cited on pages 4 and 7.)

[CVZ10] Zhenfu Cao, Ivan Visconti, and Zongyang Zhang. Constant-round concurrent non-malleable statistically binding commitments and decommitments. In *Public Key Cryptography - PKC 2010, 13th International Conference on Practice and Theory in Public Key Cryptography, Paris, France, May 26-28, 2010. Proceedings*, volume 6056 of *Lecture Notes in Computer Science*, pages 193–208. Springer, 2010. (Cited on page 10.)

[CVZ11] Zhenfu Cao, Ivan Visconti, and Zongyang Zhang. On constant-round concurrent non-malleable proof systems. *Inf. Process. Lett.*, 111(18):883–890, 2011. (Cited on page 27.)

[Dam10] Ivan Damgård. On Σ-protocol. http://www.cs.au.dk/~ivan/Sigma.pdf, 2010. (Cited on page 4.)

[DDN91] Danny Dolev, Cynthia Dwork, and Moni Naor. Non-malleable cryptography (extended abstract). In *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing, May 5-8, 1991, New Orleans, Louisiana, USA*, pages 542–552, 1991. (Cited on page 3.)

[DG03] Ivan Damgård and Jens Groth. Non-interactive and reusable non-malleable commitment schemes. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing, June 9-11, 2003, San Diego, CA, USA*, pages 426–437, 2003. (Cited on page 3.)

[DPV04] Giovanni Di Crescenzo, Giuseppe Persiano, and Ivan Visconti. Constant-round resettable zero knowledge with concurrent soundness in the bare public-key model. In Matthew K. Franklin, editor, *Advances in Cryptology - CRYPTO 2004, 24th Annual International CryptologyConference, Santa Barbara, California, USA, August 15-19, 2004, Proceedings*, volume 3152 of *Lecture Notes in Computer Science*, pages 237–253. Springer, 2004. (Cited on page 4.)

[FFS87] Uriel Feige, Amos Fiat, and Adi Shamir. Zero knowledge proofs of identity. In *Proceedings of the 19th Annual ACM Symposium on Theory of Computing, 1987, New York, New York, USA*, pages 210–217, 1987. (Cited on page 4.)

[FMNV14] Sebastian Faust, Pratyay Mukherjee, Jesper Buus Nielsen, and Daniele Venturi. Continuous non-malleable codes. In *Theory of Cryptography - 11th Theory of Cryptography Conference, TCC 2014, San Diego, CA, USA, February 24-26, 2014. Proceedings*, volume 8349 of *Lecture Notes in Computer Science*, pages 465–488. Springer, 2014. (Cited on pages 5 and 29.)

[GL89] Oded Goldreich and Leonid A. Levin. A hard-core predicate for all one-way functions. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing, May 14-17, 1989, Seattle, Washigton, USA*, pages 25–32, 1989. (Cited on page 3.)

[GLOV12] Vipul Goyal, Chen-Kuei Lee, Rafail Ostrovsky, and Ivan Visconti. Constructing non-malleable commitments: A black-box approach. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012*, pages 51–60, 2012. (Cited on pages 3 and 9.)

[GMPP16] Sanjam Garg, Pratyay Mukherjee, Omkant Pandey, and Antigoni Polychroniadou. The exact round complexity of secure computation. In *Advances in Cryptology - EUROCRYPT 2016 - 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8-12, 2016, Proceedings, Part II*, pages 448–476, 2016. (Cited on page 4.)

[Goy11] Vipul Goyal. Constant round non-malleable protocols using one way functions. In *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pages 695–704, 2011. (Cited on page 3.)

[GPR15] Vipul Goyal, Omkant Pandey, and Silas Richelson. Textbook non-malleable commitments. *IACR Cryptology ePrint Archive*, 2015, 2015. Available at http://eprint.iacr.org/2015/1178 Version 20151210:144729 (posted 10-Dec-2015 14:47:29 UTC). (Cited on pages 2, 5, 11, and 29.)

[GPR16] Vipul Goyal, Omkant Pandey, and Silas Richelson. Textbook non-malleable commitments. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing,STOC 2016, Cambridge, MA, USA, June 19 - June 21, 2016*, 2016. (Cited on pages 1, 3, 5, 11, and 31.)

[GRRV14] Vipul Goyal, Silas Richelson, Alon Rosen, and Margarita Vald. An algebraic approach to non-malleability. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 41–50, 2014. (Cited on page 3.)

[HM96] Shai Halevi and Silvio Micali. Practical and provably-secure commitment schemes from collision-free hashing. In *Advances in Cryptology - CRYPTO '96, 16th Annual International Cryptology Conference, Santa Barbara, California, USA, August 18-22, 1996, Proceedings*, pages 201–215, 1996. (Cited on page 3.)

[HV16] Carmit Hazay and Muthuramakrishnan Venkitasubramaniam. On the power of secure two-party computation. Cryptology ePrint Archive, Report 2016/074, 2016. http://eprint.iacr.org/. (Cited on page 4.)

[Kat02] Jonathan Katz. *Efficient Cryptographic Protocols Preventing "Man-in-the-Middle" Attacks*. PhD thesis, Columbia University, 2002. (Cited on pages 6 and 28.)

[KO04] Jonathan Katz and Rafail Ostrovsky. Round-optimal secure two-party computation. In *Advances in Cryptology - CRYPTO 2004, 24th Annual International Cryptology-Conference, Santa Barbara, California, USA, August 15-19, 2004, Proceedings*, pages 335–354, 2004. (Cited on page 4.)

[Lin10] Yehuda Lindell. Foundations of cryptography 89-856. http://u.cs.biu.ac.il/~lindell/89-856/complete-89-856.pdf, 2010. (Cited on page 9.)

[LP11] Huijia Lin and Rafael Pass. Constant-round non-malleable commitments from any one-way function. In *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pages 705–714, 2011. (Cited on pages 3, 8, and 16.)

[LP15] Huijia Lin and Rafael Pass. Constant-round nonmalleable commitments from any one-way function. *J. ACM*, 62(1):5:1–5:30, 2015. (Cited on page 3.)

[LPV08] Huijia Lin, Rafael Pass, and Muthuramakrishnan Venkitasubramaniam. Concurrent non-malleable commitments from any one-way function. In *Theory of Cryptography, Fifth Theory of Cryptography Conference, TCC 2008, New York, USA, March 19-21, 2008.*, pages 571–588, 2008. (Cited on pages 5, 10, and 26.)

[LPV09] Huijia Lin, Rafael Pass, and Muthuramakrishnan Venkitasubramaniam. A unified framework for concurrent security: universal composability from stand-alone non-malleability. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing,STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 179–188, 2009. (Cited on page 5.)

[LS90] Dror Lapidot and Adi Shamir. Publicly verifiable non-interactive zero-knowledge proofs. In *Advances in Cryptology - CRYPTO*, 1990. (Cited on pages 4, 8, and 11.)

[MR01] Silvio Micali and Leonid Reyzin. Soundness in the public-key model. In *Advances in Cryptology - CRYPTO 2001, 21st Annual International Cryptology Conference, Santa Barbara, California, USA, August 19-23, 2001, Proceedings*, pages 542–565, 2001. (Cited on page 5.)

[MV16] Arno Mittelbach and Daniele Venturi. Fiat-shamir for highly sound protocols is instantiable. Cryptology ePrint Archive, Report 2016/313, 2016. http://eprint.iacr.org/. (Cited on page 4.)

[Nao91] Moni Naor. Bit commitment using pseudorandomness. *J. Cryptology*, 4(2):151–158, 1991. (Cited on page 3.)

[NY89] Moni Naor and Moti Yung. Universal one-way hash functions and their cryptographic applications. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing, May 14-17, 1989, Seattle, Washigton, USA*, pages 33–43, 1989. (Cited on page 3.)

[OPV08] Rafail Ostrovsky, Giuseppe Persiano, and Ivan Visconti. Constant-round concurrent non-malleable zero knowledge in the bare public-key model. In *Automata, Languages and Programming, 35th International Colloquium, ICALP 2008, Reykjavik, Iceland, July 7-11, 2008, Proceedings, Part II - Track B: Logic, Semantics, and Theory of Programming & Track C: Security and Cryptography Foundations*, pages 548–559, 2008. (Cited on pages 5, 6, 11, 12, 26, 27, and 28.)

[OPV09] Rafail Ostrovsky, Giuseppe Persiano, and Ivan Visconti. Simulation-based concurrent non-malleable commitments and decommitments. In *Theory of Cryptography, 6th Theory of Cryptography Conference, TCC 2009, San Francisco, CA, USA, March 15-17, 2009. Proceedings*, volume 5444 of *Lecture Notes in Computer Science*, pages 91–108. Springer, 2009. (Cited on page 10.)

[ORSV13] Rafail Ostrovsky, Vanishree Rao, Alessandra Scafuro, and Ivan Visconti. Revisiting lower and upper bounds for selective decommitments. In *TCC*, pages 559–578, 2013. (Cited on page 5.)

[OV12] Rafail Ostrovsky and Ivan Visconti. Simultaneous resettability from collision resistance. *Electronic Colloquium on Computational Complexity (ECCC)*, 19:164, 2012. (Cited on page 4.)

[Pas03] Rafael Pass. Simulation in quasi-polynomial time, and its application to protocol composition. In Eli Biham, editor, *Advances in Cryptology - EUROCRYPT 2003, International Conference on the Theory and Applications of Cryptographic Techniques, Warsaw, Poland, May 4-8, 2003, Proceedings*, volume 2656 of *Lecture Notes in Computer Science*, pages 160–176. Springer, 2003. (Cited on page 27.)

[Pas04] Rafael Pass. Bounded-concurrent secure multi-party computation with a dishonest majority. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing, Chicago, IL, USA, June 13-16, 2004*, pages 232–241, 2004. (Cited on page 27.)

[Pas13] Rafael Pass. Unprovable security of perfect NIZK and non-interactive non-malleable commitments. In *TCC*, pages 334–354, 2013. (Cited on pages 1, 3, and 6.)

[Ped91] Torben P. Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In *Advances in Cryptology - CRYPTO '91, 11th Annual International Cryptology Conference, Santa Barbara, California, USA, August 11-15, 1991, Proceedings*, pages 129–140, 1991. (Cited on page 3.)

[PPV08] Omkant Pandey, Rafael Pass, and Vinod Vaikuntanathan. Adaptive one-way functions and applications. In *Advances in Cryptology - CRYPTO 2008, 28th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2008. Proceedings*, pages 57–74, 2008. (Cited on page 3.)

[PR05a] Rafael Pass and Alon Rosen. Concurrent non-malleable commitments. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2005), 23-25 October 2005, Pittsburgh, PA, USA, Proceedings*, pages 563–572, 2005. (Cited on pages 3 and 26.)

[PR05b] Rafael Pass and Alon Rosen. New and improved constructions of non-malleable cryptographic protocols. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, pages 533–542, 2005. (Cited on page 3.)

[PR08] Rafael Pass and Alon Rosen. Concurrent nonmalleable commitments. *SIAM J. Comput.*, 37(6):1891–1925, 2008. (Cited on page 3.)

[PW09] Rafael Pass and Hoeteck Wee. Black-box constructions of two-party protocols from one-way functions. In *Theory of Cryptography, 6th Theory of Cryptography Conference, TCC 2009, San Francisco, CA, USA, March 15-17, 2009. Proceedings*, pages 403–418, 2009. (Cited on page 9.)

[PW10] Rafael Pass and Hoeteck Wee. Constant-round non-malleable commitments from subexponential one-way functions. In *Advances in Cryptology - EUROCRYPT 2010, 29th Annual International Conference on the Theory and Applications of Cryptographic Techniques, French Riviera, May 30 - June 3, 2010. Proceedings*, pages 638–655, 2010. (Cited on pages 3 and 14.)

[SV12] Alessandra Scafuro and Ivan Visconti. On round-optimal zero knowledge in the bare public-key model. In *Advances in Cryptology - EUROCRYPT 2012 - 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings*, pages 153–171, 2012. (Cited on page 5.)

[Wee10] Hoeteck Wee. Black-box, round-efficient secure computation via non-malleability amplification. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*, pages 531–540. IEEE Computer Society, 2010. (Cited on page 4.)

[YYZ10] Andrew C. Yao, Moti Yung, and Yunlei Zhao. Adaptive concurrent non-malleability with bare public-keys. Cryptology ePrint Archive, Report 2010/107, 2010. `http://eprint.iacr.org/`. (Cited on page 4.)

[YZ07] Moti Yung and Yunlei Zhao. Generic and practical resettable zero-knowledge in the bare public-key model. In *Advances in Cryptology - EUROCRYPT 2007, 26th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Barcelona, Spain, May 20-24, 2007, Proceedings*, pages 129–147, 2007. (Cited on page 4.)