4-Round Concurrent Non-Malleable Commitments from One-Way Functions

MICHELE CIAMPI
DIEM, Università di Salerno
ITALY
mciampi@unisa.it

RAFAIL OSTROVSKY
UCLA
USA
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

How many rounds and which computational assumptions are needed for concurrent non-malleable commitments? The above question has puzzled researchers for several years.

Recently, Pass in [TCC 2013] proved a lower bound of 3 rounds when security is proven through black-box reductions to falsifiable assumptions. On the other side, positive results of Lin and Pass [STOC 2011], Goyal [STOC 2011] and Goyal et al. [FOCS 2012] showed that one-way functions are sufficient with a constant (at least 6) number of rounds. More recenly Ciampi et al. [CRYPTO 2016] showed that subexponentially strong one-way permutations are sufficient with just 3 rounds.

In this work we almost close the above open question by showing a 4-round concurrent non-malleable commitment scheme that only needs one-way functions. Our main technique consists in showing how to upgrade basic forms of non-malleability (i.e., non-malleability w.r.t. non-aborting adversaries) to full-fledged non-malleability without penalizing the round complexity.

Contents

1 Introduction

Commitment schemes and zero-knowledge argument systems are fundamental primitives in Cryptography. Here we consider the intriguing question of constructing round-efficient schemes that remain secure even against man-in-the-middle (MiM) attacks: non-malleable (NM) commitments and NM zero-knowledge (NMZK) argument systems [DDN91].

Non-malleable commitments. The round complexity of commitment schemes in the stand-alone setting is well understood. Non-interactive commitments exist assuming the existence of one-to-one one-way functions [GL89], and 2-round commitments exist assuming the existence of one-way functions (OWFs) only. Moreover non-interactive commitments do not exist if one relies on the black-box use of OWFs only [MP12].

Instead, the round complexity of NM commitments[1] after 25 years of research remains a fascinating open question, in particular when taking into account the required computational assumptions. The original construction of [DDN91] required a logarithmic number of rounds and the sole use of OWFs. Then, through a long sequence of very exciting positive results [Bar02, PR03, PR05b, PR05a, PR08b, PR08a, LPV08, PW10, Wee10, LP11b, LP15, Goy11, GLOV12], the above open question has been in part solved obtaining a constant[2]-round (even concurrent) NM commitment scheme by using any OWF in a black-box fashion. On the negative side, Pass recently proved that NM commitments require at least 3 rounds [Pas13][3] when security is proved through a black-box reduction to falsifiable (polynomial or subexponential time) hardness assumptions.

The basic case where the adversary plays only with one sender and one receiver has been addressed by Goyal et al. [GRRV14, GRRV16] that showed a *one-one* 4-round NM commitment scheme based on OWFs only[4]. A recent breakthrough of Goyal et al. [GPR15, GPR16] exploited the use of the NM codes in the split-state model of Aggarwal et al. [ADL14] to show a 3-round one-one NM commitment scheme based on the black-box use of any one-to-one OWF[5].

A new result of Ciampi et al. [COSV16] obtains concurrent non-malleability in 3 rounds but their security proof relies on the existence of one-way permutations (OWPs) secure against subexponential-time adversaries[6].

Non-malleable zero knowledge. The progress on NMZK arguments has tightly followed advances on NM commitments[7]. The construction of [GRRV14, GRRV16] has closed this line of

---

[1]In this paper we will consider only NM commitments w.r.t. commitments. For the case of NM w.r.t. decommitments see [PR05b, PR08b, OPV09, CVZ10, DMRV13].

[2]The construction of [GLOV12] can be squeezed to 6 rounds (see [GRRV14]).

[3]If instead one relies on non-standard assumptions or trusted setups (e.g., using trusted parameters, working in the random oracle model, relying on the existence of NM OWFs) then there exist non-interactive NM commitments [DG03, PPV08].

[4][GRRV14] claims concurrent non-malleability. Later on we have found an error in the security proof (that also applies to the construction of [BGR+15] and the claim on concurrent non-malleability has then been withdrawn in [GRRV16] where a new scheme is presented and proved one-one non-malleable.

[5]In [COSV16] Ciampi et al. showed an explicit successful concurrent man-in-the-middle for a preliminary version of their protocol [GPR15].

[6]Hardness assumptions against subexponential-time adversaries were already used [PR03, PW10, Wee10] to improve the round-complexity of NM commitments.

[7]For NMZK we will omit the case of polynomially many concurrent sessions since there is a logarithmic lower bound on the round complexity of concurrent NMZK [CKPR01] with black-box simulation. While this lower bound has been matched by Barak et al. [BPS06], no non-black-box construction is known with sub-logarithmic rounds.

research since they showed a 4-round NMZK argument of knowledge (AoK) that relies on the existence of OWFs only. This result is clearly optimal both in round complexity and in computational assumptions.

Interestingly, several years after the 4-round zero knowledge (ZK) argument system from OWFs of [BJY97], the same optimal round complexity and optimal complexity assumptions have been shown sufficient for NMZK [GRRV14, GRRV16] and resettably sound ZK [COP+14].

Delayed-input protocols. In [LS90] Lapidot and Shamir showed a 3-round witness-indistinguishable (WI) proof of knowledge (PoK) for $\mathcal{NP}$ where the instance (except its length) and the witness are not needed before playing the last round. This "delayed-input" form of completeness has been critically used in the past (e.g., [KO04, DPV04a, YZ07, Wee10]) and very recently (e.g., [CPS+16a, CPS+16b, GMPP16, COSV16, HV16a, MV16]), since it often helps in improving the round complexity of an external protocol.

1.1 Our Results

In this paper we show techniques that starting with basic non-malleability features allow us to obtain full-fledged non-malleability. By relying on such techniques, we show the first 4-round *concurrent* non-malleable commitment scheme under standard assumptions. Our commitment scheme relies on the minimal assumption of the existence of one-way functions. In light of the lower bound of Pass [Pas13], our work nearly closes the long-standing open questions of the round and computational complexities of concurrent non-malleable commitments. We stress that when just relying on one-way functions the construction of [GPR16] needs 4 rounds and achieves one-one NM only.

Our approach: non-malleability upgrades. Previous constructions for NM commitments and NMZK arguments are usually complicated and start from basic and extremely malleable tools likes regular commitment schemes and zero-knowledge proofs. Obtaining non-malleability from such basic building blocks is a well known complicated task and achieving it in a round-efficient way has always been a major challenge. Given the above indisputable difficulties, security proofs are usually very non-trivial and can be difficult to study and re-use. Here we take a different approach that consists of starting with a very basic and limited form of non-malleability, to then upgrade it to the desired notions.

Informally, we say that a commitment scheme is weak non-malleable if it is non-malleable w.r.t. adversaries that never commit to ⊥ when receiving honestly computed commitments. Moreover, we say that an adversary is synchronous if all (i.e., both left and right) sessions are played in parallel. Clearly the design of a synchronous weak one-one NM commitment scheme can be an easier task and schemes with such limited non-malleability guarantees might exist with improved round complexity, efficiency and complexity assumptions compared to previous work achieving full-fledged non-malleability. Last but not least, the security proof of a synchronous weak NM protocol is potentially much simpler to write in a robust and easy to read way than security proofs for full-fledged non-malleability.

We will sometimes assume that the last round of the receiver of the underlying (limited) NM commitment scheme (i.e., the weak non-malleable commitment scheme that we use as subprotocol) be simulatable without having the private coins used to compute the previous message of the receiver. This constraint clearly disappears when considering a 3-round protocol, and is clearly

satisfied by any public-coin protocol. Recent work on NM commitments includes 3-round and 4-round constructions that are also public coin and can be used as subprotocols in our constructions.

Improved state-of-the art. In this work we provide the following main result[8].

Starting with any 4-round public-coin weak one-many NM commitment scheme Π, we show how to obtain a 4-round concurrent NM commitment scheme by only requiring one-way functions. As a consequence of our result, by starting with the preliminary protocol Π_4 of [GRRV14] (i.e., their basic construction without the zero-knowledge argument of knowledge) we obtain 4-round concurrent NM commitments from OWFs. A comparison with positive results in the state of the art can be found in Table 1.

Paper	No. Rounds	Assumption	Concurrency
Goyal, STOC 2011	≥ 6	OWFs	Yes
Line and Pass, STOC 2011	≥ 6	OWFs	Yes
Goyal et al., FOCS 2012	≥ 6	BB OWFs	Yes
Goyal et al., FOCS 2014/ePrint 2016	4	OWFs	No
Goyal, Pandey and Richelson, STOC 2016	3	BB OWPs	No
Ciampi et al., CRYPTO 2016	3	subexp. OWPs	Yes
This work	4	OWFs	Yes

Table 1: Comparison with recent positive results.

We also show the following results on non-malleable zero knowledge. Given a 4-round public-coin one-one NM extractable commitment scheme Π, we show how to obtain a 4-round delayed-input NMZK AoK assuming CRHFs. By using the construction Π_3 of [GPR16] (in 4 rounds it needs OWFs only) we obtain a 4-round delayed-input NMZK AoK from CRHFs.

We finally give the following result on 3-round non-malleable commitments. Given a 3-round synchronous weak one-one NM commitment scheme Π, we show how to obtain a 3-round extractable one-one NM commitment scheme Π' assuming OWPs secure against subexponential-time adversaries. By using either Π_3 or Π_4 (both are based on OWPs when implemented in 3 rounds) we obtain a 3-round extractable one-one NM commitment scheme Π'. Notice that the compiler of [COSV16] on input Π' gives 3-round concurrent NM commitments from subexponentially strong OWPs. Our result improves the instantiability of the compiler of [COSV16] since we can obtain 3-round concurrent NM commitments admitting one more[9] candidate as underlying (limited) one-one NM commitment scheme.

Remark 1: on the need of public-coin protocols. We will sometimes require the underlying protocols to be public coin because in a reduction we will have to simulate the last round of the receiver without knowing the randomness he used to compute the previous round. Of course the public-coin property satisfies the above requirement and moreover the constructions [GRRV14, GPR16] that we use as subprotocol are public-coin protocols. It is straightforward to notice that any 3-round protocol would also let us conclude successfully the reduction since there is no previous round played by the receiver. Just for simplicity we state our theorems requiring the public-coin property.

---

[8]When considering a 4-round commitment scheme we always assume that the sender plays the 4th round.
[9]COSV16 can only used Π_3.

1.2 Technical Overview

Non-malleability upgrades. In contrast to previous work, we take a different approach to non-malleability. Our goal is to start with a commitment scheme Π that enjoys some partial non-malleability features only. For instance, we assume that the initial scheme is non-malleable in case the adversary never commits to $\bot$ when receiving a well formed commitment. Also we consider a limitation on the scheduling of the messages, requiring that the adversary be synchronous. We notice that this type of commitment scheme corresponds to the initial subprotocol given in [GRRV14] as well as the first subprotocol given in [GPR16]. We also require some subprotocols to be public coin (see Remark 1 for further details).

Next we show how to upgrade these limited forms of non-malleability to the desired non-malleability. Our proof approach makes use of a common structure in our protocols consisting of a subprotocol useful to extract a trapdoor from the adversary. We implement this extraction by assuming standard polynomial-time OWFs and CRHFs when 3 rounds are available to get the trapdoor (we will extract two signatures from the adversary under the same public key, following previous ideas of [DPV04b, GJO⁺13, CPS13, COP⁺14]), and by assuming subexponentially strong OWPs (similarly to [COSV16]) otherwise (indeed in this case there are only 2 rounds, therefore rewinds are useless and instead we will invert through brute-force search an element in the range of a OWP sent by the adversary).

We will make use of a delayed-input PoK where the prover proves knowledge of either a well formed message/randomness pair certifying the correctness of the execution of Π (in the case of NMZK this PoK also proves that the message is a witness) or of signatures of messages. Interestingly, through the combined use of special trapdoor commitments and special honest-verifier zero-knowledge delayed-input proofs of knowledge we can avoid issues related to an adversary that maliciously commits both to valid messages and to $\bot$ in polynomially many concurrent sessions[10].

Delayed-input NMZK. The above discussion is not sufficient for *delayed-input* NMZK. The reason is that the commitment scheme Π could require the message to commit before the last round. We address this point by instead using Π to commit to a random string s_0 and in sending in the last round a string s_1 such that $w = s_0 \oplus s_1$ is a witness. This technique was introduced in [COSV16] to obtain delayed-input NM commitments.

Efficiency. The above description of our constructions seems to indicate that our results are interesting only from a theoretical point of view, mainly because of the $\mathcal{NP}$ reductions required by the delayed-input WIPoK [LS90]. However we point out that depending on the existence of an efficient Σ-protocol for the weak/synchronous NM commitment schemes used in our constructions, the new OR-composition technique of Σ-protocols of [CPS⁺16a] could replace [LS90] avoiding $\mathcal{NP}$ reductions. Tweaking and instantiating properly our compilers for a practical scheme can be the subject of future work.

---

[10]We have designed successful aborting strategies for the one-one non-malleability of [GRRV14, BGR⁺15] and for the concurrent non-malleability of [GRRV16]. More precisely, our adversaries crucially rely on completing some commitments that correspond to $\bot$ and that therefore can not be opened anymore, therefore the adversary is implicitly aborting. In other contexts, other different abort attacks have been shown recently in literature [SV12, ORSV13, HPV15] and this makes somewhat evident that abort attacks can be very subtle and tough to detect even in the provable security framework. Recently the case of non-aborting adversaries has been considered in [HV16b], and this can be related to our notion of weak non-malleable commitments, where the adversary is not allowed to commit to $\bot$.

2 Notation and Non-Malleability Definitions

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, $x \leftarrow Q$ sampling of x from Q with uniform distribution. We use the abbreviation PPT that stays for probabilistic polynomial time. We use $\mathsf{poly}(\cdot)$ to indicate a generic polynomial function.

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 $(x, w) \in$ Rel, we call x the *instance* and w a *witness* for x. For a polynomial-time relation Rel, we define the $\mathcal{NP}$-language L_{Rel} as $L_{\mathsf{Rel}} = \{x | \exists w : (x, w) \in \mathsf{Rel}\}$. Analogously, unless otherwise specified, for an $\mathcal{NP}$-language L we denote by Rel_L the corresponding polynomial-time relation (that is, Rel_L is such that $L = L_{\mathsf{Rel}_L}$). We denote by $\hat{L}$ the language that includes both L and all well formed instances that do not have a witness. Moreover we require that membership in $\hat{L}$ can be tested in polynomial time. We implicitly assume that a PPT algorithm that is supposed to receive an instance in $\hat{L}$ will abort immediately if the instance does not belong to $\hat{L}$.

Let A and B be two interactive probabilistic algorithms. 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 (resp. B) 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.

Standard definitions and their variants w.r.t. subexponential-time adversaries can be found in App. A. We will say that a complexity assumption is $\tilde{T}$-breakable if it can be broken with overwhelming probability by running in time $\tilde{T}$.

2.1 Non-Malleable Commitments and Zero Knowledge

Here we follow [LPV08]. Let $\Pi = (\mathsf{Sen}, \mathsf{Rec})$ 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 $\mathsf{mim}_\Pi^{\mathcal{A}, m_1, \ldots, m_{\mathsf{poly}(\lambda)}}(z)$ 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 $\mathsf{sim}_\Pi^S(1^\lambda, z)$ 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 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 Rec) 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 1 (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:*
$\{\mathsf{mim}_\Pi^{\mathcal{A}, m_1, \ldots, m_{\mathsf{poly}(\lambda)}}(z)\}_{z \in \{0,1\}^\star} \approx \{\mathsf{sim}_\Pi^S(1^\lambda, z)\}_{z \in \{0,1\}^\star}.$

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* (Sen, Rec) *be a one-many NM commitment scheme. Then,* (Sen, Rec) *is also a concurrent (i.e., many-many) NM commitment scheme.*

We say that a commitment is valid or well formed if it can be decommitted to a message $m \neq \bot$. Following [LP11b] we say that a MiM is *synchronous* if it "aligns" the left and the right sessions; that is, whenever it receives message i on the left, it directly sends message i on the right, and vice versa.

Definition 2 (*synchronous* NM commitment scheme). *A commitment scheme is* synchronous *one-one (resp., one-many) non-malleable if it is one-one (resp., one-many) NM with respect to synchronous MiM adversaries.*

Definition 3 (*weak* NM commitment scheme). *A commitment scheme is* weak *one-one (resp., one-many) non-malleable if it is a one-one (resp., one-many) NM commitment scheme with respect to MiM adversaries that when receiving a well formed commitment in the left session, except with negligible probability computes well formed commitments (i.e.,* $\neq \bot$*) in the right sessions.*

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.

In the rest of the paper, following [GRRV16], we assume for simplicity that player identities are known before the protocol begins, though strictly speaking this is not necessary, as the identities do not appear in the protocol until after the first committer message. We point out that MiM can choose his identity adversarially, as long as it is not equal to Sen's. As already observed in previous work, 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 id the public key of a strong signature scheme signing the transcript. The MiM will have to use a different id or to break the signature scheme.

Delayed-input non-malleable zero knowledge. Following [LP11a] we give a definition that gives to the adversary the power of adaptive-input selection[11].

---
[11]In [LP11a] the adversary selects the instance and a Turing machines outputs the witness in exponential time. Here we slightly deviate by 1) requiring the adversary to output also the witness (similarly to [SCO+01]) and 2) allowing the adversary to make this choice at the last round.

Let $\Pi = (\mathcal{P}, \mathcal{V})$ be a delayed-input interactive argument system for a $\mathcal{NP}$-language L with witness relation Rel_L. Consider a PPT MiM adversary $\mathcal{A}$ that is simultaneously participating in one left session and one right session. Before the execution starts, both $\mathcal{P}, \mathcal{V}$ and $\mathcal{A}$ receive as a common input the security parameter in unary 1^λ, and $\mathcal{A}$ receives as auxiliary input $z \in \{0,1\}^\star$.

In the left session $\mathcal{A}$ interacts with $\mathcal{P}$ using identity $\mathtt{id}$ of his choice. In the right session, $\mathcal{A}$ interacts with $\mathcal{V}$, using identity $\tilde{\mathtt{id}}$ of his choice.

Furthermore, in the left session $\mathcal{A}$, before the last round of Π, adaptively selects the statement x to be proved and the witness w, s.t $(x, w) \in \mathsf{Rel}_L$, and sends them to $\mathcal{P}$. Also, in the right session $\mathcal{A}$, during the last round of Π, adaptively selects the statement $\tilde{x}$ to be proved and sends it to $\mathcal{V}$. Let $\mathsf{View}^{\mathcal{A}}(1^\lambda, z)$ denote a random variable that describes the view of $\mathcal{A}$ in the above experiment.

Definition 4 (Delayed-input NMZK). *A delayed-input argument system $\Pi = (\mathcal{P}, \mathcal{V})$ for a $\mathcal{NP}$-language L with witness relation Rel_L is NM Zero Knowledge (NMZK) if for any MiM adversary $\mathcal{A}$ that participates in one left session and one right session, there exists a PPT machine $S(1^\lambda, z)$ such that:*

1. *The probability ensembles $\{S^1(1^\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^\star}$ and $\{\mathsf{View}^{\mathcal{A}}(1^\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^\star}$ are computationally indistinguishable over λ, where $S^1(1^\lambda, z)$ denotes the first output of $S(1^\lambda, z)$.*

2. *Let $z \in \{0,1\}^\star$, and let $(\mathsf{View}, \tilde{w})$ denote the output of $S(1^\lambda, z)$. Let $\tilde{x}$ be the right-session statement appearing in View and let $\mathtt{id}$ and $\tilde{\mathtt{id}}$ be the identities of the left and right sessions appearing in View. If the right session is accepting and $\mathtt{id} \neq \tilde{\mathtt{id}}$, then $\mathsf{Rel}_L(\tilde{x}, \tilde{w}) = 1$.*

The above definition of NMZK allows the adversary to select statements adaptively at the last round both on left and right sessions, therefore any argument system that is NMZK according to the above definition enjoys also adaptive-input argument of knowledge and adaptive-input zero knowledge.

We stress that similarly to [SCO$^+$01], in the above definition the simulator receives from the adversary only the adaptively-computed instance, in contrast to the prover that receives both instance and witness.

3   OWFs $\Rightarrow$ 4-Round Concurrent Non-Malleable Commitments

We start discussing a useful building block.

3.1   OWFs $\Rightarrow$ 2-Round Instance-Dependent Trapdoor Commitments

Definition 5. *Let 1^λ be the security parameter, L be an $\mathcal{NP}$-language and Rel_L be the corresponding $\mathcal{NP}$-relation. A triple of PPT algorithms $\mathsf{TC} = (\mathsf{Sen}, \mathsf{Rec}, \mathsf{TFake})$ is a 2-Round Instance-Dependent Trapdoor Commitment scheme if the following properties hold.*

Correctness. *In the 1st round, Rec on input 1^λ and $x \in L$ outputs ρ. In the 2nd round Sen on input the message m, 1^λ, ρ and $x \in L$ outputs $(\mathsf{com}, \mathsf{dec})$. We will refer to the pair (ρ, com) as the commitment of m. Moreover we will refer to the execution of the above two rounds including the exchange of the corresponding two messages as the commitment phase. Then Rec on input m, x, com, dec and the private coins used to generate ρ in the commitment phase outputs 1. We will refer to the execution of this last round including the exchange of*

dec *as the decommitment phase. Notice that an adversarial sender* Sen* *could deviate from the behavior of* Sen *when computing and sending* com *and* dec *for an instance* $x \in \hat{L}$. *As a consequence* Rec *could output 0 in the decommitment phase. We will say that* dec *is a valid decommitment of* (ρ, com) *to* m *for an instance* $x \in \hat{L}$, *if* Rec *outputs 1.*

Hiding. *Given a* PPT *adversary* $\mathcal{A}$, *consider the following hiding experiment* $\mathsf{ExpHiding}^b_{\mathcal{A},\mathsf{TC}}(\lambda, x)$ *for* $b = 0, 1$ *and* $x \in \hat{L}_R$:

- *On input* 1^λ *and* x, $\mathcal{A}$ *outputs a message* m, *along with* ρ.
- *The challenger on input* x, m, ρ, b *works as follows: if* $b = 0$ *then it runs* Sen *on input* m, x *and* ρ, *obtaining a pair* (com, dec), *otherwise it runs* TFake *on input* x *and* ρ, *obtaining a pair* (com, aux). *The challenger outputs* com.
- $\mathcal{A}$ *on input* com *outputs a bit* b' *and this is the output of the experiment.*

We say that hiding *holds if for any* PPT *adversary* $\mathcal{A}$ *there exist a negligible function* ν, *s.t.:*

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

Special Binding. *There exists a* PPT *algorithm that on input a commitment* (ρ, com), *the private coins used by* Rec *to compute* ρ, *and two valid decommitments* $(\text{dec}, \text{dec}')$ *of* (ρ, com) *to two different messages* m *and* m' *w.r.t. an instance* $x \in L$, *outputs* w *s.t.* $(x, w) \in \mathsf{Rel}_\mathsf{L}$ *with overwhelming probability.*

Trapdoorness. *For any* PPT *adversary* $\mathcal{A}$ *there exist a negligible function* ν, *s.t. for all* $x \in L$ *it holds that:*

$$\left| \text{Prob} \left[\, \mathsf{ExpCom}_{\mathcal{A},\mathsf{TC}}(\lambda, x) = 1 \, \right] - \text{Prob} \left[\, \mathsf{ExpTrapdoor}_{\mathcal{A},\mathsf{TC}}(\lambda, x) = 1 \, \right] \right| < \nu(\lambda)$$

where $\mathsf{ExpCom}_{\mathcal{A},\mathsf{TC}}(\lambda, x)$ *and* $\mathsf{ExpTrapdoor}_{\mathcal{A},\mathsf{TC}}(\lambda, x)$ *are defined below*[12].

$\mathsf{ExpCom}_{\mathcal{A},\mathsf{TC}}(\lambda, x)$:	$\mathsf{ExpTrapdoor}_{\mathcal{A},\mathsf{TC}}(\lambda, x)$:
-On input 1^λ *and* x, $\mathcal{A}$ *outputs* (ρ, m).	*-On input* 1^λ *and* x, $\mathcal{A}$ *outputs* (ρ, m).
*-*Sen *on input* 1^λ, x, m *and* ρ, *outputs* (com, dec).	*-*TFake *on input* 1^λ, x *and* ρ, *outputs* (com, aux).
	*-*TFake *on input* tk *s.t.* $(x, \text{tk}) \in \mathsf{Rel}_\mathsf{L}$, x, ρ, com, aux *and* m *outputs* dec.
-$\mathcal{A}$ *on input* (com, dec) *outputs a bit* b *and this is the output of the experiment.*	*-*$\mathcal{A}$ *on input* (com, dec) *outputs a bit* b *and this is the output of the experiment.*

OWFs $\Rightarrow$ 2-round instance-dependent trapdoor commitments for any $\mathcal{NP}$ language.
Here we recall the construction $(\mathsf{Sen_H}, \mathsf{Rec_H}, \mathsf{TFake_H})$ of [FS89] for Hamiltonian graphs. $\mathsf{Sen_H}$ and $\mathsf{Rec_H}$ run as follows.

- $\mathsf{Rec_H} \rightarrow \mathsf{Sen_H}$. $\mathsf{Rec_H}$ on input a graph G with n nodes computes and sends ρ to the sender, where ρ is the 1st round of a two-round statistically binding commitment scheme from OWFs of [Nao91].

---

[12]We assume wlog that $\mathcal{A}$ is stateful.

– $\mathsf{Sen_H}$ on input a bit b, ρ and a graph G with n nodes works as follows. If $b = 0$ then $\mathsf{Sen_H}$ picks a random permutation π and computes and sends the 2nd round of the statistically binding commitment using ρ as 1st round and committing one-by-one to all bits of the adjacency matrix of $\pi(G)$. If instead $b = 1$, then $\mathsf{Sen_H}$ computes and sends the 2nd round of the statistically binding commitment, using ρ as 1st round and committing to all bits of the adjacency matrix of a a graph that consists of a random cycle H of n nodes. In both cases, $(\rho, \mathtt{com} = (\mathtt{com}_1, \ldots, \mathtt{com}_{n^2}))$ corresponds to the commitment of b. In the 1st case $\mathtt{dec}$ corresponds to the randomness used by $\mathsf{Sen_H}$, while in the 2nd case $\mathtt{dec}$ corresponds to the decommitments of those n edges in the adjacency matrix that correspond to the cycle.

– $\mathsf{Rec_H}$ on input a bit b, a graph G with n nodes, $\mathtt{dec}$ and $\mathtt{com}$ works as follows. If $b = 0$ then $\mathsf{Rec_H}$ verifies that $\mathtt{com}$ is a commitment of the adjacency matrix of $\pi(G)$ where both π and the decommitments of the adjacency matrix are taken from $\mathtt{dec}$. If instead $b = 1$ then $\mathsf{Rec_H}$ verifies that the decommitted edges in $\mathtt{dec}$ correspond to a cycle that was committed in $(\rho, \mathtt{com})$.

– $\mathsf{TFake_H}$ runs $\mathsf{Sen_H}$ on input ρ, a graph G with n nodes and $b = 0$ therefore obtaining $(\mathtt{com}, \mathtt{dec})$. Then $\mathsf{TFake_H}$ on input 0 and a cycle in G outputs $\mathtt{dec}$. Instead on input 1 and a cycle in G, $\mathsf{TFake_H}$ outputs the decommitments of the edges committed in $(\rho, \mathtt{com})$ corresponding to a cycle in $\pi(G)$ where π was the permutation selected to compute $\mathtt{com}$.

It is easy to see that the above construction is a 2-round instance-dependent trapdoor commitment scheme from OWFs. While the construction can be used to commit to a bit, in the rest of the paper we will use this construction to commit to strings by implicitly assuming that the above steps are repeated in parallel for each bit of the string.

Moreover, note that since Hamiltonicity is an $\mathcal{NP}$-complete language, the above construction works for any $\mathcal{NP}$ language through $\mathcal{NP}$ reductions. For simplicity in the rest of the paper we will omit the $\mathcal{NP}$ reduction therefore assuming that the above scheme works directly on a given $\mathcal{NP}$-language L.

3.2 4-Round Concurrent NM Commitment Scheme: $(\mathsf{NM4Sen}, \mathsf{NM4Rec})$

Our construction is crucially based on the above 2-round instance-dependent trapdoor commitment scheme and on the special honest-verifier zero-knowledge property of the delayed-input adaptive-input PoK LS. The above two tools will be used along with signatures in our protocol that upgrades a 4-round public-coin weak one-many NM commitment scheme $\Pi_{\mathsf{wom}} = (\mathsf{Sen_{wom}}, \mathsf{Rec_{wom}})$ where the last round of $\mathsf{Sen_{wom}}$ is deterministic. More in details, in addition to Π_{wom} we will need the following tools:

1. a signature scheme $\Sigma = (\mathsf{Gen}, \mathsf{Sign}, \mathsf{Ver})$ from OWFs [Rom90];
2. a 2-round instance-dependent trapdoor commitment scheme $\mathsf{TC}_\Sigma = (\mathsf{Sen}_\Sigma, \mathsf{Rec}_\Sigma, \mathsf{TFake}_\Sigma)$ from OWFs for the following $\mathcal{NP}$-language

$$L_\Sigma = \big\{ \mathsf{vk} : \exists \, (\mathsf{msg}_1, \mathsf{msg}_2, \sigma_1, \sigma_2) \text{ s.t. } \mathsf{Ver}(\mathsf{vk}, \mathsf{msg}_1, \sigma_1) = 1$$
$$\text{AND } \mathsf{Ver}(\mathsf{vk}, \mathsf{msg}_2, \sigma_2) = 1 \text{ AND } \mathsf{msg}_1 \neq \mathsf{msg}_2 \big\};$$

3. a 4-round delayed-input SHVZK[13] proof system $\mathsf{LS} = (\mathcal{P}, \mathcal{V})$ from OWFs (see App. A) for

---
[13]See Def. 10.

the language

$$L = \big\{ \big(\tau = (\pi^1_{\mathsf{wom}}, \pi^2_{\mathsf{wom}}, \pi^3_{\mathsf{wom}}, \pi^4_{\mathsf{wom}}), \mathtt{id} \big) : \exists \, (m, \mathtt{dec}) \text{ s.t.}$$
$$\mathsf{Rec}_{\mathsf{wom}} \text{ on input } (\tau, m, \mathtt{dec}, \mathtt{id}) \text{ accepts } m \text{ as a decommitment of } \tau \big\}$$

that is adaptive-input PoK for the corresponding relation Rel_L.

Let m be the message that NM4Sen wants to commit and $\mathtt{id}$ be the id for this session. In the 1st round the receiver NM4Rec computes and sends the 1st round π^1_{LS} of LS and the 1st round π^1_{wom} of Π_{wom} using as input the $\mathtt{id}$. NM4Rec also computes a pair of signature and verification keys $(\mathsf{sk}, \mathsf{vk})$, sends the verification key vk to NM4Sen and computes and sends the first round of TC_Σ by running Rec_Σ on the instance $\mathsf{vk} \in L_\Sigma$.

Then NM4Sen on input $\mathtt{id}$, the message m and the received 1st round, computes the 2nd round π^2_{wom} of Π_{wom} to commit to the message m using $\mathtt{id}$. Moreover NM4Sen computes the 2nd round π^2_{LS} of LS and runs Sen_Σ on input the instance vk to compute the commitment of π^2_{LS} therefore obtaining a pair $(\mathtt{com}, \mathtt{dec})$. NM4Sen sends $\mathtt{com}$, π^2_{wom} and a random message $\mathtt{msg}$ to NM4Rec. In the 3rd round NM4Rec sends the 3rd round π^3_{wom} of Π_{wom}, the 3rd round of LS and a signature σ (computed using sk) of the message $\mathtt{msg}$. In the last round NM4Sen verifies whether or not σ is a valid signature for $\mathtt{msg}$. If σ is a valid signature, then NM4Sen computes the last round π^4_{wom} of Π_{wom}, the 4th round π^4_{LS} of LS and sends $\pi^4_{\mathsf{wom}}, \pi^4_{\mathsf{LS}}$ and $\mathtt{dec}$ to NM4Rec. At this point NM4Rec accepts the commitment (the transcript of our protocol generated so far) iff Rec_Σ on input $\mathsf{vk}, \mathtt{com}, \mathtt{dec}, \pi^2_{\mathsf{LS}}$ accepts $(\pi^2_{\mathsf{LS}}, \mathtt{dec})$ as a decommitment of $\mathtt{com}$ and the transcript for LS is accepting for $\mathcal{V}$ with respect to the instance $(\pi^1_{\mathsf{wom}}, \pi^2_{\mathsf{wom}}, \pi^3_{\mathsf{wom}}, \pi^4_{\mathsf{wom}}, \mathtt{id})$. The decommitment phase of our scheme simply corresponds to the decommitment phase of Π_{wom}.

As described before, the delayed-input SHVZK adaptive-input PoK LS is used by NM4Sen to prove knowledge of a message and randomness consistent with the transcript computed using Π_{wom}. We remark that to execute LS the instance is not needed until the last round but the instance length is required from the onset of the protocol. We will refer to the instance length as ℓ. Since ℓ can be deterministically computed on input the session $\mathtt{id}$ and the max possible length of the message to commit, we will assume w.l.o.g. that ℓ is given in input to sender and receiver.

Fig. 1 describes in details the 4-round concurrent NM commitment scheme from OWFs.

Theorem 1. *If OWFs exist, then there exists (constructively) a 4-round concurrent NM commitment scheme.*

Proof. As discussed earlier Π_{NM4Com} can be instantiated by relying on OWFs only. Therefore to prove the theorem we just need to show that Π_{NM4Com} is a concurrent NM commitment scheme.

The security proof is divided in two parts. In the 1st part we prove that Π_{NM4Com} is indeed a commitment scheme. Then we prove that Π_{NM4Com} is also concurrent non-malleable.

Lemma 1. Π_{NM4Com} *is a statistically binding computationally hiding commitment scheme.*

Proof. **Correctness.** The correctness follows directly from the completeness of LS, the correctness of Π_{wom}, from the validity of the signature scheme Σ and from the correctness of the 2-round instance-dependent trapdoor commitment scheme TC_Σ.

Statistical Binding. Observe that the message given in output in the decommitment phase of Π_{NM4Com} is the message committed using Π_{wom}. Moreover the decommitment of Π_{NM4Com} coincides with the decommitment of Π_{wom}. Since Π_{wom} is statistically binding then so is Π_{NM4Com}.

Computationally Hiding. Computational hiding follows immediately from Lemma 2.

$\square$

Common input: security parameter λ, instance length ℓ, NM4Sen's identity $\texttt{id} \in \{0,1\}^\lambda$.
Input to NM4Sen: $m \in \{0,1\}^{\mathsf{poly}\{\lambda\}}$.
Commitment phase:

1. NM4Rec $\rightarrow$ NM4Sen
    1. Run $(\mathsf{sk}, \mathsf{vk}) \leftarrow \mathsf{Gen}(1^\lambda)$.
    2. Run $\mathcal{V}$ on input 1^λ and ℓ thus obtaining the 1st round π_{LS}^1 of LS.
    3. Run $\mathsf{Rec_{wom}}$ on input $1^\lambda, \texttt{id}$ thus obtaining the 1st round π_{wom}^1 of Π_{wom}.
    4. Run $\mathsf{Rec_\Sigma}$ on input 1^λ and vk thus obtaining ρ.
    5. Send $(\mathsf{vk}, \pi_{\mathsf{LS}}^1, \pi_{\mathsf{wom}}^1, \rho)$ to NM4Sen.

2. NM4Sen $\rightarrow$ NM4Rec
    1. Run $\mathsf{Sen_{wom}}$ on input $1^\lambda, \texttt{id}, \pi_{\mathsf{wom}}^1$ and m thus obtaining the 2nd round π_{wom}^2 of Π_{wom}.
    2. Run $\mathcal{P}$ on input $1^\lambda, \ell$ and π_{LS}^1 thus obtaining the 2nd round π_{LS}^2 of LS.
    3. Pick a message $\mathsf{msg} \leftarrow \{0,1\}^\lambda$.
    4. Run $\mathsf{Sen_\Sigma}$ on input $1^\lambda, \mathsf{vk}, \rho$ and message π_{LS}^2 to compute the pair $(\mathsf{com}, \mathsf{dec})$.
    5. Send $(\pi_{\mathsf{wom}}^2, \mathsf{com}, \mathsf{msg})$ to NM4Rec.

3. NM4Rec $\rightarrow$ NM4Sen
    1. Run $\mathsf{Rec_{wom}}$ on input π_{wom}^2 thus obtaining the 3rd round π_{wom}^3 of Π_{wom}.
    2. Run $\mathcal{V}$ on input π_{LS}^2 thus obtaining the 3rd round π_{LS}^3 of LS.
    3. Run $\mathsf{Sign}(\mathsf{sk}, \mathsf{msg})$ to obtain a signature σ of the message msg.
    4. Send $(\pi_{\mathsf{wom}}^3, \pi_{\mathsf{LS}}^3, \sigma)$ to NM4Sen.

4. NM4Sen $\rightarrow$ NM4Rec
    1. If $\mathsf{Ver}(\mathsf{vk}, \mathsf{msg}, \sigma) \neq 1$ then abort, continue as follows otherwise.
    2. Run $\mathsf{Sen_{wom}}$ on input π_{wom}^3 thus obtaining the 4th round π_{wom}^4 of Π_{wom} and the decommitment information $\mathsf{dec_{wom}}$.
    3. Set $x = (\pi_{\mathsf{wom}}^1, \pi_{\mathsf{wom}}^2, \pi_{\mathsf{wom}}^3, \pi_{\mathsf{wom}}^4, \texttt{id})$ and $w = (m, \mathsf{dec_{wom}})$ with $|x| = \ell$. Run $\mathcal{P}$ on input x, w and π_{LS}^3 thus obtaining the 4th round π_{LS}^4 of LS.
    4. Send $(\pi_{\mathsf{wom}}^4, (\mathsf{dec}, \pi_{\mathsf{LS}}^2), \pi_{\mathsf{LS}}^4)$ to NM4Rec.

5. NM4Rec : Set $x = (\pi_{\mathsf{wom}}^1, \pi_{\mathsf{wom}}^2, \pi_{\mathsf{wom}}^3, \pi_{\mathsf{wom}}^4, \texttt{id})$ and accept the commitment iff the following conditions are satisfied.
    1. $\mathsf{Rec_\Sigma}$ on input $\mathsf{vk}, \mathsf{com}, \mathsf{dec}, \pi_{\mathsf{LS}}^2$ accepts $(\pi_{\mathsf{LS}}^2, \mathsf{dec})$ as a decommitment of com.
    2. $(\pi_{\mathsf{LS}}^1, \pi_{\mathsf{LS}}^2, \pi_{\mathsf{LS}}^3, \pi_{\mathsf{LS}}^4)$ is accepting for $\mathcal{V}$ with respect to the instance x.

Decommitment phase:

1. NMSen $\rightarrow$ NMRec: Send $(\mathsf{dec_{wom}}, m)$ to NMRec.

2. NMRec: accept m as the committed message if and only if $\mathsf{Rec_{wom}}$, on input $(m, \mathsf{dec_{wom}})$, accepts m as the committed message of $(\pi_{\mathsf{wom}}^1, \pi_{\mathsf{wom}}^2, \pi_{\mathsf{wom}}^3, \pi_{\mathsf{wom}}^4, \texttt{id})$.

Figure 1: Our 4-round concurrent NM commitment scheme Π_{NM4Com} from OWFs.

Lemma 2. Π_{NM4Com} *is concurrent non-malleable.*

Proof. Here we actually prove that Π_{NM4Com} is a one-many NM commitment scheme, and then we go from one-many to concurrent non-malleability by using Proposition 1.

First we give an overview of the entire non-malleability proof, and then we will give more details. We denote by $\{\mathsf{mim}_{\mathcal{H}_i^m}^{\mathcal{A}_{\mathsf{NMCom}}, m}(z)\}_{z \in \{0,1\}^\star}$ the random variable describing the view of the

MiM $\mathcal{A}_{\mathsf{NMCom}}$ combined with the values it commits in the the $\mathsf{poly}(\lambda)$ right sessions in hybrid $\mathcal{H}_i^m(z)$.

As required by the definition, we want to show that the distribution of the real game experiment (i.e., the view of the MiM $\mathcal{A}_{\mathsf{NM4Com}}$ when playing with $\mathsf{NM4Sen}$ committing m along with the messages committed in the right sessions) and the one of the output of a simulator are computationally indistinguishable.

The proof makes use of the following hybrid experiments.

– The 1st hybrid experiment is $\mathcal{H}_1^m(z)$. In this hybrid in the left session $\mathsf{NM4Sen}$ commits to m, while in the right sessions $\mathsf{NM4Rec}_i$ interacts with $\mathcal{A}_{\mathsf{NM4Com}}$ for $i = 1 \ldots \mathsf{poly}(\lambda)$. We prove that in the i-th right session $\mathcal{A}_{\mathsf{NM4Com}}$ does not commit to a message $\tilde{m}_i =\perp$ for any $i \in \{1, \ldots, \mathsf{poly}(\lambda)\}$. The proof is by contradiction; more precisely we will show that if $\mathcal{A}_{\mathsf{NM4Com}}$ commits to $\perp$ in a session i, then we can break the security of the signature scheme.

– The 2nd hybrid experiment is $\mathcal{H}_I^m(z)$ and differs from $\mathcal{H}_1^m(z)$ in the way the commitment com and the decommitment information dec are computed in the left session. More precisely, $\mathsf{NM4Sen}$ runs TFake_Σ to compute a commitment com of 0^λ, and subsequently to compute a decommitment of com to the value π_{LS}^2 (we remark that no trapdoor is needed to run TFake_Σ in order to compute com). In more details, this experiment rewinds the adversary $\mathcal{A}_{\mathsf{NM4Com}}$ from the 3rd to the 2nd round of the left session to extract two signatures σ_1, σ_2 of two different messages $(\mathsf{msg}_1, \mathsf{msg}_2)$ and uses them as trapdoor to run TFake_Σ. The indistinguishability between $\mathcal{H}_I^m(z)$ and $\mathcal{H}_1^m(z)$ comes from the hiding and the trapdoorness of TC_Σ.

– The 3rd hybrid experiment is $\mathcal{H}_2^m(z)$ and differs from $\mathcal{H}_I^m(z)$ in the way that the transcript for LS is computed. In more details the SHVZK simulator $\mathcal{S}$ of LS is used to compute the messages π_{LS}^2 and π_{LS}^4 instead of using the honest procedure $\mathcal{P}$. The indistinguishability between $\mathcal{H}_I^m(z)$ and $\mathcal{H}_2^m(z)$ comes from the SHVZK of LS.

We also consider the hybrid experiments $\mathcal{H}_1^0(z)$, $\mathcal{H}_I^0(z)$, $\mathcal{H}_2^0(z)$, that are the same hybrid experiments described above with the difference that $\mathsf{NM4Sen}$ uses Π_{wom} to commit to a message 0^λ instead of m. From the same arguments described above we have that: $\mathsf{mim}_{\mathcal{H}_1^0}^{\mathcal{A}_{\mathsf{NM4Com}},m}(z) \approx \mathsf{mim}_{\mathcal{H}_I^0}^{\mathcal{A}_{\mathsf{NM4Com}},m}(z) \approx \mathsf{mim}_{\mathcal{H}_2^0}^{\mathcal{A}_{\mathsf{NM4Com}},m}(z)$. Observe that the distribution of the view of the adversary and the committed messages in $\mathcal{H}_1^0(z)$ are indistinguishable from the output of a simulator. Now we want to show that the view of $\mathcal{A}_{\mathsf{NM4Com}}$ along with messages committed in the right sessions in $\mathcal{H}_1^m(z)$ is indistinguishable from the view of $\mathcal{A}_{\mathsf{NM4Com}}$ along with the messages committed in the right sessions in $\mathcal{H}_1^0(z)$. In order to do this, given the indistinguishability of the hybrids discussed above, it only remains to show that $\mathcal{H}_2^m(z)$ is indistinguishable from $\mathcal{H}_2^0(z)$. This is ensured by the weak one-many non-malleability of Π_{wom}. Indeed, observe that here it suffices to rely on the *weak* one-many NM commitment because we are guaranteed (from the previous arguments) that whenever $\mathcal{A}_{\mathsf{NM4Com}}$ completes a commitment in a right session, the corresponding message committed through Π_{wom} is different from $\perp$ with overwhelming probability, and this holds both in $\mathcal{H}_2^m(z)$ and in $\mathcal{H}_2^0(z)$.

We now give more details about the hybrid experiments and the security proof.

– In the 1st experiment in the left session $\mathsf{NM4Sen}$ commits to m, while in the right sessions the receivers interact with $\mathcal{A}_{\mathsf{NM4Com}}$. We refer to this hybrid experiment as $\mathcal{H}_1^m(z)$.

Claim 1. *Let $\bar{p}_i$ be the probability that for $i = 1 \ldots \mathsf{poly}(\lambda)$ in the i-th right sessions of $\mathcal{H}_1^m(z)$ $\mathcal{A}_{\mathsf{NM4Com}}$ successfully commits to a message $\tilde{m}_i = \bot$, then $\bar{p}_i < \nu(\lambda)$ for some negligible function ν.*

Proof. We prove this claim by contradiction. More specifically we assume that there exists a right session i where $\mathcal{A}_{\mathsf{NM4Com}}$ commits to $\bot$ and then we construct an adversary that breaks the signature scheme Σ. Now we show how to extract two signatures for two different messages in order to break the signature scheme Σ. First of all we observe that if $\mathcal{A}_{\mathsf{NM4Com}}$ commits to $\bot$ in the i-th right session, then the theorem proved by LS is false. This means that for every $\tilde{\pi}_{\mathsf{LS}}^2$ there exists only one $\tilde{\pi}_{\mathsf{LS}}^3$ s.t. $\mathcal{A}_{\mathsf{NM4Com}}$ can compute an accepting 4th round for LS. Since we are assuming that $\mathcal{A}_{\mathsf{NM4Com}}$ completes the i-right session with non-negligible probability we can rewind $\mathcal{A}_{\mathsf{NM4Com}}$ sending a new 3rd round message for LS $\tilde{\pi}_{\mathsf{LS}}^{3'}$ in the i-th session possibly obtaining an accepting $\tilde{\pi}_{\mathsf{LS}}^{4'}$. By contradiction the only way for $\mathcal{A}_{\mathsf{NM4Com}}$ to answer to a different 3rd round of LS is to send in the 4th round of Π_{NM4Com} a new $\tilde{\pi}_{\mathsf{LS}}^{2'}$. This means the we have two openings of com w.r.t. two different messages ($\tilde{\pi}_{\mathsf{LS}}^{2'}$ and $\tilde{\pi}_{\mathsf{LS}}^2$), therefore we can extract two different signatures for two different messages by using the special binding of TC_Σ. The only thing that remains to discuss is how to behave in the left session when the described rewind (occurring in the right session) affects also the left session. Since in the left session the theorem proved by LS is true, we simply answer to a potentially new challenge $\pi_{\mathsf{LS}}^{3'}$ by computing a new accepting $\pi_{\mathsf{LS}}^{4'}$ without changing the 2nd round of LS π_{LS}^2 between rewinds. $\square$

– We consider the experiment $\mathcal{H}_1^0(z)$ that corresponds to $\mathcal{H}_1^m(z)$ with the only difference that the message committed using Π_{wom} is 0^λ instead of m. We prove the following claim.

Claim 2. *Let $\bar{p}_i$ be the probability that for $i = 1 \ldots \mathsf{poly}(\lambda)$ in the i-th right sessions of $\mathcal{H}_1^0(z)$ $\mathcal{A}_{\mathsf{NM4Com}}$ successfully commits to a message $\tilde{m}_i = \bot$, then $\bar{p}_i < \nu(\lambda)$ for some negligible function ν.*

The security proof of this claim follows strictly the one of Claim 1.

– We now consider the 2nd hybrid experiment $\mathcal{H}_I^m(z)$ that differs from $\mathcal{H}_1^m(z)$ as follows. In the left session, by rewinding the adversary $\mathcal{A}_{\mathsf{NM4Com}}$ from the 3rd to the 2nd round, two signatures σ_1, σ_2 for two distinct messages ($\mathsf{msg}_1, \mathsf{msg}_2$) are extracted and used as trapdoor to run TFake_Σ. The probability that $\mathcal{A}_{\mathsf{NM4Com}}$ does not give a 2nd valid signature for a randomly chosen message after a polynomial number of rewinds is negligible in λ. For the above reason the above deviation increases the abort probability of the experiment by a negligible amount. From the hiding property of TC_Σ the above holds also in case com is computed by running TFake instead of Sen.

The transcript of an execution of $\mathcal{H}_I^m(z)$ differs from $\mathcal{H}_1^m(z)$ in the way the commitment com and the corresponding decommitment information are computed. Indeed now $\mathsf{NM4Sen}$ runs TFake_Σ both to compute a commitment com of 0^λ and subsequently to compute a decommitment of com to the value π_{LS}^2 (we remark that no trapdoor is needed to run TFake_Σ in order to compute com). To prove that for all $m \in \{0,1\}^{\mathsf{poly}(\lambda)}$ $\mathsf{mim}_{\mathcal{H}_1^m}^{\mathcal{A}_{\mathsf{NM4Com}},m}(z) \approx \mathsf{mim}_{\mathcal{H}_I^m}^{\mathcal{A}_{\mathsf{NM4Com}},m}(z)$ we proceed by contradiction constructing an adversary $\mathcal{A}^{\mathsf{TC}_\Sigma}$ for the trapdoorness property of the 2-round instance-dependent trapdoor commitment scheme TC_Σ. First of all we observe

that by Claim 1 in $\mathcal{H}_1^m(z)$ we can extract from LS the messages committed by the MiM adversary. This means that in the reduction to the security of TC_Σ we can extract the messages committed by $\mathcal{A}^{\mathsf{NM4Com}}$ and run the distinguisher $\mathcal{D}_{\mathsf{NM4Com}}$ (that exists by contradiction) to distinguish $\mathsf{mim}_{\mathcal{H}_1^m}^{\mathcal{A}_{\mathsf{NM4Com}},m}(z)$ from $\mathsf{mim}_{\mathcal{H}_I^m}^{\mathcal{A}_{\mathsf{NM4Com}},m}(z)$ in order to break the trapdoorness property of the 2-round instance-dependent trapdoor commitment scheme. Now we describe the reduction by showing a successful $\mathcal{A}^{\mathsf{TC}_\Sigma}$ against the challenger of TC_Σ.

1. Upon receiving the 1st round from $\mathcal{A}_{\mathsf{NM4Com}}$, $\mathcal{A}^{\mathsf{TC}_\Sigma}$ computes π_{LS}^2 and sends it as the challenge message together with ρ.

2. $\mathcal{A}^{\mathsf{TC}_\Sigma}$, upon receiving com, uses it to compute and send the 2nd round of Π_{NM4Com} to $\mathcal{A}_{\mathsf{NM4Com}}$ on the left.

3. $\mathcal{A}^{\mathsf{TC}_\Sigma}$ extracts two valid signatures of two different messages from the left therefore obtaining the trapdoor tk for TC_Σ. Then it sends tk to the challenger.

4. Upon receiving dec from the challenger, $\mathcal{A}^{\mathsf{TC}_\Sigma}$ uses it to complete the left session against $\mathcal{A}_{\mathsf{NM4Com}}$.

5. $\mathcal{A}^{\mathsf{TC}_\Sigma}$ uses the extractor of LS to extract from the $\mathsf{poly}(\lambda)$ right sessions the witnesses used by $\mathcal{A}_{\mathsf{NM4Com}}$ to compute the LS PoKs (the witnesses correspond to the randomnesses and committed messages in Π_{wom} in the right sessions).

6. If the extraction fails with non-negligible probability, then $\mathcal{A}^{\mathsf{TC}_\Sigma}$ has a non-negligible advantage. Indeed, as proved earlier, the extraction fails with negligible probability when the trapdoor commitment scheme is played by running its honest sender. Therefore we can assume that the extraction succeeds with overwhelming probability, and thus $\mathcal{A}^{\mathsf{TC}_\Sigma}$ runs the distinguisher $\mathcal{D}_{\mathsf{NM4Com}}$ (that exists by contradiction) to distinguish $\mathsf{mim}_{\mathcal{H}_1^m}^{\mathcal{A}_{\mathsf{NM4Com}},m}(z)$ from $\mathsf{mim}_{\mathcal{H}_I^m}^{\mathcal{A}_{\mathsf{NM4Com}},m}(z)$ in order to break the trapdoorness property of TC_Σ.

We observe that rewinds made on the right sessions do not affect the reduction. Indeed the only non-trivial situation that can happen in the the left session consists of requiring to play again only the 4th round of the left session having received a different 3rd round. In this case to compute the 4th round we observe that we can reuse $(\mathsf{dec}, \pi_{\mathsf{LS}}^2)$. More precisely, to compute the 4th round of Π_{wom} we use knowledge of the message m and the randomness used to compute the second round of Π_{wom}. To answer to the new 3rd round of LS we run $\mathcal{P}$ on input π_{LS}^2, the new 3rd round, the randomness and the message used to compute Π_{wom}.

Before considering the next hybrid experiment we now discuss a special procedure to extract the messages committed by the adversary. Starting with the transcript generated by an execution of $\mathcal{H}_I^m$ it is trivial to extract from a right session whenever the corresponding rewind does result in rewinding left session giving a chance to the adversary to change the 3rd round. The extractions from the remaining right sessions proceed one by one as follows. Let i be the index of one of such right sessions. The 4th round of the left session is played again by decommitting com to a different value π_{LS}^2 (we remark that this is possible because we are using TFake_Σ to compute both commitment and decommitment) that still ends up in obtaining that session i is completed successfully. Therefore multiple rewinds trying with different values for π_{LS}^2 could be required, until a successful π_{LS}^2 is found. Then the extraction is performed on session i playing always this new successful π_{LS}^2 in the left session while rewinding session i.

The special extraction procedure discussed above allows to retrieve the messages committed by the MiM in the execution of $\mathcal{H}_I^m(z)$. This follows by observing that the 4th round of the sender of Π_{wom} is deterministic and therefore the witness extracted from sessions i when playing with the successful π_{LS}^2 in the left session, is still a valid witness for the transcript obtained in the original execution of $\mathcal{H}_I^m(z)$[14].

– The 3rd hybrid experiment is $\mathcal{H}_2^m(z)$ and differs from $\mathcal{H}_I^m(z)$ in the way the transcript for LS is computed. More precisely, the SHVZK simulator $\mathcal{S}$ of LS is used to compute the messages π_{LS}^2 and π_{LS}^4 instead of using the honest procedure $\mathcal{P}$. We observe that this is possible because in this hybrid experiment the commitment com can be opened to any value π_{LS}^2 given in output by $\mathcal{S}$ that uses as input π_{LS}^1, π_{LS}^3 and x. We now prove that $\mathsf{mim}_{\mathcal{H}_I^m}^{\mathcal{A}_{\mathsf{NM4Com}},m}(z) \approx \mathsf{mim}_{\mathcal{H}_2^m}^{\mathcal{A}_{\mathsf{NM4Com}},m}(z)$ by contradiction constructing an adversary $\mathcal{A}^{\mathsf{SHVZK}}$ for the SHVZK of LS. The reduction works as follows.

1. $\mathcal{A}^{\mathsf{SHVZK}}$ runs $\mathcal{A}_{\mathsf{NM4Com}}$ (both on the left and on the right sessions) according to $\mathcal{H}_2^m(z)$ until the 3rd round of the left session is received.

2. $\mathcal{A}^{\mathsf{SHVZK}}$ extracts from the left session two valid signature for two different messages (as described before), and then sends to the challenger of SHVZK the theorem x the witness w (computed as described in Fig. 1) and the pair $(\pi_{\mathsf{LS}}^1, \pi_{\mathsf{LS}}^3)$ received in the 1st and 3rd round from $\mathcal{A}_{\mathsf{NM4Com}}$.

3. $\mathcal{A}^{\mathsf{SHVZK}}$, upon receiving the messages $(\pi_{\mathsf{LS}}^2, \pi_{\mathsf{LS}}^4)$ from the challenger of SHVZK uses them to compute and send the last round of Π_{NM4Com}.

4. $\mathcal{A}^{\mathsf{SHVZK}}$ extracts from the $\mathsf{poly}(\lambda)$ right sessions the witnesses used by $\mathcal{A}_{\mathsf{NM4Com}}$ to compute the LS transcripts (that corresponds to the randomnesses and committed messages of Π_{wom}) using the same special extraction procedure described for $\mathcal{H}_I^m$ (i.e., if the extraction from a right session requires to play again the 4th round of the left session then the extraction procedure is run by computing a modified 4th round of the left session where the honest prover of LS is used, following the special extraction procedure described for $\mathcal{H}_1^m(z)$).

5. If the extraction fails in providing the messages committed in the original transcript (i.e., the transcript generated using the challenger in the first 3 steps) with non-negligible probability, then $\mathcal{A}^{\mathsf{SHVZK}}$ has a non-negligible advantage. Indeed, as proved earlier, the extraction fails with negligible probability when the transcript of LS is generated by running its honest prover. Therefore we can assume that the extraction succeeds with overwhelming probability, and thus $\mathcal{A}^{\mathsf{SHVZK}}$ runs the distinguisher $\mathcal{D}_{\mathsf{NM4Com}}$ (that exists by contradiction) to distinguish $\mathsf{mim}_{\mathcal{H}_I^m}^{\mathcal{A}_{\mathsf{NM4Com}},m}(z)$ from $\mathsf{mim}_{\mathcal{H}_2^m}^{\mathcal{A}_{\mathsf{NM4Com}},m}(z)$ in order to break the SHVZK of LS.

---

[14]The above argument has similarities with the adaptive-input PoK of LS where even though through rewinds the adversarial prover can change the proved instance, still the information extracted is sufficient for the original instance. Indeed the crucial point both in [LS90] and in our proof is that by making use of rewinds one gets some critical information that remained unchanged through the rewinds. For the case of [LS90] the critical information consists of the committed cycles sent in the 2nd round by the prover of [LS90], while in our case the critical information consists of the message and randomness used in the 2nd round by the sender (i.e., the MiM adversary in case of right sessions) of Π_{wom}.

– The 4th hybrid is $\mathcal{H}_2^0(m)$. The only differences between this hybrid and the previous one is that NM4Sen commits using Π_{wom} to a message 0^λ instead of m. Even in this case the proof that for all $m \in \{0,1\}^{\mathsf{poly}(\lambda)}$ it holds that $\mathsf{mim}_{\mathcal{H}_2^m}^{\mathcal{A}_{\mathsf{NM4Com}},m}(z) \approx \mathsf{mim}_{\mathcal{H}_2^0}^{\mathcal{A}_{\mathsf{NM4Com}},m}(z)$ from the concurrent weak non-malleability of Π_{wom}. Indeed observe that here it suffices to rely on the *weak* one-many NM commitment because we are guaranteed (from the previous arguments) that whenever $\mathcal{A}_{\mathsf{NM4Com}}$ completes a commitment in a right session, the corresponding message committed through Π_{wom} is different from $\perp$ with overwhelming probability, and this holds both in $\mathcal{H}_2^m(z)$ and in $\mathcal{H}_2^0(z)$. To prove the indistinguishability between those two hybrids experiments we proceed by contradiction constructing an adversary $\mathcal{A}^{\mathsf{wom}}$ that breaks the *weak* one-many non-malleability property of Π_{wom}.

Loosely speaking Π_{wom} acts as NM4Sen with $\mathcal{A}_{\mathsf{NM4Com}}$ with the following differences: 1) NM4Sen plays as proxy between $\mathcal{C}_{\mathsf{wom}}$ and $\mathcal{A}_{\mathsf{NM4Com}}$ w.r.t. messages of Π_{wom} in the main thread; 2) a second signature is extracted from the left session through rewinds; 3) random strings are played to simulate the receiver of Π_{wom} during rewinds. Then $\mathcal{A}_{\mathsf{wom}}$ runs $\mathcal{D}_{\mathsf{wom}}$ on input the messages $\tilde{m}_1, \ldots, \tilde{m}_{\mathsf{poly}}(n)$ committed by $\mathcal{A}_{\mathsf{wom}}$ and his randomness. Therefore $\mathcal{D}_{\mathsf{wom}}$ reconstructs the view of $\mathcal{A}_{\mathsf{NM4Com}}$ (by using the randomness received as input) and uses it along with the messages $\tilde{m}_1, \ldots, \tilde{m}_{\mathsf{poly}}(n)$ as inputs of $\mathcal{D}_{\mathsf{NM4Com}}$ giving in output what $\mathcal{D}_{\mathsf{NM4Com}}$ outputs. Since by contradiction $\mathcal{D}_{\mathsf{NM4Com}}$ distinguishes between $\mathsf{mim}_{\mathcal{H}_2^m}^{\mathcal{A}_{\mathsf{NM4Com}},m}(z)$ and $\mathsf{mim}_{\mathcal{H}_2^0}^{\mathcal{A}_{\mathsf{NM4Com}},m}(z)$ also $\mathcal{D}_{\mathsf{wom}}$ can tell apart which message has been committed by the MiM adversary $\mathcal{A}_{\mathsf{wom}}$. We stress that to complete the reduction we need to extract two signatures for two distinct messages in the left session. This is done by rewinding the MiM adversary $\mathcal{A}_{\mathsf{NM4Com}}$ from the third to the second round of the left session. When the rewind occurs $\mathcal{A}_{\mathsf{NM4Com}}$ also rewinds the receiver of the right session, rewinding also the receiver of Π_{wom} involved in the security reduction. To avoid this issue in the reduction we answer as a receiver of Π_{wom} would have done (we remark that this can be done because Π_{wom} is public coin) for all rewinds that occur in the right session, so that the reduction does not rewind the receiver of Π_{wom}.

The proof ends by observing that for all $m \in \{0,1\}^{\mathsf{poly}(\lambda)}$ the following holds:
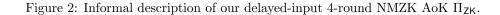
$$\mathsf{mim}_{\mathcal{H}_1^m}^{\mathcal{A}_{\mathsf{NM4Com}},m}(z) \approx \mathsf{mim}_{\mathcal{H}_I^m}^{\mathcal{A}_{\mathsf{NM4Com}},m}(z) \approx \mathsf{mim}_{\mathcal{H}_2^m}^{\mathcal{A}_{\mathsf{NM4Com}},m}(z) \approx$$
$$\approx \mathsf{mim}_{\mathcal{H}_2^0}^{\mathcal{A}_{\mathsf{NM4Com}},m}(z) \approx \mathsf{mim}_{\mathcal{H}_I^0}^{\mathcal{A}_{\mathsf{NM4Com}},m}(z) \approx \mathsf{mim}_{\mathcal{H}_1^0}^{\mathcal{A}_{\mathsf{NM4Com}},m}(z).$$

$\square$

$\square$

4 4-Round Delayed-Input NMZK from CRHFs

Our construction is based on a compiler that takes as input any 4-round public-coin extractable one-one NM commitment scheme $\Pi_{\mathsf{ext}} = (\mathsf{Sen}_{\mathsf{ext}}, \mathsf{Rec}_{\mathsf{ext}})$, a delayed-input adaptive-input statistical WI adaptive-input AoK $\mathsf{sLS} = (\mathcal{P}, \mathcal{V})$, a signature scheme, and outputs a delayed-input 4-round NMZK AoK $\Pi_{\mathsf{ZK}} = (\mathcal{P}_{\mathsf{ZK}}, \mathcal{V}_{\mathsf{ZK}})$ for the $\mathcal{NP}$-language L and corresponding relation Rel_L.

The high-level idea of our compiler is depicted in Fig. 2. In the 1st round $\mathcal{V}_{\mathsf{ZK}}$ computes and sends the 1st round π_{sLS}^1 of sLS and the 1st round π_{ext}^1 of Π_{ext} to $\mathcal{P}_{\mathsf{ZK}}$. Also $\mathcal{V}_{\mathsf{ZK}}$ computes a pair

$$\mathcal{P}_{\mathsf{ZK}}(\mathtt{id}) \qquad\qquad\qquad \mathcal{V}_{\mathsf{ZK}}(\mathtt{id})$$

$$\pi^1_{\mathsf{ext}}(s_0), \pi^1_{\mathsf{sLS}}, \mathsf{vk}$$

$$\pi^2_{\mathsf{ext}}(s_0), \pi^2_{\mathsf{sLS}}, \mathtt{msg}$$

$$\pi^3_{\mathsf{ext}}(s_0), \pi^3_{\mathsf{sLS}}, \sigma$$

Upon receiving x, w s.t. $\qquad\quad$ $s_1, \pi^4_{\mathsf{ext}}(s_0), \pi^4_{\mathsf{sLS}}$
$(x, w) \in \mathsf{Rel}_\mathsf{L}$ set $s_1 = w \oplus s_0$

- vk is a verification key of a signature scheme.
- $\tau = (\pi^1_{\mathsf{ext}}, \pi^2_{\mathsf{ext}}, \pi^3_{\mathsf{ext}}, \pi^4_{\mathsf{ext}})$ is the transcript of $\langle \mathsf{Sen}_{\mathsf{ext}}(m), \mathsf{Rec}_{\mathsf{ext}} \rangle(\mathtt{id})$.
- $(\pi^1_{\mathsf{sLS}}, \pi^2_{\mathsf{sLS}}, \pi^3_{\mathsf{sLS}}, \pi^4_{\mathsf{sLS}})$ is the transcript of sLS proving knowledge of either the decommitment of τ to a message s_0 s.t. $(x, w = s_0 \oplus s_1) \in \mathsf{Rel}_\mathsf{L}$ or of two valid signatures of two different messages w.r.t vk.

Figure 2: Informal description of our delayed-input 4-round NMZK AoK Π_{ZK}.

of signature and verification keys $(\mathsf{sk}, \mathsf{vk})$ and sends vk to $\mathcal{P}_{\mathsf{ZK}}$. $\mathcal{P}_{\mathsf{ZK}}$ input the session-id $\mathtt{id}$, picks a random string s_0, then computes and sends to $\mathcal{V}_{\mathsf{ZK}}$ the 2nd round π^2_{ZK} of Π_{ext} to commit to the message s_0 using $\mathtt{id}$ as session-id. Moreover $\mathcal{P}_{\mathsf{ZK}}$ computes the 2nd round π^2_{sLS} of sLS and sends it along with a random message $\mathtt{msg}$ to $\mathcal{V}_{\mathsf{ZK}}$. In the 3rd round $\mathcal{V}_{\mathsf{ZK}}$ sends the 3rd round π^3_{ext} of Π_{ext}, the 3rd round of sLS and a signature σ (computed using sk) of $\mathtt{msg}$ to $\mathcal{P}_{\mathsf{ZK}}$. In the last round upon receiving x, w s.t. $(x, w) \in \mathsf{Rel}_\mathsf{L}$, $\mathcal{P}_{\mathsf{ZK}}$ verifies whether or not σ is a valid signature for $\mathtt{msg}$. If σ is a valid signature, then $\mathcal{P}_{\mathsf{ZK}}$ computes the last round π^4_{ext} of Π_{ext} and the 4th round π^4_{sLS} of sLS. Finally, $\mathcal{P}_{\mathsf{ZK}}$ sets $s_1 = s_0 \oplus w$ and sends $(\pi^4_{\mathsf{ext}}, \pi^4_{\mathsf{sLS}}, s_1)$ to $\mathcal{V}_{\mathsf{ZK}}$. The delayed-input statistical WIAoK protocol sLS is used by $\mathcal{P}_{\mathsf{ZK}}$ to prove either 1) knowledge of a message s_0 and randomness that are consistent with the transcript computed using Π_{ext} and s.t. $(x, s_1 \oplus s_0) \in \mathsf{Rel}_\mathsf{L}$ or 2) knowledge of signatures of two different messages w.r.t vk.

For constructing our $\Pi_{\mathsf{ZK}} = (\mathcal{P}_{\mathsf{ZK}}, \mathcal{V}_{\mathsf{ZK}})$ for the $\mathcal{NP}$-language L and corresponding relation Rel_L we need the following tools:

1. a 4-round public-coin extractable one-one NM commitment scheme $\Pi_{\mathsf{ext}} = (\mathsf{Sen}_{\mathsf{ext}}, \mathsf{Rec}_{\mathsf{ext}})$;
2. a signature scheme $\Sigma = (\mathsf{Gen}, \mathsf{Sign}, \mathsf{Ver})$;
3. a delayed-input adaptive-input statistical WIAoK protocol $\mathsf{sLS} = (\mathcal{P}, \mathcal{V})$ for the language

$$\Lambda = \big\{ \big(\tau = (\pi^1_{\mathsf{ext}}, \pi^2_{\mathsf{ext}}, \pi^3_{\mathsf{ext}}, \pi^4_{\mathsf{ext}}), \mathtt{id}, \mathsf{vk}, x, s_1 \big) : \exists\, (s_0, \mathtt{dec}, \mathtt{msg}_1, \mathtt{msg}_2, \sigma_1, \sigma_2) \text{ s.t.}$$
$$\big((\mathsf{Rec}_{\mathsf{ext}} \text{ on input } (\tau, s_0, \mathtt{dec}, \mathtt{id}) \text{ accepts } s_0 \text{ as a decommitment of } \tau \text{ AND } (x, s_0 \oplus s_1) \in \mathsf{Rel}_\mathsf{L}) \text{ OR}$$
$$\big(\mathsf{Ver}(\mathsf{vk}, \mathtt{msg}_1, \sigma_1) = 1 \text{ AND } \mathsf{Ver}(\mathsf{vk}, \mathtt{msg}_2, \sigma_2) = 1 \text{ AND } \mathtt{msg}_1 \neq \mathtt{msg}_2 \big) \big) \big\}$$

that is adaptive-input statistical WI and adaptive-input AoK for the corresponding relation Rel_Λ.

Lemma 3. Π_{ZK} *enjoys delayed-input completeness.*

Proof. First we observe that completeness follows directly from the completeness of sLS, the correctness of Π_{ext} and the validity of the signature scheme Σ. Delayed-input completeness follows from the delayed-input completeness of sLS and from the observation that $\mathcal{P}_{\mathsf{ZK}}$ does not need to know the witness to run Π_{ext}. We stress that Π_{ext} is not required to enjoy a delayed-input property. $\square$

19

Common input: security parameter λ, the instance length ℓ of sLS and $\mathcal{P}_{\text{ZK}}$'s identity $\text{id} \in \{0,1\}^\lambda$, and the instance x available only at the last round.
Private input of $\mathcal{P}_{\text{ZK}}$: w s.t. $(x,w) \in \text{Rel}_\text{L}$, available only at the last round.

1. $\mathcal{V}_{\text{ZK}} \to \mathcal{P}_{\text{ZK}}$
   1. Run $(\text{sk}, \text{vk}) \leftarrow \text{Gen}(1^\lambda)$.
   2. Run $\mathcal{V}$ on input 1^λ and ℓ thus obtaining the 1st round π^1_{sLS} of sLS.
   3. Run Rec_{ext} on input $1^\lambda, \text{id}$ thus obtaining the 1st round π^1_{ext} of Π_{ext}.
   4. Send $(\text{vk}, \pi^1_{\text{sLS}}, \pi^1_{\text{ext}})$ to $\mathcal{P}_{\text{ZK}}$.

2. $\mathcal{P}_{\text{ZK}} \to \mathcal{V}_{\text{ZK}}$
   1. Pick at random s_0 s.t. $|s_0|$ is the witness length for an instance of L.
   2. Run Sen_{ext} on input $1^\lambda, \text{id}, \pi^1_{\text{ext}}$ and s_0 thus obtaining the 2nd round π^2_{ext} of Π_{ext}.
   3. Run $\mathcal{P}$ on input $1^\lambda, \ell$ and π^1_{sLS} thus obtaining the 2nd round π^2_{sLS} of sLS.
   4. Pick a message $\text{msg} \leftarrow \{0,1\}^\lambda$.
   5. Send $(\pi^2_{\text{ext}}, \pi^2_{\text{sLS}}, \text{msg})$ to $\mathcal{V}_{\text{ZK}}$.

3. $\mathcal{V}_{\text{ZK}} \to \mathcal{P}_{\text{ZK}}$
   1. Run Rec_{ext} on input π^2_{ext} thus obtaining the 3rd round π^3_{ext} of Π_{ext}.
   2. Run $\mathcal{V}$ on input π^2_{sLS} thus obtaining the 3rd round π^3_{sLS} of sLS.
   3. Run $\text{Sign}(\text{sk}, \text{msg})$ to obtain a signature σ of the message msg.
   4. Send $(\pi^3_{\text{ext}}, \pi^3_{\text{sLS}}, \sigma)$ to $\mathcal{P}_{\text{ZK}}$.

4. $\mathcal{P}_{\text{ZK}} \to \mathcal{V}_{\text{ZK}}$
   1. If $\text{Ver}(\text{vk}, \text{msg}, \sigma) \neq 1$ then abort, continue as follows otherwise.
   2. Set $s_1 = s_0 \oplus w$.
   3. Run Sen_{ext} on input π^3_{ext} thus obtaining the 4th round π^4_{ext} of Π_{ext} and the decommitment information dec_{ext}.
   4. Set $x_{\text{sLS}} = (\pi^1_{\text{ext}}, \pi^2_{\text{ext}}, \pi^3_{\text{ext}}, \pi^4_{\text{ext}}, \text{id}, \text{vk}, x, s_1)$ and $w_{\text{sLS}} = (s_0, \text{dec}_{\text{ext}}, \bot, \bot, \bot, \bot)$ with $|x_{\text{sLS}}| = \ell$. Run $\mathcal{P}$ on input $x_{\text{sLS}}, w_{\text{sLS}}$ and π^3_{sLS} thus obtaining the forth round π^4_{sLS} of sLS.
   5. Send $(\pi^4_{\text{ext}}, \pi^4_{\text{sLS}}, s_1)$ to $\mathcal{V}_{\text{ZK}}$.

5. $\mathcal{V}_{\text{ZK}}$: Set $x_{\text{sLS}} = (\pi^1_{\text{ext}}, \pi^2_{\text{ext}}, \pi^3_{\text{ext}}, \pi^4_{\text{ext}}, \text{id}, \text{vk}, x, s_1)$ and accept iff $(\pi^1_{\text{sLS}}, \pi^2_{\text{sLS}}, \pi^3_{\text{sLS}}, \pi^4_{\text{sLS}})$ is accepting for $\mathcal{V}$ with respect to x_{sLS}.

Figure 3: Our 4-round delayed-input NMZK argument of knowledge Π_{ZK}.

Theorem 2. *If Π_{ext} is a 4-round public-coin extractable one-one NM commitment scheme and CRHFs exist then Π_{ZK} is a delayed-input NMZK AoK for $\mathcal{NP}$.*

We will refer to the simulated experiment as the experiment where Sim_{ZK} interacts with the adversary emulating both a prover and a verifier. The simulator works in a pretty straight-forward way. It commits to a random message, it extracts a second signature from the left session and completes the execution generating the first output according to Def. 4. Then it extracts the witness from the extractable commitment Π_{ext} played by the adversary in the right session (see Fig. 6 for a detailed description of Sim_{ZK}). Obviously we will have to show that the probability that the message extracted is not a witness for the statement proved by the adversary in the right session is negligible.

We will give a lemma for each of the two properties of Def. 4.

Lemma 4. $\{\text{Sim}^1_{\text{ZK}}(1^\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*} \approx \{\text{View}^{\mathcal{A}}(1^\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*}$, *where $\text{Sim}^1_{\text{ZK}}(1^\lambda, z)$ denotes the 1st output of Sim_{ZK}.*

In order to prove the above lemma we consider an hybrid experiment $\mathcal{H}_1(1^\lambda, z)$. $\mathcal{H}_1(1^\lambda, z)$ differs from the real execution of Π_{ZK} in the witness used to compute messages of sLS. In more details in $\mathcal{H}_1(1^\lambda, z)$ $\mathcal{P}_{\mathsf{ZK}}$ extracts (through rewinds), two signatures of different messages[15] that are used as witness for sLS. Let $\{\mathsf{View}^{\mathcal{A}_{\mathsf{ZK}}}_{\mathcal{H}_1}(1^\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*}$ denote a random variable that describes the view of $\mathcal{A}_{\mathsf{ZK}}$ in $\mathcal{H}_1(1^\lambda, z)$. The adaptive-input statistical WI of sLS and the negligible probability of failing in extracting a second signature guarantee that $\{\mathsf{View}^{\mathcal{A}_{\mathsf{ZK}}}_{\mathcal{H}_1}(1^\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*}$ and $\{\mathsf{View}^{\mathcal{A}_{\mathsf{ZK}}}(1^\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*}$ are statistically close.

Observe now that the only difference between $\mathcal{H}_1(1^\lambda, z)$ and the simulated execution is the message committed using Π_{ext}. In more details, let x be the adaptively chosen statement proved by $\mathcal{P}_{\mathsf{ZK}}$. In $\mathcal{H}_1(1^\lambda, z)$ $\mathcal{P}_{\mathsf{ZK}}$ commits using Π_{ext} to a value s_0 s.t. $s_1 = w \oplus s_0$ (where $(x, w) \in \mathsf{Rel}_\mathsf{L}$). Instead in the simulated experiment $\mathsf{Sim}_{\mathsf{ZK}}$ commits to a random string. Now we can claim that $\{\mathsf{Sim}^1_{\mathsf{ZK}}(1^\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*}$ and $\{\mathsf{View}^{\mathcal{A}_{\mathsf{ZK}}}_{\mathcal{H}_1}(1^\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*}$ are computationally indistinguishable by using the computationally-hiding property of Π_{ext}. Informally, suppose by contradiction that there exist an adversary $\mathcal{A}_{\mathsf{ZK}}$ and a distinguisher $\mathcal{D}_{\mathsf{ZK}}$ such that $\mathcal{D}_{\mathsf{ZK}}$ distinguishes $\{\mathsf{Sim}^1_{\mathsf{ZK}}(1^\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*}$ from $\{\mathsf{View}^{\mathcal{A}_{\mathsf{ZK}}}_{\mathcal{H}_1}(1^\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*}$. Then we can construct an adversary $\mathcal{A}_{\mathsf{Hiding}}$ that breaks the computationally hiding of Π_{ext} in the following way. $\mathcal{A}_{\mathsf{Hiding}}$ sends to the challenger of the hiding game $\mathcal{C}_{\mathsf{Hiding}}$ two random messages (m_0, m_1). Then $\mathcal{A}_{\mathsf{Hiding}}$ acts as $\mathcal{P}_{\mathsf{ZK}}$ except for messages of Π_{ext} for which he acts as proxy between $\mathcal{C}_{\mathsf{Hiding}}$ and $\mathcal{A}_{\mathsf{ZK}}$. When $\mathcal{A}_{\mathsf{Hiding}}$ computes the last round of the left session $\mathcal{A}_{\mathsf{Hiding}}$ sets $s_1 = m_0 \oplus w$. At the end of the execution $\mathcal{A}_{\mathsf{Hiding}}$ runs $\mathcal{D}_{\mathsf{ZK}}$ and outputs what $\mathcal{D}_{\mathsf{ZK}}$ outputs. It easy to see that if $\mathcal{C}_{\mathsf{Hiding}}$ commits to m_0 then, $\mathcal{A}_{\mathsf{ZK}}$ acts as in $\mathcal{H}_1(1^\lambda, z)$, otherwise he acts as in the simulated experiment. Note that the reduction to the hiding property of Π_{ext} is possible because the rewinds to extract a second signature do not affect the execution with the challenger of Π_{ext} that remains straight-line. Thus we have proved that $\{\mathsf{View}^{\mathcal{A}_{\mathsf{ZK}}}(1^\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*} \equiv_s \{\mathsf{View}^{\mathcal{A}_{\mathsf{ZK}}}_{\mathcal{H}_1}(1^\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*} \approx \{\mathsf{Sim}^1_{\mathsf{ZK}}(1^\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*}$.

Lemma 5. *Let $\tilde{x}$ be the right-session statement appearing in $\mathsf{View} = \mathsf{Sim}^1_{\mathsf{ZK}}(1^\lambda, z)$ and let $\mathtt{id}$ and $\tilde{\mathtt{id}}$ be the identities of the left and right sessions appearing in View. If the right session is accepting and $\mathtt{id} \neq \tilde{\mathtt{id}}$, then except with negligible probability, the second output of $\mathsf{Sim}_{\mathsf{ZK}}(1^\lambda, z)$ is $\tilde{w}$ such that $\mathsf{Rel}_\mathsf{L}(\tilde{x}, \tilde{w}) = 1$.*

The formal proof can be found in App. B.1. Here we give an overview. The proof relies on hybrid experiments to prove that $\mathcal{A}_{\mathsf{ZK}}$ commits to $\tilde{s}_0$ s.t. $(\tilde{x}, \tilde{s}_0 \oplus \tilde{s}_1) \in \mathsf{Rel}_\mathsf{L}$[16] through Π_{ext} in the simulated experiment.

- The 1st hybrid is $\mathcal{H}_1(z)$ in which in the left session $\mathcal{P}_{\mathsf{ZK}}$ interacts with $\mathcal{A}_{\mathsf{ZK}}$ and in the right session $\mathcal{A}_{\mathsf{ZK}}$ interacts with $\mathcal{V}_{\mathsf{ZK}}$. We refer to this hybrid experiment as $\mathcal{H}_1(z)$. Now we prove that in the right session of $\mathcal{H}_1(z)$ the MiM adversary $\mathcal{A}_{\mathsf{ZK}}$ commits to the witness. Let $\tilde{x}$ be the adaptively chosen theorem proved by $\mathcal{A}_{\mathsf{ZK}}$. By contradiction if $\mathcal{A}_{\mathsf{ZK}}$ commits to a message s'_0 s.t. $(\tilde{x}, \tilde{s}'_0 \oplus \tilde{s}_1) \notin \mathsf{Rel}_\mathsf{L}$, then the witness used to complete an accepting transcript for sLS consists of two valid signatures of two different messages. Then, by using the adaptive-input AoK property of sLS we can reach a contradiction by breaking the security of Σ. Note that this hybrid corresponds to the real experiment where $\mathcal{A}_{\mathsf{ZK}}$ interacts with $\mathcal{P}_{\mathsf{ZK}}$ in the left session.

---

[15]In the proof of Lemma 5 we show that the extraction fails with negligible probability. The same analysis applies here.

[16]For simplicity in the rest of the proof we say that a player commits to a witness when he commits to s_0 and sends s_1 in the last round s.t. $(x, s_0 \oplus s_1) \in \mathsf{Rel}_\mathsf{L}$.

- The 2nd hybrid is $\mathcal{H}_2(z)$ and differs from $\mathcal{H}_1(z)$ only in the witness used to compute messages of sLS in the left session. In more details, we rewind the adversary $\mathcal{A}_{\mathsf{ZK}}$ from the 3rd to the 2nd round of the left session to extract two signatures σ_1, σ_2 of two different messages $(\mathtt{msg}_1, \mathtt{msg}_2)$ and we use them as witness to execute sLS in the left session. From the adaptive-statistical WI of sLS it follows that the distribution of the message committed by $\mathcal{A}_{\mathsf{ZK}}$ does not change when moving from $\mathcal{H}_1(z)$ to $\mathcal{H}_2(z)$.

- The 3rd hybrid is $\mathcal{H}_3(z)$. The only difference between this hybrid and the previous one is that both s_0 and s_1 are random strings. From the non-malleability property of Π_{ext} it follows that the distribution of the message committed by $\mathcal{A}_{\mathsf{ZK}}$ does not change when switching from $\mathcal{H}_2(z)$ to $\mathcal{H}_3(z)$. This is again a delicate reduction because it requires to show a successful MiM for Π_{ext} that is supposed to work in straight-line. However the above experiment requires to rewind the adversary in order to extract a second signature. As already discussed in the previous section, this is the place where the public-coin property is used. Indeed this allows us to simulate the additional answers of the honest receiver of Π_{ext} that are needed because of the rewinds performed to extract a second signature.

Note that $\mathcal{H}_3(z)$ corresponds to the the simulated experiment, this implies that also in the simulated game $\mathcal{A}_{\mathsf{ZK}}$ commits to the witness. Therefore, our simulator can use the extractor of Π_{ext} to get the witness $\tilde{w}$ s.t. $(\tilde{x}, \tilde{w}) \in \mathsf{Rel}_\mathsf{L}$, where $\tilde{x}$ is the adaptively chosen theorem proved by $\mathcal{A}_{\mathsf{ZK}}$. Now we can claim the following corollary.

Corollary 1. *If CRHFs exist, then there exists (constructively) a 4-round concurrent adaptive-input non-malleable zero knowledge.*

The corollary proof follows from the fact that: Π_{ext} can be constructed from OWFs if it is instantiate with the construction provided in [GPR16], sLS can be constructed from CRHFs (see App. A) and Σ can be constructed from OWFs [Rom90].

5 3-Round NM Commitments from Strong OWPs

Our construction is based on a compiler that takes as input a 3-round synchronous weak one-one NM commitment scheme $\Pi_{\mathsf{wsyn}} = (\mathsf{Sen}_{\mathsf{wsyn}}, \mathsf{Rec}_{\mathsf{wsyn}})$, a OWP f, an LS WIPoK for $\mathcal{NP}$ LS, and outputs a 3-round one-one NM commitment scheme $\Pi_{\mathsf{NMCom}} = (\mathsf{NMSen}, \mathsf{NMRec})$.

Let m be the message that NMSen wants to commit. The high-level idea of our compiler is depicted in Fig. 4. The sender NMSen, on input the session-id $\mathtt{id}$ and the message m, computes the 1st round of the protocol by sending the 1st round a_{LS} of LS and the 1st round $\mathsf{a}_{\mathsf{wsyn}}$ of Π_{wsyn} (to commit to the message m using $\mathtt{id}$ as session-id). In the 2nd round the receiver NMRec sends challenges $\mathsf{c}_{\mathsf{wsyn}}$ and c_{LS} of Π_{wsyn} and LS, also picks and sends an element Y in the range of f. In the 3rd round NMSen computes the 3rd round of Π_{wsyn} and completes the transcript for LS by sending $\mathsf{z}_{\mathsf{wsyn}}$ and z_{LS}. Let $\tau = (\mathsf{a}_{\mathsf{wsyn}}, \mathsf{c}_{\mathsf{wsyn}}, \mathsf{z}_{\mathsf{wsyn}})$ be the transcript of the execution of Π_{wsyn}. LS is used by NMSen to prove knowledge of either a decommitment of τ to a message $\neq \perp$ or of a preimage of Y.

Our compiler needs the following tools:
1. a OWP f that is secure against PPT adversaries and that is $\tilde{T}_f$-breakable;
2. a 3-round one-one synchronous weak NM commitment scheme $\Pi_{\mathsf{wsyn}} = (\mathsf{Sen}_{\mathsf{wsyn}}, \mathsf{Rec}_{\mathsf{wsyn}})$ that is T_{wsyn}-hiding/NM, and $\tilde{T}_{\mathsf{wsyn}}$-breakable;

NMSen(m, id)   NMRec(id)

$\mathsf{a_{wsyn}}, \mathsf{a_{LS}}$ →

← $\mathsf{c_{wsyn}}, \mathsf{c_{LS}}, Y$

$\mathsf{z_{wsyn}}, \mathsf{z_{LS}}$ →

- Y is an element taken from the range of the OWP f.

- $\tau = (\mathsf{a_{wsyn}}, \mathsf{c_{wsyn}}, \mathsf{z_{wsyn}})$ is the transcript of $\langle \mathsf{Sen_{wsyn}}(m), \mathsf{Rec_{wsyn}} \rangle(\mathtt{id})$.

- $(\mathsf{a_{LS}}, \mathsf{c_{LS}}, \mathsf{z_{LS}})$ is the transcript of LS for proving knowledge of either the decommitment of τ to a message $\neq \bot$ or of the preimage of Y.
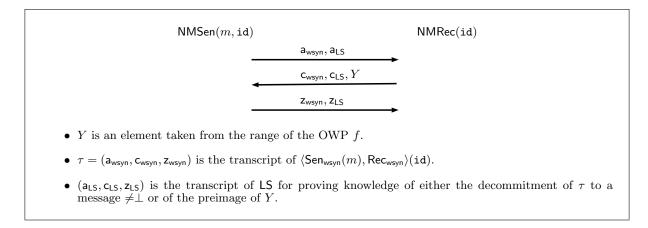
Figure 4: Informal description of our 3-round NM commitment scheme Π_{NMCom}.

3. the LS PoK $\mathsf{LS} = (\mathcal{P}, \mathcal{V})$ for the language

$$L = \big\{ (a, c, z, Y, \mathtt{id}) : \exists \, (m, \mathtt{dec}, y) \text{ s.t. } \big(\mathsf{Rec_{wsyn}} \text{ on input } (a, c, z, m, \mathtt{dec}, \mathtt{id})$$
$$\text{accepts } m \neq \bot \text{ as a decommitment of } (a, c, z, \mathtt{id}) \text{ OR } Y = f(y) \big) \big\}$$

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

Let λ be the security parameter of our scheme. We use w.l.o.g. λ also as security parameter for the one-wayness of f with respect to polynomial-time adversaries. We consider the following hierarchy of security levels: $\tilde{T}_f << T_{\mathsf{wsyn}} << \tilde{T}_{\mathsf{wsyn}} = \sqrt{T_{\mathsf{LS}}} << T_{\mathsf{LS}}$ where by "$T << T'$" we mean that "$T \cdot \mathsf{poly}(\lambda) < T'$".

Now, similarly to [PW10, COSV16], we define different security parameters, one for each tool involved in the security proof to be consistent with the hierarchy of security levels defined above. Given the security parameter λ of our scheme, we will make use of the following security parameters: 1) λ for the OWP f; 2) λ_{wsyn} for the synchronous weak one-one NM commitment scheme; 3) λ_{LS} for LS.

All of them are polynomially related to λ and they are such that the above hierarchy of security levels holds. In the construction we assume for simplicity to have a function Params that on input λ outputs $(\lambda_{\mathsf{wsyn}}, \lambda_{\mathsf{LS}}, \ell)$ where ℓ is the length of the theorem to be proved using LS.[17] The detailed scheme is described in Fig. 5 and a compact version is depicted in Fig. 4.

Theorem 3. *Suppose there exist a synchronous weak one-one NM commitment scheme and OWPs, both secure against subexponential-time adversaries, then Π_{NMCom} is a NM commitment scheme.*

The proof is divided in two parts. First we prove that Π_{NMCom} is a commitment scheme. Then we prove that Π_{NMCom} is a NM commitment scheme.

Lemma 6. Π_{NMCom} *is a statistically-binding computationally-hiding commitment scheme.*

Proof. **Correctness.** The correctness of Π_{NMCom} follows immediately from the completeness of LS, and the correctness of Π_{wsyn}.

Statistical Binding. Observe that the message given in output in the decommitment phase of Π_{NMCom} is the message committed using Π_{wsyn}. Moreover the decommitment phase of Π_{NMCom}

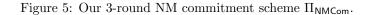[17]To compute 1st and 2nd round of LS only the length ℓ of the instance is required.

Figure 5: Our 3-round NM commitment scheme Π_{NMCom}.

coincides with the decommitment phase of Π_{wsyn}. Since Π_{wsyn} is binding we have that the same holds for Π_{NMCom}.

Hiding. Following Def. 13 to prove the hiding of Π_{NMCom} we have to show that the experiment $\mathsf{ExpHiding}^0_{\mathcal{A}, \Pi_{\mathsf{NMCom}}}(\lambda)$ in which NMSen commits to a message m_0 is computationally indistinguishable from the experiment $\mathsf{ExpHiding}^1_{\mathcal{A}, \Pi_{\mathsf{NMCom}}}(\lambda)$ in which NMSen commits to a message m_1. In order to prove this indistinguishability we consider the following hybrid experiments.

- The 1st hybrid experiment $\mathcal{H}^0(\lambda)$ is equal to the real game experiment $\mathsf{ExpHiding}^0_{\mathcal{A}, \Pi_{\mathsf{NMCom}}}(\lambda)$, with the difference that a value y s.t. $Y = f(y)$ is computed and used as a witness for LS. Observe that in order to compute y the commitment phase takes time $\tilde{T}_f$. The indistinguishability between $\mathcal{H}^0(\lambda)$ and $\mathsf{ExpHiding}^0_{\mathcal{A}, \Pi_{\mathsf{NMCom}}}(\lambda)$ comes from the adaptive-input WI of LS, that holds against adversaries with running time bounded by $T_{\mathsf{LS}} >> \tilde{T}_f$.

- The 2nd hybrid $\mathcal{H}^1(\lambda)$ differs from $\mathcal{H}^0(\lambda)$ in the message committed by the adversary using Π_{wsyn}. More precisely, Π_{wsyn} is used by NMSen to commit to the message m_1 instead of m_0. The indistinguishability between $\mathcal{H}^0(\lambda)$ and $\mathcal{H}^1(\lambda)$ comes from the hiding of Π_{wsyn} and noticing that the hiding of Π_{wsyn} still holds against adversaries with running time bounded by $T_{\mathsf{wsyn}} >> \tilde{T}_f$.

The proof ends with the observation that $\mathcal{H}^1(\lambda) \approx \mathsf{ExpHiding}^1_{\mathcal{A}, \Pi_{\mathsf{NMCom}}}(\lambda)$. The indistinguisha-

bility between $\mathcal{H}^1(\lambda)$ and $\mathsf{ExpHiding}^1_{\mathcal{A},\Pi_{\mathsf{NMCom}}}(\lambda)$ comes from the adaptive-WI property of LS and from the observation that, as before, the adaptive-input WI of LS still holds against adversaries with running time bounded by $T_{\mathsf{LS}} >> \tilde{T}_f$.

□

The full proof of non-malleability can be found in App. B.2. Here we give an overview of the proof. The proof is divided in two cases, in the first case we consider an adversarial MiM $\mathcal{A}_{\mathsf{NMCom}}$ that acts in a synchronized way, while in the second case $\mathcal{A}_{\mathsf{NMCom}}$ is non-synchronized. In both cases we want to show that the committed value (and the view) of $\mathcal{A}_{\mathsf{NMCom}}$ when interacting with a prover NMSen that commits to a message m is indistinguishable from the committed value (and the view) of a simulator. The proof for the synchronous case goes through a series of hybrid experiments listed below.

- We consider the real game experiment $\mathcal{H}^m_1(z)$ in which in the left session NMSen commits to m, while in the right session NMRec interacts with $\mathcal{A}_{\mathsf{NMCom}}$. Now we prove that in the right session the MiM adversary $\mathcal{A}_{\mathsf{NMCom}}$ does not commit to a message $\tilde{m} = \perp$. By contradiction if $\mathcal{A}_{\mathsf{NMCom}}$ commits to $\tilde{m} = \perp$ then the witness used to complete an accepting transcript for LS is a value $\tilde{y}$ s.t. $f(\tilde{Y}) = \tilde{y}$. Then, by using the adaptive PoK property of LS we can reach a contradiction by inverting f in polynomial time.

- The 2nd hybrid is $\mathcal{H}^m_2(z)$ and it differs from $\mathcal{H}^m_1(z)$ only in the witness used to compute the sLS transcript. The adversary $\mathcal{A}_{\mathsf{NMCom}}$, running in sub-exponential time, computes a value y s.t. $f(y) = Y$, and uses it as witness for the execution of LS. From the adaptive-input WI (that is stronger than inverting the OWP and of breaking Π_{wsyn}) of sLS, the view and the committed message of $\mathcal{A}_{\mathsf{NMCom}}$ do not change between $\mathcal{H}^m_2(z)$ and $\mathcal{H}^m_1(z)$.

- We now consider the hybrid experiment $\mathcal{H}^0_1(z)$ that differs from the first hybrid experiment that we have considered $\mathcal{H}^m_1(z)$ in the committed message. Indeed in this case, the message committed in the left session is 0^λ. We observe that $\mathcal{H}^0_1(z)$ actually is the simulated game. As for the hybrid experiment $\mathcal{H}^m_1(z)$ we need to prove that in the right session the MiM adversary $\mathcal{A}_{\mathsf{NMCom}}$ does not commit to a message $\tilde{m} = \perp$. By contradiction if $\mathcal{A}_{\mathsf{NMCom}}$ commits to $\tilde{m} = \perp$ then the witness used to complete an accepting transcript for LS is a value $\tilde{y}$ s.t. $f(\tilde{Y}) = \tilde{y}$. Then, by using the PoK property of LS we can reach a contradiction by inverting f in polynomial time.

- The last hybrid experiment that we consider is $\mathcal{H}^0_2(z)$ and it differs from $\mathcal{H}^0_1(z)$ only in the witness used to compute the sLS transcript. In more details the adversary $\mathcal{A}_{\mathsf{NMCom}}$, running in sub-exponential time, computes a value y s.t. $f(y) = Y$, and uses it as witness for the execution of LS. From the adaptive-input WI (that is stronger than inverting the OWP and of breaking Π_{wsyn}) of sLS, the view and the committed message of $\mathcal{A}_{\mathsf{NMCom}}$ do not change between $\mathcal{H}^0_2(z)$ and $\mathcal{H}^0_1(z)$.

To conclude this proof we show that the view and the committed message of $\mathcal{A}_{\mathsf{NMCom}}$ acting in $\mathcal{H}^m_1(z)$ are indistinguishable from the view and the committed message of $\mathcal{A}_{\mathsf{NMCom}}$ acting in $\mathcal{H}^0_1(z)$. For what has been argued above, it remains to show that the view and the committed message of $\mathcal{H}^m_2(z)$ are indistinguishable from the view and the committed message of $\mathcal{H}^0_2(z)$. This is ensured by the synchronous weak non-malleability of Π_{wsyn}. Here we need only to use a *weak* synchronous one-one NM commitment since we are guaranteed, from the above arguments, that whenever $\mathcal{A}_{\mathsf{NMCom}}$

completes a commitment in a right session the underlying commitment computed through Π_{wsyn} corresponds to $\bot$ with negligible probability only both in $\mathcal{H}_2^m(z)$ and in $\mathcal{H}_2^0(z)$.

The proof for the asynchronous case is much simpler and relies on the hiding of Π_{NMCom}. More precisely we observe that in case of asynchronous scheduling it is possible to rewind the adversary $\mathcal{A}_{\mathsf{NMCom}}$ without rewinding the sender in the left session. This allows us to extract (in polynomial time) the witness used by the adversary in the execution of LS, that with overwhelming probability corresponds to the committed message. Therefore we contradict the hiding of Π_{NMCom}.

6   Acknowledgments

References

[ADL14]   Divesh Aggarwal, Yevgeniy Dodis, and Shachar Lovett. Non-malleable codes from additive combinatorics. In David B. Shmoys, editor, *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 774–783. ACM, 2014.

[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.

[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.

[BJY97]   Mihir Bellare, Markus Jakobsson, and Moti Yung. Round-optimal zero-knowledge arguments based on any one-way function. In *Advances in Cryptology - EUROCRYPT '97, International Conference on the Theory and Application of Cryptographic Techniques, Konstanz, Germany, May 11-15, 1997, Proceeding*, volume 1233 of *Lecture Notes in Computer Science*, pages 280–305. Springer, 1997.

[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.

[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.

[COP$^+$14] Kai-Min Chung, Rafail Ostrovsky, Rafael Pass, Muthuramakrishnan Venkitasubramaniam, and Ivan Visconti. 4-round resettably-sound zero knowledge. In Yehuda Lindell, editor, *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 192–216. Springer, 2014.

[COSV16] Michele Ciampi, Rafail Ostrovsky, Luisa Siniscalchi, and Ivan Visconti. Concurrent non-malleable commitments (and more) in 3 rounds, 2016.

[CPS13] Kai-Min Chung, Rafael Pass, and Karn Seth. Non-black-box simulation from one-way functions and applications to resettable security. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 231–240. ACM, 2013.

[CPS$^+$16a] Michele Ciampi, Giuseppe Persiano, Alessandra Scafuro, Luisa Siniscalchi, and Ivan Visconti. Improved or-composition of sigma-protocols. In Eyal Kushilevitz and Tal Malkin, editors, *Theory of Cryptography - 13th International Conference, TCC 2016-A, Tel Aviv, Israel, January 10-13, 2016, Proceedings, Part II*, volume 9563 of *Lecture Notes in Computer Science*, pages 112–141. Springer, 2016.

[CPS$^+$16b] Michele Ciampi, Giuseppe Persiano, Alessandra Scafuro, Luisa Siniscalchi, and Ivan Visconti. Online/offline OR composition of sigma protocols. In Marc Fischlin and Jean-Sébastien Coron, editors, *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*, volume 9666 of *Lecture Notes in Computer Science*, pages 63–92. Springer, 2016.

[CVZ10] Zhenfu Cao, Ivan Visconti, and Zongyang Zhang. Constant-round concurrent non-malleable statistically binding commitments and decommitments. In Phong Q. Nguyen and David Pointcheval, editors, *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.

[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.

[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.

[DMRV13] Dana Dachman-Soled, Tal Malkin, Mariana Raykova, and Muthuramakrishnan Venki-tasubramaniam. Adaptive and concurrent secure computation from new adaptive, non-malleable commitments. In *Advances in Cryptology - ASIACRYPT 2013 - 19th International Conference on the Theory and Application of Cryptology and Information Security, Bengaluru, India, December 1-5, 2013, Proceedings, Part I*, pages 316–336, 2013.

[DPV04a] Giovanni Di Crescenzo, Giuseppe Persiano, and Ivan Visconti. Constant-round re-settable 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.

[DPV04b] Giovanni Di Crescenzo, Giuseppe Persiano, and Ivan Visconti. Improved setup assump-tions for 3-round resettable zero knowledge. In *Advances in Cryptology - ASIACRYPT 2004, 10th International Conference on the Theory and Application of Cryptology and Information Security, Jeju Island, Korea, December 5-9, 2004, Proceedings*, volume 3329 of *Lecture Notes in Computer Science*, pages 530–544. Springer, 2004.

[Fei90] Uriel Feige. Alternative models for zero knowledge interactive proofs. Master's thesis, Weizmann Institute of Science, Rehovot, Israel, 1990. Ph.D. thesis.

[FS89] Uriel Feige and Adi Shamir. Zero knowledge proofs of knowledge in two rounds. In *Advances in Cryptology - CRYPTO '89, 9th Annual International Cryptology Confer-ence, Santa Barbara, California, USA, August 20-24, 1989, Proceedings*, pages 526–544, 1989.

[GJO+13] Vipul Goyal, Abhishek Jain, Rafail Ostrovsky, Silas Richelson, and Ivan Visconti. Constant-round concurrent zero knowledge in the bounded player model. In *Advances in Cryptology - ASIACRYPT 2013, 19th International Conference on the Theory and Application of Cryptology and Information Security, 2004, Proceedings*, volume 8279 of *Lecture Notes in Computer Science*, pages 21–40. Springer, 2013.

[GK96] Oded Goldreich and Hugo Krawczyk. On the composition of zero-knowledge proof systems. *SIAM J. Comput.*, 25(1):169–192, 1996.

[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.

[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.

[GMPP16] Sanjam Garg, Pratyay Mukherjee, Omkant Pandey, and Antigoni Polychroniadou. The exact round complexity of secure computation. In Marc Fischlin and Jean-Sébastien

Coron, editors, *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*, volume 9666 of *Lecture Notes in Computer Science*, pages 448–476. Springer, 2016.

[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.

[GPR15]   Vipul Goyal, Omkant Pandey, and Silas Richelson. Textbook non-malleable commitments. Cryptology ePrint Archive, Report 2015/1178, 2015. http://eprint.iacr.org/2015/1178.

[GPR16]   Vipul Goyal, Omkant Pandey, and Silas Richelson. Textbook non-malleable commitments. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 1128–1141, 2016.

[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.

[GRRV16]   Vipul Goyal, Silas Richelson, Alon Rosen, and Margarita Vald. An algebraic approach to non-malleability. Cryptology ePrint Archive, Report 2014/586, June 2016. http://eprint.iacr.org/.

[HPV15]   Carmit Hazay, Antigoni Polychroniadou, and Muthuramakrishnan Venkitasubramaniam. Composable security in the tamper proof model under minimal complexity. Cryptology ePrint Archive, Report 2015/887, 2015. http://eprint.iacr.org/2015/887.

[HV16a]   Carmit Hazay and Muthuramakrishnan Venkitasubramaniam. On the power of secure two-party computation. In Matthew Robshaw and Jonathan Katz, editors, *Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part II*, volume 9815 of *Lecture Notes in Computer Science*, pages 397–429. Springer, 2016.

[HV16b]   Carmit Hazay and Muthuramakrishnan Venkitasubramaniam. What security can we achieve in less than 4-rounds? In *Security and Cryptography for Networks - 10th International Conference, SCN 2016, Amalfi, Italy, August 31 - September 2, 2016. Proceedings*, 2016.

[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.

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

[LP11a]   Huijia Lin and Rafael Pass. Concurrent non-malleable zero knowledge with adaptive inputs. In Yuval Ishai, editor, *Theory of Cryptography - 8th Theory of Cryptography Conference, TCC 2011, Providence, RI, USA, March 28-30, 2011. Proceedings*, volume 6597 of *Lecture Notes in Computer Science*, pages 274–292. Springer, 2011.

[LP11b]   Huijia Lin and Rafael Pass. Constant-round non-malleable commitments from any one-way function. In Lance Fortnow and Salil P. Vadhan, editors, *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pages 705–714. ACM, 2011.

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

[LPV08]   Huijia Lin, Rafael Pass, and Muthuramakrishnan Venkitasubramaniam. Concurrent non-malleable commitments from any one-way function. In Ran Canetti, editor, *Theory of Cryptography, Fifth Theory of Cryptography Conference, TCC 2008, New York, USA, March 19-21, 2008.*, volume 4948 of *Lecture Notes in Computer Science*, pages 571–588. Springer, 2008.

[LS90]   Dror Lapidot and Adi Shamir. Publicly verifiable non-interactive zero-knowledge proofs. In *Advances in Cryptology - CRYPTO*, 1990.

[MP12]   Mohammad Mahmoody and Rafael Pass. The curious case of non-interactive commitments - on the power of black-box vs. non-black-box use of primitives. In Reihaneh Safavi-Naini and Ran Canetti, editors, *Advances in Cryptology - CRYPTO 2012 - 32nd Annual Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2012. Proceedings*, volume 7417 of *Lecture Notes in Computer Science*, pages 701–718. Springer, 2012.

[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/.

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

[OPV09]   Rafail Ostrovsky, Giuseppe Persiano, and Ivan Visconti. Simulation-based concurrent non-malleable commitments and decommitments. In Omer Reingold, editor, *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.

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

[Pas13]   Rafael Pass. Unprovable security of perfect NIZK and non-interactive non-malleable commitments. In *TCC*, pages 334–354, 2013.

[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.

[PR03]   Rafael Pass and Alon Rosen. Bounded-concurrent secure two-party computation in a constant number of rounds. In *44th Symposium on Foundations of Computer Science (FOCS 2003), 11-14 October 2003, Cambridge, MA, USA, Proceedings*, pages 404–413. IEEE Computer Society, 2003.

[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.

[PR05b]   Rafael Pass and Alon Rosen. New and improved constructions of non-malleable cryptographic protocols. In Harold N. Gabow and Ronald Fagin, editors, *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, pages 533–542. ACM, 2005.

[PR08a]   Rafael Pass and Alon Rosen. Concurrent nonmalleable commitments. *SIAM J. Comput.*, 37(6):1891–1925, 2008.

[PR08b]   Rafael Pass and Alon Rosen. New and improved constructions of nonmalleable cryptographic protocols. *SIAM J. Comput.*, 38(2):702–752, 2008.

[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.

[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.

[Rom90]   John Rompel. One-way functions are necessary and sufficient for secure signatures. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing, May 13-17, 1990, Baltimore, Maryland, USA*, pages 387–394, 1990.

[SCO+01]   Alfredo De Santis, Giovanni Di Crescenzo, Rafail Ostrovsky, Giuseppe Persiano, and Amit Sahai. Robust non-interactive zero knowledge. In Joe Kilian, editor, *Advances in Cryptology - CRYPTO 2001, 21st Annual International Cryptology Conference, Santa Barbara, California, USA, August 19-23, 2001, Proceedings*, volume 2139 of *Lecture Notes in Computer Science*, pages 566–598. Springer, 2001.

[SV12]   Alessandra Scafuro and Ivan Visconti. On round-optimal zero knowledge in the bare public-key model. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology - EUROCRYPT 2012 - 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings*, volume 7237 of *Lecture Notes in Computer Science*, pages 153–171. Springer, 2012.

[Wee10]   Hoeteck Wee. Black-box, round-efficient secure computation via non-malleability am-
           plification. 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.

[YZ07]   Moti Yung and Yunlei Zhao. Generic and practical resettable zero-knowledge in the bare
           public-key model. In Moni Naor, editor, *Advances in Cryptology - EUROCRYPT 2007,
           26th Annual International Conference on the Theory and Applications of Cryptographic
           Techniques, Barcelona, Spain, May 20-24, 2007, Proceedings*, volume 4515 of *Lecture
           Notes in Computer Science*, pages 129–147. Springer, 2007.

A   Standard Definitions and Tools

Definition 6 (One-way function (OWF)). *A function $f : \{0,1\}^\star \to \{0,1\}^\star$ is called one way if the following two conditions hold:*

- *there exists 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 $z \in \{0,1\}^{\mathsf{poly}(\lambda)}$:*

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

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 7 (Strong Signatures [CPS13]). *A triple of PPT algorithms $(\mathsf{Gen}, \mathsf{Sign}, \mathsf{Ver})$ is called a signature scheme if it satisfies the following properties.*

Validity: *For every pair $(s,v) \leftarrow \mathsf{Gen}(1^\lambda)$, and every $m \in \{0,1\}^\lambda$, we have that*

$$\mathsf{Ver}(v,m,\mathsf{Sign}(s,m)) = 1.$$

Security: *For every PPT $\mathcal{A}$, there exists a negligible function ν, such that for all auxiliary input $z \in \{0,1\}^\star$ it holds that:*

$$\Pr[(s,v) \leftarrow \mathsf{Gen}(1^\lambda); (m,\sigma) \leftarrow \mathcal{A}^{\mathsf{Sign}(s,\cdot)}(z,v) \wedge \mathsf{Ver}(v,m,\sigma) = 1 \wedge (m,\sigma) \notin Q] < \nu(\lambda)$$

where Q denotes the list of query-answer pairs for all queries asked by $\mathcal{A}$ to the oracle $\mathsf{Sign}(s,\cdot)$.

Definition 8 (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 exists a negligible function ν such that for sufficiently large $\lambda \in \mathbb{N}$,*

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

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 8 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.

Definition 9 (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 $\mathcal{NP}$-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:*

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

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:*

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

A proof/argument system $\Pi = (\mathcal{P}, \mathcal{V})$ for an $\mathcal{NP}$-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].

An interactive protocol $\Pi = (\mathcal{P}, \mathcal{V})$ is *public coin* if, at every round, $\mathcal{V}$ simply tosses a predetermined number of coins (random challenge) and sends the outcome to the prover.

We say that the transcript τ of an execution $b = \langle \mathcal{P}(z), \mathcal{V}\rangle(x)$ is *accepting* if $b = 1$.

Definition 10 (Special Honest-Verifier Zero Knowledge (SHVZK)). *Consider a public-coin proof/argument system $\Pi = (\mathcal{S}, \mathcal{P}, \mathcal{V})$ for an $\mathcal{NP}$-language L where the verifier sends m messages of length $\ell_1, \ldots, \ell_m$. We say that Π is SHVZK if there exists a PPT simulator algorithm $\mathcal{S}$ that on input any $x \in L$, security parameter 1^λ and any $c_1 \in \{0,1\}^{\ell_1}, \ldots, c_m \in \{0,1\}^{\ell_m}$, outputs a transcript for proving $x \in L$ where $c_1, \ldots, c_m$ are the messages of the verifier, such that the distribution of the output of $\mathcal{S}$ is computationally indistinguishable from the distribution of a transcript obtained when $\mathcal{V}$ sends $c_1, \ldots, c_m$ as challenges and $\mathcal{P}$ runs on common input x and any w such that $(x, w) \in \mathsf{Rel_L}$.*

Witness indistinguishability. Let $\mathsf{View}^{\mathcal{P}(w)}_{\mathcal{V}^\star(z)}(x)$ be the random variable that denotes $\mathcal{V}^\star$'s view in an interaction with $\mathcal{P}$ when $\mathcal{V}^\star$ is given auxiliary input z, $\mathcal{P}$ is given witness w, and both parties are given common input x.

Definition 11 (Witness Indistinguishability (WI)). *An argument/proof system $\Pi = (\mathcal{P}, \mathcal{V})$ for $\mathcal{NP}$-language L, is* Witness Indistinguishable (WI) *for the corresponding relation $\mathsf{Rel_L}$ if, for every malicious PPT verifier $\mathcal{V}^\star$, for all auxiliary input $z \in \{0,1\}^\star$ and for all x, w, w' such that $(x, w) \in \mathsf{Rel_L}$ and $(x, w') \in \mathsf{Rel_L}$, the following ensembles are computationally indistinguishable:*

$$\{\mathsf{View}^{\mathcal{P}(w)}_{\mathcal{V}^\star(z)}(x)\} \approx \{\mathsf{View}^{\mathcal{P}(w')}_{\mathcal{V}^\star(z)}(x)\}.$$

The notion of a *statistically* WI proof/argument system is obtained by requiring that the two ensembles $\{\text{View}^{\mathcal{P}(w)}_{\mathcal{V}^{\star}(z)}(x)\}$ and $\{\text{View}^{\mathcal{P}(w')}_{\mathcal{V}^{\star}(z)}(x)\}$ are statistically indistinguishable.

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

In this paper we also consider a definition where the adaptive-WI property of the argument/proof system 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 WI proof system is T-Witness Indistinguishable (T-WI).

Definition 12 (Proof of Knowledge [LP11b])**.** *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 ν, every statement $x \in \{0,1\}^{\lambda}$, every randomness $r \in \{0,1\}^{\star}$ and every auxiliary input $z \in \{0,1\}^{\star}$,*

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

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-input* PoK/AoK property. Adaptive-input PoK/AoK ensures that the PoK/AoK property still holds when a malicious prover can choose the statement adaptively at the last round. In this case, to be consistent with Definition 12 of PoK/AoK 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 $\mathcal{V}$, the randomness r of $\mathcal{P}^{\star}$ 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 PoK (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. LS also enjoys adaptive-input PoK and adaptive-input WI. We also use a 4-round delayed-input, adaptive-input AoK, and adaptive-input statistical WI argument of knowledge (WIAoK), that is a variant of LS [Fei90]. More in details, the WI of LS relies on the hiding property of the underlying commitment scheme, therefore if the prover of LS uses a 2-round statistically hiding commitment scheme, then we obtain adaptive-input statistical WIAoK. Note that this variation of LS requires an additional round from verifier to prover in order to send the first round of the statistically hiding commitment scheme. Finally, in this work we also use another 4-round variant of LS that relies on OWFs only. The additional round is indeed needed to instantiate the commitment scheme used in LS under any OWF. We finally stress that all variants of LS that we consider in this paper are SHVZK.

A.1 Commitment Schemes

Definition 13 (Commitment Scheme)**.** *Given a security parameter 1^{λ}, a commitment scheme* $\mathsf{CS} = (\mathsf{Sen}, \mathsf{Rec})$ *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 com. *In the decommitment phase,* Sen *sends to* Rec *a decommitment information* d *such that* Rec *accepts* m *as the decommitment of* com.

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

Correctness:

- *Commitment phase. Let* com *be the commitment of the message* m *given as output of an execution of* CS = (Sen, Rec) *where* Sen *runs on input a message* m. *Let* d *be the private output of* Sen *in this phase.*

- *Decommitment phase[18].* Rec *on input* m *and* d *accepts* m *as decommitment of* com.

Statistical (resp. Computational) Hiding([Lin10]): *for any adversary (resp.* 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 adversary (resp. PPT *adversary)* $\mathcal{A}$, *there exist a negligible function* ν, *s.t.:*

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

Statistical (resp. Computational) Binding: *for every commitment* com *generated during the commitment phase by a possibly malicious unbounded (resp. malicious* PPT*) sender* Sen* *there exists a negligible function* ν *such that* Sen*, *with probability at most* $\nu(\lambda)$, *outputs two decommitments* (m_0, d_0) *and* (m_1, d_1), *with* $m_0 \neq m_1$, *such that* Rec *accepts both decommitments.*

We also say that a commitment scheme is perfectly binding *iff* $\nu(\lambda) = 0$.

We also consider the definition of a commitment scheme where computational hiding 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.

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

Definition 14 (Extractable Commitment Scheme [PW09]). *A perfectly (resp. statistically) binding commitment scheme* ExCS = (ExSen, ExRec) *is an* extractable commitment scheme *if there exists an expected* PPT *extractor* ExtCom *that given oracle access to any malicious* PPT *sender* ExSen*, *outputs a pair* $(\tau, \sigma^\star)$ *such that the following two properties hold:*

---

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

- **Simulatability:** τ *is identically distributed to the view of* ExSen* *(when interacting with an honest* ExRec*) in the commitment phase.*
- **Extractability:** *the probability that there exists a decommitment of* τ *to* σ*, where* $\sigma \neq \sigma^\star$ *is 0 (resp. negligible).*

B Formal Proofs

B.1 Last Part of the Proof of 4-Round NMZK

The security proof goes through a sequence of hybrid experiments that prove that $\mathcal{A}_{\mathsf{ZK}}$ commits to $\tilde{s}_0$ s.t. $(\tilde{x}, \tilde{s}_0 \oplus \tilde{s}_1) \in \mathsf{Rel_L}$ during the simulated experiment. Once we have ensured that in all the hybrids the distribution of the message committed by $\mathcal{A}_{\mathsf{ZK}}$ does not change, we show that if the right session is accepting and $\mathsf{id} \neq \tilde{\mathsf{id}}$ we can recover the witness used by $\mathcal{A}_{\mathsf{ZK}}$ (that is internally executed by $\mathsf{Sim}_{\mathsf{ZK}}$).

Let p be the probability that in the real game $\mathcal{A}_{\mathsf{ZK}}$ concludes the left session. We start considering the hybrid $\mathcal{H}_1$ in which in the left session $\mathcal{P}_{\mathsf{ZK}}$ interacts with $\mathcal{A}_{\mathsf{ZK}}$ and in the right session $\mathcal{V}_{\mathsf{ZK}}$ interacts with $\mathcal{A}_{\mathsf{ZK}}$. We refer to this hybrid experiment as $\mathcal{H}_1(z)$. Details follow below.

$\mathcal{H}_1(z)$.

Left session:

1. Second round, upon receiving $(\mathsf{vk}, \pi^1_{\mathsf{sLS}}, \pi^1_{\mathsf{ext}})$ from $\mathcal{A}_{\mathsf{ZK}}$.

    1.1. Pick at random s_0.
    1.2. Run $\mathsf{Sen}_{\mathsf{ext}}$ on input 1^λ, id, π^1_{ext} and s_0 thus obtaining the 2nd round π^2_{ext} of Π_{ext}.
    1.3. Run $\mathcal{P}$ on input 1^λ, ℓ and π^1_{sLS} thus obtaining the 2nd round π^2_{sLS} of sLS.
    1.4. Pick a message $\mathtt{msg} \leftarrow \{0,1\}^\lambda$.
    1.5. Send $(\pi^2_{\mathsf{ext}}, \pi^2_{\mathsf{sLS}}, \mathtt{msg})$ to $\mathcal{A}_{\mathsf{ZK}}$.

2. Fourth round, upon receiving $(\pi^3_{\mathsf{ext}}, \pi^3_{\mathsf{sLS}}, \sigma, x, w)$ from $\mathcal{A}_{\mathsf{ZK}}$.

    2.1. If $\mathsf{Ver}(\mathsf{vk}, \mathtt{msg}, \sigma) \neq 1$ then abort, continue as follows otherwise.
    2.2. Set $s_1 = s_0 \oplus w$.
    2.3. Run $\mathsf{Sen}_{\mathsf{ext}}$ on input $(\pi^1_{\mathsf{ext}}, \pi^3_{\mathsf{ext}})$ thus obtaining the 4th round π^4_{ext} of Π_{ext} and the decommitment information $\mathtt{dec}_{\mathsf{ext}}$.
    2.4. Set $x_{\mathsf{sLS}} = (\pi^1_{\mathsf{ext}}, \pi^2_{\mathsf{ext}}, \pi^3_{\mathsf{ext}}, \pi^4_{\mathsf{ext}}, \mathsf{id}, \mathsf{vk}, x, s_1)$ and $w_{\mathsf{sLS}} = (s_0, \mathtt{dec}_{\mathsf{ext}}, \bot, \bot, \bot, \bot)$ with $|x_{\mathsf{sLS}}| = \ell$. Run $\mathcal{P}$ on input x_{sLS}, w_{sLS}, π^1_{sLS} and π^3_{sLS} thus obtaining the 4th round π^4_{sLS} of sLS.
    2.5. Send $(\pi^4_{\mathsf{ext}}, \pi^4_{\mathsf{sLS}}, s_1)$ to $\mathcal{A}_{\mathsf{ZK}}$.

Right session: act as a proxy between $\mathcal{A}_{\mathsf{ZK}}$ and $\mathcal{V}_{\mathsf{ZK}}$.

We now prove that in the right session of $\mathcal{H}_1(z)$ the MiM adversary $\mathcal{A}_{\mathsf{ZK}}$ does not complete successfully the right session committing to a message s'_0 s.t. $(\tilde{x}, \tilde{s}'_0 \oplus \tilde{s}_1) \notin \mathsf{Rel_L}$. More formally we want to prove the following claim.

Claim 3. *Let $\bar{p}$ be the probability that in the right session of $\mathcal{H}_1(z)$ $\mathcal{A}_{\mathsf{ZK}}$ successfully commits to a message s'_0 s.t. $(\tilde{x}, \tilde{s}'_0 \oplus \tilde{s}_1) \notin \mathsf{Rel_L}$, and the verifier outputs 1. Then $\bar{p} < \nu(\lambda)$ for some negligible function ν.*

The highl-level idea of the proof of this claim follows below. Suppose by contradiction that the claim does not hold, then we can construct an adversary $\mathcal{A}_\Sigma$ that breaks the security of the signature scheme Σ. Let vk be the challenge verification key. The idea of the security proof is to create an adversary $\mathcal{A}_\Sigma$ that interacts against the MiM adversary $\mathcal{A}_{\mathsf{ZK}}$ sending vk in the 1st round of the right session and extracting the witness used by $\mathcal{A}_{\mathsf{ZK}}$ to execute sLS. Because by contradiction we are assuming that $\mathcal{A}_{\mathsf{ZK}}$ does not commit to a witness then, with non-negligible probability, the witness extracted by sLS will be a pair of signatures (σ_1, σ_2) for a pair of different messages $(\mathsf{msg}_1, \mathsf{msg}_2)$ s.t. $\mathsf{Ver}(\mathsf{vk}, \mathsf{msg}_1, \sigma_1) = 1$ and $\mathsf{Ver}(\mathsf{vk}, \mathsf{msg}_2, \sigma_2) = 1$. From the above informal description one can notice that the formal proof strictly follows the one of Claim **??**, and for this reason we omit further details.

The 2nd hybrid that we consider is $\mathcal{H}_2(z)$ and it differs from $\mathcal{H}_1(z)$ only in the way the transcript of sLS is computed. In more details, by rewinding the adversary $\mathcal{A}_{\mathsf{ZK}}$ from the 3rd to the 2nd round it is possible to extract two signatures σ_1, σ_2 of two different messages $(\mathsf{msg}_1, \mathsf{msg}_2)$ and use them as a witness to execute the WIAoK sLS. As discussed earlier, after λ/p rewinds a second signature is obtained with overwhelming probability. For the above reason we can claim that the probability that in $\mathcal{H}_2(z)$ the output of the experiment is abort is statistically close to the probability that in $\mathcal{H}_1(z)$ the output of the experiment is abort. The formal description of $\mathcal{H}_2(z)$ is the following experiment.

$\mathcal{H}_2(z)$.
Left session:

1. Second round, upon receiving $(\mathsf{vk}, \pi^1_{\mathsf{sLS}}, \pi^1_{\mathsf{ext}})$ from $\mathcal{A}_{\mathsf{ZK}}$.

   1.1. Pick at random s_0.
   1.2. Run $\mathsf{Sen}_{\mathsf{ext}}$ on input $1^\lambda, \mathsf{id}, \pi^1_{\mathsf{ext}}$ and s_0 thus obtaining the 2nd round π^2_{ext} of Π_{ext}.
   1.3. Run $\mathcal{P}$ on input $1^\lambda, \ell$ and π^1_{sLS} thus obtaining the 2nd round π^2_{sLS} of sLS.
   1.4. Pick a message $\mathsf{msg}_1 \leftarrow \{0,1\}^\lambda$.
   1.5. Send $(\pi^2_{\mathsf{ext}}, \pi^2_{\mathsf{sLS}}, \mathsf{msg}_1)$ to $\mathcal{A}_{\mathsf{ZK}}$.

2. Fourth round, upon receiving $(\pi^3_{\mathsf{ext}}, \pi^3_{\mathsf{sLS}}, \sigma_1, x, w)$ from $\mathcal{A}_{\mathsf{ZK}}$.

   2.1. If $\mathsf{Ver}(\mathsf{vk}, \mathsf{msg}_1, \sigma) \neq 1$ then abort, continue as follows otherwise.
   2.2. Repeat Step 1.4, 1.5 and follow-up right-session messages up to λ/p times in order to obtain a signature σ_2 of a random message $\mathsf{msg}_2 \neq \mathsf{msg}_1$. Abort if case of failure in obtaining σ_2.
   2.3. Run $\mathsf{Sen}_{\mathsf{ext}}$ on input $(\pi^1_{\mathsf{ext}}, \pi^3_{\mathsf{ext}})$ thus obtaining the 4th round π^4_{ext} of Π_{ext}.
   2.4. Set $x_{\mathsf{sLS}} = (\pi^1_{\mathsf{ext}}, \pi^2_{\mathsf{ext}}, \pi^3_{\mathsf{ext}}, \pi^4_{\mathsf{ext}}, \mathsf{id}, \mathsf{vk}, x, s_1)$ and $w_{\mathsf{sLS}} = (\bot, \bot, \mathsf{msg}_1, \mathsf{msg}_2, \sigma_1, \sigma_2)$ with $|x_{\mathsf{sLS}}| = \ell$. Run $\mathcal{P}$ on input $x_{\mathsf{sLS}}, w_{\mathsf{sLS}}, \pi^1_{\mathsf{sLS}}$ and π^3_{sLS} thus obtaining the 4th round π^4_{sLS} of sLS.
   2.5. Set $s_1 = w \oplus s_0$.
   2.6. Send $(\pi^4_{\mathsf{ext}}, \pi^4_{\mathsf{sLS}}, s_1)$ to $\mathcal{A}_{\mathsf{ZK}}$.

Right session: act as a proxy between $\mathcal{A}_{\mathsf{ZK}}$ and $\mathcal{V}_{\mathsf{ZK}}$.

By the adaptive-input statistical WI of sLS the distribution of the message committed by $\mathcal{A}_{\mathsf{ZK}}$ does not change when moving from $\mathcal{H}_1(z)$ to $\mathcal{H}_2(z)$.

The next hybrid is $\mathcal{H}_3(z)$. The only differences between this hybrid and the previous one is that now $s_0 \oplus s_1$ is a random string. Formally $\mathcal{H}_3(z)$ is the following experiment.

$\mathcal{H}_3(z)$.
Left session:

1. Second round, upon receiving $(\mathsf{vk}, \pi^1_{\mathsf{sLS}}, \pi^1_{\mathsf{ext}})$ from $\mathcal{A}_{\mathsf{ZK}}$.

   1.1. Pick at random s_0.
   1.2. Run $\mathsf{Sen}_{\mathsf{ext}}$ on input 1^λ, $\mathtt{id}$ and s_0 thus obtaining the 2nd round π^2_{ext} of Π_{ext}.
   1.3. Run $\mathcal{P}$ on input 1^λ, ℓ and π^1_{sLS} thus obtaining the 2nd round π^2_{sLS} of sLS.
   1.4. Pick a message $\mathtt{msg}_1 \leftarrow \{0,1\}^\lambda$.
   1.5. Send $(\pi^2_{\mathsf{ext}}, \pi^2_{\mathsf{sLS}}, \mathtt{msg}_1)$ to $\mathcal{A}_{\mathsf{ZK}}$.

2. Fourth round, upon receiving $(\pi^3_{\mathsf{ext}}, \pi^3_{\mathsf{sLS}}, \sigma_1, x, w)$ from $\mathcal{A}_{\mathsf{ZK}}$.

   2.1. If $\mathsf{Ver}(\mathsf{vk}, \mathtt{msg}_1, \sigma_1) \neq 1$ then abort, continue as follows otherwise.
   2.2. Repeat Step 1.4, 1.5 and follow-up right-session messages up to λ/p times in order to obtain a signature σ_2 of a random message $\mathtt{msg}_2 \neq \mathtt{msg}_1$. Abort if case of failure in obtaining σ_2.
   2.3. Run $\mathsf{Sen}_{\mathsf{ext}}$ on input $(\pi^1_{\mathsf{ext}}, \pi^3_{\mathsf{ext}})$ thus obtaining the 4th round π^4_{ext} of Π_{ext}.
   2.4. Set $x_{\mathsf{sLS}} = (\pi^1_{\mathsf{ext}}, \pi^2_{\mathsf{ext}}, \pi^3_{\mathsf{ext}}, \pi^4_{\mathsf{ext}}, \mathtt{id}, \mathsf{vk}, x, s_1)$ and $w_{\mathsf{sLS}} = (\bot, \bot, \mathtt{msg}_1, \mathtt{msg}_2, \sigma_1, \sigma_2)$ with $|x_{\mathsf{sLS}}| = \ell$. Run $\mathcal{P}$ on input x_{sLS}, w_{sLS}, π^1_{sLS} and π^3_{sLS} thus obtaining the 4th round π^4_{sLS} of sLS.
   2.5. Pick at random s_1.
   2.6. Send $(\pi^4_{\mathsf{ext}}, \pi^4_{\mathsf{sLS}}, s_1)$ to $\mathcal{A}_{\mathsf{ZK}}$.

Right session: act as a proxy between $\mathcal{A}_{\mathsf{ZK}}$ and $\mathcal{V}_{\mathsf{ZK}}$.

Claim 4. *The distribution of the message committed by $\mathcal{A}_{\mathsf{ZK}}$ does not change between $\mathcal{H}_2(z)$ and $\mathcal{H}_3(z)$.*

Proof. Suppose by contradiction that the claim does not hold. Then $\mathcal{A}_{\mathsf{ZK}}$ in right session commits to a witness with non-negligible probability only when $\mathcal{P}_{\mathsf{ZK}}$ commits to a witness in the left session too. Based on this observation we can construct a distinguisher $\mathcal{D}_{\mathsf{ext}}$ and an adversary $\mathcal{A}_{\mathsf{ext}}$ that break the non-malleability of Π_{ext}. Let $\mathcal{C}_{\mathsf{ext}}$ be the challenger of the NM commitment scheme and let (m_0, m_1) be the two random challenge messages.

Loosely speaking $\mathcal{A}_{\mathsf{ext}}$ acts as $\mathcal{P}_{\mathsf{ZK}}$ with $\mathcal{A}_{\mathsf{ZK}}$ in the left session and as $\mathcal{V}_{\mathsf{ZK}}$ in the right session with the following differences: 1) $\mathcal{A}_{\mathsf{ext}}$ plays as proxy between $\mathcal{C}_{\mathsf{ext}}$ and $\mathcal{A}_{\mathsf{ZK}}$ w.r.t. messages of Π_{ext} in the main thread; 2) a second signature is extracted from the left session through rewinds; 3) random strings are played to simulate the receiver of Π_{ext} during rewinds. 4) $\mathcal{A}_{\mathsf{ext}}$ in the last round of the left session sends s_1 s.t. $s_1 = m_0 \oplus w$.

Then $\mathcal{D}_{\mathsf{ext}}$, on input the message $\tilde{m}$ committed by $\mathcal{A}_{\mathsf{ext}}$ and his randomness, reconstructs the view of $\mathcal{A}_{\mathsf{ZK}}$ and recovers the adaptively chosen statement $\tilde{x}$ proved by $\mathcal{A}_{\mathsf{ZK}}$ and the messages $\tilde{s}_1$ sent by $\mathcal{A}_{\mathsf{ZK}}$ in the last round. If $\tilde{s}_1 \oplus \tilde{m}$ is s.t. $(\tilde{x}, \tilde{s}_1 \oplus \tilde{m}) \in \mathsf{Rel}_{\mathsf{L}}$ then $\mathcal{D}_{\mathsf{ext}}$ outputs 0, and a random bit otherwise. Since by contradiction $\mathcal{A}_{\mathsf{ZK}}$ commits to the witness for $\tilde{x}$ with overwhelming probability only when $\mathcal{P}_{\mathsf{ZK}}$ commits to a witness for x, then $\mathcal{D}_{\mathsf{ext}}$ can tell apart which message has ben committed by the MiM adversary $\mathcal{A}_{\mathsf{ext}}$. We notice that the reduction queries to query only once the receiver of Π_{ext} involved in the reduction. Formally the adversary $\mathcal{A}_{\mathsf{ext}}$ acts as follows.

$\mathcal{A}_{\mathsf{ext}}(m_0, m_1, z)$.
   Set $\mathtt{round}_2 = \bot$, $\mathtt{round}_3 = \bot$.
   Left session:

   1. Upon receiving $(\mathsf{vk}, \pi^1_{\mathsf{sLS}}, \pi^1_{\mathsf{ext}})$ from $\mathcal{A}_{\mathsf{ZK}}$ forward π^1_{ext} to $\mathcal{C}_{\mathsf{ext}}$.

   2. Upon receiving π^2_{ext} from $\mathcal{C}_{\mathsf{ext}}$.

2.1. Run $\mathcal{P}$ on input 1^λ, ℓ and π_{sLS}^1 thus obtaining the 2nd round π_{sLS}^2 of sLS.

2.2. Pick a message $\mathsf{msg}_1 \leftarrow \{0,1\}^\lambda$.

2.3. Send $(\pi_{\mathsf{ext}}^2, \pi_{\mathsf{sLS}}^2, \mathsf{msg}_1)$ to $\mathcal{A}_{\mathsf{ZK}}$.

3. Upon receiving $(\pi_{\mathsf{ext}}^3, \pi_{\mathsf{sLS}}^3, \sigma_1, x, w)$ from $\mathcal{A}_{\mathsf{ZK}}$.

   3.1. If $\mathsf{Ver}(\mathsf{vk}, \mathsf{msg}_1, \sigma) \neq 1$ then abort, continue with the following steps otherwise.

   3.2. Repeat Step 2.2, 2.3 and follow-up right-session messages up to λ/p times in order to obtain a signature σ_2 of a random message $\mathsf{msg}_2 \neq \mathsf{msg}_1$. Abort in case of failure in obtaining σ_2.

   3.3. Set $x_{\mathsf{sLS}} = (\pi_{\mathsf{ext}}^1, \pi_{\mathsf{ext}}^2, \pi_{\mathsf{ext}}^3, \pi_{\mathsf{ext}}^4, \mathsf{id}, \mathsf{vk}, x, s_1)$ and $w_{\mathsf{sLS}} = (\bot, \bot, \mathsf{msg}_1, \mathsf{msg}_2, \sigma_1, \sigma_2)$ with $|x_{\mathsf{sLS}}| = \ell$. Run $\mathcal{P}$ on input $x_{\mathsf{sLS}}, w_{\mathsf{sLS}}, \pi_{\mathsf{sLS}}^3$ and π_{sLS}^3 thus obtaining the 4th round π_{sLS}^4 of sLS.

   3.4. Set $s_1 = m_0 \oplus w$.

   3.5. Send $(\pi_{\mathsf{ext}}^4, \pi_{\mathsf{sLS}}^4, s_1)$ to $\mathcal{A}_{\mathsf{ZK}}$.

Right session:

1. Upon receiving $\tilde{\pi}_{\mathsf{ext}}^1$ from from $\mathsf{Rec}_{\mathsf{ext}}$.

   1.1. Run $(\tilde{\mathsf{sk}}, \tilde{\mathsf{vk}}) \leftarrow \mathsf{Gen}(1^\lambda)$.

   1.2. Run $\mathcal{V}$ on input 1^λ thus obtaining the 1st round $\tilde{\pi}_{\mathsf{sLS}}^1$ of sLS.

   1.3. Send $(\tilde{\mathsf{vk}}, \tilde{\pi}_{\mathsf{sLS}}^1, \tilde{\pi}_{\mathsf{ext}}^1)$ to $\mathcal{A}_{\mathsf{ZK}}$.

2. Upon receiving $(\tilde{\pi}_{\mathsf{ext}}^2, \tilde{\pi}_{\mathsf{sLS}}^2, \tilde{\mathsf{msg}})$ from $\mathcal{A}_{\mathsf{ZK}}$.

   2.1. If there is no rewind phase on the left-session then send $\tilde{\pi}_{\mathsf{ext}}^2$ to $\mathsf{Rec}_{\mathsf{ext}}$ and execute the following steps. execute the following steps.

      i. Upon receiving $\tilde{\pi}_{\mathsf{ext}}^3$ from $\mathsf{Rec}_{\mathsf{wsyn}}$.

      ii. Run $\mathcal{V}$ on input $\tilde{\pi}_{\mathsf{sLS}}^2$ thus obtaining the 3rd round $\tilde{\pi}_{\mathsf{sLS}}^3$ of sLS.

      iii. Run $\mathsf{Sign}(\tilde{\mathsf{sk}}, \tilde{\mathsf{msg}}_1)$ to obtain a signature $\tilde{\sigma}_1$ of the message $\tilde{\mathsf{msg}}_1$.

      iv. Send $(\tilde{\pi}_{\mathsf{ext}}^3, \tilde{\pi}_{\mathsf{sLS}}^3, \tilde{\sigma}_1)$.

   2.2. Else if there is a rewind phase in the left-session then execute the following steps.

      i. Run $\mathcal{V}$ on input $\tilde{\pi}_{\mathsf{sLS}}^2$ thus obtaining the 3rd round $\tilde{\pi}_{\mathsf{sLS}}^3$ of sLS.

      ii. Run $\mathsf{Sign}(\tilde{\mathsf{sk}}, \tilde{\mathsf{msg}}_1)$ to obtain a signature $\tilde{\sigma}_1$ of the message $\tilde{\mathsf{msg}}_1$.

      iii. Compute a random $\tilde{\pi}_{\mathsf{ext}}^3$.

      iv. Send $(\tilde{\pi}_{\mathsf{ext}}^3, \tilde{\pi}_{\mathsf{sLS}}^3, \tilde{\sigma}_1)$ to $\mathcal{A}_{\mathsf{ZK}}$.

3. Upon receiving $(\tilde{\pi}_{\mathsf{ext}}^4, \tilde{\pi}_{\mathsf{sLS}}^4, \tilde{s}_1, \tilde{x})$ from $\mathcal{A}_{\mathsf{ZK}}$.

   3.1. Set $\tilde{x}_{\mathsf{sLS}} = (\tilde{\pi}_{\mathsf{ext}}^1, \tilde{\pi}_{\mathsf{ext}}^2, \tilde{\pi}_{\mathsf{ext}}^3, \tilde{\pi}_{\mathsf{ext}}^4, \tilde{\mathsf{id}}, \tilde{\mathsf{vk}}, \tilde{x}, \tilde{s}_1)$ and abort iff $(\tilde{\pi}_{\mathsf{sLS}}^1, \tilde{\pi}_{\mathsf{sLS}}^2, \tilde{\pi}_{\mathsf{sLS}}^3, \tilde{\pi}_{\mathsf{sLS}}^4)$ is not accepting for $\mathcal{V}$ with respect to $\tilde{x}$.

   3.2. Send $\tilde{\pi}_{\mathsf{ext}}^4$ to $\mathsf{Rec}_{\mathsf{ext}}$.

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

$\mathcal{D}_{\mathsf{ext}}(\mathsf{mim}^{\mathcal{A}_{\mathsf{ext}}}(z))$: Let $\tilde{m}$ be the committed message sent in the right session by $\mathcal{A}_{\mathsf{ext}}$ to $\mathcal{V}_{\mathsf{ZK}}$. Reconstruct the view of $\mathcal{A}_{\mathsf{ZK}}$ (using randomness in $\mathsf{mim}^{\mathcal{A}_{\mathsf{ext}}}(z)$) and recover the adaptively chosen statement $\tilde{x}$ proved by $\mathcal{A}_{\mathsf{ZK}}$ and the messages $\tilde{s}_1$ sent by $\mathcal{A}_{\mathsf{ZK}}$ in the last round. Since by contradiction $\mathcal{A}_{\mathsf{ZK}}$ contradicts the claim, we have that $\mathcal{A}_{\mathsf{ext}}$ breaks the non-malleability of Π_{ext} because $(\tilde{x}, \tilde{s}_1 \oplus \tilde{m}) \in \mathsf{Rel}_L$ with non-negligible probability in $\mathcal{H}_2(z)$ where $m_0 = \tilde{m}$ is committed in com,

while the same happens with negligible probability only in $\mathcal{H}_3(z)$ where m_1 is a random string. Therefore if $(\tilde{x}, \tilde{s}_1 \oplus \tilde{m}) \in \mathsf{Rel}_\mathsf{L}$ then $\mathcal{A}_\mathsf{ext}$ outputs 0 otherwise $\mathcal{A}_\mathsf{ext}$ outputs a random bit.

The proof is concluded by observing that if $\mathcal{C}_\mathsf{ext}$ commits to m_0 then the above execution of $\mathcal{A}_\mathsf{ext}$ corresponds to $\mathcal{H}_2(z)$, otherwise it corresponds to $\mathcal{H}_3(z)$. $\qquad\square$

We now describe how Sim_ZK of Figure 6 works. Let $\mathtt{ExtCom}$ be the extractor of Π_ext. Sim_ZK runs $\mathtt{ExtCom}$ in order to get the witness $\tilde{w}$ s.t. $(\tilde{x}, \tilde{w}) \in \mathsf{Rel}_\mathsf{L}$, where $\tilde{x}$ is the adaptively chosen theorem proved by $\mathcal{A}_\mathsf{ZK}$. Before formally describing Sim_ZK we need to construct an augmented machine $\mathcal{M}_\mathsf{ext}$ that is a malicious sender that will be black-box accessed by $\mathtt{ExtCom}$.

$\mathcal{M}_\mathsf{ext}(1^\lambda, z)$.

Run $\mathcal{A}_\mathsf{ZK}$ with randomness φ.

Left session: Interact with $\mathcal{A}_\mathsf{ZK}$ as in $\mathcal{H}_3(z)$.

Right session:

1. Upon receiving $\tilde{\pi}^1_\mathsf{ext}$ from from $\mathsf{Rec}_\mathsf{ext}$.
    1.1. Run $(\tilde{\mathsf{sk}}, \tilde{\mathsf{vk}}) \leftarrow \mathsf{Gen}(1^\lambda)$.
    1.2. Run $\mathcal{V}$ on input 1^λ thus obtaining the 1st round $\tilde{\pi}^1_\mathsf{sLS}$ of sLS.
    1.3. Send $(\tilde{\mathsf{vk}}, \tilde{\pi}^1_\mathsf{sLS}, \tilde{\pi}^1_\mathsf{ext})$ to $\mathcal{A}_\mathsf{ZK}$.

2. Upon receiving $(\tilde{\pi}^2_\mathsf{ext}, \tilde{\pi}^2_\mathsf{sLS}, \tilde{\mathtt{msg}})$ from $\mathcal{A}_\mathsf{ZK}$.
    2.1. If there is no rewind phase on the left-session then send $\tilde{\pi}^2_\mathsf{ext}$ to $\mathsf{Rec}_\mathsf{ext}$ and execute the following steps. execute the following steps.
        i. Upon receiving $\tilde{\pi}^3_\mathsf{ext}$ from $\mathsf{Rec}_\mathsf{wsyn}$.
        ii. Run $\mathcal{V}$ on input $\tilde{\pi}^2_\mathsf{sLS}$ thus obtaining the 3rd round $\tilde{\pi}^3_\mathsf{sLS}$ of sLS.
        iii. Run $\mathsf{Sign}(\tilde{\mathsf{sk}}, \tilde{\mathtt{msg}}_1)$ to obtain a signature $\tilde{\sigma}_1$ of the message $\tilde{\mathtt{msg}}_1$.
        iv. Send $(\tilde{\pi}^3_\mathsf{ext}, \tilde{\pi}^3_\mathsf{sLS}, \tilde{\sigma}_1)$.
    2.2. Else if there is a rewind phase in the left-session then execute the following steps.
        i. Run $\mathcal{V}$ on input $\tilde{\pi}^2_\mathsf{sLS}$ thus obtaining the 3rd round $\tilde{\pi}^3_\mathsf{sLS}$ of sLS.
        ii. Run $\mathsf{Sign}(\tilde{\mathsf{sk}}, \tilde{\mathtt{msg}}_1)$ to obtain a signature $\tilde{\sigma}_1$ of the message $\tilde{\mathtt{msg}}_1$.
        iii. Compute a random $\tilde{\pi}^3_\mathsf{ext}$.
        iv. Send $(\tilde{\pi}^3_\mathsf{ext}, \tilde{\pi}^3_\mathsf{sLS}, \tilde{\sigma}_1)$ to $\mathcal{A}_\mathsf{ZK}$.

3. Upon receiving $(\tilde{\pi}^4_\mathsf{ext}, \tilde{\pi}^4_\mathsf{sLS}, \tilde{s}_1, \tilde{x})$ from $\mathcal{A}_\mathsf{ZK}$.
    3.1. Set $\tilde{x}_\mathsf{sLS} = (\tilde{\pi}^1_\mathsf{ext}, \tilde{\pi}^2_\mathsf{ext}, \tilde{\pi}^3_\mathsf{ext}, \tilde{\pi}^4_\mathsf{ext}, \tilde{\mathtt{id}}, \tilde{\mathsf{vk}}, \tilde{x}, \tilde{s}_1)$ and abort iff $(\tilde{\pi}^1_\mathsf{sLS}, \tilde{\pi}^2_\mathsf{sLS}, \tilde{\pi}^3_\mathsf{sLS}, \tilde{\pi}^4_\mathsf{sLS})$ is not accepting for $\mathcal{V}$ with respect to $\tilde{x}$.
    3.2. Send $\tilde{\pi}^4_\mathsf{ext}$ to $\mathsf{Rec}_\mathsf{ext}$.

Similarly to the black-box simulator of [GK96] we assume w.l.o.g. that if a transcript τ appears in the final output of a black-box extractor $\mathtt{ExtCom}$, then $\mathtt{ExtCom}$ has queried the sender of the extractable commitment Π_ext on every prefix of τ. Sim_ZK, in order to reconstruct the full transcript T of the entire execution, interacts in the right session with $\mathcal{A}_\mathsf{ZK}$ by playing messages of τ. See Figure 6 for more details.

Input: Security parameters: λ, auxiliary input: z.

1. Run ExtCom using $\mathcal{M}_{\mathsf{ext}}(1^\lambda, z)$ as a sender, and let $(\tilde{w}, \mathsf{View}_{\mathsf{ext}})$ be the output of ExtCom where $\tilde{w}$ denote the extracted value and $\mathsf{View}_{\mathsf{ext}}$ is the view of $\mathcal{M}_{\mathsf{ext}}(1^\lambda, z)$ that contains the transcript $\tau = (\tilde{\pi}^1_{\mathsf{ext}}, \tilde{\pi}^2_{\mathsf{ext}}, \tilde{\pi}^3_{\mathsf{ext}}, \tilde{\pi}^4_{\mathsf{ext}})$ (see App. A.1).

2. Use the same randomness φ used by $\mathcal{M}_{\mathsf{ext}}(1^\lambda, z)$ and $\mathcal{A}_{\mathsf{ZK}}$, and reconstruct the view View of $\mathcal{A}_{\mathsf{ZK}}$ by executing the following steps.

    2.1. Run $\mathcal{A}_{\mathsf{ZK}}$.
    2.2. Interact in the left session with $\mathcal{A}_{\mathsf{ZK}}$ as in $\mathcal{H}_3(z)$.
    2.3. Run $(\tilde{\mathsf{sk}}, \tilde{\mathsf{vk}}) \leftarrow \mathsf{Gen}(1^\lambda)$.
    2.4. Run $\mathcal{V}$ on input 1^λ thus obtaining the 1st round $\tilde{\pi}^1_{\mathsf{sLS}}$ of sLS.
    2.5. Send $(\tilde{\mathsf{vk}}, \tilde{\pi}^1_{\mathsf{sLS}}, \tilde{\pi}^1_{\mathsf{ext}})$ to $\mathcal{A}_{\mathsf{ZK}}$.
    2.6. Upon receiving $(\tilde{\pi}^2_{\mathsf{ext}}, \tilde{\pi}^2_{\mathsf{sLS}}, \tilde{\mathsf{msg}})$ from $\mathcal{A}_{\mathsf{ZK}}$.

        i. If there is no rewind phase in the left session then execute the following steps.
            A. Run $\mathcal{V}$ on input $\tilde{\pi}^2_{\mathsf{sLS}}$ thus obtaining the 3rd round $\tilde{\pi}^3_{\mathsf{sLS}}$ of sLS.
            B. Run $\mathsf{Sign}(\tilde{\mathsf{sk}}, \tilde{\mathsf{msg}}_1)$ to obtain a signature $\tilde{\sigma}_1$ of the message $\tilde{\mathsf{msg}}_1$.
            C. Send $(\tilde{\pi}^3_{\mathsf{ext}}, \tilde{\pi}^3_{\mathsf{sLS}}, \tilde{\sigma}_1)$ to $\mathcal{A}_{\mathsf{ZK}}$.
        ii. Else if there is a rewind phase in the left-session then execute the following steps.
            A. Run $\mathcal{V}$ on input $\tilde{\pi}^2_{\mathsf{sLS}}$ thus obtaining the 3rd round $\tilde{\pi}^3_{\mathsf{sLS}}$ of sLS.
            B. Run $\mathsf{Sign}(\tilde{\mathsf{sk}}, \tilde{\mathsf{msg}}_1)$ to obtain a signature $\tilde{\sigma}_1$ of the message $\tilde{\mathsf{msg}}_1$.
            C. Compute a random third round $\tilde{\pi}^\star_{\mathsf{ext}}$ of Π_{ext}.
            D. Send $(\tilde{\pi}^\star_{\mathsf{ext}}, \tilde{\pi}^3_{\mathsf{sLS}}, \tilde{\sigma}_1)$ to $\mathcal{A}_{\mathsf{ZK}}$.

    2.7. Upon receiving $(\tilde{\pi}^4_{\mathsf{ext}}, \tilde{\pi}^4_{\mathsf{sLS}}, \tilde{s}_1, \tilde{x})$ from $\mathcal{A}_{\mathsf{ZK}}$, set $\tilde{x}_{\mathsf{sLS}} = (\tilde{\pi}^1_{\mathsf{ext}}, \tilde{\pi}^2_{\mathsf{ext}}, \tilde{\pi}^3_{\mathsf{ext}}, \tilde{\pi}^4_{\mathsf{ext}}, \tilde{\mathsf{id}}, \tilde{\mathsf{vk}}, \tilde{x}, \tilde{s}_1)$ and abort iff $(\tilde{\pi}^1_{\mathsf{sLS}}, \tilde{\pi}^2_{\mathsf{sLS}}, \tilde{\pi}^3_{\mathsf{sLS}}, \tilde{\pi}^4_{\mathsf{sLS}})$ is not accepting for $\mathcal{V}$ with respect to $\tilde{x}$.

3. Let T be the transcript of the main thread in the above execution. Output $(\mathsf{View} = (\varphi, T), \tilde{w})$.

Figure 6: The simulator $\mathsf{Sim}_{\mathsf{ZK}}$.

B.2 Proof of NM of the 3-Round NM Commitment Scheme

We now formally prove that the commitment scheme Π_{NMCom} is non-malleable. This security proof consists of two parts. In the first part we consider a MiM adversary $\mathcal{A}_{\mathsf{NMCom}}$ that interacts only in a synchronized way with NMSen and NMRec showing that our scheme is synchronous one-one non-malleable. In the second part we argue that the commitment scheme is non-malleable also when $\mathcal{A}$ acts in a non-synchronized way. Putting together these two arguments we are able to conclude the proof on non-malleability.

Lemma 7. Π_{NMCom} *is a synchronous one-one NM commitment scheme.*

Proof. We show that for all $m \in \{0,1\}^{\mathsf{poly}(\lambda)}$ it holds that:

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

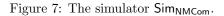where $\mathsf{Sim}_{\mathsf{NMCom}}$ is the simulator depicted in Fig. 7.

In the first experiment, in the left session NMSen commits to m playing with $\mathcal{A}_{\mathsf{NMCom}}$, while in the right session $\mathcal{A}_{\mathsf{NMCom}}$ commits on the right by playing with NMRec. We refer to this experiment as $\mathcal{H}^m_1(z)$. Details follow below.

Common input: security parameters: λ, $(\lambda_{\mathsf{wsyn}}, \lambda_{\mathsf{LS}}, \ell) = \mathsf{Params}(\lambda)$.
Identity: $\mathtt{id} \in \{0,1\}^{\lambda}$.
Internal simulation of the left session:

1. Run $\mathsf{Sen}_{\mathsf{wsyn}}$ on input $1^{\lambda_{\mathsf{wsyn}}}$, $\mathtt{id}$ and 0^{λ} thus obtaining the first round $\mathsf{a}_{\mathsf{wsyn}}$ of Π_{wsyn}.

2. Run $\mathcal{P}$ on input $1^{\lambda_{\mathsf{LS}}}$ and ℓ thus obtaining the first round a_{LS} of LS.

3. Send $(\mathsf{a}_{\mathsf{wsyn}}, \mathsf{a}_{\mathsf{LS}})$ to $\mathcal{A}_{\mathsf{NMCom}}$.

4. Upon receiving $(\mathsf{c}_{\mathsf{wsyn}}, \mathsf{c}_{\mathsf{LS}}, Y)$ from $\mathcal{A}_{\mathsf{NMCom}}$.

    4.1. Run $\mathsf{Sen}_{\mathsf{wsyn}}$ on input $\mathsf{c}_{\mathsf{wsyn}}$ thus obtaining the third round $\mathsf{z}_{\mathsf{wsyn}}$ of Π_{wsyn}.
    4.2. Run $\mathsf{Sen}_{\mathsf{wsyn}}$ thus obtaining the decommitment information $\mathsf{dec}_{\mathsf{wsyn}}$ of Π_{wsyn}.
    4.3. Set $x = \big(\mathsf{a}_{\mathsf{wsyn}}, \mathsf{c}_{\mathsf{wsyn}}, \mathsf{z}_{\mathsf{wsyn}}, Y, \mathtt{id}\big)$ and $w = (m, \mathsf{dec}_{\mathsf{wsyn}}, \perp)$ with $|x| = \ell$. Run $\mathcal{P}$ on input x, w, and c_{LS} thus obtaining the third round z_{LS} of LS.
    4.4. Send $(\mathsf{z}_{\mathsf{wsyn}}, \mathsf{z}_{\mathsf{LS}})$ to $\mathcal{A}_{\mathsf{NMCom}}$.

Stand-alone commitment:

1. $\mathsf{Sim}_{\mathsf{NMCom}}$ acts as a proxy between $\mathcal{A}_{\mathsf{NMCom}}$ and NMRec.

Figure 7: The simulator $\mathsf{Sim}_{\mathsf{NMCom}}$.

$\mathcal{H}_1^m(z)$.
  Left session:

1. First round.

    1.1. Run $\mathsf{Sen}_{\mathsf{wsyn}}$ on input $1^{\lambda_{\mathsf{wsyn}}}$, $\mathtt{id}$ and m thus obtaining the first round $\mathsf{a}_{\mathsf{wsyn}}$ of Π_{wsyn}.
    1.2. Run $\mathcal{P}$ on input $1^{\lambda_{\mathsf{LS}}}$ and ℓ thus obtaining the first round a_{LS} of LS.
    1.3. Send $(\mathsf{a}_{\mathsf{wsyn}}, \mathsf{a}_{\mathsf{LS}})$ to $\mathcal{A}_{\mathsf{NMCom}}$.

2. Third round, upon receiving $(\mathsf{c}_{\mathsf{wsyn}}, \mathsf{c}_{\mathsf{LS}}, Y)$ from $\mathcal{A}_{\mathsf{NMCom}}$, run as follows.

    2.1. Run $\mathsf{Sen}_{\mathsf{wsyn}}$ on input $\mathsf{c}_{\mathsf{wsyn}}$ thus obtaining the third round $\mathsf{z}_{\mathsf{wsyn}}$ of Π_{wsyn} and the decommitment information $\mathsf{dec}_{\mathsf{wsyn}}$.
    2.2. Set $x = (\mathsf{a}_{\mathsf{wsyn}}, \mathsf{c}_{\mathsf{wsyn}}, \mathsf{z}_{\mathsf{wsyn}}, Y, \mathtt{id})$ and $w = (m, \mathsf{dec}_{\mathsf{wsyn}}, \perp)$ with $|x| = \ell$. Run $\mathcal{P}$ on input x, w and c_{LS} thus obtaining the third round z_{LS} of LS.
    2.3. Send $(\mathsf{z}_{\mathsf{wsyn}}, \mathsf{z}_{\mathsf{LS}})$ to $\mathcal{A}_{\mathsf{NMCom}}$.

Right session: act as a proxy between $\mathcal{A}_{\mathsf{NMCom}}$ and NMRec.

The distribution of $\mathsf{mim}_{\mathcal{H}_1^m}^{\mathcal{A}_{\mathsf{NMCom}}}(z)$ clearly corresponds to the distribution of $\mathsf{mim}_{\Pi_{\mathsf{NMCom}}}^{\mathcal{A}_{\mathsf{NMCom}}, m}(z)$. We now prove that in the right session $\mathcal{A}_{\mathsf{NMCom}}$ does not commit to a message $\tilde{m} = \perp$. We can do so by proving that the LS proof of the right session is computed by $\mathcal{A}_{\mathsf{NMCom}}$ without using as witness a value $\tilde{y}$ s.t. $f(\tilde{y}) = \tilde{Y}$, where $\tilde{Y}$ is the value sent to $\mathcal{A}_{\mathsf{NMCom}}$ in the second round of the right session. Formally we want have the following claim.

Claim 5. *Let $\bar{p}$ be the probability that in the right session $\mathcal{A}_{\mathsf{NMCom}}$ successfully commits to $\tilde{m} = \perp$. Then $\bar{p} < \nu(\lambda)$ for some negligible function ν.*

Proof. Suppose by contradiction that the claim does not hold, then 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 $Y \in \{0,1\}^{\lambda}$ and sends it to $\mathcal{A}_f$. $\mathcal{A}_f$ wins if it gives as output y s.t. $Y = f(y)$.

Before describing the adversary we need to consider the augmented machine $\mathcal{M}_f$ that will be used by $\mathcal{A}_f$ to extract the witness from LS by using the extractor E (that exists from the property of adaptive-input PoK enjoyed by LS). Recall that in the case of an adaptive-input PoK, the extractor takes as input the randomnesses r of the prover and r' of the verifier of an execution of LS when theorem x has been proved by $\mathcal{P}^\star$. Now we are ready to describe how $\mathcal{M}_f$ works. $\mathcal{M}_f$ internally runs $\mathcal{A}_{\mathsf{NMCom}}$ with randomness r and interacts with him as the sender NMSen does in the left session and as the receiver NMRec does in the right session. The only difference is that all messages of LS of the right session are forwarded to the verifier $\mathcal{V}$ and vice versa. Formally $\mathcal{M}_f$ acts as follows.

$\mathcal{M}_f(z, Y, r)$.

Execute the following steps with randomness r

- Run NMSen on input m with $\mathcal{A}_{\mathsf{NMCom}}$ as in $\mathcal{H}_1^m(z)$.

- Upon receiving $(\tilde{a}_{\mathsf{wsyn}}, \tilde{a}_{\mathsf{LS}})$ from $\mathcal{A}_{\mathsf{NMCom}}$, send $\tilde{a}_{\mathsf{LS}}$ to $\mathcal{V}$.

- Upon receiving $\tilde{c}_{\mathsf{LS}}$ from $\mathcal{V}$, run as follows.

  1. Run $\mathsf{Rec}_{\mathsf{wsyn}}$ on input $\tilde{\mathsf{id}}$ and $\tilde{a}_{\mathsf{wsyn}}$ thus obtaining the second round $\tilde{c}_{\mathsf{wsyn}}$ of Π_{wsyn}.
  2. Set $\tilde{Y} = Y$.
  3. Send $(\tilde{c}_{\mathsf{wsyn}}, \tilde{c}_{\mathsf{LS}}, \tilde{Y})$ to $\mathcal{A}_{\mathsf{NMCom}}$.

- Upon receiving the 3rd round of the right session $(\tilde{z}_{\mathsf{wsyn}}, \tilde{z}_{\mathsf{LS}})$ set $\tilde{x} = (\tilde{a}_{\mathsf{wsyn}}, \tilde{c}_{\mathsf{wsyn}}, \tilde{z}_{\mathsf{wsyn}}, \tilde{Y}, \tilde{\mathsf{id}})$ and send $(\tilde{z}_{\mathsf{LS}}, \tilde{x})$ to $\mathcal{V}$.

Now we can conclude the proof of this claim by describing how $\mathcal{A}_f$ works. $\mathcal{A}_f$ runs E on input the randomness r' (used by the verifier in an execution where x has been proved) and uses $\mathcal{M}_f$ as prover with randomness r (recall that an extractor of LS plays only having access to a prover of LS). Notice that the above execution of $\mathcal{M}_f$ is distributed identically to $\mathcal{H}_1^m(z)$. Since by contradiction $\mathcal{A}_{\mathsf{NMCom}}$ is successful with non-negligible probability, we have that with non-negligible probability $\mathcal{A}_f$ in polynomial time[19] outputs the value y such that $f(y) = Y$. $\square$

The next hybrid experiment that we consider is $\mathcal{H}_1^0(z)$ that is equal to $\mathcal{H}_1^m(z)$ with the only difference that the message committed using Π_{wsyn} is 0^λ instead of m. Similarly to $\mathcal{H}_1^m(z)$, we have for $\mathcal{H}_1^0(z)$ the following claim.

Claim 6. *The probability that in the right session $\mathcal{A}_{\mathsf{NMCom}}$ successfully commits to a message $\tilde{m} = \perp$ is $p < \nu(\lambda)$ for some negligible function ν.*

Proof. The security proof strictly follows the one of Claim 5. $\square$

The next hybrid that we consider is $\mathcal{H}_2^m(z)$. $\mathcal{H}_2^m(z)$ differs from $\mathcal{H}_1^m(z)$ only in the witness used to compute the LS transcript. Formally $\mathcal{H}_2^m(z)$ is the following experiment.

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

---

[19]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 is used in another proofs but for simplicity we will not repeat this discussion.

Left session:

1. First round

   1.1. Run $\mathsf{Sen}_{\mathsf{wsyn}}$ on input $1^{\lambda_{\mathsf{wsyn}}}$, $\mathtt{id}$ and m thus obtaining the first round $\mathsf{a}_{\mathsf{wsyn}}$ of Π_{wsyn}.

   1.2. Run $\mathcal{P}$ on input $1^{\lambda_{\mathsf{LS}}}$ and ℓ thus obtaining the first round a_{LS} of LS.

   1.3. Send $(\mathsf{a}_{\mathsf{wsyn}}, \mathsf{a}_{\mathsf{LS}})$ to $\mathcal{A}_{\mathsf{NMCom}}$.

2. Third round, upon receiving $(\mathsf{c}_{\mathsf{wsyn}}, \mathsf{c}_{\mathsf{LS}}, Y)$ from $\mathcal{A}_{\mathsf{NMCom}}$, run as follows.

   2.1. Run in time $\tilde{T}_f$ to compute y s.t. $Y = f(y)$.

   2.2. Run $\mathsf{Sen}_{\mathsf{wsyn}}$ on input $\mathsf{c}_{\mathsf{wsyn}}$ thus obtaining the third round $\mathsf{z}_{\mathsf{wsyn}}$ of Π_{wsyn}.

   2.3. Set $x = (\mathsf{a}_{\mathsf{wsyn}}, \mathsf{c}_{\mathsf{wsyn}}, \mathsf{z}_{\mathsf{wsyn}}, Y, \mathtt{id})$ and $w = (\bot, \bot, y)$ with $|x| = \ell$. Run $\mathcal{P}$ on input x, w and c_{LS} thus obtaining the third round z_{LS} of LS.

   2.4. Send $(\mathsf{z}_{\mathsf{wsyn}}, \mathsf{z}_{\mathsf{LS}})$ to $\mathcal{A}_{\mathsf{NMCom}}$.

Right session: act as a proxy between $\mathcal{A}_{\mathsf{NMCom}}$ and NMRec.

Claim 7. *For all $m \in \{0,1\}^{\mathsf{poly}(\lambda)}$ it holds that $\mathsf{mim}_{\mathcal{H}_1^m}^{\mathcal{A}_{\mathsf{NMCom}}}(z) \approx \mathsf{mim}_{\mathcal{H}_2^m}^{\mathcal{A}_{\mathsf{NMCom}}}(z)$.*

Proof. Suppose by contradiction that there exist adversary $\mathcal{A}_{\mathsf{NMCom}}$ and a distinguisher $\mathcal{D}_{\mathsf{NMCom}}$ that can tell apart such two distributions. We can use this adversary and the associated distinguisher to construct ad adversary $\mathcal{A}_{\mathsf{LS}}$ for the T_{LS}-witness-indistinguishable of LS. Let $\mathcal{C}_{\mathsf{LS}}$ be the adaptive-input WI challenger. In the left session $\mathcal{A}_{\mathsf{LS}}$ acts as NMSen with $\mathcal{A}_{\mathsf{NMCom}}$ except for the messages of LS for which he acts as a proxy between $\mathcal{C}_{\mathsf{LS}}$ and $\mathcal{A}_{\mathsf{NMCom}}$. In the right session he acts as NMRec with $\mathcal{A}_{\mathsf{NMCom}}$. After the execution of the right session, $\mathcal{A}_{\mathsf{LS}}$ runs in time $\tilde{T}_{\mathsf{wsyn}}$ to obtain the message $\tilde{m}$ committed by $\mathcal{A}_{\mathsf{NMCom}}$ in the right session using Π_{wsyn}. Finally $\mathcal{A}_{\mathsf{LS}}$ gives $\tilde{m}$ and the output view of $\mathcal{A}_{\mathsf{NMCom}}$ as input to the distinguisher $\mathcal{D}_{\mathsf{NMCom}}$ and outputs what $\mathcal{D}_{\mathsf{NMCom}}$ outputs. Since by contradiction $\mathcal{D}_{\mathsf{NMCom}}$ distinguishes $\mathsf{mim}_{\mathcal{H}_1^m}^{\mathcal{A}_{\mathsf{NMCom}}}(z)$ from $\mathsf{mim}_{\mathcal{H}_2^m}^{\mathcal{A}_{\mathsf{NMCom}}}(z)$ we have that $\mathcal{A}_{\mathsf{LS}}$ can tell apart with non-negligible advantage which witness has been used to compute the transcript of LS. Formally the adversary $\mathcal{A}_{\mathsf{LS}}$ works as follows.

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

- Act as a honest receiver NMRec with $\mathcal{A}_{\mathsf{NMCom}}$, when $\mathcal{A}_{\mathsf{NMCom}}$ plays as as a sender.
- Upon receiving a_{LS} from $\mathcal{C}_{\mathsf{LS}}$, run as follows.

   1. Run $\mathsf{Sen}_{\mathsf{wsyn}}$ on input $1^{\lambda_{\mathsf{wsyn}}}$ $\mathtt{id}$ and m thus obtaining the first round $\mathsf{a}_{\mathsf{wsyn}}$ of Π_{wsyn}.

   2. Send $(\mathsf{a}_{\mathsf{wsyn}}, \mathsf{a}_{\mathsf{LS}})$ to $\mathcal{A}_{\mathsf{NMCom}}$.

- Upon receiving $(\mathsf{c}_{\mathsf{wsyn}}, \mathsf{c}_{\mathsf{LS}}, Y)$ from $\mathcal{A}_{\mathsf{NMCom}}$, run as follows.

   1. Run in time $\tilde{T}_f$ to compute y s.t. $Y = f(y)$.

   2. Run $\mathsf{Sen}_{\mathsf{wsyn}}$ on input $\mathsf{c}_{\mathsf{wsyn}}$ thus obtaining the third round $\mathsf{z}_{\mathsf{wsyn}}$ of Π_{wsyn} and the de-commitment information $\mathsf{dec}_{\mathsf{wsyn}}$.

   3. Set $x = (\mathsf{a}_{\mathsf{wsyn}}, \mathsf{c}_{\mathsf{wsyn}}, \mathsf{z}_{\mathsf{wsyn}}, Y, \mathtt{id})$, $w_0 = (m, \mathsf{dec}_{\mathsf{wsyn}}, \bot)$, $w_1 = (\bot, \bot, y)$ and send $(x, \mathsf{c}_{\mathsf{LS}}, w_0, w_1)$ to $\mathcal{C}_{\mathsf{LS}}$.

   4. Upon receiving z_{LS} from $\mathcal{C}_{\mathsf{LS}}$, send $(\mathsf{z}_{\mathsf{wsyn}}, \mathsf{z}_{\mathsf{LS}})$ to $\mathcal{A}_{\mathsf{NMCom}}$.

After the execution with $\mathcal{A}_{\mathsf{NMCom}}$, $\mathcal{A}_{\mathsf{LS}}$ computes the following steps:

1. Let $(\tilde{\mathsf{a}}_{\mathsf{wsyn}}, \tilde{\mathsf{c}}_{\mathsf{wsyn}}, \tilde{\mathsf{z}}_{\mathsf{wsyn}}, \tilde{\mathtt{id}})$ be the commitment received by NMRec when playing as in Π_{wsyn}. Run in time $\tilde{T}_{\mathsf{wsyn}}$ to compute $\tilde{m} : \exists\, \tilde{\mathsf{dec}}_{\mathsf{wsyn}}$ s.t. $\mathsf{Rec}_{\mathsf{wsyn}}$ on input $(\tilde{m}, \tilde{\mathsf{dec}}_{\mathsf{wsyn}})$ accepts $\tilde{m}$ as a decomitment of $(\tilde{\mathsf{a}}_{\mathsf{wsyn}}, \tilde{\mathsf{c}}_{\mathsf{wsyn}}, \tilde{\mathsf{z}}_{\mathsf{wsyn}}, \tilde{\mathtt{id}})$.

2. Give $\tilde{m}$ and the view of $\mathcal{A}_{\mathsf{NMCom}}$ to the distinguisher $\mathcal{D}_{\mathsf{NMCom}}$.

3. Output what $\mathcal{D}_{\mathsf{NMCom}}$ outputs.

The proof ends with the observation that if $\mathcal{C}_{\mathsf{LS}}$ has used w_0 as a witness then $\mathcal{A}_{\mathsf{NMCom}}$ acts as in $\mathcal{H}_1^m$, otherwise he acts as in $\mathcal{H}_2^m$. □

The next hybrid is $\mathcal{H}_2^0(z)$. The only differences between this hybrid and $\mathcal{H}_2^m(z)$ is that $\mathsf{Sen}_{\mathsf{wsyn}}$ commits, using Π_{wsyn}, to a message 0^λ instead of m. Formally $\mathcal{H}_2^0(z)$ is the following.

$\mathcal{H}_2^0(z)$.

Left session:

1. First round.

   1.1. Run $\mathsf{Sen}_{\mathsf{wsyn}}$ on input $1^{\lambda_{\mathsf{wsyn}}}$ $\mathtt{id}$, and 0^λ thus obtaining the first round $\mathsf{a}_{\mathsf{wsyn}}$ of Π_{wsyn}.

   1.2. Run $\mathcal{P}$ on input $1^{\lambda_{\mathsf{LS}}}$ and ℓ thus obtaining the first round a_{LS} of LS.

   1.3. Send $(\mathsf{a}_{\mathsf{wsyn}}, \mathsf{a}_{\mathsf{LS}})$ to $\mathcal{A}_{\mathsf{NMCom}}$.

2. Third round, upon receiving $(\mathsf{c}_{\mathsf{wsyn}}, \mathsf{c}_{\mathsf{LS}}, Y)$ from $\mathcal{A}_{\mathsf{NMCom}}$, run as follows.

   2.1. Run in time $\tilde{T}_f$ to compute y s.t. $Y = f(y)$.

   2.2. Run $\mathsf{Sen}_{\mathsf{wsyn}}$ on input $\mathsf{c}_{\mathsf{wsyn}}$ thus obtaining the third round $\mathsf{z}_{\mathsf{wsyn}}$ of Π_{wsyn}.

   2.3. Set $x = (\mathsf{a}_{\mathsf{wsyn}}, \mathsf{c}_{\mathsf{wsyn}}, \mathsf{z}_{\mathsf{wsyn}}, Y, \mathtt{id})$ and $w = (\bot, \bot, y)$ with $|x| = \ell$. Run $\mathcal{P}$ on input x, w and c_{LS} thus obtaining the third round z_{LS} of LS.

   2.4. Send $(\mathsf{z}_{\mathsf{wsyn}}, \mathsf{z}_{\mathsf{LS}})$ to $\mathcal{A}_{\mathsf{NMCom}}$.

Right session: act as a proxy between $\mathcal{A}_{\mathsf{NMCom}}$ and NMRec.

Claim 8. *For all $m \in \{0,1\}^{\mathsf{poly}(\lambda)}$ it holds that* $\mathsf{mim}_{\mathcal{H}_2^0}^{\mathcal{A}_{\mathsf{NMCom}}}(z) \approx \mathsf{mim}_{\mathcal{H}_1^0}^{\mathcal{A}_{\mathsf{NMCom}}}(z)$.

Proof. The security proof follows the same idea of the proof of Claim 7. □

Until now we have proved that $\mathsf{mim}_{\mathcal{H}_1^m}^{\mathcal{A}_{\mathsf{NMCom}}}(z) \approx \mathsf{mim}_{\mathcal{H}_2^m}^{\mathcal{A}_{\mathsf{NMCom}}}(z)$ and $\mathsf{mim}_{\mathcal{H}_2^0}^{\mathcal{A}_{\mathsf{NMCom}}}(z) \approx \mathsf{mim}_{\mathcal{H}_1^0}^{\mathcal{A}_{\mathsf{NMCom}}}(z)$ and that both in $\mathcal{H}_1^m(z)$ and $\mathcal{H}_1^0(z)$ the adversary $\mathcal{A}_{\mathsf{NMCom}}$ commits to $\tilde{m} = \bot$ only with negligible probability. This implies that also in $\mathcal{H}_2^m(z)$ and $\mathcal{H}_2^0(z)$ $\mathcal{A}_{\mathsf{NMCom}}$ commits to $\tilde{m} = \bot$ only with negligible probability. For this reason now we can prove the indistinguishability between $\mathsf{mim}_{\mathcal{H}_2^m}^{\mathcal{A}_{\mathsf{NMCom}}}(z)$ and $\mathsf{mim}_{\mathcal{H}_2^0}^{\mathcal{A}_{\mathsf{NMCom}}}(z)$ by relying only on the synchronous weak one-one non-malleability of Π_{wsyn}. Formally we prove the following claim.

Claim 9. *For all $m \in \{0,1\}^{\mathsf{poly}(\lambda)}$ it holds that* $\mathsf{mim}_{\mathcal{H}_2^m}^{\mathcal{A}_{\mathsf{NMCom}}}(z) \approx \mathsf{mim}_{\mathcal{H}_2^0}^{\mathcal{A}_{\mathsf{NMCom}}}(z)$.

Proof. Suppose by contradiction that there exists an adversary $\mathcal{A}_{\mathsf{NMCom}}$ and a distinguisher $\mathcal{D}_{\mathsf{NMCom}}$ that can tell apart such two distributions. We can construct a distinguisher $\mathcal{D}_{\mathsf{wsyn}}$ and an adversary $\mathcal{A}_{\mathsf{wsyn}}$ that break the synchronous weak one-one non-malleability of Π_{wsyn}. It is important to observe that we can reduce the security of our scheme to the security of a synchronous weak one-one NM commitment because the previous claims ensure that the message that $\mathcal{A}_{\mathsf{NMCom}}$ commits in the right session (using Π_{wsyn}) is valid with overwhelming probability. Let $\mathcal{C}_{\mathsf{wsyn}}$ be the challenger of the *synchronous weak* one-one NM commitment and let $(0^\lambda, m)$ be the two challenge messages given by $\mathcal{C}_{\mathsf{wsyn}}$.

Loosely speaking in the left session $\mathcal{A}_{\mathsf{wsyn}}$ acts as NMSen with $\mathcal{A}_{\mathsf{NMCom}}$ with the difference that w.r.t. to the messages of Π_{wsyn} he acts as a proxy between $\mathcal{C}_{\mathsf{wsyn}}$ and $\mathcal{A}_{\mathsf{NMCom}}$. In the right session he acts as NMRec with $\mathcal{A}_{\mathsf{NMCom}}$ and, as in the left session, acts as a proxy w.r.t. the messages of Π_{wsyn} exchanged between $\mathsf{Rec}_{\mathsf{wsyn}}$ and $\mathcal{A}_{\mathsf{wsyn}}$. Then $\mathcal{A}_{\mathsf{wsyn}}$ runs $\mathcal{D}_{\mathsf{wsyn}}$ on input the message $\tilde{m}$ committed by $\mathcal{A}_{\mathsf{wsyn}}$ and his randomness. $\mathcal{D}_{\mathsf{wsyn}}$ reconstructs the view of $\mathcal{A}_{\mathsf{NMCom}}$ (by using the same randomness) and uses it and the message $\tilde{m}$ as inputs of $\mathcal{D}_{\mathsf{NMCom}}$ giving in output what $\mathcal{D}_{\mathsf{NMCom}}$ outputs. Since by contradiction $\mathcal{D}_{\mathsf{NMCom}}$ distinguishes between $\mathsf{mim}_{\mathcal{H}_2^m}^{\mathcal{A}_{\mathsf{NM4Com}}}(z)$ and $\mathsf{mim}_{\mathcal{H}_2^0}^{\mathcal{A}_{\mathsf{NM4Com}}}(z)$, we have that $\mathcal{D}_{\mathsf{wsyn}}$ tells apart which message has ben committed by the MiM adversary $\mathcal{A}_{\mathsf{wsyn}}$.

The adversary $\mathcal{A}_{\mathsf{wsyn}}$ acts as follows (we recall that this reduction is possible only because the message scheduling that we are considering is synchronous).

$\mathcal{A}_{\mathsf{wsyn}}(0^\lambda, m, z)$.
   Left session:

      1. Upon receiving $\mathsf{a}_{\mathsf{wsyn}}$ from $\mathcal{C}_{\mathsf{wsyn}}$, run as follows.

         1.1. Run $\mathcal{P}$ on input $1^{\lambda_{\mathsf{LS}}}$ and ℓ thus obtaining the first round a_{LS} of LS.

         1.2. Send $(\mathsf{a}_{\mathsf{wsyn}}, \mathsf{a}_{\mathsf{LS}})$ to $\mathcal{A}_{\mathsf{NMCom}}$.

      2. Upon receiving $(\mathsf{c}_{\mathsf{wsyn}}, \mathsf{c}_{\mathsf{LS}}, Y)$ from $\mathcal{A}_{\mathsf{NMCom}}$, run as follows.

         2.1. Run in time $\tilde{T}_f$ to compute y s.t. $Y = f(y)$.

         2.2. Set $x = (\mathsf{a}_{\mathsf{wsyn}}, \mathsf{c}_{\mathsf{wsyn}}, \mathsf{z}_{\mathsf{wsyn}}, Y, \mathtt{id})$, $w = (\bot, \bot, y)$. Run $\mathcal{P}$ on input x, w and c_{LS} thus obtaining the third round z_{LS} of LS.

      3. Upon receiving $\mathsf{z}_{\mathsf{wsyn}}$ from $\mathcal{C}_{\mathsf{wsyn}}$, send $(\mathsf{z}_{\mathsf{wsyn}}, \mathsf{z}_{\mathsf{LS}})$ to $\mathcal{A}_{\mathsf{NMCom}}$.

   Right session:

      1. Forward $\tilde{\mathsf{a}}_{\mathsf{wsyn}}$ to $\mathsf{Rec}_{\mathsf{wsyn}}$.

      2. Upon receiving $\tilde{\mathsf{c}}_{\mathsf{wsyn}}$ from $\mathsf{Rec}_{\mathsf{wsyn}}$, run as follows.

         2.1. Pick a random $\tilde{Y}$.

         2.2. Run $\mathcal{V}$ on input $\tilde{\mathsf{a}}_{\mathsf{LS}}$ thus obtaining the second round $\tilde{\mathsf{c}}_{\mathsf{LS}}$ of Π_{LS}.

         2.3. Send $(\tilde{\mathsf{c}}_{\mathsf{wsyn}}, \tilde{\mathsf{c}}_{\mathsf{LS}}, Y)$ to $\mathcal{A}_{\mathsf{NMCom}}$.

      3. Upon receiving $\tilde{\mathsf{z}}_{\mathsf{wsyn}}, \tilde{\mathsf{z}}_{\mathsf{LS}}$ from $\mathcal{A}_{\mathsf{NMCom}}$, run as follows:

         3.1. Set $\tilde{x} = (\tilde{\mathsf{a}}_{\mathsf{wsyn}}, \tilde{\mathsf{c}}_{\mathsf{wsyn}}, \tilde{\mathsf{z}}_{\mathsf{wsyn}}, \tilde{Y}, \tilde{\mathtt{id}})$ and abort iff $(\tilde{\mathsf{a}}_{\mathsf{LS}}, \tilde{\mathsf{c}}_{\mathsf{LS}}, \tilde{\mathsf{z}}_{\mathsf{LS}})$ is not accepted by $\mathcal{V}$ for $\tilde{x} \in L$.

         3.2. Send $\tilde{\mathsf{z}}_{\mathsf{wsyn}}$ to $\mathsf{Rec}_{\mathsf{wsyn}}$.

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

$\mathcal{D}_{\mathsf{wsyn}}(\mathsf{mim}^{\mathcal{A}_{\mathsf{wsyn}}}(z))$: Let $\tilde{m}$ be the committed message sent in the right session by $\mathcal{A}_{\mathsf{wsyn}}$ to $\mathsf{Rec}_{\mathsf{wsyn}}$. Reconstruct the view of $\mathcal{A}_{\mathsf{NMCom}}$ (using the randomness given in $\mathsf{mim}^{\mathcal{A}_{\mathsf{wsyn}}}(z)$) and give it and $\tilde{m}$ to the distinguisher $\mathcal{D}_{\mathsf{NMCom}}$. Output what $\mathcal{D}_{\mathsf{NMCom}}$ outputs. We observe that the reduction could fail if $\mathcal{A}_{\mathsf{NMCom}}$ commit to $\bot$ when $\mathcal{C}_{\mathsf{wsyn}}$ commits to 0^λ (by definition Π_{wsyn} is not secure against MiM adversary that can commit to $\bot$ when the commitment on the left is honestly generated). Actually the probability that $\mathcal{A}_{\mathsf{NMCom}}$ commits to $\bot$ is negligible. This is because $\mathsf{mim}_{\mathcal{H}_2^0}^{\mathcal{A}_{\mathsf{NMCom}}}(z) \approx \mathsf{mim}_{\mathcal{H}_1^0}^{\mathcal{A}_{\mathsf{NMCom}}}(z)$ and because of Claim 6. The proof ends with the observation that if $\mathcal{C}_{\mathsf{wsyn}}$ commits to m, $\mathcal{A}_{\mathsf{NMCom}}$ acts as in $\mathcal{H}_2^m$, otherwise he acts as in $\mathcal{H}_2^0$. $\qquad\square$

Now we can conclude the security proof of Lemma 7 by observing that for all $m \in \{0,1\}^{\mathsf{poly}(\lambda)}$ the following holds:

$$\mathsf{mim}_{\Pi_{\mathsf{NMCom}}}^{\mathcal{A}_{\mathsf{NMCom}},m}(z) = \mathsf{mim}_{\mathcal{H}_1^m}^{\mathcal{A}_{\mathsf{NMCom}}}(z) \approx \mathsf{mim}_{\mathcal{H}_2^m}^{\mathcal{A}_{\mathsf{NMCom}}}(z) \approx$$
$$\mathsf{mim}_{\mathcal{H}_2^0}^{\mathcal{A}_{\mathsf{NMCom}}}(z) \approx \mathsf{mim}_{\mathcal{H}_1^0}^{\mathcal{A}_{\mathsf{NMCom}}}(z) = \mathsf{sim}_{\Pi_{\mathsf{NMCom}}}^{\mathsf{Sim}_{\mathsf{NMCom}}}(1^\lambda, z).$$

$\square$

We conclude the proof of Theorem 3 by proving the following lemma.

Lemma 8. Π_{NMCom} *is a one-one NM commitment scheme.*

Proof. The proofs starts with the observation that the only non-trivial adversary using a non-synchronizing scheduling[20] is the sequential scheduling where $\mathcal{A}_{\mathsf{NMCom}}$ lets the left interaction complete before beginning the right. Considering this scheduling we now prove again that for all $m \in \{0,1\}^{\mathsf{poly}(\lambda)}$ it holds that

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

We prove the indistinguishability through a sequence of two hybrid experiments. The first hybrid experiment that we consider is $\mathcal{H}_1^m(z)$, that corresponds to the $\mathcal{H}_1^m(z)$ showed in the proof of Lemma 7 with the only difference that $\mathcal{A}_{\mathsf{NMCom}}$ acts in a non-synchronized way. Therefore Claim 5 holds also in this case.

The second hybrid that we consider is $\mathcal{H}_1^0(z)$. The only differences between this hybrid and the previous one is that NMSen commits to the message 0^λ instead of m. We observe that Claim 6 holds also in this case.

Claim 10. *For all $m \in \{0,1\}^{\mathsf{poly}(\lambda)}$ it holds that $\{\mathsf{mim}_{\mathcal{H}_1^m}^{\mathcal{A}_{\mathsf{NMCom}}}(z)\}_{z \in \{0,1\}^\star} \approx \{\mathsf{mim}_{\mathcal{H}_1^0}^{\mathcal{A}_{\mathsf{NMCom}}}(z)\}_{z \in \{0,1\}^\star}.$*

Proof. Suppose by contradiction that there exist an adversary $\mathcal{A}_{\mathsf{NMCom}}$ and a distinguisher $\mathcal{D}_{\mathsf{NMCom}}$ that can tell apart such two distributions. We can construct an adversary $\mathcal{A}_{\mathsf{Hiding}}$ that breaks the hiding of Π_{NMCom} (recall the hiding property of Π_{NMCom} comes from Lemma 6). Let $\mathcal{C}_{\mathsf{Hiding}}$ be the challenger of the hiding game and let $(m, 0^\lambda)$ be two challenge messages of the hiding game sent by $\mathcal{A}_{\mathsf{Hiding}}$. The high-level idea of this proof is that $\mathcal{A}_{\mathsf{Hiding}}$ can break the hiding of Π_{NMCom} using the witness extracted from the LS transcript computed by $\mathcal{A}_{\mathsf{NMCom}}$ in the right session. In more details if the witness extract from the LS transcript corresponds to the message committed by $\mathcal{A}_{\mathsf{NMCom}}$ then $\mathcal{A}_{\mathsf{Hiding}}$ can win the hiding game by running $\mathcal{D}_{\mathsf{NMCom}}$. We observe that Claim 5 and Claim 6 ensure that with non-negligible probability the witness extracted from LS in $\mathcal{H}_1^m$ and also in $\mathcal{H}_1^0$ is the committed message $\tilde{m}$.

Before describing the adversary we need to consider the augmented machine $\mathcal{M}_{\mathsf{Hiding}}$ that will be used by $\mathcal{A}_{\mathsf{Hiding}}$ to extract the witness from LS by using the extractor (that exists from the property of adaptive-input PoK enjoyed by LS). Recall that the extractor takes as input a randomness r for the prover and a randomness r' of $\mathcal{V}_{r'}$ in an execution of LS where x has been proved by $\mathcal{P}_r^\star$. Therefore $\mathcal{A}_{\mathsf{Hiding}}$ runs $\mathcal{A}_{\mathsf{NMCom}}$ and interacts in the left session acting as a proxy between $\mathcal{C}_{\mathsf{Hiding}}$ and $\mathcal{A}_{\mathsf{NMCom}}$ in order to obtain the transcript $\tau_{\mathsf{NMCom}} = (\mathsf{a}_{\mathsf{NMCom}}, \mathsf{c}_{\mathsf{NMCom}}, \mathsf{z}_{\mathsf{NMCom}})$ of Π_{NMCom}. In the right session $\mathcal{A}_{\mathsf{Hiding}}$ acts as NMRec with $\mathcal{A}_{\mathsf{NMCom}}$.

---

[20] As discussed earlier, an adverary using a trivial non-syncrhonizing scheduling can be simulated by an adversary using a syncrhonizing scheduling. Therefore the security proof for the synchronizing case applies.

Then $\mathcal{A}_{\mathsf{Hiding}}$ uses $\mathcal{M}_{\mathsf{Hiding}}$ to extract the witness of the LS transcript. The augmented machine $\mathcal{M}_{\mathsf{Hiding}}$ runs $\mathcal{A}_{\mathsf{NMCom}}$ acting in the left session with $\mathcal{A}_{\mathsf{NMCom}}$ as the sender NMSen using the messages $\mathsf{a}_{\mathsf{NMCom}}, \mathsf{z}_{\mathsf{NMCom}}$ of τ_{NMCom}. In the right session $\mathcal{M}_{\mathsf{Hiding}}$ interacts with $\mathcal{A}_{\mathsf{NMCom}}$ as the receiver NMRec with the only difference that all the messages of LS received by $\mathcal{A}_{\mathsf{NMCom}}$ are forwarded to the verifier $\mathcal{V}$ and vice versa. Now we describe the augmented machine $\mathcal{M}_{\mathsf{Hiding}}$.

$\mathcal{M}_{\mathsf{Hiding}}(\tau_{\mathsf{NMCom}}, r, z)$.

Let r be the randomness used for all next steps.

- Send $\mathsf{a}_{\mathsf{NMCom}}$ to $\mathcal{A}_{\mathsf{NMCom}}$.
- Upon receiving $\mathsf{c}_{\mathsf{NMCom}}$ from $\mathcal{A}_{\mathsf{NMCom}}$, send $\mathsf{z}_{\mathsf{NMCom}}$ to $\mathcal{A}_{\mathsf{NMCom}}$.
- Upon receiving $(\tilde{\mathsf{a}}_{\mathsf{wsyn}}, \tilde{\mathsf{a}}_{\mathsf{LS}})$ from $\mathcal{A}_{\mathsf{NMCom}}$, send $\tilde{\mathsf{a}}_{\mathsf{LS}}$ to $\mathcal{V}$.
- Upon receiving $\tilde{\mathsf{c}}_{\mathsf{LS}}$ from $\mathcal{V}$, run as follows.

    1. Run $\mathsf{Rec}_{\mathsf{wsyn}}$ on input $\tilde{\mathsf{id}}$ and $\tilde{\mathsf{a}}_{\mathsf{wsyn}}$ thus obtaining the second round $\tilde{\mathsf{c}}_{\mathsf{wsyn}}$ of Π_{wsyn}.
    2. Pick a random $\tilde{Y}$.
    3. Send $(\tilde{\mathsf{c}}_{\mathsf{wsyn}}, \tilde{\mathsf{c}}_{\mathsf{LS}}, \tilde{Y})$ to $\mathcal{A}_{\mathsf{NMCom}}$.

- Upon receiving $(\tilde{\mathsf{z}}_{\mathsf{wsyn}}, \tilde{\mathsf{z}}_{\mathsf{LS}})$, set $\tilde{x} = (\tilde{\mathsf{a}}_{\mathsf{wsyn}}, \tilde{\mathsf{c}}_{\mathsf{wsyn}}, \tilde{\mathsf{z}}_{\mathsf{wsyn}}, \tilde{Y}, \tilde{\mathsf{id}})$ and send $(\tilde{\mathsf{z}}_{\mathsf{LS}}, \tilde{x})$ to $\mathcal{V}$.

Now we can conclude the proof of this claim by describing how $\mathcal{A}_{\mathsf{Hiding}}$ works. $\mathcal{A}_{\mathsf{Hiding}}$ runs the extractor of LS (on input the randomnesses r and r') with oracle access to $\mathcal{M}_{\mathsf{Hiding}}$ (recall that an extractor of LS plays having oracle access to an adversarial prover of LS). We know from Claim 5 and from Claim 6 that with overwhelming probability the witness extracted from LS in $\mathcal{H}_1^m$ and in $\mathcal{H}_1^0$ is the committed message $\tilde{m}$. Therefore, $\mathcal{A}_{\mathsf{Hiding}}$ runs the distinguisher $\mathcal{D}_{\mathsf{NMCom}}$ on input $\tilde{m}$ and the view of $\mathcal{A}_{\mathsf{NMCom}}$, and outputs what $\mathcal{D}_{\mathsf{NMCom}}$ outputs. The proof ends with the observation that if $\mathcal{C}_{\mathsf{Hiding}}$ commits to m $\mathcal{A}_{\mathsf{NMCom}}$ acts as in $\mathcal{H}_1^m(z)$, otherwise he acts as in $\mathcal{H}_1^0(z)$. $\qquad\square$

Now, observe that the distribution of $\mathsf{mim}_{\mathcal{H}_1^m}^{\mathcal{A}_{\mathsf{NMCom}}}(z)$ corresponds to the distribution of $\mathsf{mim}_{\Pi_{\mathsf{NMCom}}}^{\mathcal{A}_{\mathsf{NMCom}},m}(z)$ and that the distribution of $\mathsf{mim}_{\mathcal{H}_1^0}^{\mathcal{A}_{\mathsf{NMCom}}}(z)$ corresponds to the distribution of $\mathsf{sim}_{\Pi_{\mathsf{NMCom}}}^{\mathsf{Sim}_{\mathsf{NMCom}}}(1^\lambda, z)$. With this observation we have proved that for all $m \in \{0,1\}^{\mathsf{poly}(\lambda)}$ the following relation holds:

$$\mathsf{mim}_{\Pi_{\mathsf{NMCom}}}^{\mathcal{A}_{\mathsf{NMCom}},m}(z) = \mathsf{mim}_{\mathcal{H}_1^m}^{\mathcal{A}_{\mathsf{NMCom}}}(z) \approx \mathsf{mim}_{\mathcal{H}_1^0}^{\mathcal{A}_{\mathsf{NMCom}}}(z) = \mathsf{sim}_{\Pi_{\mathsf{NMCom}}}^{\mathsf{Sim}_{\mathsf{NMCom}}}(1^\lambda, z).$$

$\qquad\square$

The proof of Theorem 3 follows from Lemma 6, Lemma 7 and Lemma 8.