

Statistical Concurrent Non-Malleable Zero Knowledge

Claudio Orlandi*
Aarhus University, Denmark
Email: orlandi@cs.au.dk

Rafail Ostrovsky
UCLA, USA
Email: rafail@cs.ucla.edu

Vanishree Rao
UCLA, USA
Email: vanishri@cs.ucla.edu

Amit Sahai
UCLA, USA
Email: sahai@cs.ucla.edu

Ivan Visconti
University of Salerno, Italy
Email: visconti@unisa.it

Abstract

The notion of Zero Knowledge introduced by Goldwasser, Micali and Rackoff in STOC 1985 is fundamental in Cryptography. Motivated by conceptual and practical reasons, this notion has been explored under stronger definitions. We will consider the following two main strengthened notions.

Statistical Zero Knowledge: here the zero-knowledge property will last forever, even in case in future the adversary will have unlimited power.

Concurrent Non-Malleable Zero Knowledge: here the zero-knowledge property is combined with non-transferability and the adversary fails in mounting a concurrent man-in-the-middle attack aiming at transferring zero-knowledge proofs/arguments.

Besides the well-known importance of both notions, it is still unknown whether one can design a zero-knowledge protocol that satisfies both notions simultaneously.

In this work we shed light on this question in a very strong sense. We show a *statistical concurrent non-malleable* zero-knowledge argument system for \mathcal{NP} with a *black-box* simulator-extractor.

1 Introduction

The notion of zero knowledge, first introduced in [GMR85], is one of the most pivotal cryptographic constructs. Depending on both natural and real-world attack scenarios, zero knowledge has been studied considering different conceptual flavors and practical applications.

Zero knowledge and man-in-the-middle attacks. In distributed settings such as the Internet, an adversary that controls the network can play concurrently as a verifier in some proofs¹ and as

*Work done while visiting UCLA.

¹While in our general discussion, we often refer to zero-knowledge proofs, we will finally need to resort to only arguments since our goal is to achieve statistical zero-knowledge property.

a prover in the other proofs. The goal of the adversary is to exploit the proofs it receives from the provers to then generate new proofs for the verifiers. The original notion of zero knowledge does not prevent such attacks since it assumes the adversarial verifier to only play as a verifier and only in sequential sessions.

The need of providing non-transferable proofs secure against such man-in-the-middle (MiM, for short) attacks was first studied by Dolev, Dwork and Naor in [DDN91]. In [BPS06], Barak, Prabhakaran and Sahai achieved for the first time such a strong form of zero knowledge, referred to as concurrent non-malleable zero knowledge (CNMZK, for short) is possible in the plain model. They provide a $\text{poly}(\lambda)$ -round construction, for λ being the security parameter, based on one-way functions, and a $O(\log(\lambda))$ -round construction based on collision-resistant hash functions. More recent results focused on achieving round efficiency with a mild setup [OPV08], computationally efficient constructions [OPV10], security with adaptive inputs [LP11].

Zero knowledge and forward security. The zero-knowledge property says that the view of the adversarial verifier does not help her in gaining any useful information. This means that it does not include information that can be exploited by a PPT machine. However, even though the execution of a zero-knowledge protocol can be based on the current hardness of some complexity assumptions, it is quite risky to rely on the assumed resilience of such assumptions against more powerful machines of the future. What is zero knowledge in a transcript produced today could not be zero knowledge in the eyes of a distinguisher that will read the transcript in 2040.

It is therefore appealing to provide some forward security flavor so that whatever is zero knowledge today will be zero knowledge forever. Statistical zero knowledge [BMO90, SV03, Oka00, GSV98, MOSV06, GMOS07, MX13] is the notion that satisfies this requirement. It has been achieved in constant rounds using collision-resistant hash functions [HM96], and even under the sole assumption that one-way functions exist requiring more rounds [HNO⁺09].

Unfortunately, all the known constructions for CNMZK protocols strongly rely on the computational indistinguishability of the output of the simulator. Techniques so far used to design protocols that are then proved to be CNMZK require the protocol to fix a witness in a commitment, that therefore must be statistically binding and thus only computationally hiding. There is therefore no hope to prove those protocol to be statistical zero knowledge. Moreover it does not seem that minor changes can establish the statistical zero knowledge property still allowing to prove CNMZK.

The Open Problem. Given the above state-of-the-art a natural question is the following: *is it possible to design an argument system that combines the best of both worlds, namely, a statistical concurrent non-malleable zero-knowledge argument system?*

1.1 Our Contribution

In this work, we provide the first statistical concurrent non-malleable zero-knowledge argument system. Our construction is an argument of knowledge (AoK, for short) and has a black-box simulator-extractor producing a statistically indistinguishable distribution.

As mentioned earlier, Barak et al. [BPS06] presented the first CNMZKAoK protocol; we will refer to their work here as BPS. However, their construction had an inherent limitation that the simulation can only be computational, the reason being the following. In their protocol, the prover needs to commit to a valid witness via a statistically binding non-malleable commitment scheme.

The commitment scheme being statistically binding is extremely crucial in their proof of security. This implies that when the simulator cheats and commits to a non-witness, the simulated view can only be computationally indistinguishable and not statistically so.

In this work, we overcome this shortcoming with the following idea. We take the BPS argument as a starting point and modify it. Firstly, we work on the root of the problem – the non-malleable commitment. We replace it with a special kind of a commitment scheme called ‘*mixed non-malleable commitment*’ scheme. The notion of mixed commitment was first introduced by Damgård and Nielsen [DN02]. Our mixed non-malleable commitment is parameterized by a string that if sampled with uniform distribution makes the scheme statistically hiding and computationally binding. Instead, when it is taken from another (computationally indistinguishable) distribution it is a statistically binding, computationally hiding, and non-malleable. We will construct such a scheme by using as distributions non-DDH and DDH tuples.

The next idea would be to append the (modified) BPS argument to a coin-flipping phase in which the prover and the verifier generate a random string. Thus, in the real-world the above mixed commitment is statistically hiding. This thus enables us to prove statistical simulatability of our protocol. Furthermore, in order to also achieve extractability of witnesses for the arguments given by the adversary, we switch to a hybrid which biases the coin-flipping outcome to a random DDH tuple. Typically, a coin-flipping protocol would involve the verifier committing to its share of randomness, the prover sending its share of randomness in the clear, and finally, the verifier opening the commitment. However, in order to enable the simulator to bias the outcome, instead of the verifier opening the commitment to its share of randomness, it gives only the committed value in the clear and presents an AoK for the randomness used. This argument is again played by using the BPS AoK, since we would need concurrent non-malleability here.

In order to simplify our proofs, we rely on the Robust Extraction Lemma of Goyal et al. [GLP⁺12] that generalizes concurrent extractability of the PRS preamble (or concurrently extractable commitments – CECOM, for short) [PRS02] in the following sense. Consider an adversary who sends multiple CECOM commitments interleaving them arbitrarily and also interacts with an external party B in an arbitrary protocol. Then, [GLP⁺12] shows how to perform concurrent extraction of the CECOM commitments without rewinding the external party B . The extractor designed by them is called the ‘robust simulator’.

Technical Challenges. While we will encounter multiple technical challenges, which will be clear as we go ahead, we point out the core technical challenge here and the way we will solve it.

One of the main technical challenges is when we prove witness extractability of our protocol. Namely, in our hybrid argument, we will encounter two consecutive hybrids H_a and H_b , wherein a coin-flipping phase of a particular right hand session is ‘intact’ in H_a , but is biased in H_b . This results in the mixed commitment changing from statistically hiding to statistically binding. In order to finally be able to argue that the extracted values are indeed valid witnesses, we will need to argue for the hybrid H_b that the value committed in this commitment is a valid witness. Herein, we will need to reduce our claim to computational binding of a CECOM commitment in the protocol. Thus, the requirement in this reduction would be that no extraction performed should rewind the external CECOM sender. Even the Robust Extraction Lemma will not be helpful here as the Lemma requires that the external protocol have round complexity strictly less than the round complexity of CECOM commitments (on which the robust simulator performs extraction) and the external protocol in this case is a CECOM commitment itself. The condition for the Lemma thus cannot be

met. We get around this difficulty through a carefully designed sequence of hybrid arguments. A similar difficulty arises in the proof of statistical simulatability of our protocol. Here again, we rely on a carefully designed sequence of hybrids.

The second main technical challenge, still of the same flavor as the first one above, is in the proof of witness extractability. Here, we encounter a pair of hybrids: in the former hybrid, we would have a few CECOM commitments of the right session being extracted by the robust simulator; in the latter hybrid, the modification introduced would be to change the value committed in a (statistically hiding) CECOM commitment of a left session from a valid witness to a zero-string. Here again, we will not be able to argue a reduction to the hiding property of the CECOM commitment of the left session in question, just by relying on the Robust Extraction Lemma. Here, we instead present a more detailed hybrid argument. Namely, in the CECOM commitment, we change the committed value one sub-commitment at a time [PRS02]. Since every sub-commitment in the standard CECOM commitment of [PRS02] ranges over just three rounds, we are now still able to apply the Robust Extraction Lemma.

2 Background

We assume familiarity with interactive Turing machines, denoted ITM. Given a pair of ITMs, A and B , we denote by $\langle A(x), B(y) \rangle(z)$ the random variable representing the (local) output of B , on common input z and private input y , when interacting with A with private input x , when the random tape of each machine is uniformly and independently chosen. In addition, we denote $\text{view}_B^A(x, z)$ to be the random variable representing the content of the random tape of B together with the messages received by B from A during the interaction on common input x and auxiliary input z to B .

If \mathcal{D}_1 and \mathcal{D}_2 are two distributions, then we denote that they are statistically close by $\mathcal{D}_1 \approx_s \mathcal{D}_2$; we denote that they are computationally indistinguishable by $\mathcal{D}_1 \approx_c \mathcal{D}_2$; and we denote that they are identical by $\mathcal{D}_1 \equiv \mathcal{D}_2$.

Definition 1 (Pseudorandom Language). *An NP-language $L \subseteq \{0, 1\}^*$ is said to be a pseudorandom language if the following holds. For $\lambda \in \mathbb{N}$, let \mathcal{D}_λ be a uniform distribution over $L \cap \{0, 1\}^\lambda$. Then, for every distinguisher \mathcal{D} running in time polynomial in λ , there exists a negligible function $\text{negl}(\cdot)$ such that \mathcal{D} can distinguish between \mathcal{D}_λ and U_λ with probability at most $\text{negl}(\lambda)$.*

Definition 2 (Witness relation). *A witness relation for an NP-language L is a binary relation R_L that is polynomially bounded, polynomial time recognizable and characterizes L by $L = \{x : \exists w.s.t.(x, w) \in R_L\}$. We say that w is a witness for the membership $x \in L$ if $(x, w) \in R_L$ (also denoted $R_L(x, w) = 1$). We will also let $R_L(x)$ denote the set of witnesses for the membership $x \in L$, i.e., $R_L(x) = \{w : (x, w) \in R_L\}$.*

In the following, we assume a fixed witness relation R_L for each NP-language L .

Definition 3 (Statistical Witness-Indistinguishable Argument of Knowledge (sWIAoK)). *An interactive argument system $\langle \mathcal{P}, \mathcal{V} \rangle$ for an NP-language L is called a statistical witness-indistinguishable argument of knowledge if it satisfies the following properties:*

Statistical witness-indistinguishability. *For every interactive machine \mathcal{V}^* and for every two sequences $\{w_x^1\}_{x \in L}, \{w_x^2\}_{x \in L}$, such that $w_x^1, w_x^2 \in R_L(x)$, the ensembles $\{\text{view}_{\mathcal{V}^*}^{\mathcal{P}(w_x^1)}(x)\}_{x \in L}$ and $\{\text{view}_{\mathcal{V}^*}^{\mathcal{P}(w_x^2)}(x)\}_{x \in L}$ are statistically indistinguishable.*

Knowledge Soundness. *There exists a PPT ITM called the ‘extractor’ E , such that for every PPT machine \mathcal{P}^* , for every $x \in L$, auxiliary input z , and random tape r , $\Pr[E^{\mathcal{P}^*}(x, z, r) = w : (x, w) \in R_L]$ is negligibly close to $\Pr[\langle \mathcal{P}^*(z; r), \mathcal{V} \rangle(x) = 1]$.*

Definition 4 (Interactive Argument System). *A two-party game $\langle \mathcal{P}, \mathcal{V} \rangle$ is called an Interactive Argument System for a language L if \mathcal{P}, \mathcal{V} are PPT ITMs and the following two conditions hold:*

Completeness. *For every $x \in L$,*

$$\Pr[\langle \mathcal{P}, \mathcal{V} \rangle(x) = 1] = 1.$$

Soundness. *For every $x \notin L$, every PPT ITM \mathcal{P}^* , there exists a negligible function $\epsilon(\cdot)$ such that,*

$$\Pr[\langle \mathcal{P}^*, \mathcal{V} \rangle(x) = 1] \leq \epsilon(|x|)$$

The verifier’s view of an interaction consists of the common input x , followed by its random tape and the sequence of prover messages the verifier receives during the interaction. We denote by $\text{view}_{\mathcal{V}^*}^{\mathcal{P}}(x, z)$ a random variable describing $\mathcal{V}^*(z)$ ’s view of the interaction with \mathcal{P} on common input x .

We will use various forms of commitment schemes. We will denote by SB, SH, CB, CH the usual properties that can be enjoyed by classic commitment schemes, namely: statistical binding, statistical hiding, computational binding and computational hiding.

Statistical Concurrent Non-Malleable Zero Knowledge. The definition of statistical CNMZK is taken almost verbatim from [BPS06] except for the additional requirement on the simulation being statistical. Let $\langle \mathcal{P}, \mathcal{V} \rangle$ be an interactive proof for an NP-language L with witness relation R_L , and let λ be the security parameter. Consider a man-in-the-middle adversary \mathcal{M} that participates in m_L “left interactions” and m_R “right interactions” described as follows. In the left interactions, the adversary \mathcal{M} interacts with $\mathcal{P}_1, \dots, \mathcal{P}_{m_L}$, where each \mathcal{P}_i is an honest prover and proves the statement $x_i \in L$. In the right interactions, the adversary proves the validity of statements $\bar{x}_1, \dots, \bar{x}_{m_R}$. Prior to the interactions, both $\mathcal{P}_1, \dots, \mathcal{P}_{m_L}$ receive $(x_1, w_1), \dots, (x_{m_L}, w_{m_L})$, respectively, where for all i , $(x_i, w_i) \in R_L$. The adversary \mathcal{M} receives x_1, \dots, x_{m_L} and the auxiliary input z , which in particular might contain a-priori information about $(x_1, w_1), \dots, (x_{m_L}, w_{m_L})$. On the other hand, the statements proved in the right interactions $\bar{x}_1, \dots, \bar{x}_{m_R}$ are chosen by \mathcal{M} . Let $\text{view}_{\mathcal{M}}(x_1, \dots, x_{m_L}, z)$ denote a random variable that describes the view of \mathcal{M} in the above experiment. Loosely speaking, an interactive argument is statistical concurrent non-malleable zero-knowledge (sCNMZK) if for every man-in-the-middle adversary \mathcal{M} , there exists a probabilistic polynomial time machine (called the simulator-extractor) that can *statistically* simulate both the left and the right interactions for \mathcal{M} , while outputting a witness for every statement proved by the adversary in the right interactions.

Definition 5 ((Black-Box) Statistical Concurrent Non-Malleable Zero Knowledge Argument of Knowledge). *An interactive protocol $\langle \mathcal{P}, \mathcal{V} \rangle$ is said to be a (Black-Box) Statistical Concurrent Non-Malleable Zero Knowledge (sCNMZK) argument of knowledge for membership in an NP language L with witness relation R_L , if the following hold:*

1. $\langle \mathcal{P}, \mathcal{V} \rangle$ is an interactive argument system;

2. For every m_L and m_R that are polynomial in λ , for every PPT adversary \mathcal{M} launching a concurrent non-malleable attack (i.e., \mathcal{M} interacts with honest provers $\mathcal{P}_1, \dots, \mathcal{P}_{m_L}$ in “left sessions” and honest verifiers $\mathcal{V}_1, \dots, \mathcal{V}_{m_R}$ in “right sessions”), there exists an expected polynomial time simulator-extractor \mathcal{SE} such that for every set of “left inputs” x_1, \dots, x_{m_L} we have $\mathcal{SE}(x_1, \dots, x_{m_L}) = (\text{view}, \bar{w}_1, \dots, \bar{w}_{m_R})$ such that:

- **view** is the simulated joint view of \mathcal{M} and $\mathcal{V}_1, \dots, \mathcal{V}_{m_R}$. Further, for any set of witnesses (w_1, \dots, w_{m_L}) defining the provers $\mathcal{P}_1, \dots, \mathcal{P}_{m_L}$, the view **view** is distributed statistically indistinguishable from the view of \mathcal{M} , denoted $\text{view}_{\mathcal{M}}(x_1, \dots, x_{m_L}, z)$, in a real execution;
- In the view **view**, let trans_ℓ denote the transcript of ℓ -th left execution, and $\overline{\text{trans}}_t$ that of t -th right execution, $\ell \in [m_L], t \in [m_R]$. If \bar{x}_t is the common input in $\overline{\text{trans}}_t$, $\overline{\text{trans}}_t \neq \text{trans}_\ell$ (for all ℓ) and \mathcal{V}_t accepts, then $R_L(\bar{x}_t, \bar{w}_t) = 1$ except with probability negligible in λ .

The probability is taken over the random coins of \mathcal{SE} . Further, the protocol is black-box sCNMZK, if \mathcal{SE} is a universal simulator that uses \mathcal{M} only as an oracle, i.e., $\mathcal{SE} = \mathcal{SE}^{\mathcal{M}}$.

We remark here that the statistical indistinguishability is considered only against computationally unbounded distinguishers, and not against unbounded man-in-the-middle adversaries. This restriction is inherent to the definition since we require statistical zero-knowledge and thus cannot simultaneously ask for soundness against unbounded provers.

Extractable Commitment Schemes.

Definition 6 (Extractable Commitment Schemes). *An extractable commitment scheme $(\text{Sender}, \text{Receiver})$ is a commitment scheme such that given oracle access to any PPT malicious sender Sender^* , committing to a string, there exists an expected PPT extractor \mathbb{E} that outputs a pair (τ, σ^*) such that the following properties hold:*

Simulatability. *The simulated view τ is identically distributed to the view of Sender^* (when interacting with an honest Receiver) in the commitment phase.*

Extractability. *the probability that τ is accepting and σ^* correspond to \perp is at most $1/2$. Moreover if $\sigma^* \neq \perp$ then the probability that Sender^* opens τ to a value different than σ^* is negligible.*

Lemma 1. *[LP09] Com_{nm} is an extractable commitment scheme.*

As shown in [LP09], Com_{nm} is an extractable commitment scheme. This is in fact the core property of the scheme that is relied upon in proving its non-malleability in [DDN00, LP09].

Extractable Mixed Robust Non-Malleable Commitments w.r.t. 1-Round Protocols.

In our protocol we make use of a special kind of commitment scheme, that we call a *extractable mixed robust non-malleable commitment scheme*. These are basically the mixed commitment schemes introduced by Damgård and Nielsen [DN02] that are also non-malleable (or robust) not only w.r.t. themselves but also w.r.t. 1-round protocols and also extractable.

We shall first discuss how we get mixed non-malleable commitments, and then at the end, we shall discuss how we also get mixed non-malleable commitments that are also robust w.r.t. 1-round protocols.

Intuitively, a mixed non-malleable commitment scheme is a commitment scheme that is parameterized by a string srs in such a way that if srs is from some specific distribution, then the commitment scheme is SH, and if srs is from another specific indistinguishable distribution, then the scheme is non-malleable. We require that both the distributions be efficiently samplable. When srs is randomly sampled (from the domain over which both the distributions are defined), we would require that srs is such that with all but negligible probability the scheme is SH. We denote such a scheme by $\text{NMMXCom}_{\text{srs}}$. More formally:

Definition 7 (Mixed Non-Malleable Commitments). *A commitment scheme is said to be a mixed non-malleable commitment scheme if it is parameterized by a string srs and if there exist two efficiently samplable distributions $\mathcal{D}_1, \mathcal{D}_2$, such that, $\mathcal{D}_1 \approx_c \mathcal{D}_2$, and if $\text{srs} \leftarrow \mathcal{D}_1$ then the commitment scheme is SH and if $\text{srs} \leftarrow \mathcal{D}_2$ then the commitment scheme is non-malleable. Furthermore, $|\text{Supp}(\mathcal{D}_2)|/|\text{Supp}(\mathcal{D}_1)| = \text{negl}(\lambda)$.*

Below, we show how to construct such a scheme. At a high level, we achieve this by using a *mixed commitment scheme* which, roughly speaking, is a commitment scheme parameterized by a string srs in such a way that if srs is from some specific efficiently samplable distribution, then the commitment scheme is SH, and if srs is from another specific indistinguishable efficiently samplable distribution, then the scheme is SB. We denote such a scheme by $\text{MXCom}_{\text{srs}}$. More formally:

Definition 8 (Mixed Commitments). *A commitment scheme is said to be a mixed commitment scheme if it is parameterized by a string srs and if there exist two efficiently samplable distributions $\mathcal{D}_1, \mathcal{D}_2$, such that, $\mathcal{D}_1 \approx_c \mathcal{D}_2$, and if $\text{srs} \leftarrow \mathcal{D}_1$ then the commitment scheme is SH and if $\text{srs} \leftarrow \mathcal{D}_2$ then the commitment scheme is SB. Furthermore, $|\text{Supp}(\mathcal{D}_2)|/|\text{Supp}(\mathcal{D}_1)| = \text{negl}(\lambda)$.*

In [DN02], Damgård and Nielsen gave two constructions of mixed commitment schemes, one based on one based on the Paillier cryptosystem and the other based on the Okamoto-Uchiyama cryptosystem. For concreteness, we provide a construction below based on Σ -protocols and that builds on previous ideas presented in [DG03, CV05, CV07].

Constructing Mixed Commitments. Let us first describe how to construct a mixed commitment scheme. The idea is to have \mathcal{D}_1 be uniform over $\{0, 1\}^{\text{poly}(\lambda)}$ and \mathcal{D}_2 be uniform over a pseudorandom language L (as per Definition 9) with a Σ -protocol (i.e., public-coin 3-round special-sound special honest-verifier zero-knowledge proof system). Then, to commit to a value β , sender would first run the simulator of the Σ -protocol for the statement that $\text{srs} \in L$ such that the simulated proof has β as the challenge; let (α, β, γ) be the simulated proof. Then the commitment would just be α . The opening would be γ .

Observe that if $\text{srs} \notin L$, then for any β there is only one accepting (α, β, γ) , making the scheme parameterized by this srs to be SB. Furthermore, with srs sampled uniformly at random from $\{0, 1\}^* \setminus L$, we will also be able to argue that the resulting scheme is CH. On the other hand, if $\text{srs} \in L$, then, for every α (in its valid domain as defined by the Σ -protocol), there exists γ' for every β' such that $(\alpha, \beta', \gamma')$ is an accepting transcript. This implies that there exists an opening of α to any β' . This makes the scheme SH. Furthermore, with srs sampled uniformly at random from L , it shall hold for any PPT machine that it can only run the simulator and it is infeasible for the machine to open α to *also* any $\beta' \neq \beta$ (with some γ' as an opening), assuming special-soundness of the Σ -protocol (Otherwise, one could extract the witness from $(\alpha, \beta, \gamma, \beta', \gamma')$). This makes the system only computationally binding. In detail:

Mixed Commitment from Σ -protocol. Let R_L be a hard relation for a pseudorandom language L i.e., $L = \{\text{srs} \in \{0,1\}^\lambda \mid \exists w : R_L(\text{srs}, w) = 1\}$ and $L \approx_c U_\lambda$. Consider a Σ -protocol for the above language L . The special honest-verifier zero-knowledge property of the Σ -protocol implies existence of a simulator S that on input the instance srs , a string β and a randomness r , outputs a pair (α, γ) such that $(\text{srs}, \alpha, \beta, \gamma)$ is computationally indistinguishable from a transcript $(\text{srs}, \alpha, \beta, \gamma)$ played by the honest prover when receiving β as challenge.

The commitment scheme played by sender C and receiver R that we need goes as follows.

Shared Random String: A random string $\text{srs} \in \{0,1\}^\lambda$ is given as a common input to both the parties;

Commitment Phase: We denote the commitment function by $\text{MXCom}_{\text{srs}}(\cdot; \cdot)$ and to commit to a string $\beta \in \{0,1\}^\lambda$:

1. C runs the Σ -protocol simulator $S(\text{srs}, \beta, r)$ to obtain (α, γ) ;
2. C sends α to R ;

Decommitment Phase: To open α to β :

1. C sends (β, γ) to R ;
2. R accepts if $(\text{srs}, \alpha, \beta, \gamma)$ is an accepting transcript for the Σ -protocol.

If $\text{srs} \in L$, then the commitment is computationally binding (since, with two openings one gets two accepting conversations for the same α , and from the special-soundness property of the Σ -protocol one can extract the witness) and statistically hiding (which is directly implied by perfect completeness of the Σ -protocol; i.e., for any α output as the first message by the simulator – for any β as the challenge – for every β' , given the witness, one can efficiently compute a final message γ' such that the verifier accepts). If $\text{srs} \notin L$ the commitment is statistically binding (since, for any α , there exists at most one β that makes R accept the decommitment, as there is no witness for $\text{srs} \in L$ and two accepting transcripts $(\alpha, \beta, \gamma), (\alpha, \beta', \gamma')$ with $\beta \neq \beta'$ implies a witness owing to the special-soundness property of the Σ -protocol) and computationally hiding (since, if on input α , one can guess β efficiently, then this can be used to decide whether or not $\text{srs} \in L$, a contradiction).

While there are many instantiations for L , we shall work with the following simple one. Define $L = \{(g_1, g_2, g_3, g_4) \in \mathbb{G}^4 \mid \exists a, b : a \neq b \wedge g_1^a = g_2 \wedge g_3^b = g_4\}$ with \mathbb{G} being a prime order group, where DDH is believed to be hard. That is, L is the language of non-DDH triplets. Note that in this case if srs is chosen uniformly at random from \mathbb{G}^4 the commitment is statistically hiding with overwhelming probability (most strings are not DDH triplets).

Relaxing the assumption. Another example for L is the following language: let (G, E, D) be a *dense* cryptosystem (i.e., valid public keys and ciphertexts can be easily extracted from random strings). The language L is:

$$L = \{(pk_0, pk_1, c_0, c_1) \mid \exists r_0, r_1, m_0, m_1, s_0, s_1 : m_0 \neq m_1, (pk_0, sk_0) \leftarrow G(1^k, r_0), \\ c_0 = E_{pk_0}(m_0, s_0), (pk_1, sk_1) \leftarrow G(1^k, r_1), c_1 = E_{pk_1}(m_1, s_1)\}.$$

Also in this case most strings are in the language, while the simulator can choose a string not in the language (i.e., with $m_0 = m_1$).

Moreover, we can plug this mixed commitment MXCom in a zero-knowledge protocol in the SRS model NMMXCom , so that when srs is a random DDH triple, the zero-knowledge protocol is a proof (i.e., statistically sound) and computational zero-knowledge, while when the srs is a random non-DDH triple then the zero-knowledge protocol is statistical zero-knowledge (and computationally sound). For eg., an implementation of Blum’s protocol by using MXCom as commitment scheme when the prover commits to the permuted adjacency matrices gives us a computational zero-knowledge proof-of-knowledge (ZKPoK, for short) if srs of the MXCom commitment used is a random DDH tuple and a statistical zero-knowledge argument-of-knowledge (ZKAoK, for short) if the srs is a random non-DDH tuple.

Constructing Mixed Non-Malleable Commitments. As mentioned earlier, we show how to construct a mixed non-malleable commitment scheme by using a mixed commitment scheme. For concreteness, we shall work with the mixed commitment scheme MXCom described earlier. To thus recall, by the construction of MXCom , our mixed non-malleable commitment scheme will be non-malleable when srs is a random DDH tuple and, is statistically hiding and computationally binding when srs is a random non-DDH tuple.

Our scheme $\text{NMMXCom}_{\text{srs}}$ is described as follows. At a high level, our approach is to slightly modify the DDN non-malleable commitment scheme in [DDN00]. In fact, we shall describe our modification by considering the concurrent non-malleable commitment scheme that appears in [LP09] (whose analysis of non-malleability is similar to that of the DDN commitment and is simpler). The protocol in [LP09] is in fact non-malleable w.r.t. any arbitrary protocols of logarithmic round-complexity, a property that is called $\log(\lambda)$ -robust non-malleability. This is one of the properties which will be of a crucial use to us and we shall elaborate on this property shortly. In fact, we only need 1-robust non-malleability. The scheme of [LP09] is described below.

At a high level, the protocol of the sender who wishes to commit to some value v proceeds as follows. To catch the core of the intuition, we describe here a simplified version of the protocol while ignoring the currently unnecessary details (such as parallel repetitions, etc.); later in the formal description, we shall present the original protocol of [LP09]. The sender proceeds as follows. In the first stage, upon receiving an output of a one-way function from the receiver, commit to v using a statistically binding commitment scheme Com_{sb} . In the second stage, engage in $\log(\lambda)$ (special-sound) \mathcal{WI} proofs of knowledge of either the value committed to using Com_{sb} or of a pre-image of the one-way function output sent by the receiver. (The number of \mathcal{WI} proofs is logarithmic in the length of the identities of the senders; hence, it is considered to be $\log(\lambda)$ in general). We note here that a special-sound \mathcal{WI} proof can be instantiated by using Blum’s Hamiltonicity protocol, wherein the commitment sent by the \mathcal{WI} prover in this protocol is SB.

Now to construct the mixed non-malleable commitment, the idea is to replace the SB commitment Com_{sb} of the first stage and the SB commitment within the Blum’s Hamiltonicity protocol (where both the commitments are given by the sender to the receiver) with the mixed commitment $\text{MXCom}_{\text{srs}}$. We shall analyze the properties of the resulting commitment scheme, denoted by $\text{NMMXCom}_{\text{srs}}$, below.

Recall that if srs is a random DDH tuple, then $\text{MXCom}_{\text{srs}}$ is SB and CH. Under this case, the resulting scheme would have the properties identical to the original scheme of [LPV08]; namely it is SB, CH, and non-malleable. On the other hand, if srs is a random non-DDH tuple, then $\text{MXCom}_{\text{srs}}$ is SH and CB. This would render the the resulting scheme to be SH (owing to the SH property of the commitment scheme in the first phase and witness-indistinguishability of the Hamiltonicity

protocol that is instantiated with SH commitment) and CB (owing to the computational binding property of the commitment scheme in the first phase; this is due to the fact that decommitment of the scheme in [LP09] is simply an opening of the commitment of the first phase). In fact, if srs is a random string, then it is a non-DDH tuple with all but negligible probability. Hence, we also have that when srs is a random string, $\text{MXCom}_{\text{srs}}$ is SH and CB with all but negligible probability. For future reference, we shall bookmark this into the following proposition.

Proposition 1. *If srs is a uniform DDH tuple, then $\text{MXCom}_{\text{srs}}$ is SB, CH, and non-malleable. If srs is a uniform random string, then $\text{MXCom}_{\text{srs}}$ is SH and CB.*

Robustness w.r.t. 1-Round Protocols of the Mixed Non-Malleable Commitments.

Recall that we modified the [LP09] non-malleable commitment scheme that is robust w.r.t. 1-round protocols to get mixed non-malleable commitment scheme. It turns out that the modified scheme still retains robust w.r.t. 1-round protocols. Here, we only give a high-level description of the reason behind this fact as this can be easily verified. The reason is that robustness of the non-malleable commitment scheme in Figure 3 is proved in [LP09] by relying only upon the structure (the ‘designs’, in particular) of the commitment scheme in Figure 3. In particular, this proof does not rely upon the specifics of the underlying commitment scheme. Now recall that the only modification we introduced in the robust non-malleable commitment scheme of [LP09] to get a mixed non-malleable commitment scheme is the following. Instead of using any underlying commitment scheme, we used a mixed commitment scheme. Thus, the scheme continues to be non-malleable commitment scheme robust w.r.t. 1-round protocols even when the underlying commitment schemes are mixed commitments.

Non-Malleability of $\text{NMMXCom}_{\text{srs}}$ w.r.t. Com_{nm} . Another property of $\text{NMMXCom}_{\text{srs}}$ that we need is the following. Let Com_{nm} be the NMCCom commitment robust w.r.t. 1-round protocol. We shall argue below that $\text{NMMXCom}_{\text{srs}}$ is non-malleable w.r.t. Com_{nm} (as per Definition 16).

Proposition 2. *The non-malleable commitment $\text{NMMXCom}_{\text{srs}}$ is robust w.r.t. the non-malleable commitment Com_{nm} .*

Proof sketch. Essentially, the proof is exactly the same as the proof of non-malleability of the non-malleable commitment scheme of [LP09] presented in Figure 3. We argue this here next. Consider a MiM adversary against non-malleability of $\text{NMMXCom}_{\text{srs}}$ that executes a Com_{nm} session on the left by playing the role of the receiver and a $\text{NMMXCom}_{\text{srs}}$ session on the right by playing the role of a sender. The key technique in proving non-malleability in [DDN00, LPV08, LP09] is to show that, immaterial of the way a MiM adversary interleaves the left and right commitments, there exists at least one \mathcal{WI} proof (within some design) on the right session such that it is ‘safe’ to rewind the MiM adversary for this proof; by ‘safe’, we mean that rewinding the MiM adversary at this point can be done without rewinding the external sender on the left. (Recall that to rewind a \mathcal{WI} proof is to rewind to the point between the first and the second message of the proof). To then understand what \mathcal{WI} proof qualifies to be safe to rewind, we begin by giving a high level idea of when a proof *does not* qualify to be safe. Consider any \mathcal{WI} proof $(\alpha_r, \beta_r, \gamma_r)$ on the right. If it is trying to use and ‘maul’ some \mathcal{WI} proof $(\alpha_l, \beta_l, \gamma_l)$ on the left, then the right proof is positioned in time with respect to the left one as shown in Figure B. Observe that rewinding such a proof on the right with a new challenge may make the MiM adversary send a new challenge for the left proof too asking for a new response which tantamounts to rewinding the sender on the left.




Figure 1: Prefix (until the dotted line) that is not a safe point.

[DDN00, LPV08, LP09] provide a characterization for the \mathcal{WZ} proofs on the right that qualify as safe for being rewound; however, the details of this characterization itself will not be important to us; the core argument in proving non-malleability in [DDN00, LPV08, LP09] is an argument that, immaterial of the way a MiM adversary interleaves the left and right commitments, there exists a \mathcal{WZ} proof on the right that is safe to rewind. This is so owing to the fact that the adversary can use only one proof on the left for every proof on the right and to the fact that there are exactly the same number of proofs on the left and the right. This would imply that if the left and the right identities are distinct (at least at one bit position), then at proofs corresponding to this bit position, design_0 on the left ‘matches up’ with design_1 on the right, depicted in Figure B. With a closer look at this interleaving, it can be easily derived that at least one of the \mathcal{WZ} proofs within this design_1 on the right is safe to be rewound.

We first observe that the only way $\text{NMMXCom}_{\text{srs}}$ differs from Com_{nm} in Figure 3 is that a specific kind of commitment, namely, a mixed commitment is used to instantiate the underlying commitments used in building Com_{nm} in Figure 3. Next, we observe that non-malleability of the commitment scheme $\text{NMMXCom}_{\text{srs}}$ is mainly due to the structure (or designs) of the \mathcal{WZ} proofs, and the same arguments on interleaving and safety of rewinding would hold even if the left commitment is under an Com_{nm} session. \square

We remark that in fact the non-malleable commitments $\text{NMMXCom}_{\text{srs}}$ and Com_{nm} are robust w.r.t. each other by the same arguments as above. However, it suffices for us that $\text{NMMXCom}_{\text{srs}}$ is robust w.r.t. Com_{nm} .

Concurrently Extractable Commitment Schemes. Concurrently extractable commitment (CECom) schemes consist of committing using the PRS preamble, and decommitting by opening all the commitments within the preamble [PRS02]. Roughly speaking, the preamble consists of the sender committing to multiple shares of the value to be committed; then the receiver, in multiple rounds, would challenge the sender to open a subset of them in such a way that the opened shares do not reveal the committed value, but this would somehow facilitate consistency checks as shown




Figure 2: A design₀ matches up with design₁.

in [PRS02, MOSV06].

A challenge-response pair in the preamble is called a ‘slot’. [MOSV06] formalized concurrent extractability and showed that the PRS preamble satisfies it if the number of slots therein is $\omega(\log(\lambda))$. We denote a CECOM commitment that is SB by CECom_{sb} , the one that is SH by CECom_{sh} .

Robust Concurrent Extraction. In [PRS02], Prabhakaran et al. demonstrated an extraction procedure by which, for an adversary Sender^* that executes multiple concurrent sessions of CECOM commitments, commitment information (commitment value and randomness) for each session can be extracted in polynomial time before the corresponding commitment phase is completed.

In [GLP⁺12], Goyal et al. extended the technique of [PRS02] and showed how to perform efficient extractions of CECOM commitments when an adversary Sender^* , besides concurrently performing CECOM commitments, also interacts with an ‘external’ party B in some arbitrary protocol Π . This setting now additionally requires that the extraction procedure rewinds the adversary Sender^* in a way that B does not get rewound in the process. This is achieved in [GLP⁺12] by building a *robust concurrent simulator* (or just ‘robust simulator’) RobustSim that interacts with both a *robust concurrent adversary*, which commits to multiple CECOM commitments, and an external party B , with which it runs some arbitrary protocol Π . For every CECOM commitment that is successfully completed, Goyal et al. show that, the robust concurrent simulator – without rewinding the external party – extracts a commitment information, with all but negligible probability. [GLP⁺12] present this result as the *Robust Extraction Lemma* which informally states that if $\ell_{\text{external}} = \ell_{\text{external}}(\lambda)$ and $\ell_{\text{cecom}} = \ell_{\text{cecom}}(\lambda)$ denote the round complexities of Π and the CECOM commitment, respectively, the Lemma guarantees the following two properties for RobustSim :

- RobustSim outputs a view whose statistical distance from the adversary’s view is at most $2^{-(\ell_{\text{cecom}} - \ell_{\text{external}} \cdot \log(T(\lambda)))}$, where, $T(\lambda)$ is the maximum number of total CECOM commitments by the adversary.

- RobustSim outputs commitment information for every CECOM commitment sent by the adversary with an assurance that the external party B of protocol Π is not rewind.

3 Statistical Concurrent Non-Malleable Zero-Knowledge

We start by giving an intuition behind the design of our protocol. In [BPS06], Barak et al. gave a construction of a computational CNMZK argument of knowledge. The simulation for this protocol was restricted to be only computational due to the following reason. In their protocol, one of the messages sent by the prover is a non-malleable commitment to a valid witness. Since the non-malleable commitment is SB, and the simulator, unlike an honest prover, does not use a valid witness in this non-malleable commitment, the simulated view was only computationally indistinguishable from the real-world view of a MiM adversary. It will be quite relevant for us to note that the non-malleable commitment being SB was crucially used in the proof of concurrent non-malleability of their protocol, therefore it is not possible to replace the above commitment scheme with a statistically hiding non-malleable commitment. More specifically, the proof would begin with the real-world view and through a series of hybrids would move towards the simulated view. In some certain hybrid along the way there would be introduced PRS rewindings to facilitate simulation. Given such a hybrid that performs PRS rewindings, it would be difficult to establish that one can extract a value out of the non-malleable commitment and that the extracted value is a valid-witness. The difficulty here is in ensuring that the PRS rewindings would not interfere with the non-malleable commitment on which the NMCom extractor is run. The idea in their proof instead was to first prove for the real-world view itself that the value committed in the NMCom commitment is a valid witness, and then make transitions to hybrids by introducing PRS rewindings. The point to be noted here is that it was crucial in their proof that the non-malleable commitment is a *statistically binding* commitment, so that they could put forth arguments on the values committed in it. With this, since introducing PRS rewindings would only bias the distribution of the view output by at most a negligible amount, their proof boiled down to proving that the value committed in the NMCom commitment does not adversely change as we move across various hybrids. Now, since we began with a hybrid where the values committed were valid witnesses, the values committed in the NMCom commitments after the PRS rewindings too are valid witnesses by non-malleability (and in particular statistical binding) of the commitment scheme.

Our idea begins from noticing that statistical binding of the NMCom commitment is crucial in proving extractability of valid witnesses and not important in simulating the view of the adversary. So the core idea is to somehow ensure that when we prove the indistinguishability of the simulation, the commitment scheme is statistically hiding. Instead, when we need to argue that the distribution of the extracted message does not change, then the commitment should be statistically binding. With this being the crux of our idea, the way we shall execute it is via what we call ‘mixed non-malleable commitments’. Intuitively, a mixed non-malleable commitment scheme is associated with two efficiently samplable, computationally indistinguishable distributions, and every commitment is parameterized by some string. Furthermore, one of the distributions is such that if the string is uniformly sampled from this distribution then the commitment is SH and CB; on the other hand, a commitment that is parameterized by a string that is uniformly sampled from the other distribution is SB and CH. Given such a commitment scheme, our protocol basically is an instantiation of the BPS protocol except that the NMCom commitment in the BPS protocol is replaced by a mixed non-malleable commitment. Also, the string that parameterizes this commitment computed jointly

by both the prover and the verifier is the outcome of a coin-flipping protocol. Namely, in our mixed non-malleable commitment scheme, the distribution on the parameter that produces a SH, CB commitment is the uniform distribution. Hence, the parameter generated via the coin-flipping protocol is SH and CB, as required. The BPS protocol forms the **Main BPS Phase** and the coin-flipping protocol is run in the **Coin-flipping Phase** of our protocol.

A traditional coin-flipping protocol would involve the verifier committing to a random string in the first round, followed by the prover sending another random string in the clear in the second round, the verifier opening the commitment in the third round, and finally having the prover's and the verifier's strings XOR-ed as the outcome of the coin-flipping protocol. However, now that we would also like to be able to cheat and bias the outcome to another (computationally indistinguishable) distribution (so that the mixed non-malleable commitment would then be SB), we modify the third round. Namely, instead of the third round being the verifier opening the commitment by giving both the committed value and the randomness used, the verifier would only give the committed value and then give an argument that there exists a randomness that would explain the commitment to this value. However, we won't be able to work with just any argument since we are in the concurrent setting. Furthermore, we also would like to ensure that when our simulator cheats in the argument to bias the coin-flipping outcome, the MiM adversary will not get any undue advantage. Thus, the argument that we use here is a CNMZK argument. In particular, we use the BPS argument itself. This argument forms the **BPS^{CFP} Phase** in our protocol.

Furthermore, towards simplifying our proof, we introduce the following slight modification of the BPS protocol in the 'Main BPS Phase'. In the original BPS protocol, the commitment in which the prover commits the valid witness to is an NMCCom commitment; on the other hand, in the 'Main BPS Phase', besides sending the NMCCom commitment to the witness, the prover also sends a concurrently extractable (CECom) commitment to the same witness. The simplification we achieve by adding the CECom commitment is that even the extraction of the witnesses (by the simulator-extractor) can be performed just like an extraction on any other CECom commitments in the protocol. Since, for simulation, we anyway need to employ certain techniques for the extraction from the other CECom commitments, we are now able to recycle the same techniques for witness extractions too, thus letting our focus stay on the other crucial subtleties (which we shall see as we get to the proofs of security).

We will now give a formal description of the protocol.

3.1 Our sCNMZKAoK Protocol $\langle \mathcal{P}, \mathcal{V} \rangle$

Ingredients.

1. Let CECom_{sh} and CECom_{sb} be SH and SB concurrently-extractable commitment scheme, respectively. Let each of them be of k_{cecom} -slots, where $k_{\text{cecom}} \in \omega(\log \lambda)$. Let the sender's randomness space for these commitments be $\text{RandSpace}_{\text{cecom}}$.
2. Let Com_{sh} be a SH commitment scheme. Let k_{sh} be its round-complexity, where k_{sh} is a constant.
3. Let sWIAoK be a statistical WIAoK protocol. Let k_{swiaok} be its round-complexity, where k_{swiaok} is a constant.
4. Let $\text{NMMXCom}_{(\cdot)}$ be our mixed non-malleable commitment scheme. Recall that it satisfies

extractability and is robust w.r.t. 1-round protocols. Let k_{nmcom} be its round-complexity, where k_{nmcom} is $O(\log(\lambda))$.

5. Let Com_{nm} be the non-malleable commitment scheme (described in Fig. 3). Recall that it satisfies extractability and is robust w.r.t. 1-round protocols. Let k_{nmcom} be its round-complexity.

In summary, the round complexities of the sub-protocols in our protocol are as follows: $k_{\text{cecom}} \in \omega(\log \lambda)$, $k_{\text{swiaok}}, k_{\text{sh}}$ are constants, and $k_{\text{nmcom}}, k_{\text{nmcom}} \in O(\log(\lambda))$.

Coin-Flipping Phase (CFP).

cfp_1 ($\mathcal{V} \rightarrow \mathcal{P}$): Sample $r_V \leftarrow \{0, 1\}^\lambda$, $\text{rand} \leftarrow \text{RandSpace}_{\text{cecom}}$ and commit to r_V using CECom_{sh} and randomness rand .

cfp_2 ($\mathcal{P} \rightarrow \mathcal{V}$): Sample $r_P \leftarrow \{0, 1\}^\lambda$ and send r_P .

cfp_3 ($\mathcal{V} \rightarrow \mathcal{P}$): Send r_V .

BPS^{CFP} Phase.

$\text{bps}^{\text{cfp}}_1$ ($\mathcal{P} \rightarrow \mathcal{V}$): Sample $\alpha \leftarrow \{0, 1\}^\lambda$ and commit to α using CECom_{sb} .

$\text{bps}^{\text{cfp}}_2$ ($\mathcal{V} \rightarrow \mathcal{P}$): Commit to 0^λ using Com_{sh} and argue knowledge of a commitment information (i.e., a commitment value and randomness) using swIAoK .

$\text{bps}^{\text{cfp}}_3$ ($\mathcal{P} \rightarrow \mathcal{V}$): Open the commitment of Step $\text{bps}^{\text{cfp}}_1$ to α .

$\text{bps}^{\text{cfp}}_4$ ($\mathcal{V} \rightarrow \mathcal{P}$): Commit to rand (used as commitment randomness in Step cfp_1) using the NMCom commitment Com_{nm} . In the rest of the paper, we shall refer to rand as the *sub-witness*.

$\text{bps}^{\text{cfp}}_5$ ($\mathcal{V} \rightarrow \mathcal{P}$): Send swIAoK to argue knowledge of either rand or r_{comsh} such that:

1. the value committed to by \mathcal{V} with the NMCom commitment at Step $\text{bps}^{\text{cfp}}_4$ is rand and rand explains the CECom commitment at Step cfp_1 to r_V .
2. Randomness r_{comsh} explains Com_{sh} at Step $\text{bps}^{\text{cfp}}_2$ being committed to α .

Let $\text{srs} = r_P \oplus r_V$.

Main BPS Phase.

bps_1 ($\mathcal{V} \rightarrow \mathcal{P}$): Sample $\sigma \leftarrow \{0, 1\}^\lambda$ and commit to it using CECom_{sb} .

bps_2 ($\mathcal{P} \rightarrow \mathcal{V}$): Commit to 0^λ using Com_{sh} and argue knowledge of a commitment information (i.e., a commitment value and randomness) using swIAoK .

bps_3 ($\mathcal{V} \rightarrow \mathcal{P}$): Open the commitment of Step bps_1 to σ .

bps_4 ($\mathcal{P} \rightarrow \mathcal{V}$): Commit to the witness w using mixed commitment $\text{NMCom}_{\text{srs}}$.

bps_{4+} ($\mathcal{P} \rightarrow \mathcal{V}$): Commit to the witness w using CECom_{sh} ².

bps_5 ($\mathcal{P} \rightarrow \mathcal{V}$): Send sWIAoK to argue knowledge of either w, r_{nm}, r_{cecom} or r'_{comsh} such that:

1. r_{nm} and r_{cecom} explain the NMMXCom_{srs} commitment of Step bps_4 and the CECom commitment of Step bps_{4+} to w , respectively, and w is such that $R_L(x, w) = 1$,
2. Randomness r'_{comsh} explains Com_{sh} at Step bps_2 being committed to σ .

3.2 Proofs of Security

In this section, we prove that our proposed protocol $\langle \mathcal{P}, \mathcal{V} \rangle$ is a statistical concurrent non-malleable zero-knowledge argument of knowledge. In other words, we show that there exists a simulator-extractor \mathcal{SE} that, for every concurrent MiM adversary \mathcal{M} , outputs a view view that is statistically indistinguishable from the view $\text{view}_{\mathcal{M}}(x_1, \dots, x_{m_L}, z)$ of \mathcal{M} in a real execution, and also outputs valid witnesses $\bar{y}_1, \dots, \bar{y}_{m_R}$ for all accepting right sessions.

Our simulator-extractor. The simulator-extractor \mathcal{SE} runs RobustSim which is the robust concurrent simulator for a robust concurrent attack. The adversary of the robust concurrent attack is a procedure I that we describe below. \mathcal{SE} will then output the output of $\text{RobustSim}^I(z)$. Recall that RobustSim runs a given adversary that mounts a robust concurrent attack by committing to multiple CECom commitments, where the adversary also interacts with an external party B in an arbitrary external protocol. RobustSim then is guaranteed to extract commitment information from every CECom commitment sent by the adversary before the completion of its commitment phase, in such a way that the external party B does not get rewind.

Procedure $I(z)$. I incorporates the MiM adversary \mathcal{M} , initiates an execution, and simulates its view as follows. Let the m_L left sessions be ordered with some arbitrary ordering. Let the m_R right sessions be ordered as follows: Consider any two right sessions, the i -th and the j -th; $i \leq j$ if and only if the CECom_{sb} commitment at Step bps_1 of the i -th session begins earlier to the CECom_{sb} commitment at Step bps_1 of the j -th session.

For every right session: Run the code of the verifier except isolate CECom_{sh} at Step bps_{4+} and relay it to external receiver. Let value y'_t be received from the outside (RobustSim) at the end of the CECom_{sh} commitment.

For every left session: When \mathcal{M} initiates an ℓ -th new session on the left, I proceeds as follows.

- Run the coin-flipping phase and the BPS^{CFP} phase honestly. Let srs be the outcome.
- Isolate CECom_{sb} at Step bps_1 and relay it to an external receiver. Let σ' be the value received from the outside (RobustSim) at the end of the CECom_{sb} commitment.
- Then commit to σ' using Com_{sh} at Step bps_2 ; also, use the same extracted value as the witness in executing the sWIAoK of Step bps_2 .

²In order to make the difference from the BPS protocol more easily noticeable, the five steps here that are common to the BPS protocol are numbered in sequence from bps_1 through bps_5 , while this ‘extra’ step is given a distinctive notation, bps_{4+} .

- In Step bps_3 , let \mathcal{M} opens its CECom_{sb} (of Step bps_1) to σ . Abort if $\sigma \neq \sigma'$.
- Commit to 0^λ using the mixed non-malleable commitment $\text{NMMXCom}_{\text{srs}}$ in Step bps_4 .
- Commit to 0^λ using the CECom_{sh} commitment in Step bps_{4+} .
- Use σ' committed to in Step bps_2 as the witness in executing sWIAoK of Step bps_5 .

When \mathcal{M} halts, I outputs the view of \mathcal{M} together with y'_1, \dots, y'_{m_R} , and halts.

Statistical simulation. We shall prove that the view output by \mathcal{SE} is distributed statistically close to the real-world view of the MiM adversary \mathcal{M} .

Theorem 1. *For every PPT adversary \mathcal{M} , $\{\text{view}_{\mathcal{M}}(x_1, \dots, x_{m_L})\}_{x_1, \dots, x_{m_L} \in L} \approx_s \{\text{view}\}_{x_1, \dots, x_{m_L} \in L}$.*

We only provide an intuition to the proof here below. Full proof appears in the full version of the paper.

Proof sketch. To prove the indistinguishability, we first take note of the ways in which the view generated by the simulator differs from the real-world view of the MiM adversary. Basically, the differences are that: for left sessions, the simulator does not use valid witnesses but tries to get ‘fake’ witnesses via the robust simulator; and for the right sessions, the simulator tries to extract witnesses via the robust simulator. While we know that using the robust simulator can incur at most negligible distance, what still remains to be shown is that the simulator using fake-witnesses for the left sessions also creates at most negligible distance from the real-view. For this, we simply rely on the statistical properties of the sub-protocols in which the simulator uses different values; namely, we rely upon SH of Com_{sh} of Step bps_2 , sWI property of sWIAoK of Step bps_2 , SH of the mixed non-malleable commitment of Step bps_4 , and sWI of sWIAoK of Step bps_5 – the steps at which the simulator uses different values in left sessions. Except for SH of the mixed non-malleable commitment of Step bps_4 , all the above properties are already guaranteed by the corresponding primitives themselves; however, on the other hand, to ensure that the mixed non-malleable commitment – parameterized by srs which is the outcome of the coin-flipping protocol – is SH, we need to ensure that srs is uniformly random with all but negligible probability. Before we proceed, we thus prove that in the real-world view srs is uniform in every left session with all but negligible probability.

Claim 1. *In the real-world view $\text{view}_{\mathcal{M}}(x_1, \dots, x_{m_L})$, for every left session, srs is uniformly random with all but negligible probability.*

Proof sketch. We begin by outlining the structure of the proof.

1. First, we show that, there exists a PPT algorithm that can extract a value r'_V from CECom_{sh} of Step cfp_1 of every left session *before* Step cfp_2 of that session is reached. Thus, since r_P is sent to the adversary after r'_V is extracted, r'_V is independent of r_P , and since r_P is uniformly random, $r_P \oplus r'_V$ is also uniformly random with all but negligible probability.
2. Then, we show that, in every left session, with all but negligible probability, $r'_V = r_V$, where, r_V is the value sent by \mathcal{M} in Step cfp_3 .

The above items together imply that $\text{srs} = r_P \oplus r_V$ is uniformly random, with all but negligible probability.

We prove the first step above by relying upon the Robust Extraction Lemma. Basically, the PPT algorithm (mentioned in the first step above) just emulates honest provers and honest verifiers to \mathcal{M} except that it relays the CECom_{sh} of Step cfp_1 of every left session to RobustSim for extraction. We establish the second step as follows. Recall that a commitment information for r'_V of CECom_{sh} of Step cfp_1 in question is extractable as shown for the first step. Furthermore, from the witness-extractability of the BPS protocol in BPS^{CFP} phase, we can extract a witness – that we call sub-witness – for r_V being committed in the same CECom_{sh} commitment. Thus, if $r_V \neq r'_V$, we break CB of CECom_{sh} .

However, the proof is still not complete. The reason is for an implicit assumption in proving the second step above that the BPS argument given by the adversary in BPS^{CFP} phase of the left session is sound. To prove this, we establish the following claim.

Sub-Claim 1. *In the real world view, if BPS^{CFP} phase of the l -th left session is accepted by the prover \mathcal{P}_ℓ , then the value committed to by \mathcal{M} in Com_{nm} at Step $\text{bps}^{\text{cfp}}_4$ of the l -th left session is a valid sub-witness.*

Proof sketch. Intuitively, Com_{nm} at Step $\text{bps}^{\text{cfp}}_4$ of the l -th left session contains a valid sub-witness owing to

computational hiding of CECom_{sb} – to argue that \mathcal{M} does not learn α , committed to by the prover in CECom_{sb} , and use it in its commitment Com_{sh} and sWIAoK at Step $\text{bps}^{\text{cfp}}_2$,

knowledge-soundness of sWIAoK in Step $\text{bps}^{\text{cfp}}_2$ – to extract knowledge of commitment information (i.e., commitment value and randomness) for Com_{sh} in Step $\text{bps}^{\text{cfp}}_2$ and to verify that the extracted value will not be α ,

knowledge-soundness of sWIAoK in Step $\text{bps}^{\text{cfp}}_5$ – to argue that either the value committed to in Com_{nm} at Step $\text{bps}^{\text{cfp}}_4$ is a valid sub-witness or to argue knowledge of a commitment information for Com_{sh} in Step $\text{bps}^{\text{cfp}}_2$ with commitment value as α ,

and finally, computational binding of Com_{sh} at Step $\text{bps}^{\text{cfp}}_2$ to show the knowledge extracted is not α as a commitment value.

We prove each of the above steps by carefully designing interfaces that launch robust concurrent attacks and by crucially relying upon the Robust Extraction Lemma for extraction of commitment information out of these interfaces. \square

With this, we continue with a hybrid argument by moving from the real-world view to the simulated view. This is facilitated by the already established facts that the messages where the simulator deviates in its behavior from the real-world are statistically hiding (in some sense). \square

Witness extractability. We shall prove that the values y'_1, \dots, y'_{m_R} extracted by the simulator-extractor \mathcal{SE} are valid witnesses for the statements of the corresponding right sessions.

Theorem 2. *For every PPT adversary \mathcal{M} , the output of the simulator $\mathcal{SE}(x_1, \dots, x_{m_L}, z) = (\text{view}, \bar{y}_1, \dots, \bar{y}_{m_R})$ is such that, $\forall i \in [m_R]$, $(\bar{x}_i, \bar{y}_i) \in R_L$.*

We discuss some of the core technical difficulties of the proof together with a high-level proof structure. Full proof appears in the full version of the paper

Proof sketch. Recall that in our protocol, the prover commits to a valid witness in $\text{NMMXCom}_{\text{srs}}$ at Step bps_4 and also commits to the same valid witness in CECom_{sh} at Step bps_{4+} (accompanied by a sWIAoK later in Step bps_5 for correctness of behavior). Note that both of these commitments are extractable. However, we cannot in a straight-forward manner employ the proof techniques of [BPS06] or [LPTV10] to prove that the values extracted from these commitments by the simulator are indeed valid witnesses.

We begin by pointing out the reason why we are not able to simply make use of the proofs of [BPS06] or [LPTV10]. In both [BPS06] and [LPTV10], the prover commits to the witness with a non-malleable commitment. Thus, the commitment is *statistically binding*. Their proofs essentially proceed in the following manner: First, prove that the values committed to in the non-malleable commitments are valid witnesses. Secondly, move to a hybrid where extractions are performed to extract ‘trapdoors’ for cheating in the left sessions and to extract witnesses of the right sessions. Although cheating by the simulator on the left sessions may adversely change the values committed by \mathcal{M} in the commitments of the right sessions, one can argue that the values committed to in the commitments of the right sessions are still valid witnesses owing to non-malleability of the commitment schemes.

Indeed, the statistically binding NMCCom commitments are the reason why the protocols of [BPS06] and [LPTV10] are not statistical CNMZK , but only computationally so. Our approach, to recall, is to use a mixed NMCCom commitment which is parameterized by a string that is output of the coin-flipping phase that precedes the main argument phase. Thus, in the real-world, as proven earlier for Theorem 1, the parameter is a uniform random string rendering the mixed NMCCom commitment to be SH . (Recall that the commitment being SH was crucial in proving statistical simulation in Theorem 1). Thus, it is not clear how to solely rely on the proof techniques of [BPS06, LPTV10] for our proof.

Our proof technique instead is as follows. We begin with the real-world experiment where the outcome of the coin-flipping protocol is a uniform random string and thus the commitment scheme at Step bps_4 is a SH commitment. Then we start moving towards the hybrid which cheats in right sessions by biasing the outcome of the coin-flipping protocol to a uniform DDH tuple. The technical challenge will be the following. Fix any right session. Let H_a and H_b be the two hybrids in our hybrid sequence such that, the commitment at Step bps_4 in H_a is SH while the same commitment is SB in H_b (due to cheating in the coin-flipping protocol). Here, we need to establish that in H_b , the committed value in the commitment at Step bps_4 is a valid witness. We establish this through a careful design of hybrids and their sequence. We expand on our techniques and the whole high-level structure of the proof here below. We shall discuss the further multiple technical difficulties in the full proof in the full version of the paper.

We begin with a hybrid that is identical to the real-world view. Then we gradually modify the behavior of the hybrid for the right sessions towards biasing the coin-flipping protocol outcome to a random DDH tuple (from a uniform random string). Here, we will also prove that the values committed to by the MiM adversary in the mixed commitment at Step bps_4 is a valid witness (note that, with the outcome of coin-flipping being a random DDH tuple, this commitment scheme is now SB , thus allowing us to put forth arguments on the values committed in it). Next, we further move to hybrids which also behave differently in the left sessions by using ‘trapdoors’ (or fake-witnesses) extracted from the adversary itself (instead of valid witnesses). Here, we argue that such deviation

in the hybrids' behavior for the left sessions does not adversely change the values committed to in the mixed NMCom commitments of the right sessions. Finally, we thereby reach a hybrid that behaves the same as our simulator-extractor, thus proving that the values extracted by \mathcal{SE} are indeed valid witnesses.

Observe that it is easy to prove indistinguishability of hybrids as we change hybrids' behavior for the left sessions. The reason is that the left sessions will still have the outcome of coin-flipping to be uniformly random and thus the corresponding mixed commitment is SH. Thus, hybrids using fake-witnesses instead of the real ones will only introduce negligible statistical distance. However, the challenging part would be to argue indistinguishability of hybrids as they deviate in their behavior on the right sessions. We expand on the difficulty and our techniques briefly here below.

In order for hybrids to start cheating in coin-flipping phases of the right sessions, it is crucial that the hybrids are ordered carefully. Note that, we cannot at once move to a hybrid which changes the outcome of the coin-flipping phase due to soundness of the BPS protocol in BPS^{CFP} phase. Thus, we first *simulate* this BPS protocol. We do so by extracting a trapdoor from the adversary in a way similar to [BPS06]. Then, the next hybrid would be 'free' to bias the coin-flipping outcome to a random DDH tuple. However, note that this change is not statistically indistinguishable but only computationally so. Hence, this may adversely change the values committed to in the NMCom commitments in the protocol. However, with a careful sequence of arguments, we will be able to obtain a reduction to robustness w.r.t. 1-round protocols. Here it will be crucial to ensure that the other rewindings performed by the hybrids would not rewind the external NMCom receiver of the reduction.

Let us now consider the first hybrid that biases the coin-flipping outcome of the i -th right session. By this hybrid, we will already have biased coin-flipping outcomes of the first $i - 1$ sessions. We thus need to make sure that this biasing will also not adversely change the values committed to in the mixed NMCom commitments at Step bps_4 of the first $i - 1$ right sessions. Here again we rely on w.r.t. 1-round protocols for these NMCom commitments too.

A major technical difficulty would be the following. Fix any right session. Consider the first hybrid that biases the coin-flipping outcome of this session. Note that the previous hybrid had coin-flipping outcome to be a random string and thus the mixed commitment at Step bps_4 of the right session here to be SH. But in the current hybrid, due to the bias, the commitment scheme is SB. Here we need to argue that the committed value is a valid witness. As shown in the full proof, this would entail proving computational binding of a CECom_{sh} commitment. Here, we are no longer able to rely only upon the Robust Extraction Lemma to ensure us of successful extractions for the following reason. In Robust Extraction Lemma, it is essential that the external protocol whose party is not supposed to be rewound is such that its round complexity is strictly less than the number of slots of the CECom commitments extracted from. However, in the current case, the external protocol itself is a CECom commitment and hence this condition can not be met. We get around this difficulty again with a careful sequencing of hybrid arguments.

Furthermore, the above technical difficulty arises at another juncture in the proof of witness extractability. Namely, we encounter a hybrid where coin-flippings of all right sessions are biased, and in the subsequent hybrid we start changing the values committed in CECom_{sh} commitments of the left sessions. Here, we are still able to rely on the robustness of the concurrent extraction as follows. Although one cannot use the Robust Extraction Lemma for a reduction to statistical hiding of the entire left CECom_{sh} commitment, we can consider intermediate hybrids where, at a time, only one sub-commitment of the CECom_{sh} commitment is changed. Thus, we are still able

to use robustness of the concurrent extraction since the sub-protocol in question is only of three rounds (as per the standard CECOM commitment of [PRS02]).

Then, once we ensure that the commitments at Step bps_4 of right sessions contain valid witnesses, we proceed to argue that the values extracted from the CECom_{sh} commitments are valid witnesses with the following argument. We, along the way, show that the adversary cannot have a trapdoor, namely, $r'_{com_{sh}}$ that explains Com_{sh} at Step bps_2 being committed to σ . This implies that, for every right session, the witness that is extractable from the sWIAoK argument at Step bps_5 of is an opening of the CECom_{sh} commitment (together with the opening of the $\text{NMMXCom}_{\text{srs}}$ commitment of Step bps_4) to a valid witness.

With this, we finally are at a hybrid that extracts valid witnesses from the right sessions. Furthermore, this hybrid is identical to our simulator-extractor, thus proving witness extractability of our protocol $\langle \mathcal{P}, \mathcal{V} \rangle$. □

Acknowledgments

Work supported in part by NSF grants 0830803, 09165174, 1065276, 1118126 and 1136174, US-Israel BSF grant 2008411, OKAWA Foundation Research Award, IBM Faculty Research Award, Xerox Faculty Research Award, B. John Garrick Foundation Award, Teradata Research Award, and Lockheed-Martin Corporation Research Award. This material is based upon work supported by the Defense Advanced Research Projects Agency through the U.S. Office of Naval Research under Contract N00014 – 11 – 1 – 0392. The views expressed are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

References

- [BMO90] Mihir Bellare, Silvio Micali, and Rafail Ostrovsky. The (true) complexity of statistical zero knowledge. In *STOC*, pages 494–502, 1990.
- [BPS06] Boaz Barak, Manoj Prabhakaran, and Amit Sahai. Concurrent non-malleable zero knowledge. In *FOCS*, pages 345–354, 2006. Full version available on eprint archive.
- [CV05] Dario Catalano and Ivan Visconti. Hybrid trapdoor commitments and their applications. In *Automata, Languages and Programming, 32nd International Colloquium, ICALP 2005, Lisbon, Portugal, July 11-15, 2005, Proceedings*, volume 3580 of *Lecture Notes in Computer Science*, pages 298–310. Springer, 2005.
- [CV07] Dario Catalano and Ivan Visconti. Hybrid commitments and their applications to zero-knowledge proof systems. *Theor. Comput. Sci.*, 374(1-3):229–260, 2007.
- [DDN91] Danny Dolev, Cynthia Dwork, and Moni Naor. Non-malleable cryptography (extended abstract). In *STOC*, pages 542–552, 1991.
- [DDN00] Danny Dolev, Cynthia Dwork, and Moni Naor. Nonmalleable cryptography. *SIAM Journal on Computing*, 30(2):391–437 (electronic), 2000. Preliminary version in *STOC* 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. ACM, 2003.
- [DN02] Ivan Damgård and Jesper Buus Nielsen. Perfect hiding and perfect binding universally composable commitment schemes with constant expansion factor. In *CRYPTO*, pages 581–596, 2002.
- [GLP⁺12] Vipul Goyal, Huijia Lin, Omkant Pandey, Rafael Pass, and Amit Sahai. Round-efficient concurrently composable secure computation via a robust extraction lemma. *IACR Cryptology ePrint Archive*, 2012:652, 2012.
- [GMOS07] Vipul Goyal, Ryan Moriarty, Rafail Ostrovsky, and Amit Sahai. Concurrent statistical zero-knowledge arguments for np from one way functions. In *ASIACRYPT*, pages 444–459, 2007.
- [GMR85] S. Goldwasser, S. Micali, and C. Rackoff. The knowledge complexity of interactive proof-systems. In *Proc. 17th STOC*, pages 291–304, 1985.
- [GSV98] Oded Goldreich, Amit Sahai, and Salil P. Vadhan. Honest-verifier statistical zero-knowledge equals general statistical zero-knowledge. In *STOC*, pages 399–408, 1998.
- [HM96] Shai Halevi and Silvio Micali. Practical and provably-secure commitment schemes from collision-free hashing. In *CRYPTO*, pages 201–215, 1996.
- [HNO⁺09] Iftach Haitner, Minh-Huyen Nguyen, Shien Jin Ong, Omer Reingold, and Salil P. Vadhan. Statistically hiding commitments and statistical zero-knowledge arguments from any one-way function. *SIAM J. Comput.*, 39(3):1153–1218, 2009.
- [LP09] Huijia Lin and Rafael Pass. Non-malleability amplification. In *STOC*, pages 189–198, 2009.
- [LP11] Huijia Lin and Rafael Pass. Concurrent non-malleable zero knowledge with adaptive inputs. In *Theory of Cryptography Conference (TCC)*, volume 6597, pages 274–292, 2011.
- [LPTV10] Huijia Lin, Rafael Pass, Wei-Lung Dustin Tseng, and Muthuramakrishnan Venkatasubramanian. Concurrent non-malleable zero knowledge proofs. In *CRYPTO '2010*, pages 429–446, 2010.
- [LPV08] Huijia Lin, Rafael Pass, and Muthuramakrishnan Venkatasubramanian. Concurrent non-malleable commitments from any one-way function. In *TCC*, pages 571–588, 2008.
- [MOSV06] Daniele Micciancio, Shien Jin Ong, Amit Sahai, and Salil P. Vadhan. Concurrent zero knowledge without complexity assumptions. In *TCC*, pages 1–20, 2006.
- [MX13] Mohammad Mahmoody and David Xiao. Languages with efficient zero-knowledge pcps are in szk. In *TCC*, pages 297–314, 2013.
- [Oka00] Tatsuaki Okamoto. On relationships between statistical zero-knowledge proofs. *J. Comput. Syst. Sci.*, 60(1):47–108, 2000.

- [OPV08] Rafail Ostrovsky, Giuseppe Persiano, and Ivan Visconti. Constant-round concurrent non-malleable zero knowledge in the bare public-key model. In *Automata, Languages and Programming, 35th International Colloquium, ICALP 2008, Reykjavik, Iceland, July 7-11, 2008, Proceedings, Part II*, volume 5126 of *Lecture Notes in Computer Science*, pages 548–559. Springer, 2008.
- [OPV10] Rafail Ostrovsky, Omkant Pandey, and Ivan Visconti. Efficiency preserving transformations for concurrent non-malleable zero knowledge. In *Theory of Cryptography, 7th Theory of Cryptography Conference, TCC 2010, Zurich, Switzerland, February 9-11, 2010. Proceedings*, volume 5978 of *Lecture Notes in Computer Science*, pages 535–552. Springer, 2010.
- [PR05] Rafael Pass and Alon Rosen. Concurrent non-malleable commitments. In *FOCS*, pages 563–572, 2005.
- [PRS02] Manoj Prabhakaran, Alon Rosen, and Amit Sahai. Concurrent zero knowledge with logarithmic round-complexity. In *Proc. 43rd FOCS*, 2002.
- [SV03] Amit Sahai and Salil P. Vadhan. A complete problem for statistical zero knowledge. *J. ACM*, 50(2):196–249, 2003.

A Notations

The basic notational conventions used in the paper are listed below.

For a fixed $\lambda \in \mathbb{N}$, let $\text{dom}B$ denote the domain of valid input for an algorithm B . Although the set $\text{dom}B$ is a function of λ , we skip mentioning λ explicitly for simplicity of notation. Let S be a set. Often we let S also denote a uniform distribution on S , whenever it is clear from the context whether it is the set or a distribution that is in question. We assume familiarity with interactive Turing machines, denoted ITM. Given a pair of ITMs, A and B , we denote by $\langle A(x), B(y) \rangle(z)$ the random variable representing the (local) output of B , on common input z and private input y , when interacting with A with private input x , when the random tape of each machine is uniformly and independently chosen. In addition, we denote $\text{view}_B^A(x, z)$ to be the random variable representing the content of the random tape of B together with the messages received by B from A during the interaction on common input x and auxiliary input z to B .

If \mathcal{D}_1 and \mathcal{D}_2 are two distributions, then we denote that they are statistically close by $\mathcal{D}_1 \approx_s \mathcal{D}_2$; we denote that they are computationally indistinguishable by $\mathcal{D}_1 \approx_c \mathcal{D}_2$; and we denote that they are identical by $\mathcal{D}_1 \equiv \mathcal{D}_2$.

We often refer to computational binding (respectively, hiding) of a commitment scheme in short as CB and CH, respectively. Also, we refer to statistical binding (respectively, hiding) of a commitment scheme in short as SB and SH, respectively. We refer to a protocol $\langle A, B \rangle$ as a dummy protocol if the parties A, B do not exchange any messages (i.e., the number of rounds of the protocol is zero). For a commitment, by a valid commitment information (or just commitment information), we mean a value and a randomness such that the randomness explains the commitment to be a commitment to that value.

B Preliminaries

Definition 9 (Pseudorandom Language). *An NP-language $L \subseteq \{0, 1\}^*$ is said to be a pseudorandom language if the following holds. For $\lambda \in \mathbb{N}$, let \mathcal{D}_λ be a uniform distribution over $L \cap \{0, 1\}^\lambda$. Then, for every distinguisher \mathcal{D} running in time polynomial in λ , there exists a negligible function $\text{negl}(\cdot)$ such that \mathcal{D} can distinguish between \mathcal{D}_λ and U_λ with probability at most $\text{negl}(\lambda)$.*

Definition 10 (Witness relation). *A witness relation for an NP-language L is a binary relation R_L that is polynomially bounded, polynomial time recognizable and characterizes L by $L = \{x : \exists w.s.t.(x, w) \in R_L\}$. We say that w is a witness for the membership $x \in L$ if $(x, w) \in R_L$ (also denoted $R_L(x, w) = 1$). We will also let $R_L(x)$ denote the set of witnesses for the membership $x \in L$, i.e., $R_L(x) = \{w : (x, w) \in R_L\}$.*

In the following, we assume a fixed witness relation R_L for each NP-language L .

Definition 11 (Statistical Witness-Indistinguishable Argument of Knowledge (sWIAoK)). *An interactive argument system $\langle \mathcal{P}, \mathcal{V} \rangle$ for an NP-language L is called a statistical witness-indistinguishable argument of knowledge if it satisfies the following properties:*

Statistical witness-indistinguishability. *For every interactive machine \mathcal{V}^* and for every two sequences $\{w_x^1\}_{x \in L}, \{w_x^2\}_{x \in L}$, such that $w_x^1, w_x^2 \in R_L(x)$, the ensembles $\{\text{view}_{\mathcal{V}^*}^{\mathcal{P}(w_x^1)}(x)\}_{x \in L}$ and $\{\text{view}_{\mathcal{V}^*}^{\mathcal{P}(w_x^2)}(x)\}_{x \in L}$ are statistically indistinguishable.*

Knowledge Soundness. *There exists a PPT ITM called the ‘extractor’ E , such that for every PPT machine \mathcal{P}^* , for every $x \in L$, auxiliary input z , and random tape r , $\Pr[E^{\mathcal{P}^*}(x, z, r) = w : (x, w) \in R_L]$ is negligibly close to $\Pr[\langle \mathcal{P}^*(z; r), \mathcal{V} \rangle(x) = 1]$.*

Definition 12 (Interactive Argument System). *A two-party game $\langle \mathcal{P}, \mathcal{V} \rangle$ is called an Interactive Argument System for a language L if \mathcal{P}, \mathcal{V} are PPT ITMs and the following two conditions hold:*

Completeness. *For every $x \in L$,*

$$\Pr[\langle \mathcal{P}, \mathcal{V} \rangle(x) = 1] = 1.$$

Soundness. *For every $x \notin L$, every PPT ITM \mathcal{P}^* , there exists a negligible function $\epsilon(\cdot)$ such that,*

$$\Pr[\langle \mathcal{P}^*, \mathcal{V} \rangle(x) = 1] \leq \epsilon(|x|)$$

The verifier’s view of an interaction consists of the common input x , followed by its random tape and the sequence of prover messages the verifier receives during the interaction. We denote by $\text{view}_{\mathcal{V}^*}^{\mathcal{P}}(x, z)$ a random variable describing $\mathcal{V}^*(z)$ ’s view of the interaction with \mathcal{P} on common input x .

(Black-Box) Statistical Concurrent Non-Malleable Zero Knowledge Argument of Knowledge. The definition of statistical CNMZK is taken almost verbatim from [BPS06] except for the additional requirement on the simulation being statistical. Let $\langle \mathcal{P}, \mathcal{V} \rangle$ be an interactive proof for an NP-language L with witness relation R_L , and let λ be the security parameter. Consider a man-in-the-middle adversary \mathcal{M} that participates in m_L “left interactions” and m_R “right interactions” described as follows. In the left interactions, the adversary \mathcal{M} interacts with $\mathcal{P}_1, \dots, \mathcal{P}_{m_L}$, where

each \mathcal{P}_i is an honest prover and proves the statement $x_i \in L$. In the right interactions, the adversary proves the validity of statements $\bar{x}_1, \dots, \bar{x}_{m_R}$. Prior to the interactions, both $\mathcal{P}_1, \dots, \mathcal{P}_{m_L}$ receive $(x_1, w_1), \dots, (x_{m_L}, w_{m_L})$, respectively, where for all i , $(x_i, w_i) \in R_L$. The adversary \mathcal{M} receives x_1, \dots, x_{m_L} and the auxiliary input z , which in particular might contain a-priori information about $(x_1, w_1), \dots, (x_{m_L}, w_{m_L})$. On the other hand, the statements proved in the right interactions $\bar{x}_1, \dots, \bar{x}_{m_R}$ are chosen by \mathcal{M} . Let $\text{view}_{\mathcal{M}}(x_1, \dots, x_{m_L}, z)$ denote a random variable that describes the view of \mathcal{M} in the above experiment. Loosely speaking, an interactive argument is statistical concurrent non-malleable zero-knowledge (sCNMZK) if for every man-in-the-middle adversary \mathcal{M} , there exists a probabilistic polynomial time machine (called the simulator-extractor) that can *statistically* simulate both the left and the right interactions for \mathcal{M} , while outputting a witness for every statement proved by the adversary in the right interactions.

Definition 13 ((Black-Box) Statistical Concurrent Non-Malleable Zero Knowledge Argument of Knowledge). *An interactive protocol $\langle \mathcal{P}, \mathcal{V} \rangle$ is said to be a (Black-Box) Statistical Concurrent Non-Malleable Zero Knowledge (sCNMZK) argument of knowledge for membership in an NP language L with witness relation R_L , if the following hold:*

1. $\langle \mathcal{P}, \mathcal{V} \rangle$ is an interactive argument system;
2. For every m_L and m_R that are polynomial in λ , for every PPT adversary \mathcal{M} launching a concurrent non-malleable attack (i.e., \mathcal{M} interacts with honest provers $\mathcal{P}_1, \dots, \mathcal{P}_{m_L}$ in “left sessions” and honest verifiers $\mathcal{V}_1, \dots, \mathcal{V}_{m_R}$ in “right sessions”), there exists an expected polynomial time simulator-extractor \mathcal{SE} such that for every set of “left inputs” x_1, \dots, x_{m_L} we have $\mathcal{SE}(x_1, \dots, x_{m_L}) = (\text{view}, \bar{w}_1, \dots, \bar{w}_{m_R})$ such that:
 - view is the simulated joint view of \mathcal{M} and $\mathcal{V}_1, \dots, \mathcal{V}_{m_R}$. Further, for any set of witnesses (w_1, \dots, w_{m_L}) defining the provers $\mathcal{P}_1, \dots, \mathcal{P}_{m_L}$, the view view is distributed statistically indistinguishable from the view of \mathcal{M} , denoted $\text{view}_{\mathcal{M}}(x_1, \dots, x_{m_L}, z)$, in a real execution;
 - In the view view , let trans_ℓ denote the transcript of ℓ -th left execution, and $\overline{\text{trans}}_t$ that of t -th right execution, $\ell \in [m_L], t \in [m_R]$. If \bar{x}_t is the common input in $\overline{\text{trans}}_t$, $\overline{\text{trans}}_t \neq \text{trans}_\ell$ (for all ℓ) and \mathcal{V}_t accepts, then $R_L(\bar{x}_t, \bar{w}_t) = 1$ except with probability negligible in λ .

The probability is taken over the random coins of \mathcal{SE} . Further, the protocol is black-box sCNMZK, if \mathcal{SE} is a universal simulator that uses \mathcal{M} only as an oracle, i.e., $\mathcal{SE} = \mathcal{SE}^{\mathcal{M}}$.

Non-Malleable Commitment Schemes. We recall the definition of non-malleability from [LPV08] (which builds upon the definition of [DDN00, PR05]). Let $\langle \text{Sender}, \text{Receiver} \rangle$ be a tag-based statistically binding commitment scheme. Consider a man-in-the-middle adversary \mathcal{M} that, on auxiliary input z , participates in one left and one right interaction simultaneously. In the left interaction, the man-in-the-middle adversary \mathcal{M} interacts with Sender, receiving a commitment to value v , using identity id of its choice. In the right interaction \mathcal{M} interacts with Receiver attempting to commit to a related value \tilde{v} , again using identity $\tilde{\text{id}}$ of its choice. If the right commitment is invalid, or undefined, its value is set to \perp . Furthermore, if $\tilde{\text{id}} = \text{id}$, \tilde{v} is also set to \perp – i.e., a commitment where the adversary copies the identity of the left interaction is considered invalid. Let $\text{nmc}_{\langle \text{Sender}, \text{Receiver} \rangle}^{\mathcal{M}}(v, z)$ denote a random variable that describes the value \tilde{v} and the view of \mathcal{M} , in the above experiment.

Definition 14 (Non-Malleable Commitment Schemes). *A statistically binding commitment scheme $\langle \text{Sender}, \text{Receiver} \rangle$ is said to be non-malleable (with respect to itself) if for every polynomial $p(\cdot)$, and every probabilistic polynomial-time man-in-the-middle adversary \mathcal{M} , the following ensembles are computationally indistinguishable.*

$$\begin{aligned} & \{\text{nmc}_{\langle \text{Sender}, \text{Receiver} \rangle}^{\mathcal{M}}(v, z)\}_{\lambda \in \mathbb{N}, v \in \{0,1\}^\lambda, v' \in \{0,1\}^\lambda, z \in \{0,1\}^*} \\ & \{\text{nmc}_{\langle \text{Sender}, \text{Receiver} \rangle}^{\mathcal{M}}(v', z)\}_{\lambda \in \mathbb{N}, v \in \{0,1\}^\lambda, v' \in \{0,1\}^\lambda, z \in \{0,1\}^*} \end{aligned}$$

Robust Non-Malleable Commitment Schemes.

Definition 15 (Robust Non-Malleable Commitment Schemes). *Let $\langle \text{Sender}, \text{Receiver} \rangle$ be a commitment scheme, and B a PPT ITM. We say the commitment scheme $\langle \text{Sender}, \text{Receiver} \rangle$ is non-malleable w.r.t. B , if for every two sequences $\{y_\lambda^1\}_{\lambda \in \mathbb{N}}$ and $\{y_\lambda^2\}_{\lambda \in \mathbb{N}}$, $y_\lambda^1, y_\lambda^2 \in \{0,1\}^\lambda$, such that, for all PPT ITM A^* , it holds that*

$$\{\langle B(y_\lambda^1), A^*(z) \rangle(1^\lambda)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*} \approx \{\langle B(y_\lambda^2), A^*(z) \rangle(1^\lambda)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*}$$

then it also holds that, for every PPT man-in-the-middle adversary \mathcal{M} ,

$$\{\text{mim}_{\langle \text{Sender}, \text{Receiver} \rangle}^{B, \mathcal{M}}(y_\lambda^1, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*} \approx \{\text{mim}_{\langle \text{Sender}, \text{Receiver} \rangle}^{B, \mathcal{M}}(y_\lambda^2, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*}$$

We say that $\langle \text{Sender}, \text{Receiver} \rangle$ is non-malleable w.r.t k -round protocols if $\langle \text{Sender}, \text{Receiver} \rangle$ is non-malleable w.r.t any machine B that interacts with the man-in-the-middle adversary in k rounds.

[LP09] show how to construct a robust non-malleable commitment scheme w.r.t. ℓ -round protocols, where ℓ is logarithmic in the length of the identifiers and hence is $\log(\lambda)$ in general. In fact, roughly speaking, they show that any commitment scheme that is ‘extractable’ and has more than k ‘rewinding slots’ is itself robust non-malleable w.r.t. k -round protocols. In this work, we build upon the robust non-malleable commitment scheme constructed in [LP09] based on the techniques from [DDN00] to prepare the ingredients for our final scheme. The [LP09] robust non-malleable commitment scheme w.r.t. ℓ rounds is thus described below in Figure 3.

A specific case of robustness of a non-malleable commitment scheme that we will consider in this work is robustness w.r.t. a different non-malleable commitment scheme.

Definition 16 (Robust Non-Malleability w.r.t. Distinct Commitment Schemes). *Let Com_{nm}^L and Com_{nm}^R be two non-malleable commitment schemes. Let Sender_R denote the sender of Com_{nm}^L commitment scheme. We say that the scheme Com_{nm}^R is robust w.r.t. the scheme Com_{nm}^L , if for every polynomial $p(\cdot)$, and every probabilistic polynomial-time man-in-the-middle adversary \mathcal{M} , for every pair of messages $y_\lambda^1, y_\lambda^2 \in \{0,1\}^\lambda$, the following holds.*

$$\{\text{mim}_{\text{Com}_{\text{nm}}^R}^{\text{Com}_{\text{nm}}^L, \mathcal{M}}(y_\lambda^1, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*} \approx \{\text{mim}_{\langle \text{Sender}, \text{Receiver} \rangle}^{\text{Com}_{\text{nm}}^L, \mathcal{M}}(y_\lambda^2, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*}$$

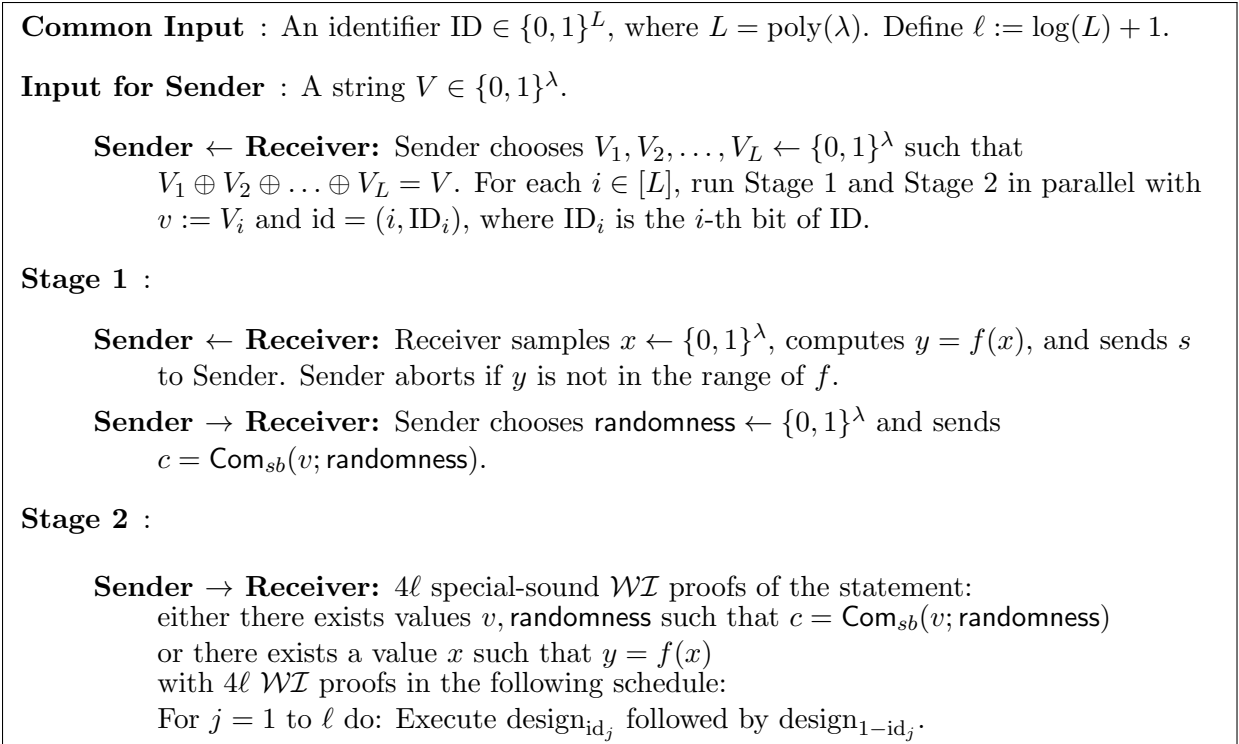


Figure 3: $O(\log(\lambda))$ -round Non-Malleable Commitment of [LP09]

Extractable Commitment Schemes.

Definition 17 (Extractable Commitment Schemes). *An extractable commitment scheme $(\text{Sender}, \text{Receiver})$ is a commitment scheme such that given oracle access to any PPT malicious sender Sender^* , committing to a string, there exists an expected PPT extractor E that outputs a pair (τ, σ^*) such that the following properties hold:*

Simulatability. *The simulated view τ is identically distributed to the view of Sender^* (when interacting with an honest Receiver) in the commitment phase.*

Extractability. *the probability that τ is accepting and σ^* correspond to \perp is at most $1/2$. Moreover if $\sigma^* \neq \perp$ then the probability that Sender^* opens τ to a value different than σ^* is negligible.*

Lemma 2. *[LP09] Com_{nm} is an extractable commitment scheme.*

As shown in [LP09], Com_{nm} is an extractable commitment scheme. This is in fact the core property of the scheme that is relied upon in proving its non-malleability in [DDN00, LP09].

Extractable Mixed Robust Non-Malleable Commitments w.r.t. 1-Round Protocols.

In our protocol we make use of a special kind of commitment scheme, that we call a *extractable mixed robust non-malleable commitment scheme*. These are basically the mixed commitment schemes introduced by Damgård and Nielsen [DN02] that are also non-malleable (or robust) not only w.r.t. themselves but also w.r.t. 1-round protocols and also extractable.

We shall first discuss how we get mixed non-malleable commitments, and then at the end, we shall discuss how we also get mixed non-malleable commitments that are also robust w.r.t. 1-round protocols.

Intuitively, a mixed non-malleable commitment scheme is a commitment scheme that is parameterized by a string srs in such a way that if srs is from some specific distribution, then commitment scheme is SH, and if srs is from another specific indistinguishable distribution, then the scheme is non-malleable. We require that both the distributions be efficiently samplable. When srs is randomly sampled (from the dominion over which both the distributions are defined), we would require that srs is such that with all but negligible probability the scheme is SH. We denote such a scheme by $\text{NMMXCom}_{\text{srs}}$. More formally:

Definition 18 (Mixed Non-Malleable Commitments). *A commitment scheme is said to be a mixed non-malleable commitment scheme if it is parameterized by a string srs and if there exist two efficiently samplable distributions $\mathcal{D}_1, \mathcal{D}_2$, such that, $\mathcal{D}_1 \approx_c \mathcal{D}_2$, and if $\text{srs} \leftarrow \mathcal{D}_1$ then the commitment scheme is SH and if $\text{srs} \leftarrow \mathcal{D}_2$ then the commitment scheme is non-malleable. Furthermore, $|\text{Supp}(\mathcal{D}_2)|/|\text{Supp}(\mathcal{D}_1)| = \text{negl}(\lambda)$.*

Below, we show how to construct such a scheme. At a high level, we achieve this by using a *mixed commitment scheme* which, roughly speaking, is a commitment scheme parameterized by a string srs in such a way that if srs is from some specific efficiently samplable distribution, then commitment scheme is SH, and if srs is from another specific indistinguishable efficiently samplable distribution, then the scheme is SB. We denote such a scheme by $\text{MXCom}_{\text{srs}}$. More formally:

Definition 19 (Mixed Commitments). *A commitment scheme is said to be a mixed commitment scheme if it is parameterized by a string srs and if there exist two efficiently samplable distributions $\mathcal{D}_1, \mathcal{D}_2$, such that, $\mathcal{D}_1 \approx_c \mathcal{D}_2$, and if $\text{srs} \leftarrow \mathcal{D}_1$ then the commitment scheme is SH and if $\text{srs} \leftarrow \mathcal{D}_2$ then the commitment scheme is SB. Furthermore, $|\text{Supp}(\mathcal{D}_2)|/|\text{Supp}(\mathcal{D}_1)| = \text{negl}(\lambda)$.*

In [DN02], Damgård and Nielsen gave two constructions of mixed commitment schemes, one based on one based on the Paillier cryptosystem and the other based on the Okamoto-Uchiyama cryptosystem. For concreteness, we provide a construction below based on Σ -protocols.

Constructing Mixed Commitments. Let us first describe how to construct a mixed commitment scheme. The idea is to have \mathcal{D}_1 be uniform over $\{0, 1\}^{\text{poly}(\lambda)}$ and \mathcal{D}_2 be uniform over a pseudorandom language L (as per Definition 9) with a Σ -protocol (i.e., public-coin 3-round special-sound special honest-verifier zero-knowledge proof system). Then, to commit to a value β , sender would first run the simulator of the Σ -protocol for the statement that $\text{srs} \in L$ such that the simulated proof has β as the challenge; let (α, β, γ) be the simulated proof. Then the commitment would just be α . The opening would be γ .

Observe that if $\text{srs} \notin L$, then for any β there is only one accepting (α, β, γ) , making the scheme parameterized by this srs to be SB. Furthermore, with srs sampled uniformly at random from $\{0, 1\}^* \setminus L$, we will also be able to argue that the resulting scheme is CH. On the other hand, if $\text{srs} \in L$, then, for every α (in its valid domain as defined by the Σ -protocol), there exists γ' for every β' such that $(\alpha, \beta', \gamma')$ is an accepting transcript. This implies that there exists an opening of α to any β' . This makes the scheme SH. Furthermore, with srs sampled uniformly at random from L , it shall hold for any PPT machine that it can only run the simulator and it is infeasible for the machine to open α to *also* any $\beta' \neq \beta$ (with some γ' as an opening), assuming special-soundness of the Σ -protocol (Otherwise, one could extract the witness from $(\alpha, \beta, \gamma, \beta', \gamma')$). This makes the system only computationally binding. In detail:

Mixed Commitment from Σ -protocol. Let R_L be a hard relation for a pseudorandom language L i.e., $L = \{\text{srs} \in \{0,1\}^\lambda \mid \exists w : R_L(\text{srs}, w) = 1\}$ and $L \approx_c U_\lambda$. Consider a Σ -protocol for the above language L . The special honest-verifier zero-knowledge property of the Σ -protocol implies existence of a simulator S that on input the instance srs , a string β and a randomness r , outputs a pair (α, γ) such that $(\text{srs}, \alpha, \beta, \gamma)$ is computationally indistinguishable from a transcript $(\text{srs}, \alpha, \beta, \gamma)$ played by the honest prover when receiving β as challenge.

The commitment scheme played by sender C and receiver R that we need goes as follows.

Shared Random String: A random string $\text{srs} \in \{0,1\}^\lambda$ is given as a common input to both the parties;

Commitment Phase: We denote the commitment function by $\text{MXCom}_{\text{srs}}(\cdot; \cdot)$ and to commit to a string $\beta \in \{0,1\}^\lambda$:

1. C runs the Σ -protocol simulator $S(\text{srs}, \beta, r)$ to obtain (α, γ) ;
2. C sends α to R ;

Decommitment Phase: To open α to β :

1. C sends (β, γ) to R ;
2. R accepts if $(\text{srs}, \alpha, \beta, \gamma)$ is an accepting transcript for the Σ -protocol.

If $\text{srs} \in L$, then the commitment is computationally binding (since, with two openings one gets two accepting conversations for the same α , and from the special-soundness property of the Σ -protocol one can extract the witness) and statistically hiding (which is directly implied by perfect completeness of the Σ -protocol; i.e., for any α output as the first message by the simulator – for any β as the challenge – for every β' , given the witness, one can efficiently compute a final message γ' such that the verifier accepts). If $\text{srs} \notin L$ the commitment is statistically binding (since, for any α , there exists at most one β that makes R accept the decommitment, as there is no witness for $\text{srs} \in L$ and two accepting transcripts $(\alpha, \beta, \gamma), (\alpha, \beta', \gamma')$ with $\beta \neq \beta'$ implies a witness owing to the special-soundness property of the Σ -protocol) and computationally hiding (since, if on input α , one can guess β efficiently, then this can be used to decide whether or not $\text{srs} \in L$, a contradiction).

While there are many instantiations for L , we shall work with the following simple one. Define $L = \{(g_1, g_2, g_3, g_4) \in \mathbb{G}^4 \mid \exists a, b : a \neq b \wedge g_1^a = g_2 \wedge g_3^b = g_4\}$ with \mathbb{G} being a prime order group, where DDH is believed to be hard. That is, L is the language of non-DDH triplets. Note that in this case if srs is chosen uniformly at random from \mathbb{G}^4 the commitment is statistically hiding with overwhelming probability (most strings are not DDH triplets).

Relaxing the Assumption. Another example for L is the following language: let (G, E, D) be a *dense* cryptosystem (i.e., valid public keys and ciphertexts can be easily extracted from random strings). The language L is:

$$L = \{(pk_0, pk_1, c_0, c_1) \mid \begin{array}{l} \exists r_0, r_1, m_0, m_1, s_0, s_1 : \\ m_0 \neq m_1, \\ (pk_0, sk_0) \leftarrow G(1^k, r_0), c_0 = E_{pk_0}(m_0, s_0), \\ (pk_1, sk_1) \leftarrow G(1^k, r_1), c_1 = E_{pk_1}(m_1, s_1) \end{array}\}$$

Also in this case most strings are in the language, while the simulator can choose a string not in the language (i.e., with $m_0 = m_1$).

Moreover, we can plug this mixed commitment MXCom in a zero-knowledge protocol in the SRS model NMMXCom , so that when srs is a random DDH triple, the zero-knowledge protocol is a proof (i.e., statistically sound) and computational zero-knowledge, while when the srs is a random non-DDH triple then the zero-knowledge protocol is statistical zero-knowledge (and computationally sound). For eg., an implementation of Blum’s protocol by using MXCom as commitment scheme when the prover commits to the permuted adjacency matrices gives us a computational zero-knowledge proof-of-knowledge (ZKPoK, for short) if srs of the MXCom commitment used is a random DDH tuple and a statistical zero-knowledge argument-of-knowledge (ZKAoK, for short) if the srs is a random non-DDH tuple.

Constructing Mixed Non-Malleable Commitments. As mentioned earlier, we show how to construct a mixed non-malleable commitment scheme by using a mixed commitment scheme. For concreteness, we shall work with the mixed commitment scheme MXCom described earlier. To thus recall, by the construction of MXCom , our mixed non-malleable commitment scheme will be non-malleable when srs is a random DDH tuple and, is statistically hiding and computationally binding when srs is a random non-DDH tuple.

Our scheme $\text{NMMXCom}_{\text{srs}}$ is described as follows. At a high level, our approach is to slightly modify the DDN non-malleable commitment scheme in [DDN00]. In fact, we shall describe our modification by considering the concurrent non-malleable commitment scheme that appears in [LP09] (whose analysis of non-malleability is similar to that of the DDN commitment and is simpler). The protocol in [LP09] is in fact non-malleable w.r.t. any arbitrary protocols of logarithmic round-complexity, a property that is called $\log(\lambda)$ -robust non-malleability. This is one of the properties which will be of a crucial use to us and we shall elaborate on this property shortly. In fact, we only need 1-robust non-malleability. The scheme of [LP09] is described below.

At a high level, the protocol of the sender who wishes to commit to some value v proceeds as follows. To catch the core of the intuition, we describe here a simplified version of the protocol while ignoring the currently unnecessary details (such as parallel repetitions, etc.); later in the formal description, we shall present the original protocol of [LP09]. The sender proceeds as follows. In the first stage, upon receiving an output of a one-way function from the receiver, commit to v using a statistically binding commitment scheme Com_{sb} . In the second stage, engage in $\log(\lambda)$ (special-sound) \mathcal{WI} proofs of knowledge of either the value committed to using Com_{sb} or of a pre-image of the one-way function output sent by the receiver. (The number of \mathcal{WI} proofs is logarithmic in the length of the identities of the senders; hence, it is considered to be $\log(\lambda)$ in general). We note here that a special-sound \mathcal{WI} proof can be instantiated by using Blum’s Hamiltonicity protocol, wherein the commitment sent by the \mathcal{WI} prover in this protocol is SB.

Now to construct the mixed non-malleable commitment, the idea is to replace the SB commitment Com_{sb} of the first stage and the SB commitment within the Blum’s Hamiltonicity protocol (where both the commitments are given by the sender to the receiver) with the mixed commitment $\text{MXCom}_{\text{srs}}$. We shall analyze the properties of the resulting commitment scheme, denoted by $\text{NMMXCom}_{\text{srs}}$, below.

Recall that if srs is a random DDH tuple, then $\text{MXCom}_{\text{srs}}$ is SB and CH. Under this case, the resulting scheme would have the properties identical to the original scheme of [LPV08]; namely it is SB, CH, and non-malleable. On the other hand, if srs is a random non-DDH tuple, then $\text{MXCom}_{\text{srs}}$

is SH and CB. This would render the the resulting scheme to be SH (owing to the SH property of the commitment scheme in the first phase and witness-indistinguishability of the Hamiltonicity protocol that is instantiated with SH commitment) and CB (owing to the computational binding property of the commitment scheme in the first phase; this is due to the fact that decommitment of the scheme in [LP09] is simply an opening of the commitment of the first phase). In fact, if srs is a random string, then it is a non-DDH tuple with all but negligible probability. Hence, we also have that when srs is a random string, $\text{MXCom}_{\text{srs}}$ is SH and CB with all but negligible probability. For future reference, we shall bookmark this into the following proposition.

Proposition 3. *If srs is a uniform DDH tuple, then $\text{MXCom}_{\text{srs}}$ is SB, CH, and non-malleable. If srs is a uniform random string, then $\text{MXCom}_{\text{srs}}$ is SH and CB.*

Robustness w.r.t. 1-Round Protocols of the Mixed Non-Malleable Commitments.

Recall that we modified the [LP09] non-malleable commitment scheme that is robust w.r.t. 1-round protocols to get mixed non-malleable commitment scheme. It turns out that the modified scheme still retains robust w.r.t. 1-round protocols. Here, we only give a high-level description of the reason behind this fact as this can be easily verified. The reason is that robustness of the non-malleable commitment scheme in Figure 3 is proved in [LP09] by relying only upon the structure (the ‘designs’, in particular) of the commitment scheme in Figure 3. In particular, this proof does not rely upon the specifics of the underlying commitment scheme. Now recall that the only modification we introduced in the robust non-malleable commitment scheme of [LP09] to get a mixed non-malleable commitment scheme is the following. Instead of using any underlying commitment scheme, we used a mixed commitment scheme. Thus, the scheme continues to be non-malleable commitment scheme robust w.r.t. 1-round protocols even when the underlying commitment schemes are mixed commitments.

Non-Malleability of $\text{NMMXCom}_{\text{srs}}$ w.r.t. Com_{nm} . Another property of $\text{NMMXCom}_{\text{srs}}$ that we need is the following. Let Com_{nm} be the NMCCom commitment robust w.r.t. 1-round protocol. We shall argue below that $\text{NMMXCom}_{\text{srs}}$ is non-malleable w.r.t. Com_{nm} (as per Definition 16).

Proposition 4. *The non-malleable commitment $\text{NMMXCom}_{\text{srs}}$ is robust w.r.t. the non-malleable commitment Com_{nm} .*

Proof sketch. Essentially, the proof is exactly the same as the proof of non-malleability of the non-malleable commitment scheme of [LP09] presented in Figure 3. We argue this here next. Consider a MiM adversary against non-malleability of $\text{NMMXCom}_{\text{srs}}$ that executes a Com_{nm} session on the left by playing the role of the receiver and a $\text{NMMXCom}_{\text{srs}}$ session on the right by playing the role of a sender. The key technique in proving non-malleability in [DDN00, LPV08, LP09] is to show that, immaterial of the way a MiM adversary interleaves the left and right commitments, there exists at least one \mathcal{WI} proof (within some design) on the right session such that it is ‘safe’ to rewind the MiM adversary for this proof; by ‘safe’, we mean that rewinding the MiM adversary at this point can be done without rewinding the external sender on the left. (Recall that to rewind a \mathcal{WI} proof is to rewind to the point between the first and the second message of the proof). To then understand what \mathcal{WI} proof qualifies to be safe to rewind, we begin by giving a high level idea of when a proof *does not* qualify to be safe. Consider any \mathcal{WI} proof $(\alpha_r, \beta_r, \gamma_r)$ on the right. If it is trying to use and ‘maul’ some \mathcal{WI} proof $(\alpha_l, \beta_l, \gamma_l)$ on the left, then the right proof is positioned in time with respect to the left one as shown in Figure B. Observe that rewinding such a proof




Figure 4: Prefix (until the dotted line) that is not a safe point.

on the right with a new challenge may make the MiM adversary send a new challenge for the left proof too asking for a new response which tantamounts to rewinding the sender on the left. [DDN00, LPV08, LP09] provide a characterization for the \mathcal{WI} proofs on the right that qualify as safe for being rewound; however, the details of this characterization itself will not be important to us; the core argument in proving non-malleability in [DDN00, LPV08, LP09] is an argument that, immaterial of the way a MiM adversary interleaves the left and right commitments, there exists a \mathcal{WI} proof on the right that is safe to rewind. This is so owing to the fact that the adversary can use only one proof on the left for every proof on the right and to the fact that there are exactly the same number of proofs on the left and the right. This would imply that if the left and the right identities are distinct (at least at one bit position), then at proofs corresponding to this bit position, design_0 on the left ‘matches up’ with design_1 on the right, depicted in Figure B. With a closer look at this interleaving, it can be easily derived that at least one of the \mathcal{WI} proofs within this design_1 on the right is safe to be rewound.

We first observe that the only way $\text{NMMXCom}_{\text{srs}}$ differs from Com_{nm} in Figure 3 is that a specific kind of commitment, namely, a mixed commitment is used to instantiate the underlying commitments used in building Com_{nm} in Figure 3. Next, we observe that non-malleability of the commitment scheme $\text{NMMXCom}_{\text{srs}}$ is mainly due to the structure (or designs) of the \mathcal{WI} proofs, and the same arguments on interleaving and safety of rewinding would hold even if the left commitment is under an Com_{nm} session. \square

We remark that in fact the non-malleable commitments $\text{NMMXCom}_{\text{srs}}$ and Com_{nm} are robust w.r.t. each other by the same arguments as above. However, it suffices for us that $\text{NMMXCom}_{\text{srs}}$ is robust w.r.t. Com_{nm} .

Concurrently Extractable Commitment Schemes. Concurrently extractable commitment (CECom, for short,) schemes consist of committing using the PRS preamble, and decommitting by opening all the commitments within the preamble [PRS02]. More specifically, the preamble uses an underlying commitment scheme Com , and roughly speaking, the sender first commits to




Figure 5: A design_0 matches up with design_1 .

many shares of the value v to be committed using Com ; this is followed by several rounds of interaction where in each round, the receiver sends a random challenge, and the sender responds with appropriate decommitments. A challenge-response pair is called a ‘slot’.

This commitment scheme can be either statistically binding or statistically hiding depending on whether the underlying commitment protocol (used within the preamble) is either statistically binding or statistically hiding, respectively. We denote a SH concurrently extractable commitment scheme by CECom_{sh} and a SB one by CECom_{sb} .

We will not require details of the CECom scheme itself, but only rely on certain properties of it established in [GLP⁺12]; we thus provide an informal definition of a CECom scheme as defined in [MOSV06].

Definition 20 (Concurrently Extractable Commitment Schemes (Informal)). *A commitment scheme $\langle \text{Sender}, \text{Receiver} \rangle$ is said to be concurrently extractable if there exists a $\text{CEC} - \text{Sim}$ whose output has two parts and that satisfies the following two properties.*

- *For every adversarial sender Sender^* that interacts with multiple receivers concurrently only in the commitment phase, the first part of the output of $\text{CEC} - \text{Sim}$, $\text{CEC} - \text{Sim}_1^{\text{Sender}^*}$ is distributed statistically close to $\langle \text{Receiver}, \text{Sender}^* \rangle$.*
- *For every session s in the output $\text{CEC} - \text{Sim}_1^{\text{Sender}^*}$, there exists a message $M(s)$ in the second part of $\text{CEC} - \text{Sim}$ ’s output $\text{CEC} - \text{Sim}_2^{\text{Sender}^*}$ such that no adversary (efficient if the commitment scheme is only computationally binding) having generated the commitment phase transcript s could have opened s to value different than $M(s)$.*

In [PRS02], Prabhakaran et al. demonstrated an extraction procedure by which, for an adversary Sender^* that executes multiple concurrent sessions of CECom commitments, commitment information (commitment value and randomness) can be extracted in polynomial time in such a way that the extraction outputs the commitment information for a CECom commitment before the commitment phase is completed.

Robust Concurrent Extraction. In [GLP⁺12], Goyal et al. extended the technique of [PRS02] and showed how to perform efficient extractions of CECOM commitments when an adversary A^* , besides concurrently performing CECOM commitments, also interacts with an ‘external’ party B in some arbitrary protocol Π . This setting now additionally requires that the extraction procedure rewinds the adversary A^* in a way that B does not get rewound in the process. This is achieved in [GLP⁺12] by building a *robust concurrent simulator* (or just ‘robust simulator’, for short). **RobustSim** interacts with both a *robust concurrent adversary*, which commits to multiple CECOM commitments, and an external party B , with which it runs some arbitrary protocol Π . For every CECOM commitment that is successfully completed, Goyal et al. show that, the robust concurrent simulator – without rewinding the external party – extracts a value (together with its randomness) that can explain the commitment, with all but negligible probability.

The requirements for such a concurrent extraction is formalized by considering an online extractor which is allowed to run in exponential time and which extracts the commitment information of the CECOM commitment, before the completion of the commitment phase (and no later). The online extractor also outputs the view of the adversary in the main-thread. With this, it would suffice to show that the output by the robust concurrent simulator (i.e., the view and the commitment information) is statistically close to the view output by the online extractor. This indistinguishability for the robust simulator in [GLP⁺12] is established in their *Robust Extraction Lemma* that we recall below.

PROTOCOL Π . Let $\Pi = \langle B, A \rangle$ be an arbitrary two-party computation protocol. Let $\text{dom}B$ denote the domain of valid inputs for algorithm B , and let $\ell_{\text{external}} = \ell_{\text{external}}(\lambda)$ denote the round complexity of Π .

The Robust Concurrent Attack. Let A^* be an algorithm, and $\beta \in \text{dom}B$ an input. In the robust-concurrent attack, A^* interacts with a special, not necessarily polynomial time, party \mathcal{E} , called the ‘online extractor’. Party \mathcal{E} simultaneously participates in one execution of the protocol Π , and several executions of CECOM commitments, all with A^* . Party \mathcal{E} follows the (honest) algorithm $B(\beta)$ in the execution of Π with A^* . Further, it follows the (honest) receiver algorithm in each execution of the CECOM commitments. If A^* successfully completes a CECOM commitment s , \mathcal{E} sends a string α_s to A^* . The scheduling of all messages in all sessions – Π as well as CECOM commitments – is controlled by A^* including starting new sessions and finishing or aborting existing sessions. At some point, A^* halts. We say that A^* launches the robust concurrent attack.

For $\beta \in \text{dom}B$, $z \in \{0, 1\}^*$, let $\text{REAL}_{\mathcal{E}, \Pi}^{A^*}(\beta, z)$ denote the output of the following probabilistic experiment: on input an auxiliary input z , the experiment starts an execution of A^* . Adversary A^* launches the robust-concurrent attack by interacting with the special party \mathcal{E} throughout the experiment, as described above. When A^* halts, the experiment outputs the view of A^* which includes: all messages sent/received by A^* , the auxiliary input z , the randomness of A^* .

We are now ready to present the Robust Extraction Lemma. Informally speaking, the lemma states that there exists an interactive PPT machine **RobustSim**, a.k.a the robust (concurrent) simulator, whose output is statistically close to $\text{REAL}_{\mathcal{E}, \Pi}^{A^*}(\beta, z)$ even when given that the final response of \mathcal{E} at the end of a successful CECOM commitment session is actually a valid commitment information for that commitment. Further, the robust simulator does not rewind B , and runs in time polynomial in total sessions opened by A^* .

Lemma 3 ([GLP⁺12]). *There exists an interactive Turing machine **RobustSim**, called the ‘robust simulator’, such that, for every PPT A^* , for every $\Pi = \langle B, A \rangle$, there exists a party \mathcal{E} , called the*

“online extractor”, for every $\beta \in \text{dom}B$, and every $z \in \{0, 1\}^*$, the following conditions hold:

1. **Validity constraint.** For every output ν of $\text{REAL}_{\mathcal{E}, \Pi}^{A^*}(\beta, z)$, we have:

- (a) for every statistically-binding CECOM commitment s (appearing in ν) with transcript τ_s , if there exists a unique value $v \in \{0, 1\}^\lambda$ in the commitment-transcript τ_s , then $\alpha_s = v$,
- (b) for every statistically-hiding CECOM commitment s (appearing in ν) with transcript τ_s , if there exists a valid opening (v_s, rand_s) in the view ν , then $\alpha_s = v_s$,

where α_s is the value \mathcal{E} sends at the completion of s .

2. **Statistical simulation.** If $\ell_{\text{external}} = \ell_{\text{external}}(\lambda)$ and $\ell_{\text{cecom}} = \ell_{\text{cecom}}(\lambda)$ denote the round complexities of Π and the CECOM commitment respectively, then the statistical distance between distributions $\text{REAL}_{\mathcal{E}, \Pi}^A(\lambda, \beta, z)$ and $\text{OUT}_s[B(\beta) \leftrightarrow \text{RobustSim}^{A^*}(z)]$ is given by

$$\Delta(\lambda) \leq 2^{-(\ell_{\text{cecom}} - \ell_{\text{external}} \cdot \log(T(\lambda)))}$$

where $T(\lambda)$ is the maximum number of total CECOM commitments between A^* and \mathcal{E} . Further, the running time of RobustSim is $\text{poly}(\lambda) \cdot T(\lambda)^2$.

Corollary 3 (Identical simulation if no CECOM played with the online extractor [GLP⁺12]). *If the robust concurrent adversary A^* sends no CECOM commitments to the online extractor \mathcal{E} then the view output by the robust simulator is an identical simulation of the real-world view of A^* . That is:*

$$\text{REAL}_{\mathcal{E}, \Pi}^A(\lambda, \beta, z) \equiv \text{OUT}_s[B(\beta) \leftrightarrow \text{RobustSim}^{A^*}(z)].$$

Remark 1. *We note that in fact the value returned by the robust simulator is not just the commitment value but also the commitment randomness. However, wherever not necessary we avoid explicitly mentioning the randomness returned by the robust simulator.*

C Proofs of Security

In this section, we prove that our proposed protocol $\langle \mathcal{P}, \mathcal{V} \rangle$ is a statistical concurrent non-malleable zero-knowledge argument of knowledge. Recall from Definition 13 that $\langle \mathcal{P}, \mathcal{V} \rangle$ is a statistical concurrent non-malleable zero-knowledge argument of knowledge protocol, if it satisfies the following properties;

1. $\langle \mathcal{P}, \mathcal{V} \rangle$ is an interactive argument system,
2. **Simulatability and Extractability:** for every m_L, m_R that are polynomial in λ , there exists a PPT simulator-extractor \mathcal{SE} that, for every concurrent man-in-the-middle adversary \mathcal{M} with some auxiliary information z , outputs a view, view , and also outputs $\bar{y}_1, \dots, \bar{y}_{m_R}$ for all accepting right sessions except for those right sessions that are just copied off from some left session, such that the outputs satisfy the following properties:
 - (a) **Statistical Simulation:** view is statistically indistinguishable from the view $\text{view}_{\mathcal{M}}(x_1, \dots, x_{m_L}, z)$ of \mathcal{M} in a real execution;

- (b) **Witness Extractability:** $\bar{y}_1, \dots, \bar{y}_{m_R}$ are valid witnesses for the statements of the corresponding sessions.

That is, we would like that $\mathcal{SE}(x_1, \dots, x_{m_L}, z) = (\text{view}, \bar{y}_1, \dots, \bar{y}_{m_R})$.

Although witness extractability implies that the protocol is an interactive argument system, we provide a separate proof for the latter for completeness; this may also serve as a warm-up for the techniques coming up ahead.

C.1 $\langle \mathcal{P}, \mathcal{V} \rangle$ is an Interactive Argument System

We shall now prove that $\langle \mathcal{P}, \mathcal{V} \rangle$ is an interactive argument system.

Lemma 4. *$\langle \mathcal{P}, \mathcal{V} \rangle$ is an interactive argument system.*

Proof. We shall show that our protocol $\langle \mathcal{P}, \mathcal{V} \rangle$ is an interactive argument system by establishing its completeness and soundness (as defined in Definition 12). Completeness directly follows from that of the sub-protocols.

It remains to show that any PPT adversarial prover \mathcal{P}^* can make \mathcal{V} accept any $x \notin L$ with at most negligible probability. To prove this, looking ahead, we would need $\text{NMMXCom}_{\text{srs}}$ in Step bps_4 to be computationally binding; before we go further, let us first ensure this. We ensure this by arguing that, if \mathcal{V} accepts then srs is uniformly random with all but negligible probability, and from Proposition 3, such an $\text{NMMXCom}_{\text{srs}}$ is CB.

Now consider an execution that is accepted by the verifier (and hence not aborted by either the prover or the verifier). We observe the following.

- r_V is statistically hidden in CECom_{sh} of Step cfp_1 .
- r_V is revealed at Step cfp_3 only after \mathcal{P}^* sends r_P (in Step cfp_2).
- r_V is uniformly random.

Thus, we have that $\text{srs} = r_P \oplus r_V$ is uniformly random (for any adversarially chosen r_P) with all but negligible probability.

Now, given that the $\text{NMMXCom}_{\text{srs}}$ in Step bps_4 is computationally binding, at a high level, the soundness of our protocol reduces to

- computational hiding of CECom_{sb} – to argue that \mathcal{P}^* does not learn σ , committed to by the prover in CECom_{sb} , and use it in its commitment Com_{sh} and sWIAoK at Step bps_2 ,
- knowledge-soundness of sWIAoK in Step bps_2 – to extract knowledge of commitment information (i.e., commitment value and randomness) for Com_{sh} in Step bps_2 and to verify that the extracted value will not be σ ,
- knowledge-soundness of sWIAoK in Step bps_5 – to argue knowledge of a commitment information for $\text{NMMXCom}_{\text{srs}}$ at Step bps_4 with commitment value as a valid witness or knowledge of a commitment information for Com_{sh} in Step bps_2 with commitment value as σ ,
- and finally, computational binding of Com_{sh} at Step bps_2 to show the knowledge extracted is not σ as a commitment value.

Using CH of CECom_{sb} and knowledge-soundness of sWIAoK in Step bps_2 . We begin by showing that one can extract a commitment information for Com_{sh} at Step bps_2 from sWIAoK at the same step, and owing to the computational hiding of CECom_{sb} at Step bps_1 , the value will not be σ .

Consider an adversarial prover P_1^* against knowledge-soundness of sWIAoK which behaves as follows.

- P_1^* runs $\text{RobustSim}^{I_{\text{sound}}^{(1)}}(z)$, where $I_{\text{sound}}^{(1)}$ is described as follows. $I_{\text{sound}}^{(1)}$ incorporates \mathcal{P}^* in a black-box way and interacts with it by playing the code of the honest verifier, except that $I_{\text{sound}}^{(1)}$ isolates sWIAoK in Step bps_2 and forwards it to an external sWIAoK verifier.
- Upon completion of this sWIAoK protocol, if the external sWIAoK verifier accepts, then it runs sWIAoK extractor on $\text{RobustSim}^{I_{\text{sound}}^{(1)}}(z)$.

Since the sWIAoK argument is isolated and relayed to an external sWIAoK verifier, (with no other messages being isolated by $I_{\text{sound}}^{(1)}$), knowledge-soundness of sWIAoK implies that the sWIAoK extractor extracts a valid opening – a commitment value and randomness – for the Com_{sh} commitment of Step bps_2 . This also implies that the value is not σ (committed to in CECom_{sb} of Step bps_1) with all but negligible probability, as otherwise we can build an adversary \mathcal{A}_{CH} that breaks computational hiding of CECom_{sb} of Step bps_1 with the same probability as P_1^* extracting σ . \mathcal{A}_{CH} runs $\text{RobustSim}^{I_{\text{sound}}^{(2)}}(z)$, where $I_{\text{sound}}^{(2)}$ behaves the same as $I_{\text{sound}}^{(1)}$ except that, besides isolating the sWIAoK argument of Step bps_2 , also isolates CECom_{sb} of Step bps_1 . While the CECom_{sb} commitment is forwarded to an external CECom receiver, \mathcal{A}_{CH} itself runs the honest verifier code of the isolated sWIAoK argument. \mathcal{A}_{CH} also runs the sWIAoK extractor on the isolated sWIAoK argument. Since neither $I_{\text{sound}}^{(2)}$ nor $I_{\text{sound}}^{(1)}$ isolate any CECom commitments, applying Corollary 3 of the Robust Extraction Lemma, the view of the adversary \mathcal{P}^* when run by \mathcal{A}_{CH} is identical to its view when run by P_1^* . Furthermore, since the view of the sWIAoK extractor also remains identical, we have that the probability that sWIAoK extractor extracts σ when run by P_1^* is equal to that when run by \mathcal{A}_{CH} , thus breaking computational hiding of CECom_{sb} of Step bps_1 with the same probability.

Using knowledge-soundness of sWIAoK in Step bps_5 and CB of Com_{sh} at Step bps_2 .

Now, we show that one can extract a witness for sWIAoK of Step bps_5 , and from its knowledge-soundness, we have that either we extract a commitment information (i.e., a commitment value and randomness) in NMMXCom_{sr5} at Step bps_4 such that this value is a valid witness that $x \in L$ or we extract an opening of Com_{sh} at Step bps_2 to σ . Finally, we will see that CB of Com_{sh} at Step bps_2 implies that the extracted value is not an opening to σ , which implies extraction of a valid witness, and existence of a valid witness in turn implies soundness.

Consider an adversarial prover P_2^* against knowledge-soundness of sWIAoK which behaves as follows.

- P_2^* runs $\text{RobustSim}^{I_{\text{sound}}^{(3)}}(z)$, where $I_{\text{sound}}^{(3)}$ behaves the same as $I_{\text{sound}}^{(1)}$, except for the following modification. Recall that $I_{\text{sound}}^{(1)}$ isolated sWIAoK of Step bps_2 and forwarded it to an external sWIAoK verifier; here $I_{\text{sound}}^{(3)}$ runs the sWIAoK verifier’s task of sWIAoK of Step bps_2 by itself.
- $I_{\text{sound}}^{(3)}$ instead isolates sWIAoK of Step bps_5 and forwards it to an external sWIAoK verifier.

- Upon completion of this sWIAoK protocol, if the sWIAoK verifier accepts, then it runs sWIAoK extractor on $\text{RobustSim}_{\text{sound}}^{I^{(3)}}(z)$.

Since the sWIAoK argument is isolated, knowledge-soundness of sWIAoK implies that the sWIAoK extractor extracts a valid sWIAoK witness – *either* y such that $(x, y) \in R_L$, *or* an opening of Com_{sh} at Step bps_2 to σ . We shall shortly show that owing to CB of Com_{sh} of Step bps_2 , the extracted value is not a Com_{sh} opening to σ . Given this, we have that the extracted value is some valid witness y . Existence of a valid witness thus implies soundness.

Now it remains to show that the extracted output of sWIAoK extractor above is not an opening of Com_{sh} to σ , with all but negligible probability. Assume for contradiction that the the extracted output is an opening of Com_{sh} to σ with some non-negligible probability ε . Then we construct an adversary \mathcal{A}_{CB} that breaks CB of Com_{sh} with probability $\varepsilon - \text{negl}(\lambda)$.

\mathcal{A}_{CB} is described as follows. \mathcal{A}_{CB} runs $\text{RobustSim}_{\text{sound}}^{I^{(4)}}(z)$, where $I_{\text{sound}}^{(4)}$ behaves the same as $I_{\text{sound}}^{(3)}$, except for the following modification.

- Unlike $I_{\text{sound}}^{(3)}$ (or $I_{\text{sound}}^{(1)}$) which isolates only one of the two sWIAoK sub-protocols present in our protocol, \mathcal{A}_{CB} isolates both the sWIAoK protocols, one at Step bps_2 and and the other at Step bps_5 .
- Furthermore, it also isolates Com_{sh} of Step bps_2 and forwards it to an external Com_{sh} receiver.

However, the sWIAoK verifiers' roles for both the isolated sWIAoK arguments are played by \mathcal{A}_{CB} itself. As proven earlier, the extracted output of sWIAoK at Step bps_2 is an opening of Com_{sh} , with all but negligible probability; furthermore, the extracted value, however, is *not* σ , with all but negligible probability. Thus, we have that the extracted output obtained by \mathcal{A}_{CB} out of this sWIAoK is $(\delta, \text{rand}_\delta)$, such that $\delta \neq \sigma$ and $(\delta, \text{rand}_\delta)$ is a valid opening of Com_{sh} . Furthermore, as also proven above, the extracted output of sWIAoK at Step bps_2 is *either* an opening of $\text{NMMXCom}_{\text{srs}}$ to a valid witness y , *or* an opening of Com_{sh} at Step bps_2 to σ . If the extracted output is the latter, i.e., an opening of Com_{sh} to σ , then \mathcal{A}_{CB} has openings of Com_{sh} (which is isolated and given to an external Com_{sh} receiver) to two distinct values δ and σ . In that event, \mathcal{A}_{CB} breaks computational binding of Com_{sh} . Since neither $I_{\text{sound}}^{(4)}$ nor $I_{\text{sound}}^{(3)}$ isolate any CECOM commitments, applying Corollary 3 of the Robust Extraction Lemma, the view of the adversary \mathcal{P}^* when run by \mathcal{A}_{CB} is identical to its view when run by P_2^* . Thus, \mathcal{A}_{CB} breaks computational binding of Com_{sh} with probability $\varepsilon - \text{negl}(\lambda)$ (where, $\text{negl}(\lambda)$ corresponds to the event that \mathcal{A}_{CB} fails in the sWIAoK extraction from sWIAoK at Step bps_2).

Since Com_{sh} is computationally binding, we have proven that the value extracted from sWIAoK at Step bps_5 is an opening of $\text{NMMXCom}_{\text{srs}}$ to y with all but negligible probability. From the existence of a valid witness, we have that $x \in L$, thus proving soundness. □

C.2 $\langle \mathcal{P}, \mathcal{V} \rangle$ Satisfies Simulatability and Extractability

Here, we prove that there exists a PPT simulator-extractor \mathcal{SE} that, for every concurrent man-in-the-middle adversary \mathcal{M} with some auxiliary information z , outputs a view, view , that is statistically indistinguishable from the view $\text{view}_{\mathcal{M}}(x_1, \dots, x_{m_L}, z)$ of \mathcal{M} in a real execution, and also outputs valid witnesses $\bar{y}_1, \dots, \bar{y}_{m_R}$ for all accepting right sessions except for those right sessions

that are just copied off from some left session. That is, we would like that $\mathcal{SE}(x_1, \dots, x_{m_L}, z) = (\text{view}, \bar{y}_1, \dots, \bar{y}_{m_R})$. In the following we describe our simulator-extractor \mathcal{SE} .

Our simulator-extractor. Let \mathcal{M} be a concurrent man-in-the-middle adversary and let $x_1, \dots, x_{m_L} \in L$ be the statements of the left sessions. The simulator-extractor \mathcal{SE} outputs the output of $\text{RobustSim}^I(z)$, where procedure I – which has black-box access to adversary \mathcal{M} – is described shortly. Before we proceed, to briefly recall, the robust simulator RobustSim interacts with an adversary I which mounts a robust concurrent attack by committing to multiple CECom commitments and interacting with an external party B in a protocol Π . Under such an attack, RobustSim is guaranteed to extract commitment information from every CECom commitment sent by the adversary I before the completion of its commitment phase, in such a way that the external party B does not get rewound.

Procedure $I(z)$. Procedure I launches the robust-concurrent attack by committing in several CECom commitments to external receivers. At the end of each of those CECom commitments, it expects to receive a string. I incorporates the MiM adversary \mathcal{M} internally as a black-box. I initiates an execution of \mathcal{M} , simulating its view as follows. Let the m_L left sessions be ordered with some arbitrary ordering. Let the m_R right sessions be ordered as follows: Consider any two right sessions, the i -th and the j -th; $i \leq j$ if and only if the CECom_{sb} commitment at Step bps_1 of the i -th session begins earlier to the CECom_{sb} commitment at Step bps_1 of the j -th session.

For right sessions: When \mathcal{M} initiates an t -th new session on the right, I runs the code of the honest verifier of $\langle \mathcal{P}, \mathcal{V} \rangle$ except for the following modification.

- Initiate a new CECom commitment with an external CECom receiver and upon \mathcal{M} initiating CECom_{sh} at Step bps_{4+} , relay messages between \mathcal{M} and the external receiver. Let value y'_t be received from the outside at the end of the CECom_{sh} commitment.
- Include y'_t in the output (of I).

For left sessions: When \mathcal{M} initiates an ℓ -th new session on the left, I proceeds as follows.

- Run the coin-flipping phase and the BPS^{CFP} phase honestly. Let srs be the outcome.
- Initiate a new CECom commitment with an external CECom receiver and upon \mathcal{M} reaching the BPS phase and initiating CECom_{sb} at Step bps_1 , relay messages between \mathcal{M} and the external receiver. Let value σ' be received from the outside at the end of the CECom_{sb} commitment.
- Then commit to σ' using Com_{sh} at Step bps_2 ; also, use the same extracted value as the witness in executing the statistical WIAoK of Step bps_2 .
- Let σ be the value that \mathcal{M} opens its CECom_{sb} commitment (of Step bps_1) to in Step bps_3 . Abort if $\sigma \neq \sigma'$.
- Commit to 0^λ using the mixed non-malleable commitment $\text{NMMXCom}_{\text{srs}}$ in Step bps_4 .
- Commit to 0^λ using the CECom_{sh} commitment in Step bps_{4+} .

- Use the value, σ' , committed to in Step bps_2 as the witness in executing sWIAoK of Step bps_5 .

When \mathcal{M} halts, I outputs the view of \mathcal{M} together with y'_1, \dots, y'_{m_R} , and halts.

Running Time of \mathcal{SE} . Notice that except for the extraction of committed values from the CECOM commitments, all steps of the simulator in every session take only as much time as they would for the honest prover in the real world. Furthermore, recall that the robust concurrent simulator runs in time $\text{poly}(\lambda) \cdot T(\lambda)^2$, where $T(\lambda)$ is the maximum number of total CECOM commitments isolated and forwarded to external CECOM verifiers while interacting with \mathcal{M} . Hence, the running time of the simulator as far as the simulator's extraction step is concerned is also polynomial in λ . Hence, overall, the running time of the simulator is polynomial in the running time of the adversary \mathcal{M} , λ , and $|x_1|, \dots, |x_{m_L}|$.

Statistical Simulation: We shall prove that the view output by \mathcal{SE} is distributed statistically close to the real-world view of the MiM adversary \mathcal{M} .

Theorem 4.

$$\{\text{view}_{\mathcal{M}}(x_1, \dots, x_{m_L})\}_{x_1, \dots, x_{m_L} \in L} \approx_s \{\text{view}\}_{x_1, \dots, x_{m_L} \in L},$$

where, $\text{view}_{\mathcal{M}}(x_1, \dots, x_{m_L})$ is the view of the adversary \mathcal{M} in the real-world and view is the view output by the simulator-extractor \mathcal{SE} .

Proof. We begin by providing an intuition to the proof. To prove the indistinguishability, we firstly take note of the ways in which the view generated by the simulator differs from the real-world view of the MiM adversary. Basically, the differences are that, for left sessions, the simulator does not use valid witnesses but tries to get ‘fake-witnesses’ (which we also sometimes refer to as ‘trapdoors’) via the robust simulator; and for the right sessions, the simulator tries to extract witnesses via the robust simulator. While we know that using the robust simulator can incur at most negligible distance, what still remains to be shown is that the simulator using fake-witnesses for the left sessions also creates at most negligible distance from the real-view. For this, we simply rely on the statistical properties of the sub-protocols in which the simulator uses different values; namely, we prove statistical indistinguishability between the distributions of the real and simulated views by relying upon SH of Com_{sh} of Step bps_2 , sWI property of sWIAoK of Step bps_2 , SH of the mixed non-malleable commitment of Step bps_4 , and sWI of sWIAoK of Step bps_5 – the steps at which the simulator uses different values in left sessions. Except for SH of the mixed non-malleable commitment of Step bps_4 , all the above properties are already guaranteed by the corresponding primitives themselves; however, on the other hand, to ensure that the mixed non-malleable commitment – parameterized by srs which is the outcome of the coin-flipping protocol – is SH, we need to ensure that srs is uniformly random with all but negligible probability (see Proposition 3). Let us thus begin proving the lemma by arguing that in the real-world view srs is uniform in every left session with all but negligible probability. We thus first establish the following claim.

Claim 2. *In the real-world view $\text{view}_{\mathcal{M}}(x_1, \dots, x_{m_L})$, for every left session, srs is uniformly random with all but negligible probability.*

Proof. We begin by outlining the structure of the proof.

1. Firstly, we show that, there exists a PPT algorithm that can extract a value r'_V from CECom_{sh} of Step cfp_1 of every left session *before* Step cfp_2 of that session is reached. Thus, since r_P is sent to the adversary after r'_V is extracted, r'_V is independent of r_P , and since r_P is uniformly random, $r_P \oplus r'_V$ is also uniformly random with all but negligible probability.
2. Then, we show that, in every left session, with all but negligible probability, $r'_V = r_V$, where, r_V is the value sent by \mathcal{M} in Step cfp_3 .

The above items together imply that $\text{srs} = r_P \oplus r_V$ is uniformly random, with all but negligible probability. In the rest of the proof, we shall refer to the above two steps as the ‘*steps of our proof*’.

Towards proving these steps, we shall present an equivalent description of the real-world experiment and then proceed with our proof by using this description as the base. More specifically, this equivalent description consists of an interface I_{real} that incorporates the MiM adversary \mathcal{M} as a black-box; then, the output of the real-world experiment will be the output of the robust simulator run on I_{real} . I_{real} , intuitively will be a dummy interface that does not isolate any sub-protocols. That is, it invokes \mathcal{M} and runs the code of the honest provers and honest verifiers. We shall later build upon this interface while slowly isolating various sub-protocols.

Procedure $I_{\text{real}}(z)$. Let \mathcal{M} be a concurrent man-in-the-middle adversary and let $x_1, \dots, x_{m_L} \in L$ be the statements of the left sessions. For every $\ell \in [m_L]$, I_{real} receives y_ℓ such that $(x_\ell, y_\ell) \in R_L$. I_{real} incorporates the MiM adversary \mathcal{M} internally, as a black-box. I_{real} invokes \mathcal{M} . For every left session initiated by \mathcal{M} , I_{real} runs the code of the honest prover of $\langle \mathcal{P}, \mathcal{V} \rangle$. For every right session, it runs the code of the honest verifier of $\langle \mathcal{P}, \mathcal{V} \rangle$. This basic interface does not isolate any CECom commitments. When \mathcal{M} halts, I_{real} outputs the view of \mathcal{M} , and halts.

Extracting r'_V before Step cfp_2 for any left session. Consider any $\ell \in [m_L]$. We shall show that one can efficiently extract r'_V before Step cfp_2 of the ℓ -th left session.

We pursue this step of extracting r'_V from CECom_{sh} of Step cfp_1 by modifying the interface I_{real} to $I_{\text{real}}^{(1)}$ which isolates the CECom_{sh} commitment of Step cfp_1 of the ℓ -th left session and forwards it to an external receiver. And then we run the robust simulator on $I_{\text{real}}^{(1)}$.

In detail, we construct a PPT algorithm hyb that interacts with \mathcal{M} , outputs a view distributed statistically close to $\text{view}_{\mathcal{M}}(x_1, \dots, x_{m_L})$, and also extracts a value r'_V from CECom_{sh} of Step cfp_1 of the ℓ -th left session before Step cfp_2 of that session is reached. hyb simply runs $\text{RobustSim}^{I_{\text{real}}^{(1)}}(z)$, where $I_{\text{real}}^{(1)}$ is a modified version of I_{real} and is described as follows. $I_{\text{real}}^{(1)}$ – with a black-box access to adversary \mathcal{M} – behaves the same way as I_{real} except that it also isolates CECom_{sh} commitment of Step cfp_1 of the ℓ -th left session and forwards it to an external CECom receiver. $I_{\text{real}}^{(1)}$ additionally outputs the value r'_V that RobustSim gives to $I_{\text{real}}^{(1)}$ after extracting it from the isolated CECom_{sh} commitment.

Now to argue extractability of r'_V and to argue statistical indistinguishability of the view output by hyb from the real-world view, we need to invoke the Robust Extraction Lemma of [GLP⁺12], for which we consider the following: CECom_{sh} of Step cfp_1 is of k_{cecom} -slots and the external protocol that the robust simulator is participating in, here, is the empty protocol (i.e., $\ell_{\text{cecom}} = k_{\text{cecom}}$ and $\ell_{\text{external}} = 0$). Now by applying the Robust Extraction Lemma, we have that statistical distance between the outputs of the simulator and hyb is at most:

$$\Delta(\lambda) \leq 2^{-(k_{\text{cecom}} - 0 \cdot \log(T(\lambda)))} \leq \text{negl}(\lambda),$$

since, $k_{\text{cecom}} \in \omega(\log(\lambda))$.

Also by applying the Robust Extraction Lemma, if BPS^{CFP} phase of the ℓ -th left session is successfully completed, RobustSim , and hence hyb , fail to extract r'_V with at most negligible probability. Since $\text{RobustSim}^{I_{\text{real}}^{(1)}}(z)$ runs in polynomial time (from the Robust Extraction Lemma stated in Lemma 3), we have proven that one can extract r'_V from the Step cfp_1 CECom commitment of the ℓ -th left session before Step cfp_2 is reached. Hence, in Step cfp_2 , when $I_{\text{real}}^{(1)}$ sends uniformly chosen r_P , we have that $r_P \oplus r'_V$ is uniformly random, with all but negligible probability. This establishes the first step of our proof.

Showing that $r'_V = r_V$. Now it remains to show that, in every left session, with all but negligible probability, $r'_V = r_V$, where, r_V is the value given by \mathcal{M} in Step cfp_3 . Assume for contradiction that there exists a left session ℓ such that $r'_V \neq r_V$ with some non-negligible probability ϵ . Then we construct an adversary \mathcal{A}_{CB} that breaks computational binding property of CECom_{sh} (of Step cfp_1) with probability $\epsilon - \text{negl}(\lambda)$. That is \mathcal{A}_{CB} should be able to open a random CECom_{sh} commitment to two distinct values with probability $\epsilon - \text{negl}(\lambda)$. Intuitively, recall that we have shown how to extract an opening of this commitment to r'_V in the description of hyb ; it remains to show how \mathcal{A}_{CB} can get an opening (of the same commitment) to r_V too. Observe that in our protocol the verifier would *never* open the CECom_{sh} commitment, but would only argue knowledge of the committed value in the BPS^{CFP} phase. In fact, we can show that the value committed to in the NMCom commitment Com_{nm} – which is SB since it is parameterized by a random DDH tuple ddh – is the randomness rand which is an opening of the CECom_{sh} commitment to r_V . (Indeed, referring back to the description of our protocol, rand is what an honest verifier commits to using Com_{nm} ; the idea thus will be to show that the verifier cannot cheat here). Then \mathcal{A}_{CB} would obtain this rand by running the NMCom extractor. (Recall that we refer to rand as the sub-witness for the session).

Towards maintaining the flow of the proof, let us for now proceed in proving that \mathcal{A}_{CB} breaks computational binding property of CECom_{sh} with probability $\epsilon - \text{negl}(\lambda)$, under the assumption that the value committed to in the NMCom commitment Com_{nm} of BPS^{CFP} phase is rand ; at the end of the proof of this lemma, in Sub-claim 2, we shall establish that this assumption is true with all but negligible probability. Description of \mathcal{A}_{CB} follows.

\mathcal{A}_{CB} behaves the same way as hyb , with the only two differences being the following.

1. While hyb ran $\text{RobustSim}^{I_{\text{real}}^{(1)}}(z)$, \mathcal{A}_{CB} runs $\text{RobustSim}^{I_{\text{real}}^{(2)}}(z)$, where $I_{\text{real}}^{(2)}$ differs from $I_{\text{real}}^{(1)}$ as follows.
 - $I_{\text{real}}^{(2)}$ isolates the NMCom commitment Com_{nm} at Step $\text{bps}^{\text{cfp}}_4$ of the ℓ -th left session. (Recall that the only message that $I_{\text{real}}^{(1)}$ isolates is the CECom commitment CECom_{sh} of Step cfp_1 of the same left session.)
2. \mathcal{A}_{CB} runs the code of the honest NMCom receiver for the isolated NMCom commitment.
3. \mathcal{A}_{CB} runs the NMCom extractor on $\text{RobustSim}^{I_{\text{real}}^{(2)}}(z)$ for the isolated NMCom commitment.

Observe that, since the NMCom commitment is isolated, rewindings by the robust simulator do not interfere with the NMCom commitment. (An alternative and a more intuitive reason, especially here in this part of the proof, is that the only two messages isolated are the CECom commitment at Step cfp_1 and the NMCom commitment at Step $\text{bps}^{\text{cfp}}_4$ of the *same* session. Thus, we have that the

isolated CECOM commitment completes before the beginning of the NMCOM commitment, thus ensuring us that the rewindings on the CECOM commitment will not interfere with the NMCOM commitment).

Thus, extractability of the NMCOM commitment ensures that \mathcal{A}_{CB} succeeds in extracting a valid sub-witness. Also, since the CECOM_{sh} of Step cfp_1 is isolated and forwarded to an external CECOM receiver, \mathcal{A}_{CB} outputting two openings to two distinct values for this commitment amounts to breaking binding of this commitment. Moreover, to ensure extractability from the isolated CECOM commitments and to ensure that the view of \mathcal{M} in its interaction with \mathcal{A}_{CB} is statistically close to the view output by hyb , we apply the Robust Extraction Lemma of [GLP⁺12] (stated in Lemma 3), which roughly says that the robust simulator which avoids rewinding any external parties outputs a view that is statistically close to the view output by an online extractor which (not necessarily running in polynomial time) with all but negligible probability provides a valid commitment information for every CECOM commitment relayed to the online extractor.

In order to apply the Robust Extraction Lemma, we will first create two hybrids hyb_A^* and hyb_B^* , whose outputs are identical, and the output of the former is statistically close to the view of \mathcal{M} when run by hyb and the output of the latter is statistically close to the view of the MiM adversary \mathcal{M} when run by \mathcal{A}_{CB} .

hyb_A^* is described as follows. It simply outputs the view output by the online extractor, namely, $\text{REAL}_{\mathcal{E}, \Pi}^{I_{\text{real}}^{(1)}}(\beta, z)$, where the external protocol Π here is the empty protocol. Now, we can invoke the Robust Extraction Lemma, for which we consider the following: the CECOM commitments, that are isolated and relayed to external CECOM receivers, have k_{cecom} -slots and the external protocol that the robust simulator is participating in, here, is the empty protocol. Thus, we have that $\ell_{\text{cecom}} = k_{\text{cecom}}$ and $\ell_{\text{external}} = 0$. Now by applying the Robust Extraction Lemma, we have that statistical distance between the views $\text{REAL}_{\mathcal{E}, \Pi}^{I_{\text{real}}^{(1)}}(\beta, z)$ and the view of the adversary when run by hyb is at most:

$$\Delta(\lambda) \leq 2^{-(k_{\text{cecom}} - 0 \cdot \log(T(\lambda)))} \leq \text{negl}(\lambda),$$

since, $k_{\text{cecom}} \in \omega(\log(\lambda))$ and T is at most a polynomial.

Next, we describe another hybrid hyb_B^* whose output is identical to that of hyb_A^* . For this consider an interface, $I_{\text{extr}}^{(B)}$ behaves the same way as $I_{\text{real}}^{(1)}$ except that it also isolates the Com_{nm} commitment at Step $\text{bps}^{\text{cfp}}_4$ of the ℓ -th left session and forwards it to an external NMCOM receiver.

hyb_B^* simply outputs the view output by the online extractor, namely, $\text{REAL}_{\mathcal{E}, \Pi}^{I_{\text{extr}}^{(B)}}(\beta, z)$, where the external protocol Π here consists of the Com_{nm} commitment at Step $\text{bps}^{\text{cfp}}_4$ of the ℓ -th left session and the external party runs the code of NMCOM receiver for the isolated commitments. Again, we can invoke the Robust Extraction Lemma, for which we consider the following: the CECOM commitments, that are isolated and relayed to external CECOM receivers, have k_{cecom} -slots and the external protocol Π that the robust simulator is participating in, here, is the Com_{nm} commitment at Step $\text{bps}^{\text{cfp}}_4$ of the ℓ -th left session. Thus, we have that $\ell_{\text{cecom}} = k_{\text{cecom}}$ and $\ell_{\text{external}} = k_{\text{nmcom}}$. Now by applying the Robust Extraction Lemma, we have that statistical distance between the views $\text{REAL}_{\mathcal{E}, \Pi}^{I_{\text{extr}}^{(B)}}(\beta, z)$ and the view of the \mathcal{M} during its interaction with \mathcal{A}_{CB} is at most:

$$\Delta(\lambda) \leq 2^{-(k_{\text{cecom}} - k_{\text{nmcom}} \cdot \log(T(\lambda)))} \leq \text{negl}(\lambda),$$

since, $k_{\text{cecom}} \in \omega(\log(\lambda))$, $k_{\text{nmcom}} \in O(\log(\lambda))$, and T is at most a polynomial.

Thus, we have proven that the view of the MiM adversary \mathcal{M} during its interaction with \mathcal{A}_{CB} is statistically close to its view during its interaction with hyb . Also, by invoking the Robust Extraction Lemma, we have that the CECOM extraction is successful with all but negligible probability.

We have thus proven that, with all but negligible probability, \mathcal{A}_{CB} extracts (r'_V, rand') through robust simulator and rand through the NMCom extractor for r_V , with these values being such that rand explains r_V being committed to CECom_{sh} commitment of Step cfp_1 and rand' explains r'_V being committed to in the same CECom_{sh} commitment.

Moreover, we have also proven that the view of the MiM adversary \mathcal{M} when run by \mathcal{A}_{CB} is statistically close to the view of \mathcal{M} when run by hyb . Since $r_V \neq r'_V$ with probability $\epsilon - \text{negl}(\lambda)$, \mathcal{A}_{CB} thus breaks computational binding of CECom_{sh} commitment with probability $\epsilon - \text{negl}(\lambda)$.

Finally, once we prove validity of our assumption (made earlier in the proof) that the value committed to in the NMCom commitment Com_{nm} at Step $\text{bps}^{\text{cfp}}_4$ of the ℓ -th left session is indeed a valid sub-witness with all but negligible probability, we will have established the second (and the final) step of the proof (listed at the beginning of the proof). We now set out to prove validity of this assumption.

Sub-Claim 2. *In the real world view $\text{view}_{\mathcal{M}}(x_1, \dots, x_{m_L})$, if BPS^{CFP} phase of the ℓ -th left session is accepted by the ℓ -th prover \mathcal{P}_ℓ , then the value committed to by \mathcal{M} in the NMCom commitment Com_{nm} at Step $\text{bps}^{\text{cfp}}_4$ of the ℓ -th left session is a valid sub-witness.*

Proof. Intuitively, Com_{nm} at Step $\text{bps}^{\text{cfp}}_4$ of the ℓ -th left session contains a valid sub-witness owing to:

- computational hiding of CECom_{sb} – to argue that \mathcal{M} does not learn α , committed to by the prover in CECom_{sb} , and use it in its commitment Com_{sh} and sWIAoK at Step $\text{bps}^{\text{cfp}}_2$,
- knowledge-soundness of sWIAoK in Step $\text{bps}^{\text{cfp}}_2$ – to extract knowledge of commitment information (i.e., commitment value and randomness) for Com_{sh} in Step $\text{bps}^{\text{cfp}}_2$ and to verify that the extracted value will not be α ,
- knowledge-soundness of sWIAoK in Step $\text{bps}^{\text{cfp}}_5$ – to argue that either the value committed to in Com_{nm} at Step $\text{bps}^{\text{cfp}}_4$ is a valid sub-witness or to argue knowledge of a commitment information for Com_{sh} in Step $\text{bps}^{\text{cfp}}_2$ with commitment value as α ,
- and finally, computational binding of Com_{sh} at Step $\text{bps}^{\text{cfp}}_2$ to show the knowledge extracted is not α as a commitment value.

Using CH of CECom_{sb} and knowledge-soundness of sWIAoK in Step $\text{bps}^{\text{cfp}}_2$. We begin by showing that one can extract a commitment information for Com_{sh} at Step $\text{bps}^{\text{cfp}}_2$ from sWIAoK at the same Step, and by computational hiding of CECom_{sb} at Step $\text{bps}^{\text{cfp}}_1$, the value will not be α .

Consider an adversarial prover P_1^* against knowledge-soundness of sWIAoK which behaves as follows. Here again we consider the dummy interface I_{real} and build upon it. P_1^* runs $\text{RobustSim}^{I_{\text{real}}^{(2')}}(z)$, where $I_{\text{real}}^{(2')}$ where $I_{\text{real}}^{(2')}$ differs from I_{real} in the following sense.

- $I_{\text{real}}^{(2')}$ isolates sWIAoK in Step $\text{bps}^{\text{cfp}}_2$ of the ℓ -th left session and forwards it to an external sWIAoK verifier.

Upon completion of this sWIAoK protocol, if the sWIAoK verifier accepts, then it runs sWIAoK extractor on $\text{RobustSim}^{I_{\text{real}}^{(2')}}(z)$.

Observe that the only sub-protocol isolated by $I_{\text{real}}^{(2')}$ is the sWIAoK protocol. Hence, there are no other rewindings that could interfere with his sWIAoK protocol. Thus, since the sWIAoK argument is isolated and relayed to an honest external sWIAoK verifier, knowledge-soundness of sWIAoK implies that the sWIAoK extractor extracts a valid sWIAoK witness – the value committed to together with the randomness used – of the Com_{sh} commitment of Step $\text{bps}^{\text{cfp}}_2$. This also implies that the value is not α (committed to in CECom_{sb} of Step $\text{bps}^{\text{cfp}}_1$) with all but negligible probability, as otherwise we can build an adversary \mathcal{A}_{CH} that breaks computational hiding of CECom_{sb} of Step $\text{bps}^{\text{cfp}}_1$ as follows.

Assume for contradiction that the value extracted by P_1^* is α with some non-negligible probability ϵ . Then we shall show that \mathcal{A}_{CH} breaks hiding with probability $\epsilon - \text{negl}(\lambda)$. \mathcal{A}_{CH} is described as follows. \mathcal{A}_{CH} runs $\text{RobustSim}^{I_{\text{real}}^{(2')}}(z)$, where $I_{\text{real}}^{(2')}$ behaves the same as $I_{\text{real}}^{(2')}$ except that, besides isolating the sWIAoK argument of Step $\text{bps}^{\text{cfp}}_2$, also isolates CECom_{sb} of Step bps_1 . While the CECom_{sb} commitment is forwarded to an external CECOM sender, \mathcal{A}_{CH} itself runs the honest verifier code of the isolated sWIAoK argument. If the sWIAoK verifier (run by \mathcal{A}_{CH}) accepts the sWIAoK argument, then \mathcal{A}_{CH} also runs the sWIAoK extractor on the isolated sWIAoK argument. Furthermore, \mathcal{A}_{CH} *does not continue* the interaction with \mathcal{M} after the sWIAoK argument. Crucially, note that the isolated sWIAoK argument (at Step $\text{bps}^{\text{cfp}}_2$) begins strictly after the completion of the isolated CECom_{sb} commitment (at Step $\text{bps}^{\text{cfp}}_1$) as they both belong to the same session. Hence, the sWIAoK rewindings do not interfere with the isolated CECom_{sb} commitment thus ensuring that the external CECOM sender will not be rewound.

Finally, note that neither $I_{\text{real}}^{(2')}$ nor $I_{\text{real}}^{(2')}$ (which are the interfaces used by P_1^* and \mathcal{A}_{CH} , respectively,) neither $I_{\text{sound}}^{(2)}$ nor $I_{\text{sound}}^{(1)}$ isolate any CECOM commitments, applying Corollary 3 of the Robust Extraction Lemma, the view of the adversary \mathcal{M} when run by \mathcal{A}_{CH} is identical to its view when run by P_1^* until the sWIAoK argument at Step $\text{bps}^{\text{cfp}}_2$ of the ℓ -th left session, (after which \mathcal{A}_{CH} aborts). Furthermore, since the view of the sWIAoK extractor also remains identical, we have that the probability that sWIAoK extractor extracts α when run by P_1^* is equal to that when run by \mathcal{A}_{CH} , thus breaking computational hiding of CECom_{sb} of Step bps_1 with the same probability.

Using knowledge-soundness of sWIAoK in Step $\text{bps}^{\text{cfp}}_5$ and CB of Com_{sh} at Step $\text{bps}^{\text{cfp}}_2$. Now, we show that one can extract a witness for sWIAoK of Step $\text{bps}^{\text{cfp}}_5$, and from its knowledge-soundness, we have that either we extract a commitment information (i.e., a commitment value and randomness) in Com_{nm} at Step $\text{bps}^{\text{cfp}}_4$ such that this value is a valid sub-witness or we extract an opening of Com_{sh} at Step $\text{bps}^{\text{cfp}}_2$ to α . Finally, we will see that CB of Com_{sh} at Step $\text{bps}^{\text{cfp}}_2$ implies that the extracted value is not an opening of Com_{sh} to α . Putting it all together, we will have established soundness of the argument proved in the BPS^{CFP} phase. Finally, since Com_{nm} is statistically binding, we will have that the value committed in it is a valid sub-witness, rand , with all but negligible probability.

Consider an adversarial prover P_2^* against knowledge-soundness of sWIAoK which behaves as follows. Recall that P_1^* executed $\text{RobustSim}^{I_{\text{real}}^{(2')}}(z)$, where $I_{\text{real}}^{(2')}$ differs from I_{real} in that it isolated sWIAoK in Step $\text{bps}^{\text{cfp}}_2$ of the ℓ -th left session. P_2^* instead runs $\text{RobustSim}^{I_{\text{real}}^{(2'')}}(z)$, where $I_{\text{real}}^{(2'')}$ instead isolates sWIAoK in Step $\text{bps}^{\text{cfp}}_5$ of the same session. That is, $I_{\text{real}}^{(2'')}$ differs from I_{real} as

follows.

- $I_{\text{real}}^{(2')}$ isolates sWIAoK in Step $\text{bps}^{\text{cfp}}_5$ of the ℓ -th left session and forwards it to an external sWIAoK verifier.

Upon completion of this sWIAoK protocol, if the sWIAoK verifier accepts, then it runs sWIAoK extractor on $\text{RobustSim}_{\text{real}}^{I_{\text{real}}^{(2')}}(z)$.

Like for P_1^* , the only sub-protocol isolated by $I_{\text{real}}^{(2')}$ is the sWIAoK protocol. Hence, there are no other rewindings that could interfere with this sWIAoK protocol. Thus, since the sWIAoK argument is isolated and relayed to an honest external sWIAoK verifier, knowledge-soundness of sWIAoK implies that the sWIAoK extractor extracts a valid sWIAoK witness – namely, *either* rand such that rand explains CECom_{sh} at Step cfp_1 to be a commitment to r_V , *or* an opening of Com_{sh} at Step $\text{bps}^{\text{cfp}}_2$ to α .

We shall shortly show that owing to CB of Com_{sh} , the extracted value is not a Com_{sh} opening to α . With this, we will have proven that the extracted value is rand . Since Com_{nm} is statistically binding, *the value committed in it should be rand itself*, with all but negligible probability.

Now it remains to be shown that the extracted output of sWIAoK extractor above is not an opening of Com_{sh} to α with all but negligible probability. Assume for contradiction that the the extracted output is an opening of Com_{sh} to α with probability ϵ' . Then we construct an adversary \mathcal{A}'_{CB} that breaks CB of Com_{sh} with probability $\epsilon' - \text{negl}(\lambda)$.

\mathcal{A}'_{CB} runs $\text{RobustSim}_{\text{real}}^{I_{\text{real}}^{(2')}}(z)$, where $I_{\text{real}}^{(2')}$ differs from I_{real} in the following sense.

- $I_{\text{real}}^{(2')}$ isolates Com_{sh} of Step $\text{bps}^{\text{cfp}}_2$ of the ℓ -th session and forwards it to an external Com_{sh} receiver.
- $I_{\text{real}}^{(2')}$ also isolates sWIAoK in Step $\text{bps}^{\text{cfp}}_2$ and sWIAoK in Step $\text{bps}^{\text{cfp}}_5$ of the ℓ -th left session. (Recall that in contrast the interfaces for P_1^* , P_2^* isolated only one of the two sWIAoK arguments).

\mathcal{A}'_{CB} itself runs the code of the honest sWIAoK verifiers for these sWIAoK arguments. Furthermore, \mathcal{A}'_{CB} checks if both the sWIAoK protocols are accepting. If so, then it runs the sWIAoK extractor on $\text{RobustSim}_{\text{real}}^{I_{\text{real}}^{(2')}}(z)$ once over the sWIAoK protocol at Step $\text{bps}^{\text{cfp}}_2$ and again over the sWIAoK protocol at Step $\text{bps}^{\text{cfp}}_5$. If either extraction fails, it aborts.

Observe that the only sub-protocols isolated by \mathcal{A}'_{CB} are the sWIAoK arguments at Step $\text{bps}^{\text{cfp}}_2$ and at Step $\text{bps}^{\text{cfp}}_5$ and the Com_{sh} commitment of Step $\text{bps}^{\text{cfp}}_2$, all of the same ℓ -th session. Hence we are assured that the sWIAoK extractions do not interfere with the isolated Com_{sh} commitment thus implying that the external Com_{sh} receiver will not be rewound. This in turn implies that \mathcal{A}'_{CB} outputting two openings to two distinct values amounts to breaking binding of Com_{sh} .

As proven earlier, the extracted output of sWIAoK at Step $\text{bps}^{\text{cfp}}_2$ is an opening of Com_{sh} with all but negligible probability; as also proven earlier, the extracted value, however, is *not* α with all but negligible probability. Thus, we have the extracted output obtained by \mathcal{A}'_{CB} out of this sWIAoK is $(\delta, \text{rand}_\delta)$, such that $\delta \neq \alpha$ and $(\delta, \text{rand}_\delta)$ is a valid opening of Com_{sh} . Furthermore, as also proven above, the extracted output of sWIAoK at Step $\text{bps}^{\text{cfp}}_5$ is *either* an opening of Com_{nm} to a value rand such that rand explains CECom_{sh} at Step cfp_1 to be a commitment to r_V , *or* an opening of Com_{sh} at Step $\text{bps}^{\text{cfp}}_2$ to α . If the extracted output is the latter, i.e., an opening of Com_{sh} to α , then \mathcal{A}'_{CB} has openings of Com_{sh} (which is isolated) to two distinct values δ and α . Thus, \mathcal{A}'_{CB}

breaks computational binding of Com_{sh} with probability $\epsilon' - \text{negl}(\lambda)$, under the assumption that the view of the MiM adversary \mathcal{M} when run by \mathcal{A}'_{CB} is statistically close to its view when run by P_2^* . Now to prove right the assumption: we make the following observations. The interface in \mathcal{A}'_{CB} , namely $I_{\text{real}}^{(2')}$, behaves the same way as I_{real} except that it also isolates certain sub-protocols, and even the isolated protocols are run honestly. Also, neither interfaces of \mathcal{A}'_{CB} and P_2^* isolate any CECOM commitments, applying Corollary 3 of the Robust Extraction Lemma, the view of the adversary \mathcal{M} when run by \mathcal{A}'_{CB} is identical to its view when run by P_2^* , under the event that \mathcal{A}'_{CB} does not abort due to failure in sWIAoK extractions. Furthermore, since the sWIAoK extractor fails in the extraction with only negligible probability, we have proven that \mathcal{A}'_{CB} breaks computational binding of Com_{sh} with probability $\epsilon' - \text{negl}(\lambda)$ (where, $\text{negl}(\lambda)$ corresponds to the event that \mathcal{A}'_{CB} fails in either of the two sWIAoK extractions).

Thus, we have that the value extracted from sWIAoK at Step $\text{bps}^{\text{cfp}}_5$ is an opening of Com_{nm} to rand with all but negligible probability. Since Com_{nm} is SB, the value committed in it should be rand itself, with all but negligible probability, thus proving Sub-claim 2. □

This concludes the proof of Claim 2. □

We have thus proven that in the real-world view of the adversary, for every left session, srs is uniformly random with all but negligible probability. With this we are now ready to prove statistical indistinguishability between the real-world view of the adversary and the simulated view output by the simulator-extractor \mathcal{SE} . We begin with the real-world game and reach \mathcal{SE} through a series of hybrids $\text{hyb}_0, \dots, \text{hyb}_7$ as follows. Also, let us denote the view output by a hybrid hyb_i by $\text{view}^{(i)}$.

hyb₀: This is identical to the real-world experiment. To understand the upcoming hybrids easily, we choose to explain this hybrid again in terms of the dummy interface I_{real} , defined in the proof of Claim 2. Recall that I_{real} receives valid witnesses for all left sessions, invokes the MiM adversary \mathcal{M} , and simply runs the code of honest provers and honest verifiers in the left and right sessions, respectively. hyb_0 outputs $\text{RobustSim}^{I_{\text{real}}}(z)$. For the sake of notations, we shall rename I_{real} as $I_{\text{real}}^{(0)}$.

hyb₁: hyb_1 runs $\text{RobustSim}^{I_{\text{real}}^{(1)}}(z)$, where $I_{\text{real}}^{(1)}$ behaves the same way as $I_{\text{real}}^{(0)}$ except for the following modification.

- For every right session, $I_{\text{real}}^{(1)}$ initiates a new CECOM commitment with an external CECOM receiver and upon \mathcal{M} initiating CECOM_{sh} at Step bps_{4+} , relays messages between \mathcal{M} and the external receiver. Let value y'_t be received from the outside at the end of the CECOM_{sh} commitment.

Claim 3.

$$\text{view}^{(1)} \approx_s \text{view}^{(0)}.$$

Proof. We prove this by applying the Robust Extraction Lemma. Since in hyb_0 , $I_{\text{real}}^{(0)}$ does not isolate any CECOM commitments, to apply the Robust Extraction Lemma, we will first create an

intermediate hybrid hyb_1^* . hyb_1^* simply outputs the view output by the online extractor, namely, $\text{REAL}_{\mathcal{E}, \Pi}^{I_{\text{real}}^{(1)}}(\beta, z)$, where the external protocol Π here is the empty protocol. Note that this view is identical to the view of the MiM adversary \mathcal{M} when run by hyb_0 , since $I_{\text{real}}^{(1)}$ behaves the same way as $I_{\text{real}}^{(0)}$ except that it also isolates certain CECOM commitments and relays them to external CECOM receivers. That is,

$$\text{REAL}_{\mathcal{E}, \Pi}^{I_{\text{real}}^{(1)}}(\beta, z) \equiv \text{view}^{(0)}.$$

Now, we can invoke the Robust Extraction Lemma, for which we consider the following: the CECOM commitments, that are isolated and relayed to external CECOM receivers, have k_{cecom} -slots and the external protocol that the robust simulator is participating in, here, is the empty protocol. Thus, we have that $\ell_{\text{cecom}} = k_{\text{cecom}}$ and $\ell_{\text{external}} = 0$. Now by applying the Robust Extraction Lemma, we have that statistical distance between the views $\text{REAL}_{\mathcal{E}, \Pi}^{I_{\text{real}}^{(1)}}(\beta, z)$ and $\text{view}^{(1)}$ output by \mathcal{H}_1^* and \mathcal{H}_1 , respectively, is at most:

$$\Delta(\lambda) \leq 2^{-(k_{\text{cecom}} - 0 \cdot \log(T(\lambda)))} \leq \text{negl}(\lambda),$$

since, $k_{\text{cecom}} \in \omega(\log(\lambda))$ and T is at most a polynomial. Thus,

$$\text{view}^{(1)} \approx_s \text{REAL}_{\mathcal{E}, \Pi}^{I_{\text{real}}^{(1)}}(\beta, z).$$

□

hyb₂: hyb_2 runs $\text{RobustSim}^{I_{\text{real}}^{(2)}}(z)$, where $I_{\text{real}}^{(2)}$ behaves the same way as $I_{\text{real}}^{(1)}$ except for the following modification.

- For every left session, $I_{\text{real}}^{(2)}$ initiates a new CECOM commitment with an external CECOM receiver and upon \mathcal{M} initiating CECOM_{sb} at Step bps_1 , relays messages between \mathcal{M} and the external receiver. Let value σ' be received from the outside at the end of the CECOM_{sb} commitment.

Claim 4.

$$\text{view}^{(2)} \approx_s \text{view}^{(1)}.$$

Proof. Recall that hyb_1 and hyb_2 run $\text{RobustSim}^{I_{\text{real}}^{(1)}}(z)$ and $\text{RobustSim}^{I_{\text{real}}^{(2)}}(z)$, respectively, where $I_{\text{real}}^{(1)}$ and $I_{\text{real}}^{(2)}$ differ as follows. While $I_{\text{real}}^{(1)}$ isolates only some CECOM commitments of the right sessions, $I_{\text{real}}^{(2)}$ also isolates some CECOM commitments of the left sessions. Now to establish statistical indistinguishability between the views output by hyb_1 and hyb_2 , we apply the Robust Extraction Lemma. Since $I_{\text{real}}^{(1)}$ does not isolate certain CECOM commitments that are isolated by $I_{\text{real}}^{(2)}$, to apply the Robust Extraction Lemma, we will first create an intermediate hybrid hyb_2^* . hyb_2^* simply outputs the view output by the online extractor, namely, $\text{REAL}_{\mathcal{E}, \Pi}^{I_{\text{real}}^{(2)}}(\beta, z)$, where the external protocol Π here is the empty protocol. Note that this view is statistically close to the view of the MiM adversary \mathcal{M} when run by hyb_1 , since $I_{\text{real}}^{(2)}$ behaves the same way as $I_{\text{real}}^{(1)}$ except that it also

isolates certain CECOM commitments that were originally not isolated by $I_{\text{real}}^{(1)}$ and relays them to external CECOM receivers. That is,

$$\text{REAL}_{\mathcal{E}, \Pi}^{I_{\text{real}}^{(2)}}(\beta, z) \approx_s \text{view}^{(1)}.$$

Now, we can invoke the Robust Extraction Lemma, for which we consider the following: the CECOM commitments, that are isolated and relayed to external CECOM receivers, have k_{cecom} -slots and the external protocol that the robust simulator is participating in, here, is the empty protocol. Thus, we have that $\ell_{\text{cecom}} = k_{\text{cecom}}$ and $\ell_{\text{external}} = 0$. Now by applying the Robust Extraction Lemma, we have that statistical distance between the views $\text{REAL}_{\mathcal{E}, \Pi}^{I_{\text{real}}^{(2)}}(\beta, z)$ and $\text{view}^{(2)}$ output by \mathcal{H}_2^* and \mathcal{H}_2 , respectively, is at most:

$$\Delta(\lambda) \leq 2^{-(k_{\text{cecom}} - 0 \cdot \log(T(\lambda)))} \leq \text{negl}(\lambda),$$

since, $k_{\text{cecom}} \in \omega(\log(\lambda))$ and T is at most a polynomial. Thus,

$$\text{view}^{(2)} \approx_s \text{REAL}_{\mathcal{E}, \Pi}^{I_{\text{real}}^{(2)}}(\beta, z).$$

□

hyb₃. **hyb₃** differs from **hyb₂** in that, while **hyb₂** ran $\text{RobustSim}^{I_{\text{real}}^{(2)}}(z)$, **hyb₃** runs $\text{RobustSim}^{I_{\text{real}}^{(3)}}(z)$, where $I_{\text{real}}^{(3)}$ differs from $I_{\text{real}}^{(2)}$ for every left session as follows.

- Recall that $I_{\text{real}}^{(2)}$ (among other messages) isolated the CECOM_{sb} commitments at Step **bps₁** of the left session.
- Let value σ' be received from the outside at the end of the CECOM_{sb} commitment. Then commit to σ' using Com_{sh} at Step **bps₂**; also, use the same extracted value as the witness in proving sWIAoK at Step **bps₂**.

Claim 5.

$$\text{view}^{(3)} \approx_s \text{view}^{(2)}.$$

Proof. Since Com_{sh} at Step **bps^{cfp}₂** is a statistically hiding commitment scheme, and sWIAoK is statistical witness-indistinguishability, the claim follows. □

hyb₄. **hyb₄** differs from **hyb₃** in that, while **hyb₃** ran $\text{RobustSim}^{I_{\text{real}}^{(3)}}(z)$, **hyb₄** runs $\text{RobustSim}^{I_{\text{real}}^{(4)}}(z)$, where $I_{\text{real}}^{(4)}$ differs from $I_{\text{real}}^{(3)}$ in every left session as follows.

- Recall that $I_{\text{real}}^{(3)}$ isolated the CECOM_{sb} commitment at Step **bps₁** of the left session. Let the extracted value received from the outside be σ' . Also, in Step **bps₃**, let σ be the value that \mathcal{M} opens the CECOM_{sb} commitment to. If $\sigma \neq \sigma'$, then abort.

Claim 6.

$$\text{view}^{(4)} \approx_s \text{view}^{(3)}.$$

Proof. Assume for contradiction that there exists $\ell \in [m_L]$ such that for the ℓ -th left session, $\sigma \neq \sigma'$, with some non-negligible probability ϵ . Then we construct an adversary that breaks CB of the CECom_{sb} commitment.

Recall that hyb_4 isolates the CECom_{sb} commitment at Step $\text{bps}_{\text{cfp}_1}$ of the ℓ -th left session (among other messages). It thus receives an opening for it, say $(\sigma', \text{rand}_{\sigma'})$ from the outside. Furthermore, \mathcal{M} provides an opening to the same CECom_{sb} commitment at Step bps_3 ; call it $(\sigma, \text{rand}_\sigma)$. From the assumption in the proof that $\alpha \neq \alpha'$ with some non-negligible probability ϵ , we can construct an adversary that breaks CB of the CECom_{sb} commitment with the same probability. This clearly follows from the fact that the CECom_{sb} commitment in question is already isolated and relayed to an external CECom receiver. Thus, hyb_4 itself can be deemed our adversary against CB of the CECom_{sb} commitment, a contradiction. Thus, $\sigma = \sigma'$, with all but negligible probability, and the Claim follows. \square

hyb_5 . hyb_5 differs from hyb_4 in that, while hyb_4 ran $\text{RobustSim}^{I_{\text{real}}^{(4)}}(z)$, hyb_5 runs $\text{RobustSim}^{I_{\text{real}}^{(5)}}(z)$, where $I_{\text{real}}^{(5)}$ differs from $I_{\text{real}}^{(4)}$ in every ℓ -th left session as follows.

- Recall that, in the ℓ -th left session in question, $I_{\text{real}}^{(4)}$ committed to a valid witness y_ℓ in $\text{NMMXCom}_{\text{srs}}$ at Step bps_4 and in CECom_{sh} at Step bps_{4+} . Furthermore, it uses this committed value y_ℓ and the randomnesses used in the $\text{NMMXCom}_{\text{srs}}$ and CECom_{sh} commitments as the witness in proving sWIAoK at Step bps_5 . Here, the modification is that $I_{\text{real}}^{(5)}$ uses the commitment information from Com_{sh} at Step bps_2 where it committed to σ as the witness for the sWIAoK argument.

Claim 7.

$$\text{view}^{(5)} \approx_s \text{view}^{(4)}.$$

Proof. Since sWIAoK at Step bps_5 is statistical witness-indistinguishability, the claim follows. \square

hyb_6 . hyb_6 differs from hyb_5 in that, while hyb_5 ran $\text{RobustSim}^{I_{\text{real}}^{(5)}}(z)$, hyb_6 runs $\text{RobustSim}^{I_{\text{real}}^{(6)}}(z)$, where $I_{\text{real}}^{(6)}$ differs from $I_{\text{real}}^{(5)}$ in every ℓ -th left session as follows.

- Recall that, in the ℓ -th left session in question, $I_{\text{real}}^{(5)}$ committed to a valid witness y_ℓ in $\text{NMMXCom}_{\text{srs}}$ at Step bps_4 . Here, the modification is that $I_{\text{real}}^{(6)}$ commits to 0^λ in $\text{NMMXCom}_{\text{srs}}$ at Step bps_4 .

Claim 8.

$$\text{view}^{(6)} \approx_s \text{view}^{(5)}.$$

Proof. Recall that we have proven in Claim 2 that for every left session srs is uniformly random with all but negligible probability. Thus, from Proposition 3, $\text{NMMXCom}_{\text{srs}}$ at Step bps_4 of left sessions are SH with all but negligible probability. The claim thus immediately follows. \square

hyb₇. hyb₇ differs from hyb₆ in that, while hyb₆ ran $\text{RobustSim}^{I_{\text{real}}^{(6)}}(z)$, hyb₇ runs $\text{RobustSim}^{I_{\text{real}}^{(7)}}(z)$, where $I_{\text{real}}^{(7)}$ differs from $I_{\text{real}}^{(6)}$ in every ℓ -th left session as follows.

- Recall that, in the ℓ -th left session in question, $I_{\text{real}}^{(6)}$ committed to a valid witness y_ℓ in CECom_{sh} at Step bps_{4+} . Here, the modification is that $I_{\text{real}}^{(7)}$ commits to 0^λ in CECom_{sh} at Step bps_{4+} .

Claim 9.

$$\text{view}^{(7)} \approx_s \text{view}^{(6)}.$$

Proof. Since CECom_{sh} is SH, the claim follows. □

Note that hyb₇ does not use any witness for the left sessions. By construction, hyb₇ is identical to our simulator-extractor \mathcal{SE} .

We have thus proven Theorem 4. □

Witness Extractability: We shall prove that the values y'_1, \dots, y'_{m_R} extracted by the simulator-extractor \mathcal{SE} are valid witnesses for the statements of the corresponding right sessions.

Theorem 5. *For every PPT adversary \mathcal{M} , the output of the simulator $\mathcal{SE}(x_1, \dots, x_{m_L}, z) = (\text{view}, \bar{y}_1, \dots, \bar{y}_{m_R})$ is such that, $\forall i \in [m_R], (\bar{x}_i, \bar{y}_i) \in R_L$.*

Proof. We shall prove this theorem through a series of hybrids \mathcal{H}_i , for $i = 1, \dots, 7$. Every \mathcal{H}_i outputs a view, $\text{view}_{\text{extr}}^{(i)}$. We shall show that the final hybrid outputs a view that is identical to the view output by our simulator-extractor \mathcal{SE} . Furthermore, every \mathcal{H}_i will also output a list of values $y_1^{(i)}, \dots, y_{m_R}^{(i)}$. We shall show for one of the hybrids that these values are valid witnesses for the statements being proved in the right sessions. Finally, we shall prove that, not only the views but also these values are (computationally) indistinguishable across the above hybrids, implying that the values output by our simulator-extractor \mathcal{SE} are also valid witnesses, thus establishing ‘witness extractability’.

We define $2m_R$ random variables $\{b_t^{(i)}, \beta_t^{(i)}\}_{t=1}^{m_R}$, where $b_t^{(i)}$ is a bit denoting whether according to $\text{view}_{\text{extr}}^{(i)}$ verifier \mathcal{V}_t accepted the proof from the adversary or not, and $\beta_t^{(i)}$ is the value contained in the $\text{NMMXCom}_{\text{srs}}$ commitment received by \mathcal{V}_t in Step bps_4 ; if there is no unique value, then $\beta_t^{(i)}$ is defined to be \perp .

hyb_{real}: This is just the real-world experiment. Like in the proof of Claim 2, we interpret the real-world view as output of the robust simulator RobustSim which interacts with an interface that launches a robust concurrent attack by incorporating the MiM adversary \mathcal{M} and by playing the role of honest provers and honest verifiers in left and right sessions, respectively. To recall, I_{real} is described as follows.

hyb_0 : This is similar to the real-world experiment hyb_{real} except for a slight modification.

Recall that the real-world view is just the output of $\text{RobustSim}^{I_{\text{real}}}(z)$, where, I_{real} is a dummy interface, defined in the proof of Claim 2. Recall that I_{real} receives valid witnesses for all left sessions, invokes the MiM adversary \mathcal{M} , and simply runs the code of honest provers and honest verifiers in the left and right sessions, respectively. The modification is that hyb_0 runs $\text{RobustSim}^{I_{\text{extr}}^{(0)}}(z)$, where $I_{\text{extr}}^{(0)}$ differs from I_{real} as follows.

- For every t -th right session, isolate CECom_{sh} commitment at Step bps_{4+} and forward it to an external CECom receiver. Let the value received from the outside be $y_t^{(1)}$.

Besides outputting the view $\text{view}_{\text{extr}}^{(0)}$, \mathcal{H}_0 also outputs $y_1^{(0)}, \dots, y_{m_R}^{(0)}$.

Remark 2. From now onwards, all the upcoming hybrids shall isolate CECom_{sh} commitment at Step bps_{4+} of every right session. We shall denote the m_R values extracted in any hybrid by $y_1^{(ind)}, \dots, y_{m_R}^{(ind)}$, where ind is the index of the corresponding hybrid. In fact, these are the values that every hybrid here shall output besides the view of the MiM adversary \mathcal{M} .

Sub-Claim 3.

$$\text{view}_{\text{extr}}^{(0)} \approx_s \text{view}_{\mathcal{M}}(x_1, \dots, x_{m_L}),$$

where, $\text{view}_{\mathcal{M}}(x_1, \dots, x_{m_L})$ is the real-world view of \mathcal{M} .

Proof. We prove this by applying the Robust Extraction Lemma. Since I_{real} does not isolate any CECom commitments, to apply the Robust Extraction Lemma, we will first create an intermediate hybrid hyb_0^* . hyb_0^* simply outputs the view output by the online extractor, namely, $\text{REAL}_{\mathcal{E}, \Pi}^{I_{\text{extr}}^{(0)}}(\beta, z)$, where the external protocol Π here is the empty protocol. Note that this view is identical to the view of the MiM adversary \mathcal{M} when run by the real-world experiment, since $I_{\text{extr}}^{(0)}$ behaves the same way as I_{real} except that it also isolates certain CECom commitments and relays them to external CECom receivers. That is,

$$\text{REAL}_{\mathcal{E}, \Pi}^{I_{\text{real}}^{(0)}}(\beta, z) \equiv \text{view}_{\mathcal{M}}(x_1, \dots, x_{m_L}).$$

Now, we can invoke the Robust Extraction Lemma, for which we consider the following: the CECom commitments, that are isolated and relayed to external CECom receivers, have k_{cecom} -slots and the external protocol that the robust simulator is participating in, here, is the empty protocol. Thus, we have that $\ell_{\text{cecom}} = k_{\text{cecom}}$ and $\ell_{\text{external}} = 0$. Now by applying the Robust Extraction Lemma, we have that statistical distance between the views $\text{REAL}_{\mathcal{E}, \Pi}^{I_{\text{extr}}^{(0)}}(\beta, z)$ and $\text{view}_{\text{extr}}^{(0)}$ output by \mathcal{H}_0^* and hyb_0 , respectively, is at most:

$$\Delta(\lambda) \leq 2^{-(k_{\text{cecom}} - 0 \cdot \log(T(\lambda)))} \leq \text{negl}(\lambda),$$

since, $k_{\text{cecom}} \in \omega(\log(\lambda))$ and T is at most a polynomial. Thus,

$$\text{view}_{\text{extr}}^{(0)} \approx_s \text{REAL}_{\mathcal{E}, \Pi}^{I_{\text{real}}^{(1)}}(\beta, z).$$

Hence, the claim. □

Claim 10. For every right session, with all but negligible probability, srs is uniformly random.

Proof. Consider any t -th right session. We observe the following.

- r_V is statistically hidden in CECom_{sh} of Step cfp_1 .
- r_V is revealed at Step cfp_3 only after \mathcal{M} sends r_P (in Step cfp_2).
- r_V is uniformly random.

Hence, we have that $\text{srs} = r_P \oplus r_V$ (for any adversarially chosen r_P) is uniformly random with all but negligible probability. \square

\mathcal{H}_1 : \mathcal{H}_1 runs $\text{RobustSim}^{I_{\text{extr}}^{(1)}}(z)$, where we define $I_{\text{extr}}^{(1)}$ to be identical to $I_{\text{extr}}^{(0)}$, except that it biases the coin-flipping phase of every right session to a random DDH tuple by ‘cheating’ in the BPS argument at the BPS^{CFP} phase. It does so by proceeding as follows for every right session.

1. In Step cfp_1 , (like in $I_{\text{extr}}^{(0)}$), choose a random string and commit to it using CECom_{sh} .
2. In Step cfp_2 , let r_P be the value sent by \mathcal{M} .
3. In Step cfp_3 , sample a random DDH tuple srs . Define $r_V := r_P \oplus \text{srs}$. Send r_V .
4. Isolate the CECom commitment CECom_{sb} at Step $\text{bps}^{\text{cfp}}_1$ and forward it to an external CECom receiver.
5. Let value α' be received from outside at the end of the CECom_{sb} commitment. Then commit to α' using Com_{sh} at Step $\text{bps}^{\text{cfp}}_2$; also, use the same extracted value as the witness in proving the statistical ZKAoK of Step $\text{bps}^{\text{cfp}}_2$.
6. Let α be the value that \mathcal{M} opens its CECom_{sb} (of Step $\text{bps}^{\text{cfp}}_1$) to in Step $\text{bps}^{\text{cfp}}_3$. Abort if $\alpha \neq \alpha'$.
7. Commit to 0^λ using the mixed non-malleable commitment Com_{nm} in Step $\text{bps}^{\text{cfp}}_4$.
8. Use the value, α' , committed to in Step $\text{bps}^{\text{cfp}}_2$ as the witness in proving sZKAoK of Step $\text{bps}^{\text{cfp}}_5$.
9. Isolate CECom_{sh} commitment at Step bps_{4+} and forward it to an external CECom receiver. Let the value received from the outside be $y_t^{(1)}$.
10. Execute the rest of the steps of any right session honestly, like in $I_{\text{extr}}^{(0)}$.

For every left session, proceed exactly the same way as $I_{\text{extr}}^{(0)}$. Namely, run the code of honest prover (using valid witness) for every left session.

Besides outputting the view $\text{view}_{\text{extr}}^{(1)}$, \mathcal{H}_1 also outputs $y_1^{(1)}, \dots, y_{m_R}^{(1)}$.

Note that $\text{view}_{\text{extr}}^{(1)}$ is *not* statistically indistinguishable from $\text{view}_{\text{extr}}^{(0)}$, the reason being the following: While in $\text{view}_{\text{extr}}^{(0)}$, for every right session, srs is uniformly random with all but negligible probability, as proven in Claim 10, in $\text{view}_{\text{extr}}^{(1)}$, $I_{\text{extr}}^{(1)}$ biases the outcome of the coin-flippings of the right sessions to random DDH tuples. Thus, it maybe the case that \mathcal{M} is now also able to bias the coin-flippings of some left sessions. We would first like to prove that this is not the case; i.e., we shall prove that for every left session srs is still uniformly random with all but negligible probability.

Claim 11. *For every left session, with all but negligible probability, \mathbf{srs} is uniformly random.*

Proof. We begin by sketching the outline of the proof. Recall that until now our approach in proving uniform distribution of \mathbf{srs} in left sessions was to first construct a PPT algorithm that extracts r'_V from CECom_{sh} of Step cfp_1 and then argue that the value r_V given by \mathcal{M} is such that $r_V = r'_V$, with all but negligible probability. That is the statement that \mathcal{M} gives a BPS argument for in the BPS^{CFP} phase of the left session is for the same r'_V that was extracted, with all but negligible probability. Since r_P would then be sent only after the extraction of r'_V , we could argue that r_P being uniformly random implies $\mathbf{srs} = r_P \oplus r_V$ being uniformly random too, with all but negligible probability. However, we cannot naïvely proceed with these steps any more for the following reason. Note that $I_{\text{extr}}^{(2)}$ differs from $I_{\text{extr}}^{(1)}$ (at least) in the following two ways.

1. In every right session, while $I_{\text{extr}}^{(1)}$ sent the same r_V that it committed to in Step cfp_1 at Step cfp_3 , $I_{\text{extr}}^{(2)}$ cheats in cfp_3 by sampling a random DDH tuple \mathbf{srs} and by setting r_V to be $r_V := r_P \oplus \mathbf{srs}$.
2. In every right session, while $I_{\text{extr}}^{(1)}$ committed to a valid sub-witness rand – used as the randomness in the CECom_{sh} commitment of Step cfp_1 – in Com_{nm} in Step $\text{bps}^{\text{cfp}}_4$, $I_{\text{extr}}^{(2)}$ commits to 0^λ using Com_{nm} in Step $\text{bps}^{\text{cfp}}_4$.

Note that both the changes are *not statistically indistinguishable*, but only computationally so. Hence, it maybe the case that \mathcal{M} also is now able to prove a false statement in BPS^{CFP} phase of a left session thus biasing the distribution of \mathbf{srs} on the left far from uniform, although computationally indistinguishable from the latter. Our task thus would be to ensure that the computational changes introduced by $I_{\text{extr}}^{(2)}$ do not lead \mathbf{srs} on left sessions to be distributed statistically far from uniform.

The proof shall proceed by defining a sequence of ‘sub-hybrids’ from \mathcal{H}_0 to \mathcal{H}_1 . Consider the sequence of sub-hybrids, $\mathcal{H}_{0,0}, \mathcal{H}_{0,1}, \dots, \mathcal{H}_{0,m_R}$, and further ‘intermediate-hybrids’ $\mathcal{H}_{0,i:1}, \dots, \mathcal{H}_{0,i:7}$ that interpolate between $\mathcal{H}_{0,i}$ and $\mathcal{H}_{0,i+1}$, defined as follows. The sequence of the sub-hybrids, as we shall shortly describe, correspond to the following ordering of the right sessions: Consider any two right sessions, i -th and j -th; $i \leq j$ if and only if the CECom_{sb} commitment at Step bps_1 of the i -th session begins earlier to the CECom_{sb} commitment at Step bps_1 of the j -th session. The hybrids are defined as below.

Sub-hybrid $\mathcal{H}_{0,0}$. $\mathcal{H}_{0,0}$ is identical to \mathcal{H}_0 . Recall that \mathcal{H}_0 ran $\text{RobustSim}^{I_{\text{extr}}^{(0)}}(z)$. For notational purposes, we shall define $I_{\text{extr}}^{(0,0)}$ to be identical to $I_{\text{extr}}^{(0)}$.

Sub-Claim 4.

$$\text{view}_{\text{extr}}^{(0,0)} \equiv \text{view}_{\text{extr}}^{(0)}.$$

Proof. This follows immediately from the fact that $\mathcal{H}_{0,0}$ is identical to \mathcal{H}_0 . □

Intermediate-hybrid $\mathcal{H}_{0,i:1}$. $\mathcal{H}_{0,i:1}$ differs from $\mathcal{H}_{0,i}$ in that, while $\mathcal{H}_{0,i}$ ran $\text{RobustSim}^{I_{\text{extr}}^{(0,i)}}(z)$, $\mathcal{H}_{0,i:1}$ runs $\text{RobustSim}^{I_{\text{extr}}^{(0,i:1)}}(z)$, where $I_{\text{extr}}^{(0,i:1)}$ differs from $I_{\text{extr}}^{(0,i)}$ as follows.

- In the $i + 1$ -th right session, isolate the CECom_{sb} commitment at Step $\text{bps}^{\text{cfp}}_1$ and forward it to an external CECOM receiver.

Sub-Claim 5.

$$\text{view}_{\text{extr}}^{(0,i:1)} \approx_s \text{view}_{\text{extr}}^{(0,i)}.$$

Proof. Recall that $\mathcal{H}_{0,i}$ and $\mathcal{H}_{0,i:1}$ run $\text{RobustSim}^{I_{\text{extr}}^{(0,i)}}(z)$ and $\text{RobustSim}^{I_{\text{extr}}^{(0,i:1)}}(z)$, respectively, where $I_{\text{extr}}^{(0,i)}$ and $I_{\text{extr}}^{(0,i:1)}$ differ in that $I_{\text{extr}}^{(0,i:1)}$ isolates one more CECOM commitment than $I_{\text{extr}}^{(0,i)}$ and relays it to an external CECOM receiver.

Now to establish statistical indistinguishability between the views output by $\mathcal{H}_{0,i}$ and $\mathcal{H}_{0,i:1}$, we apply the Robust Extraction Lemma. Since $I_{\text{extr}}^{(0,i)}$ does not isolate certain CECOM commitments that are isolated by $I_{\text{extr}}^{(0,i:1)}$, to apply the Robust Extraction Lemma, we will first create an intermediate hybrid $\text{hyb}_{0,i:1}^*$. $\text{hyb}_{0,i:1}^*$ simply outputs the view output by the online extractor, namely, $\text{REAL}_{\mathcal{E},\Pi}^{I_{\text{extr}}^{(0,i:1)}}(\beta, z)$, where the external protocol Π here is the empty protocol. Note that this view is statistically close to the view of the MiM adversary \mathcal{M} when run by $\mathcal{H}_{0,i}$, since $I_{\text{extr}}^{(0,i:1)}$ behaves the same way as $I_{\text{extr}}^{(0,i)}$ except that it also isolates certain CECOM commitments that were originally not isolated by $I_{\text{extr}}^{(0,i)}$ and relays them to external CECOM receivers. That is,

$$\text{REAL}_{\mathcal{E},\Pi}^{I_{\text{extr}}^{(0,i:1)}}(\beta, z) \approx_s \text{view}_{\text{extr}}^{(0,i)}.$$

Now, we can invoke the Robust Extraction Lemma, for which we consider the following: the CECOM commitments, that are isolated and relayed to external CECOM receivers, have k_{cecom} -slots and the external protocol that the robust simulator is participating in, here, is the empty protocol. Thus, we have that $\ell_{\text{cecom}} = k_{\text{cecom}}$ and $\ell_{\text{external}} = 0$. Now by applying the Robust Extraction Lemma, we have that statistical distance between the views $\text{REAL}_{\mathcal{E},\Pi}^{I_{\text{extr}}^{(0,i:1)}}(\beta, z)$ and $\text{view}_{\text{extr}}^{(0,i)}$ output by $\mathcal{H}_{0,i:1}^*$ and $\mathcal{H}_{0,i:1}$, respectively, is at most:

$$\Delta(\lambda) \leq 2^{-(k_{\text{cecom}} - 0 \cdot \log(T(\lambda)))} \leq \text{negl}(\lambda),$$

since, $k_{\text{cecom}} \in \omega(\log(\lambda))$ and T is at most a polynomial. Thus,

$$\text{view}_{\text{extr}}^{(0,i:1)} \approx_s \text{REAL}_{\mathcal{E},\Pi}^{I_{\text{extr}}^{(0,i:1)}}(\beta, z).$$

□

Intermediate-hybrid $\mathcal{H}_{0,i:2}$. $\mathcal{H}_{0,i:2}$ differs from $\mathcal{H}_{0,i:1}$ in that, while $\mathcal{H}_{0,i:1}$ ran $\text{RobustSim}^{I_{\text{extr}}^{(0,i:1)}}(z)$, $\mathcal{H}_{0,i:2}$ runs $\text{RobustSim}^{I_{\text{extr}}^{(0,i:2)}}(z)$, where $I_{\text{extr}}^{(0,i:2)}$ differs from $I_{\text{extr}}^{(0,i:1)}$ as follows.

- Recall that $I_{\text{extr}}^{(0,i:1)}$ isolated CECom_{sb} commitment at Step $\text{bps}^{\text{cfp}}_1$ of the $i + 1$ -th right session.
- Let value α' be received from outside at the end of the CECom_{sb} commitment. Then commit to α' using Com_{sh} at Step $\text{bps}^{\text{cfp}}_2$; also, use the same extracted value as the witness in proving sWIAoK at Step $\text{bps}^{\text{cfp}}_2$.

Sub-Claim 6.

$$\text{view}_{\text{extr}}^{(0,i:2)} \approx_s \text{view}_{\text{extr}}^{(0,i:1)}.$$

Proof. Since Com_{sh} at Step $\text{bps}^{\text{cfp}}_2$ is a statistically hiding commitment scheme, and sWIAoK is statistically witness-indistinguishable, the sub-claim follows. \square

Intermediate-hybrid $\mathcal{H}_{0,i:3}$. $\mathcal{H}_{0,i:3}$ differs from $\mathcal{H}_{0,i:2}$ in that, while $\mathcal{H}_{0,i:2}$ ran $\text{RobustSim}^{I_{\text{extr}}^{(0,i:2)}}(z)$, $\mathcal{H}_{0,i:3}$ runs $\text{RobustSim}^{I_{\text{extr}}^{(0,i:3)}}(z)$, where $I_{\text{extr}}^{(0,i:3)}$ differs from $I_{\text{extr}}^{(0,i:2)}$ as follows.

- Recall that $I_{\text{extr}}^{(0,i:2)}$ isolated the CECom_{sb} commitment at Step $\text{bps}^{\text{cfp}}_1$ of the $i + 1$ -th right session. Let the extracted value received from the outside be α' . Also, in Step $\text{bps}^{\text{cfp}}_3$, let α be the value that \mathcal{M} opens the CECom_{sb} commitment to. If $\alpha \neq \alpha'$, then abort.

Sub-Claim 7.

$$\text{view}_{\text{extr}}^{(0,i:3)} \approx_s \text{view}_{\text{extr}}^{(0,i:2)}.$$

Proof. Recall that $\mathcal{H}_{0,i:3}$ isolates the CECom_{sb} commitment at Step $\text{bps}^{\text{cfp}}_1$ of the $i + 1$ -th right session (among other messages). It thus receives an opening for it, say $(\alpha', \text{rand}_{\alpha'})$ from the outside. Furthermore, \mathcal{M} provides an opening to the same CECom_{sb} commitment at Step $\text{bps}^{\text{cfp}}_3$; call it $(\alpha, \text{rand}_{\alpha})$. Since CECom_{sb} is SB, $\alpha = \alpha'$ with all but negligible probability, and hence the sub-claim follows. \square

Intermediate-hybrid $\mathcal{H}_{0,i:4}$. $\mathcal{H}_{0,i:4}$ differs from $\mathcal{H}_{0,i:3}$ in that, while $\mathcal{H}_{0,i:3}$ ran $\text{RobustSim}^{I_{\text{extr}}^{(0,i:3)}}(z)$, $\mathcal{H}_{0,i:4}$ runs $\text{RobustSim}^{I_{\text{extr}}^{(0,i:4)}}(z)$, where $I_{\text{extr}}^{(0,i:4)}$ differs from $I_{\text{extr}}^{(0,i:3)}$ as follows.

- Recall that, in the $i + 1$ -th session, $I_{\text{extr}}^{(0,i:3)}$ committed to the value rand – used as randomness in committing r_V at Step cfp_1 – in Com_{nm} at Step $\text{bps}^{\text{cfp}}_4$. Furthermore, it uses this committed value rand as the witness in proving sWIAoK at Step $\text{bps}^{\text{cfp}}_5$. Here, the modification is that $I_{\text{extr}}^{(0,i:4)}$ uses the commitment information from Com_{sh} at Step $\text{bps}^{\text{cfp}}_2$, where it committed to α , as the witness in proving sWIAoK at Step $\text{bps}^{\text{cfp}}_5$.

Sub-Claim 8.

$$\text{view}_{\text{extr}}^{(0,i:4)} \approx_s \text{view}_{\text{extr}}^{(0,i:3)}.$$

Proof. Since sWIAoK at Step $\text{bps}^{\text{cfp}}_5$ is statistically witness-indistinguishable, the sub-claim follows. \square

Intermediate-hybrid $\mathcal{H}_{0,i:5}$. $\mathcal{H}_{0,i:5}$ differs from $\mathcal{H}_{0,i:4}$ in that, while $\mathcal{H}_{0,i:4}$ ran $\text{RobustSim}^{I_{\text{extr}}^{(0,i:4)}}(z)$, $\mathcal{H}_{0,i:5}$ runs $\text{RobustSim}^{I_{\text{extr}}^{(0,i:5)}}(z)$, where $I_{\text{extr}}^{(0,i:5)}$ differs from $I_{\text{extr}}^{(0,i:4)}$ as follows.

- Recall that, in the $i + 1$ -th right session, $I_{\text{extr}}^{(0,i:4)}$ committed to the value rand – used as randomness in committing r_V at Step cfp_1 – in Com_{nm} at Step $\text{bps}^{\text{cfp}}_4$. Here, the modification is that $I_{\text{extr}}^{(0,i:5)}$ commits to 0^λ in Com_{nm} at Step $\text{bps}^{\text{cfp}}_4$.

Sub-Claim 9.

$$\text{view}_{\text{extr}}^{(0,i:5)} \approx_c \text{view}_{\text{extr}}^{(0,i:4)}.$$

Proof. Recall that the NMCCom commitment Com_{nm} at Step $\text{bps}^{\text{cfp}}_4$ is computational hiding. Since the only modification introduced in $\mathcal{H}_{0,i:5}$ is in the value committed to in the NMCCom commitment Com_{nm} at Step $\text{bps}^{\text{cfp}}_4$, the sub-claim follows. \square

Sub-Claim 10. *In $\mathcal{H}_{0,i:5}$, $\forall \ell \in [m_L]$, if \mathcal{P}_ℓ accepts the BPS argument of BPS^{CFP} phase, then the value committed to by \mathcal{M} in the Com_{nm} commitment at Step $\text{bps}^{\text{cfp}}_4$ of the ℓ -th left session is a valid sub-witness.*

Proof. Observe that the change introduced by $\mathcal{H}_{0,i:5}$ is only computationally indistinguishable, but not statistically so. Furthermore, the value committed to in the SB NMCCom commitment Com_{nm} at Step $\text{bps}^{\text{cfp}}_4$ of any left session is not revealed by \mathcal{M} in any part of the protocol. Hence, computational hiding of the Com_{nm} commitment itself would not suffice in arguing that the value committed to by the adversary \mathcal{M} in the Com_{nm} commitment at Step $\text{bps}^{\text{cfp}}_4$ of any left session does not change adversely. For this, we shall rely upon non-malleability of the Com_{nm} commitment.

Assume for contradiction that there exists $\ell \in [m_L]$ such that, with some non-negligible probability ϵ , \mathcal{P}_ℓ accepts the BPS argument of BPS^{CFP} phase, but the value committed to by \mathcal{M} in the Com_{nm} commitment at Step $\text{bps}^{\text{cfp}}_4$ of the ℓ -th left session is not a valid sub-witness. Then we construct an adversary $\mathcal{A}_{\text{nmcom}}$ against non-malleability of the NMCCom commitment that wins with probability $\epsilon - \text{negl}(\lambda)$.

$\mathcal{A}_{\text{nmcom}}$ behaves exactly the same way as $\mathcal{H}_{0,i:4}$ except that it also isolates the Com_{nm} commitment at Step $\text{bps}^{\text{cfp}}_4$ of the ℓ -th left session and Com_{nm} commitment at Step $\text{bps}^{\text{cfp}}_4$ of the $i + 1$ -th right session, and forwards them to an external NMCCom receiver and to an external NMCCom sender, respectively. Furthermore, $\mathcal{A}_{\text{nmcom}}$ sends the sub-witness rand (i.e., the value committed to in the isolated Com_{nm} commitment of the $i + 1$ -th right session by $\mathcal{H}_{0,i:4}$) and 0^λ (i.e., the value committed to in the isolated Com_{nm} commitment of the $i + 1$ -th right session by $\mathcal{H}_{0,i:5}$) to the external NMCCom sender, who chooses one of the values to commit to.

Before we proceed we shall first prove that the view of the MiM adversary \mathcal{M} during its interaction with $\mathcal{A}_{\text{nmcom}}$ in the case where the value committed to in the isolated Com_{nm} commitment of the $i + 1$ -th right session is rand is statistically close to the view output by $\mathcal{H}_{0,i:4}$. For this, we apply the Robust Extraction Lemma. Since $\mathcal{A}_{\text{nmcom}}$ also isolates two NMCCom commitments and relays them to external parties, to apply the Robust Extraction Lemma, we will first create two intermediate hybrids $\text{hyb}_{0,i:4,A}^*$ and $\text{hyb}_{0,i:4,B}^*$, whose outputs are identical, and the output of the former is statistically close to the view of the MiM adversary \mathcal{M} when run by $\mathcal{H}_{0,i:4}$ and the output of the latter is statistically close to the view of the MiM adversary \mathcal{M} when run by $\mathcal{A}_{\text{nmcom}}$ in the

case where the value committed to in the isolated Com_{nm} commitment of the $i + 1$ -th right session is rand .

$\text{hyb}_{0,i:4,A}^*$ is described as follows. It simply outputs the view output by the online extractor, namely, $\text{REAL}_{\mathcal{E},\Pi}^{I_{\text{extr}}^{(0,i:4)}}(\beta, z)$, where the external protocol Π here is the empty protocol. Now, we can invoke the Robust Extraction Lemma, for which we consider the following: the CECOM commitments, that are isolated and relayed to external CECOM receivers, have k_{cecom} -slots and the external protocol that the robust simulator is participating in, here, is the empty protocol. Thus, we have that $\ell_{\text{cecom}} = k_{\text{cecom}}$ and $\ell_{\text{external}} = 0$. Now by applying the Robust Extraction Lemma, we have that statistical distance between the views $\text{REAL}_{\mathcal{E},\Pi}^{I_{\text{extr}}^{(0,i:4)}}(\beta, z)$ and $\text{view}_{\text{extr}}^{(0,i:4)}$ is at most:

$$\Delta(\lambda) \leq 2^{-(k_{\text{cecom}} - 0 \cdot \log(T(\lambda)))} \leq \text{negl}(\lambda),$$

since, $k_{\text{cecom}} \in \omega(\log(\lambda))$ and T is at most a polynomial.

Next, we describe an intermediate hybrid $\text{hyb}_{0,i:4,B}^*$ whose output is identical to that of $\text{hyb}_{0,i:4,A}^*$. For this consider an interface, $I_{\text{extr}}^{(0,i:4,B)}$ behaves the same way as $I_{\text{extr}}^{(0,i:4)}$ except that it also isolates the Com_{nm} commitment at Step $\text{bps}^{\text{cfp}}_4$ of the ℓ -th left session and Com_{nm} commitment at Step $\text{bps}^{\text{cfp}}_4$ of the $i + 1$ -th right session, and forwards them to an external NMCOM receiver and to an external NMCOM sender, respectively. $\text{hyb}_{0,i:4,B}^*$ simply outputs the view output by the online extractor, namely, $\text{REAL}_{\mathcal{E},\Pi}^{I_{\text{extr}}^{(0,i:4,B)}}(\beta, z)$, where the external protocol Π here consists of the Com_{nm} commitment at Step $\text{bps}^{\text{cfp}}_4$ of the ℓ -th left session and Com_{nm} commitment at Step $\text{bps}^{\text{cfp}}_4$ of the $i + 1$ -th right session, and the external party running the codes of NMCOM receiver and NMCOM sender (committing to rand), respectively, for these isolated commitments. Again, we can invoke the Robust Extraction Lemma, for which we consider the following: the CECOM commitments, that are isolated and relayed to external CECOM receivers, have k_{cecom} -slots and the external protocol that the robust simulator is participating in, here, Π here consists of the Com_{nm} commitment at Step $\text{bps}^{\text{cfp}}_4$ of the ℓ -th left session and Com_{nm} commitment at Step $\text{bps}^{\text{cfp}}_4$ of the $i + 1$ -th right session, and the external party running the codes of NMCOM receiver and NMCOM sender (committing to rand), respectively, for these isolated commitments. Thus, we have that $\ell_{\text{cecom}} = k_{\text{cecom}}$ and $\ell_{\text{external}} = 2k_{\text{nmcom}}$. Now by applying the Robust Extraction Lemma, we have that statistical distance between the views $\text{REAL}_{\mathcal{E},\Pi}^{I_{\text{extr}}^{(0,i:4,B)}}(\beta, z)$ and the view of the \mathcal{M} during its interaction with $\mathcal{A}_{\text{nmcom}}$ in the case where the value committed to in the isolated Com_{nm} commitment of the $i + 1$ -th right session is 0^{rand} is at most:

$$\Delta(\lambda) \leq 2^{-(k_{\text{cecom}} - 2k_{\text{nmcom}} \cdot \log(T(\lambda)))} \leq \text{negl}(\lambda),$$

since, $k_{\text{cecom}} \in \omega(\log(\lambda))$, $k_{\text{nmcom}} \in O(\log(\lambda))$, and T is at most a polynomial.

Thus, we have proven that the view of the MiM adversary \mathcal{M} during its interaction with $\mathcal{A}_{\text{nmcom}}$ in the case where the value committed to in the isolated Com_{nm} commitment of the $i + 1$ -th right session is rand is statistically close to the view output by $\mathcal{H}_{0,i:4}$.

Also, similarly, we have that the view of the MiM adversary \mathcal{M} during its interaction with $\mathcal{A}_{\text{nmcom}}$ in the case where the value committed to in the isolated Com_{nm} commitment of the $i + 1$ -th right session is 0^λ is statistically close to the view output by $\mathcal{H}_{0,i:5}$.

Hence, by the construction of $\mathcal{A}_{\text{nmcom}}$, $\mathcal{H}_{0,i:4}$, and $\mathcal{H}_{0,i:5}$, we have that if the external NMCOM sender commits to the valid sub-witness rand in Com_{nm} , then the view of the adversary \mathcal{M} in its

interaction with \mathcal{A}_{nmcom} is a statistical simulation of the view $\text{view}_{\text{extr}}^{(0,i:4)}$; on the other hand, if the external NMCCom sender commits to 0^λ , then the view of the adversary \mathcal{M} in its interaction with \mathcal{A}_{nmcom} is a statistical simulation of the view $\text{view}_{\text{extr}}^{(0,i:5)}$. Putting this all together, from the assumption that, for $\mathcal{H}_{0,i:5}$, with some non-negligible probability ϵ , for the ℓ -th left session \mathcal{P}_ℓ accepts the BPS argument of BPS^{CFP} phase but the value committed to by \mathcal{M} in Com_{nm} commitment at Step $\text{bps}_4^{\text{CFP}}$ is not a valid sub-witness, and from induction that this is not the case in $\mathcal{H}_{0,i:4}$, we have that \mathcal{A}_{nmcom} breaks non-malleability of the NMCCom commitment with probability $\epsilon - \text{negl}(\lambda)$. \square

Sub-Claim 11. In $\mathcal{H}_{0,i:5}$, $\forall i, j \leq i$, $\beta_j^{(0,i:4)} \approx_c \beta_j^{(0,i:5)}$.

Proof. We shall see shortly that in $\mathcal{H}_{0,i:5}$, for every $j \leq i$ NMMXCom_{srs} commitment at Step bps_4 of the j -th right session is statistically binding. In particular, this will be clear in the hybrid $\mathcal{H}_{0,i:7}$. We now argue that changing the value committed in Com_{nm} at Step $\text{bps}_4^{\text{CFP}}$ of the $i+1$ -th right session does not adversely change the value committed to by the adversary in NMMXCom_{srs} commitment at Step bps_4 of the j -th right session.

Observe that the change introduced by $\mathcal{H}_{0,i:5}$ is only computationally indistinguishable, but not statistically so. Furthermore, the value committed to in the SB NMCCom commitment NMMXCom_{srs} at Step bps_4 of any right session is not revealed by \mathcal{M} in any part of the protocol. Hence, computational hiding of the NMMXCom_{srs} commitment itself would not suffice in arguing that the value committed to by the adversary \mathcal{M} in the NMMXCom_{srs} commitment at Step bps_4 of any left session does not change adversely. For this, we shall rely upon non-malleability of the NMMXCom_{srs} commitment w.r.t. Com_{nm} commitment (as per Definition 16).

Assume for contradiction that there exists $i \in [m_R], j \in [1, i]$ such that, with some non-negligible probability ϵ , the values committed to in NMMXCom_{srs} at Step bps_4 of the j -th right session in the hybrids $\mathcal{H}_{0,i:4}$ and $\mathcal{H}_{0,i:5}$ can be distinguished with probability ϵ ; that is, $\beta_j^{(0,i:4)}$ and $\beta_j^{(0,i:5)}$ can be distinguished with probability ϵ . Then we construct an adversary \mathcal{A}_{nmcom} against non-malleability of NMMXCom_{srs} w.r.t. Com_{nm} that wins with probability $\epsilon - \text{negl}(\lambda)$.

\mathcal{A}_{nmcom} behaves exactly the same way as $\mathcal{H}_{0,i:4}$ except that it also isolates the NMMXCom_{srs} commitment at Step bps_4 of the $i+1$ -th right session and Com_{nm} commitment at Step $\text{bps}_4^{\text{CFP}}$ of the $i+1$ -th right session, and forwards them to an external NMCCom receiver and to an external NMCCom sender, respectively. Furthermore, \mathcal{A}_{nmcom} sends the sub-witness rand (i.e., the value committed to in the isolated Com_{nm} commitment of the $i+1$ -th right session by $\mathcal{H}_{0,i:4}$) and 0^λ (i.e., the value committed to in the isolated Com_{nm} commitment of the $i+1$ -th right session by $\mathcal{H}_{0,i:5}$) to the external NMCCom sender, who chooses one of the values to commit to.

Before we proceed we shall first prove that the view of the MiM adversary \mathcal{M} during its interaction with \mathcal{A}_{nmcom} in the case where the value committed to in the isolated Com_{nm} commitment of the $i+1$ -th right session is rand is statistically close to the view output by $\mathcal{H}_{0,i:4}$. For this, we apply the Robust Extraction Lemma. Since \mathcal{A}_{nmcom} also isolates two NMCCom commitments and relays them to external parties, to apply the Robust Extraction Lemma, we will first create two intermediate hybrids $\text{hyb}_{0,i:4,A}^*$ and $\text{hyb}_{0,i:4,B}^*$, whose outputs are identical, and the output of the former is statistically close to the view of the MiM adversary \mathcal{M} when run by $\mathcal{H}_{0,i:4}$ and the output of the latter is statistically close to the view of the MiM adversary \mathcal{M} when run by \mathcal{A}_{nmcom} in the case where the value committed to in the isolated Com_{nm} commitment of the $i+1$ -th right session is rand .

$\text{hyb}_{0,i:4,A}^*$ is described as follows. It simply outputs the view output by the online extractor,

namely, $\text{REAL}_{\mathcal{E}, \Pi}^{I_{\text{extr}}^{(0,i:4)}}(\beta, z)$, where the external protocol Π here is the empty protocol. Now, we can invoke the Robust Extraction Lemma, for which we consider the following: the CECOM commitments, that are isolated and relayed to external CECOM receivers, have k_{cecom} -slots and the external protocol that the robust simulator is participating in, here, is the empty protocol. Thus, we have that $\ell_{\text{cecom}} = k_{\text{cecom}}$ and $\ell_{\text{external}} = 0$. Now by applying the Robust Extraction Lemma, we have that statistical distance between the views $\text{REAL}_{\mathcal{E}, \Pi}^{I_{\text{extr}}^{(0,i:4)}}(\beta, z)$ and $\text{view}_{\text{extr}}^{(0,i:4)}$ is at most:

$$\Delta(\lambda) \leq 2^{-(k_{\text{cecom}} - 0 \cdot \log(T(\lambda)))} \leq \text{negl}(\lambda),$$

since, $k_{\text{cecom}} \in \omega(\log(\lambda))$ and T is at most a polynomial.

Next, we describe an intermediate hybrid $\text{hyb}_{0,i:4,B}^*$ whose output is identical to that of $\text{hyb}_{0,i:4,A}^*$. For this consider an interface, $I_{\text{extr}}^{(0,i:4,B)}$ behaves the same way as $I_{\text{extr}}^{(0,i:4)}$ except that it also isolates the $\text{NMMXCom}_{\text{srs}}$ commitment at Step bps_4 of the j -th right session and Com_{nm} commitment at Step $\text{bps}_{\text{cfp}_4}$ of the $i+1$ -th right session, and forwards them to an external NMCom receiver and to an external NMCom sender, respectively. $\text{hyb}_{0,i:4,B}^*$ simply outputs the view output by the online extractor, namely, $\text{REAL}_{\mathcal{E}, \Pi}^{I_{\text{extr}}^{(0,i:4,B)}}(\beta, z)$, where the external protocol Π here consists of the $\text{NMMXCom}_{\text{srs}}$ commitment at Step bps_4 of the j -th right session and Com_{nm} commitment at Step $\text{bps}_{\text{cfp}_4}$ of the $i+1$ -th right session, and the external party running the codes of NMCom receiver and NMCom sender (committing to rand), respectively, for these isolated commitments. Again, we can invoke the Robust Extraction Lemma, for which we consider the following: the CECOM commitments, that are isolated and relayed to external CECOM receivers, have k_{cecom} -slots and the external protocol that the robust simulator is participating in, here, Π here consists of the two isolated NMCom commitments, and the external party running the codes of NMCom receiver and NMCom sender (committing to rand), respectively, for these isolated commitments. Thus, we have that $\ell_{\text{cecom}} = k_{\text{cecom}}$ and $\ell_{\text{external}} = k_{\text{nmcom}} + k_{\text{nmmxcom}}$. Now by applying the Robust Extraction Lemma, we have that statistical distance between the views $\text{REAL}_{\mathcal{E}, \Pi}^{I_{\text{extr}}^{(0,i:4,B)}}(\beta, z)$ and the view of the \mathcal{M} during its interaction with $\mathcal{A}_{\text{nmcom}}$ in the case where the value committed to in the isolated Com_{nm} commitment of the $i+1$ -th right session is 0^{rand} is at most:

$$\Delta(\lambda) \leq 2^{-(k_{\text{cecom}} - (k_{\text{nmcom}} + k_{\text{nmmxcom}}) \cdot \log(T(\lambda)))} \leq \text{negl}(\lambda),$$

since, $k_{\text{cecom}} \in \omega(\log(\lambda))$, $k_{\text{nmcom}}, k_{\text{nmmxcom}} \in O(\log(\lambda))$, and T is at most a polynomial.

Thus, we have proven that the view of the MiM adversary \mathcal{M} during its interaction with $\mathcal{A}_{\text{nmcom}}$ in the case where the value committed to in the isolated Com_{nm} commitment of the $i+1$ -th right session is rand is statistically close to the view output by $\mathcal{H}_{0,i:4}$.

Also, similarly, we have that the view of the MiM adversary \mathcal{M} during its interaction with $\mathcal{A}_{\text{nmcom}}$ in the case where the value committed to in the isolated Com_{nm} commitment of the $i+1$ -th right session is 0^λ is statistically close to the view output by $\mathcal{H}_{0,i:5}$.

Hence, by the construction of $\mathcal{A}_{\text{nmcom}}$, $\mathcal{H}_{0,i:4}$, and $\mathcal{H}_{0,i:5}$, we have that if the external NMCom sender commits to the valid sub-witness rand in Com_{nm} , then the view of the adversary \mathcal{M} in its interaction with $\mathcal{A}_{\text{nmcom}}$ is a statistical simulation of the view $\text{view}_{\text{extr}}^{(0,i:4)}$; on the other hand, if the external NMCom sender commits to 0^λ , then the view of the adversary \mathcal{M} in its interaction with $\mathcal{A}_{\text{nmcom}}$ is a statistical simulation of the view $\text{view}_{\text{extr}}^{(0,i:5)}$. Putting this all together, from the assumption that, the values committed to in $\text{NMMXCom}_{\text{srs}}$ at Step bps_4 of the j -th right session in the hybrids $\mathcal{H}_{0,i:4}$ and $\mathcal{H}_{0,i:5}$ can be distinguished with probability ϵ , we have that $\mathcal{A}_{\text{nmcom}}$ breaks

non-malleability of the $\text{NMMXCom}_{\text{srs}}$ commitment w.r.t. Com_{nm} commitment with probability $\epsilon - \text{negl}(\lambda)$. \square

Intermediate-hybrid $\mathcal{H}_{0,i:6}$. $\mathcal{H}_{0,i:6}$ differs from $\mathcal{H}_{0,i:5}$ in that, while $\mathcal{H}_{0,i:5}$ ran $\text{RobustSim}^{I_{\text{extr}}^{(0,i:5)}}(z)$, $\mathcal{H}_{0,i:6}$ runs $\text{RobustSim}^{I_{\text{extr}}^{(0,i:6)}}(z)$, where $I_{\text{extr}}^{(0,i:6)}$ differs from $I_{\text{extr}}^{(0,i:5)}$ as follows.

- Isolate CECom_{sh} commitment at Step bps_{4+} of the $i+1$ -th right session and forward it to an external CECom receiver. Let the value received from the outside be $y_{i+1}^{(0,i:6)}$.

Besides outputting the view $\text{view}_{\text{extr}}^{(0,i:6)}$, $\mathcal{H}_{0,i:6}$ also outputs $y_1^{(0,i:6)}, \dots, y_{i+1}^{(0,i:6)}$.

Sub-Claim 12.

$$\text{view}_{\text{extr}}^{(0,i:6)} \approx_s \text{view}_{\text{extr}}^{(0,i:5)}.$$

Proof. Recall that $\mathcal{H}_{0,i:5}$ and $\mathcal{H}_{0,i:6}$ run $\text{RobustSim}^{I_{\text{extr}}^{(0,i:5)}}(z)$ and $\text{RobustSim}^{I_{\text{extr}}^{(0,i:6)}}(z)$, respectively, where $I_{\text{extr}}^{(0,i:5)}$ and $I_{\text{extr}}^{(0,i:6)}$ differ in that $I_{\text{extr}}^{(0,i:6)}$ isolates one more CECom commitment than $I_{\text{extr}}^{(0,i:5)}$ and relays it to an external CECom receiver.

Now to establish statistical indistinguishability between the views output by $\mathcal{H}_{0,i:5}$ and $\mathcal{H}_{0,i:6}$, we apply the Robust Extraction Lemma. Since $I_{\text{extr}}^{(0,i:5)}$ does not isolate certain CECom commitments that are isolated by $I_{\text{extr}}^{(0,i:6)}$, to apply the Robust Extraction Lemma, we will first create an intermediate hybrid $\text{hyb}_{0,i:6}^*$. $\text{hyb}_{0,i:6}^*$ simply outputs the view output by the online extractor, namely, $\text{REAL}_{\mathcal{E},\Pi}^{I_{\text{extr}}^{(0,i:6)}}(\beta, z)$, where the external protocol Π here is the empty protocol. Note that this view is statistically close to the view of the MiM adversary \mathcal{M} when run by $\mathcal{H}_{0,i:5}$, since $I_{\text{extr}}^{(0,i:6)}$ behaves the same way as $I_{\text{extr}}^{(0,i:5)}$ except that it also isolates certain CECom commitments that were originally not isolated by $I_{\text{extr}}^{(0,i:5)}$ and relays them to external CECom receivers. That is,

$$\text{REAL}_{\mathcal{E},\Pi}^{I_{\text{extr}}^{(0,i:6)}}(\beta, z) \approx_s \text{view}_{\text{extr}}^{(0,i:5)}.$$

Now, we can invoke the Robust Extraction Lemma, for which we consider the following: the CECom commitments, that are isolated and relayed to external CECom receivers, have k_{cecom} -slots and the external protocol that the robust simulator is participating in, here, is the empty protocol. Thus, we have that $\ell_{\text{cecom}} = k_{\text{cecom}}$ and $\ell_{\text{external}} = 0$. Now by applying the Robust Extraction Lemma, we have that statistical distance between the views $\text{REAL}_{\mathcal{E},\Pi}^{I_{\text{extr}}^{(0,i:6)}}(\beta, z)$ and $\text{view}_{\text{extr}}^{(0,i:6)}$ output by $\mathcal{H}_{0,i:6}^*$ and $\mathcal{H}_{0,i:6}$, respectively, is at most:

$$\Delta(\lambda) \leq 2^{-(k_{\text{cecom}} - 0 \cdot \log(T(\lambda)))} \leq \text{negl}(\lambda),$$

since, $k_{\text{cecom}} \in \omega(\log(\lambda))$ and T is at most a polynomial. Thus,

$$\text{view}_{\text{extr}}^{(0,i:6)} \approx_s \text{REAL}_{\mathcal{E},\Pi}^{I_{\text{extr}}^{(0,i:6)}}(\beta, z).$$

Thus, we have that $\text{view}_{\text{extr}}^{(0,i:6)} \approx_s \text{view}_{\text{extr}}^{(0,i:5)}$. \square

Intermediate-hybrid $\mathcal{H}_{0,i:7}$. $\mathcal{H}_{0,i:7}$ differs from $\mathcal{H}_{0,i:6}$ in that, while $\mathcal{H}_{0,i:6}$ ran $\text{RobustSim}^{I_{\text{extr}}^{(0,i:6)}}(z)$, $\mathcal{H}_{0,i:7}$ runs $\text{RobustSim}^{I_{\text{extr}}^{(0,i:7)}}(z)$, where $I_{\text{extr}}^{(0,i:7)}$ differs from $I_{\text{extr}}^{(0,i:6)}$ as follows.

- Recall that in Step cfp_3 of the $i + 1$ -th right session, $I_{\text{extr}}^{(0,i:6)}$ gives the value r_V that was committed to in CECom_{sh} of Step cfp_1 of the same session. The modification now is that $I_{\text{extr}}^{(0,i:7)}$ samples a random DDH tuple srs , defines $r_V := r_P \oplus \text{srs}$, and sends r_V .
- $I_{\text{extr}}^{(0,i:7)}$ isolates CECom_{sh} commitment at Step bps_{4+} of the $i + 1$ -th right session and forwards it to an external CECom receiver. Let the value received from the outside be $y'_{i+1}{}^{(0,i:7)}$.

Besides outputting the view $\text{view}_{\text{extr}}^{(0,i:7)}$, $\mathcal{H}_{0,i:7}$ also outputs $y'_1{}^{(0,i:7)}, \dots, y'_{i+1}{}^{(0,i:7)}$.

Sub-Claim 13.

$$\text{view}_{\text{extr}}^{(0,i:7)} \approx_c \text{view}_{\text{extr}}^{(0,i:6)}.$$

Proof. This immediately follows from the DDH assumption. □

Sub-Claim 14. *In $\mathcal{H}_{0,i:7}$, $\forall \ell \in [m_L]$, if \mathcal{P}_ℓ accepts the BPS argument of BPS^{CFP} phase, then the value committed to by \mathcal{M} in Com_{nm} commitment at Step $\text{bps}^{\text{CFP}}_4$ of the ℓ -th left session is a valid sub-witness.*

Proof. Observe that the change introduced by $\mathcal{H}_{0,i:7}$ is only computationally indistinguishable, but not statistically so. Furthermore, as mentioned earlier, the value committed to in the SB NMCom commitment Com_{nm} at Step $\text{bps}^{\text{CFP}}_4$ of any left session is not revealed by \mathcal{M} in any part of the protocol. Hence, computational indistinguishability of DDH tuples from uniform strings is insufficient to argue that the value committed to by the adversary \mathcal{M} in the Com_{nm} commitment at Step $\text{bps}^{\text{CFP}}_4$ of any left session does not change adversely. For this, we shall rely upon robust non-malleability of the Com_{nm} commitment scheme w.r.t. 1-round protocols (see Definition 15).

Assume for contradiction that there exists $\ell \in [m_L]$ such that \mathcal{P}_ℓ accepts the BPS argument of BPS^{CFP} phase, but the value committed to by \mathcal{M} in Com_{nm} commitment at Step $\text{bps}^{\text{CFP}}_4$ of the ℓ -th left session is not a valid sub-witness, with some non-negligible probability ϵ . Then, by relying on the DDH assumption, we construct an adversary $\mathcal{A}_{\text{rob-nmcom}}$ against robust non-malleability of the NMCom commitment scheme w.r.t. 1-round protocols that succeeds with probability $\epsilon - \text{negl}(\lambda)$.

$\mathcal{A}_{\text{rob-nmcom}}$ behaves exactly the same way as $\mathcal{H}_{0,i:6}$ except that it also isolates the Com_{nm} commitment at Step $\text{bps}^{\text{CFP}}_4$ of the ℓ -th left session and the Step cfp_3 message of the $i + 1$ -th right session, and forwards them to an external NMCom receiver and to an external DDH challenger, respectively. The DDH challenger chooses srs which is either a random DDH tuple or a uniform random string (and sets r_V accordingly).

Before we proceed to analyze the success probability of $\mathcal{A}_{\text{rob-nmcom}}$, we need to show that the view of the MiM adversary \mathcal{M} during its interaction with $\mathcal{A}_{\text{rob-nmcom}}$ in the case where the string received from the external DDH challenger is a uniform random string (in the isolated message of Step cfp_3 of the $i + 1$ -th right session) is statistically close to the view output by $\mathcal{H}_{0,i:6}$. For this, we apply the Robust Extraction Lemma. Since $\mathcal{A}_{\text{rob-nmcom}}$ also isolates the Com_{nm} commitment at Step $\text{bps}^{\text{CFP}}_4$ of the ℓ -th left session and the Step cfp_3 message of the $i + 1$ -th right session, and relays them to external parties, to apply the Robust Extraction Lemma, we will first create two

intermediate hybrids $\text{hyb}_{0,i:6,A}^*$ and $\text{hyb}_{0,i:6,B}^*$, whose outputs are identical, and the output of the former is statistically close to the view of the MiM adversary \mathcal{M} when run by $\mathcal{H}_{0,i:6}$ and the output of the latter is statistically close to the view of the MiM adversary \mathcal{M} when run by $\mathcal{A}_{rob-nmcom}$ in the case where the string received from the external DDH challenger is a uniform random string.

$\text{hyb}_{0,i:6,A}^*$ is described as follows. It simply outputs the view output by the online extractor, namely, $\text{REAL}_{\mathcal{E},\Pi}^{I_{\text{extr}}^{(0,i:6)}}(\beta, z)$, where the external protocol Π here is the empty protocol. Now, we can invoke the Robust Extraction Lemma, for which we consider the following: the CECOM commitments, that are isolated and relayed to external CECOM receivers, have k_{cecom} -slots and the external protocol that the robust simulator is participating in, here, is the empty protocol. Thus, we have that $\ell_{\text{cecom}} = k_{\text{cecom}}$ and $\ell_{\text{external}} = 0$. Now by applying the Robust Extraction Lemma, we have that statistical distance between the views $\text{REAL}_{\mathcal{E},\Pi}^{I_{\text{extr}}^{(0,i:6)}}(\beta, z)$ and $\text{view}_{\text{extr}}^{(0,i:6)}$ is at most:

$$\Delta(\lambda) \leq 2^{-(k_{\text{cecom}} - 0 \cdot \log(T(\lambda)))} \leq \text{negl}(\lambda),$$

since, $k_{\text{cecom}} \in \omega(\log(\lambda))$ and T is at most a polynomial.

Next, we describe an intermediate hybrid $\text{hyb}_{0,i:6,B}^*$ whose output is identical to that of $\text{hyb}_{0,i:6,A}^*$. For this consider an interface, $I_{\text{extr}}^{(0,i:6,B)}$ behaves the same way as $I_{\text{extr}}^{(0,i:6)}$ except that it also isolates the Com_{nm} commitment at Step $\text{bps}^{\text{cfp}}_4$ of the ℓ -th left session and Com_{nm} commitment at Step $\text{bps}^{\text{cfp}}_4$ of the $i+1$ -th right session, and forwards them to an external NMCom receiver and to an external NMCom sender, respectively. $\text{hyb}_{0,i:6,B}^*$ simply outputs the view output by the online extractor, namely, $\text{REAL}_{\mathcal{E},\Pi}^{I_{\text{extr}}^{(0,i:6,B)}}(\beta, z)$, where the external protocol Π here consists of the Com_{nm} commitment at Step $\text{bps}^{\text{cfp}}_4$ of the ℓ -th left session and the Step cfp_3 message of the $i+1$ -th right session, and the role of the external party is to run the code of the NMCom receiver and the DDH challenger, respectively, for the isolated sub-protocols. Again, we can invoke the Robust Extraction Lemma, for which we consider the following: the CECOM commitments, that are isolated and relayed to external CECOM receivers, have k_{cecom} -slots and the external protocol that the robust simulator is participating in, here, Π here consists of the consists of two sub-protocols: namely, one Com_{nm} commitment of k_{nmcom} rounds and one 1-round protocol. Thus, we have that $\ell_{\text{cecom}} = k_{\text{cecom}}$ and $\ell_{\text{external}} = k_{\text{nmcom}} + 1$. Now by applying the Robust Extraction Lemma, we have that statistical distance between the views $\text{REAL}_{\mathcal{E},\Pi}^{I_{\text{extr}}^{(0,i:6,B)}}(\beta, z)$ and the view of the \mathcal{M} during its interaction with $\mathcal{A}_{rob-nmcom}$ in the case where the string received from the external DDH challenger is a uniform random string is at most:

$$\Delta(\lambda) \leq 2^{-(k_{\text{cecom}} - (k_{\text{nmcom}} + 1) \cdot \log(T(\lambda)))} \leq \text{negl}(\lambda),$$

since, $k_{\text{cecom}} \in \omega(\log(\lambda))$, $k_{\text{nmcom}} \in O(\log(\lambda))$, and T is at most a polynomial.

Thus, we have proven that the view of the MiM adversary \mathcal{M} during its interaction with $\mathcal{A}_{rob-nmcom}$ in the case where the string received from the external DDH challenger is a uniform random string is statistically close to the view output by $\mathcal{H}_{0,i:6}$.

Also, similarly, we have that the view of the MiM adversary \mathcal{M} during its interaction with $\mathcal{A}_{rob-nmcom}$ in the case where the string received from the external DDH challenger is a random DDH tuple is statistically close to the view output by $\mathcal{H}_{0,i:7}$.

We have thus proven that if the external DDH challenger sends a uniform random string, then the view of the adversary \mathcal{M} in its interaction with $\mathcal{A}_{rob-nmcom}$ is a statistical simulation of the

view $\text{view}_{\text{extr}}^{(0,i:6)}$; on the other hand, if the external DDH challenger sends a random DDH tuple, then the view of the adversary \mathcal{M} in its interaction with $\mathcal{A}_{\text{rob-nmcom}}$ is a statistical simulation of the view $\text{view}_{\text{extr}}^{(0,i:7)}$. Thus, from the assumption that, for $\mathcal{H}_{0,i:7}$, for the ℓ -th left session, \mathcal{P}_ℓ accepts the BPS argument of BPS^{CFP} phase, but the value committed to by \mathcal{M} in Com_{nm} commitment at Step $\text{bps}^{\text{cfp}}_4$ is not a valid sub-witness, with some non-negligible probability ϵ , and from induction that this is not the case in $\mathcal{H}_{0,i:6}$, by relying on the DDH assumption, we have that $\mathcal{A}_{\text{rob-nmcom}}$ breaks robust non-malleability of the NMCom commitment with probability $\epsilon - \text{negl}(\lambda)$. \square

Note that the last intermediate hybrid $\mathcal{H}_{0,mR:7}$ is identical to \mathcal{H}_1 . Hence, if for any left session, if the BPS^{CFP} phase is accepted by the prover, then the value committed to by \mathcal{M} in Com_{nm} commitment at Step $\text{bps}^{\text{cfp}}_4$ of every left session is a valid sub-witness rand ; i.e., the committed value is rand such that (r_V, rand) form a valid opening to the CECom_{sh} commitment at Step cfp_1 of that session.

With this, we can argue that for every left session, $r'_V = r_V$, with all but negligible probability. Intuitively, this *now* follows from computational binding of CECom_{sh} commitment at Step cfp_1 , since we have ensured that the value committed to in Com_{nm} commitment at Step $\text{bps}^{\text{cfp}}_4$ of every left session is a valid sub-witness rand , where (r_V, rand) forms a valid opening to the CECom_{sh} commitment. To formalize this argument, assume for contradiction that for some ℓ -th left session $r'_V \neq r_V$ with some non-negligible probability ϵ . Then we can construct an adversary \mathcal{A}_{CB} against computational binding of CECom_{sh} commitment at Step cfp_1 such that \mathcal{A}_{CB} succeeds with probability $\epsilon - \text{negl}(\lambda)$ as follows.

\mathcal{A}_{CB} behaves the same way as \mathcal{H}_1 except for a few modifications. Recall that \mathcal{H}_1 ran $\text{RobustSim}^{I_{\text{extr}}^{(1)}}(z)$, where the only sub-protocols that $I_{\text{extr}}^{(1)}$ isolated are certain CECom commitments. \mathcal{A}_{CB} runs $\text{RobustSim}^{I_{\text{extr}}^{(1*)}}(z)$ where $I_{\text{extr}}^{(1*)}$ differs from $I_{\text{extr}}^{(1)}$ in the following sense.

- $I_{\text{extr}}^{(1*)}$ (besides the CECom commitments isolated by $I_{\text{extr}}^{(1)}$) also isolates CECom_{sh} commitment at Step cfp_1 of the ℓ -th left session and forwards it to an external CECom receiver.
- $I_{\text{extr}}^{(1*)}$ also isolates the NMCom commitment Com_{nm} commitment at Step $\text{bps}^{\text{cfp}}_4$ of the ℓ -th left session.

\mathcal{A}_{CB} itself runs the code of the honest NMCom receiver for the isolated NMCom commitment. Furthermore, \mathcal{A}_{CB} also runs the NMCom extractor on $\text{RobustSim}^{I_{\text{extr}}^{(1*)}}(z)$ for the isolated NMCom commitment.

Now recall that we had earlier proven that the value committed to in Com_{nm} commitment at Step $\text{bps}^{\text{cfp}}_4$ of any left session whose BPS^{CFP} phase is accepted by the prover is a valid sub-witness rand (i.e., a randomness that explains the CECom_{sh} commitment at Step cfp_1 of that session to r_V that is given by \mathcal{M} at Step cfp_3). Hence, \mathcal{A}_{CB} has two openings of the CECom commitment in question, one being (r'_V, rand') obtained from the robust simulator, and the other being (r_V, rand) , where r_V is the value given by \mathcal{M} at Step cfp_3 in the corresponding session, and rand is an opening of the same CECom commitment to r_V obtained via the NMCom extractor. Hence, if $r'_V \neq r_V$, then \mathcal{A}_{CB} has two openings of the same CECom commitment (that is isolated and forwarded to an external CECom receiver) to two distinct values.

Now it remains to show that the view of the adversary \mathcal{M} when run by \mathcal{A}_{CB} is statistically close to the view $\text{view}_{\text{extr}}^{(1)}$ output by \mathcal{H}_1 . Note that the only difference is that \mathcal{A}_{CB} additionally isolates

a CECOM commitment and a NMCOM commitment. Like before, this indistinguishability can again be proven by applying the Robust Extraction Lemma. Since \mathcal{A}_{CB} additionally also isolates a CECOM commitment and a NMCOM commitment, to apply the Robust Extraction Lemma, we will first create two intermediate hybrids $\text{hyb}_{0,A}^*$ and $\text{hyb}_{0,B}^*$, whose outputs are identical, and the output of the former is statistically close to the view of the MiM adversary \mathcal{M} when run by \mathcal{H}_1 and the output of the latter is statistically close to the view of the MiM adversary \mathcal{M} when run by \mathcal{A}_{CB} .

$\text{hyb}_{0,A}^*$ is described as follows. It simply outputs the view output by the online extractor, namely, $\text{REAL}_{\mathcal{E},\Pi}^{I_{\text{extr}}^{(1)}}(\beta, z)$, where the external protocol Π here is the empty protocol. Now, we can invoke the Robust Extraction Lemma, for which we consider the following: the CECOM commitments, that are isolated and relayed to external CECOM receivers, have k_{cecom} -slots and the external protocol that the robust simulator is participating in, here, is the empty protocol. Thus, we have that $\ell_{\text{cecom}} = k_{\text{cecom}}$ and $\ell_{\text{external}} = 0$. Now by applying the Robust Extraction Lemma, we have that statistical distance between the views $\text{REAL}_{\mathcal{E},\Pi}^{I_{\text{extr}}^{(0)}}(\beta, z)$ and $\text{view}_{\text{extr}}^{(0)}$ is at most:

$$\Delta(\lambda) \leq 2^{-(k_{\text{cecom}} - 0 \cdot \log(T(\lambda)))} \leq \text{negl}(\lambda),$$

since, $k_{\text{cecom}} \in \omega(\log(\lambda))$ and T is at most a polynomial.

Next, we describe an intermediate hybrid $\text{hyb}_{0,B}^*$ whose output is identical to that of $\text{hyb}_{0,A}^*$. For this consider an interface, $I_{\text{extr}}^{(0,B)}$ behaves the same way as $I_{\text{extr}}^{(0)}$ except that it also isolates also isolates CECOM_{sh} commitment at Step cfp_1 of the ℓ -th left session and forwards them to an external CECOM receiver, and the NMCOM commitment Com_{nm} commitment at Step $\text{bps}^{\text{cfp}}_4$ of the ℓ -th left session and forwards it to an external NMCOM receiver. $\text{hyb}_{0,B}^*$ simply outputs the view output by the online extractor, namely, $\text{REAL}_{\mathcal{E},\Pi}^{I_{\text{extr}}^{(0,B)}}(\beta, z)$, where the external protocol Π here is the NMCOM commitment Com_{nm} commitment at Step $\text{bps}^{\text{cfp}}_4$ of the ℓ -th left session, and the role of the external party is to run the code of the NMCOM receiver. Again, we can invoke the Robust Extraction Lemma, for which we consider the following: the CECOM commitments, that are isolated and relayed to external CECOM receivers, have k_{cecom} -slots and the external protocol that the robust simulator is participating in, here, Π here is an Com_{nm} commitment of k_{nmcom} rounds. Thus, we have that $\ell_{\text{cecom}} = k_{\text{cecom}}$ and $\ell_{\text{external}} = k_{\text{nmcom}}$. Now by applying the Robust Extraction Lemma, we have that statistical distance between the views $\text{REAL}_{\mathcal{E},\Pi}^{I_{\text{extr}}^{(0,B)}}(\beta, z)$ and the view of the \mathcal{M} during its interaction with \mathcal{A}_{CB} is at most:

$$\Delta(\lambda) \leq 2^{-(k_{\text{cecom}} - k_{\text{nmcom}} \cdot \log(T(\lambda)))} \leq \text{negl}(\lambda),$$

since, $k_{\text{cecom}} \in \omega(\log(\lambda))$, $k_{\text{nmcom}} \in O(\log(\lambda))$, and T is at most a polynomial.

Thus, we have that the adversary \mathcal{M} when run by \mathcal{A}_{CB} is statistically close to the view $\text{view}_{\text{extr}}^{(1)}$ output by \mathcal{H}_1 .

Thus, \mathcal{A}_{CB} breaks computational binding of the CECOM_{sh} commitment with probability $\epsilon - \text{negl}(\lambda)$.

Finally, we thus have that, $\text{srs} = r_P \oplus r_V$ is uniformly random, with all but negligible probability. Thus, Claim 11. \square

Claim 12.

$$(b_t^{(1)} = 1) \implies R_L(\bar{x}_t, \beta_t^{(1)}) = 1.$$

Proof. We begin by providing some intuition to the proof. Before we present a sketch of the proof, it will be helpful to recall certain aspects. Recall that the modification introduced as we move from $\mathcal{H}_{0,(i-1):7}$ to $\mathcal{H}_{0,i:7}$ is that, in the $i + 1$ -th right session while $\mathcal{H}_{0,(i-1):7}$ sent that random r_V in Step cfp_3 that it committed to in Step cfp_1 , $\mathcal{H}_{0,i:7}$ chooses an r_V such that the resulting $\text{srs} = r_P \oplus r_V$ is a random DDH tuple. This modification thus renders $\text{NMMXCom}_{\text{srs}}$ at Step bps_4 of the $i + 1$ -th right session to be SB and CH (see Proposition 3).

At a high level, the structure of our proof would be to establish the following two claims:

1. We shall first prove that, in $\text{view}_{\text{extr}}^{(0,i:7)}$, if the $i + 1$ -th verifier accepts the $i + 1$ -th right session, then the value committed to in $\text{NMMXCom}_{\text{srs}}$ at Step bps_4 of the $i + 1$ -th right session is a valid witness for the statement of that session.
2. The modification of biasing the srs of the $i + 1$ -th right session from uniformly random to a random DDH tuple does not adversely affect the values committed to in $\text{NMMXCom}_{\text{srs}}$ at Step bps_4 of the j -th right session for every $j \leq i$.

Clearly, putting these two claims together would give us that in $\text{view}_{\text{extr}}^{(0,i:7)}$, the values committed to in $\text{NMMXCom}_{\text{srs}}$ at Step bps_4 of every accepted right session is a valid witness, as required.

We shall now proceed to establish each of these claims.

Sub-Claim 15. $\forall i$, in $\text{view}_{\text{extr}}^{(0,i:7)}$, $(b_{i+1}^{(0,i:7)} = 1) \implies R_L(\bar{x}_{i+1}, \beta_{i+1}^{(0,i:7)}) = 1$.

Proof. Here, we need to prove that $\forall i$, in $\text{view}_{\text{extr}}^{(0,i:7)}$, if \mathcal{V}_{i+1} accepts, then the value committed to in $\text{NMMXCom}_{\text{srs}}$ in Step bps_4 of the $(i + 1)$ -th right session is a valid witness.

Before we proceed to present the formal proof, here follows a high-level sketch of the proof for an intuition. The idea would be to follow our earlier proof strategy of proving soundness of our argument system. (Namely, the idea to prove soundness would be to reduce soundness to knowledge-soundness of sWIAoK at Step bps_2 and sWIAoK at Step bps_5 while making use of certain properties like CH, and CB of certain other sub-protocols in the Main BPS Phase.) A particular aspect of such a proof that will be crucial to us is the fact that we will need a reduction to CH property of CECom_{sb} at Step bps_1 of the $i + 1$ -th right session. For this, in our reduction, we will need to isolate this sub-protocol and delegate the task of sending \mathcal{M} the CECom commitment to an external CECom sender. Although, until now, our proofs had crucially used Robust Extraction Lemma, for this reduction we will not be able to work our way through with just this Lemma. The reason here is that the Robust Extraction Lemma can be successfully invoked only when the isolated protocol is of round-complexity strictly smaller than that of the other CECom commitments. This is necessary to render the view output by the robust simulator to be statistically close and to render the extraction successful with all but negligible probability. What comes to our rescue here is a *careful ordering of our hybrids*. The ordering of hybrids is such that at this point in the sequence of hybrids, delegating the task of the CECom sender for CECom_{sb} at Step bps_1 of the $i + 1$ -th right session can be done in such a way that all other CECom commitments that are isolated and forwarded to external CECom receivers are completed strictly earlier to the beginning of this CECom_{sb} commitment, ignoring the isolated CECom_{sh} commitments at Step bps_{4+} of the right sessions, since in the reductions these commitments need not be isolated as the values extracted from them are not used by the hybrids in any part of the execution. This proof direction is as

against assuring that the isolated sub-protocol will not be rewound just by the property of the robust simulator.

We are now ready to present the formal proof.

Let $t \in [m_R]$. To gain further insight into the proof, intuitively, $\text{NMMXCom}_{\text{srs}}$ at Step bps_4 of the t -th right session contains a valid witness owing to

- computational hiding of CECom_{sb} – to argue that \mathcal{M} does not learn σ , committed to by the verifier in CECom_{sb} , and use it in its Com_{sh} commitment and sWIAoK at Step bps_2 ,
- knowledge-soundness of sWIAoK in Step bps_2 – to extract knowledge of commitment information (i.e., committed value and randomness) for Com_{sh} in Step bps_2 and to verify that the extracted value will not be σ ,
- knowledge-soundness of sWIAoK in Step bps_5 – to argue that either the value committed to in $\text{NMMXCom}_{\text{srs}}$ at Step bps_4 is a valid (sub-)witness or the value committed to by the adversary in Step bps_2 is σ ,
- and finally, computational binding of Com_{sh} at Step bps_2 to show that the value committed to in this commitment is not σ .

Using CH of CECom_{sb} and knowledge-soundness of sWIAoK in Step bps_2 . We begin by showing that one can extract the committed value and the randomness in Com_{sh} at Step bps_2 from sWIAoK at the same Step, and by computational hiding of CECom_{sb} at Step bps_1 , the value will not be σ .

Consider an adversarial prover P_1^* against knowledge-soundness of sWIAoK which behaves as follows. P_1^* runs $\text{RobustSim}^{\tilde{I}_1}(z)$, where \tilde{I}_1 behaves the same way as $I_{\text{extr}}^{(0,i:7)}$ except for the following modifications.

- Recall that the only sub-protocols isolated by $I_{\text{extr}}^{(0,i:7)}$ are the CECom_{sb} commitments at Step bps_1^{cp} and CECom_{sh} commitments at Step bps_{4+} of the first $i+1$ right sessions. On the other hand, \tilde{I}_2 additionally isolates sWIAoK in Step bps_2 of the t -th right session and forwards it to an external sWIAoK verifier.

Upon completion of this sWIAoK protocol, if the sWIAoK verifier accepts, then it runs sWIAoK extractor on $\text{RobustSim}^{\tilde{I}_2}(z)$.

To argue a reduction to knowledge-soundness of the sWIAoK argument in question, we will need to argue that the view of the adversary \mathcal{M} when run by P_1^* is statistically close to the view of the adversary \mathcal{M} when run by $I_{\text{extr}}^{(0,i:7)}$. For this, we apply the Robust Extraction Lemma. Since P_1^* also isolates sWIAoK protocol and relays it to an external party, to apply the Robust Extraction Lemma, we will first create two intermediate hybrids $\text{hyb}_{0,i:7,A}^*$ and $\text{hyb}_{0,i:7,B}^*$, whose outputs are identical, and the output of the former is statistically close to the view of the MiM adversary \mathcal{M} when run by $\mathcal{H}_{0,i:7}$ and the output of the latter is statistically close to the view of the MiM adversary \mathcal{M} when run by P_1^* .

$\text{hyb}_{0,i:7,A}^*$ is described as follows. It simply outputs the view output by the online extractor, namely, $\text{REAL}_{\mathcal{E},\Pi}^{I_{\text{extr}}^{(0,i:7)}}(\beta, z)$, where the external protocol Π here is the empty protocol. Now, we can invoke the Robust Extraction Lemma, for which we consider the following: the CECom commitments, that are isolated and relayed to external CECom receivers, have k_{cecom} -slots and the external

protocol that the robust simulator is participating in, here, is the empty protocol. Thus, we have that $\ell_{cecom} = k_{cecom}$ and $\ell_{external} = 0$. Now by applying the Robust Extraction Lemma, we have that statistical distance between the views $\text{REAL}_{\mathcal{E}, \Pi}^{(0,i:7)}(\beta, z)$ and $\text{view}_{\text{extr}}^{(0,i:7)}$ is at most:

$$\Delta(\lambda) \leq 2^{-(k_{cecom} - 0 \cdot \log(T(\lambda)))} \leq \text{negl}(\lambda),$$

since, $k_{cecom} \in \omega(\log(\lambda))$ and T is at most a polynomial.

Next, we describe an intermediate hybrid $\text{hyb}_{0,i:7,B}^*$ whose output is identical to that of $\text{hyb}_{0,i:7,A}^*$. For this consider an interface, $I_{\text{extr}}^{(0,i:7,B)}$ behaves the same way as $I_{\text{extr}}^{(0,i:7)}$ except that it also isolates the sWIAoK protocol in Step bps_2 of the t -th right session. $\text{hyb}_{0,i:7,B}^*$ simply outputs the view output by the online extractor, namely, $\text{REAL}_{\mathcal{E}, \Pi}^{(0,i:7,B)}(\beta, z)$, where the external protocol Π here is the sWIAoK protocol in Step bps_2 of the t -th right session, and the external party running the code of the sWIAoK verifier for the isolated sWIAoK argument. Again, we can invoke the Robust Extraction Lemma, for which we consider the following: the CECOM commitments, that are isolated and relayed to external CECOM receivers, have k_{cecom} -slots and the external protocol that the robust simulator is participating in, here, Π here is the sWIAoK protocol in Step bps_2 of the t -th right session, and the external party running the code of the sWIAoK verifier for the isolated sWIAoK argument. Thus, we have that $\ell_{cecom} = k_{cecom}$ and $\ell_{external} = k_{swiaok}$. Now by applying the Robust Extraction Lemma, we have that statistical distance between the views $\text{REAL}_{\mathcal{E}, \Pi}^{(0,i:7,B)}(\beta, z)$ and the view of the \mathcal{M} during its interaction with P_1^* is at most:

$$\Delta(\lambda) \leq 2^{-(k_{cecom} - k_{swiaok} \cdot \log(T(\lambda)))} \leq \text{negl}(\lambda),$$

since, $k_{cecom} \in \omega(\log(\lambda))$, k_{swiaok} is a constant, and T is at most a polynomial.

Thus, we have proven that the view of the MiM adversary \mathcal{M} during its interaction with P_1^* is statistically close to the view output by $\mathcal{H}_{0,i:7}$.

Moreover, since the sWIAoK argument is isolated and relayed to an honest external sWIAoK verifier, knowledge-soundness of sWIAoK implies that the sWIAoK extractor extracts a valid sWIAoK witness – the value committed to together with the randomness used – of the Com_{sh} commitment of Step bps_2 . This also implies that the value is not σ (committed to in CECom_{sb} of Step bps_1) with all but negligible probability, as otherwise we can build an adversary \mathcal{A}_{CH} that breaks computational hiding of CECom_{sb} of Step bps_1 as follows.

Assume for contradiction that the value extracted by P_1^* is σ with some non-negligible probability ϵ . Then we shall show that \mathcal{A}_{CH} breaks hiding with probability $\epsilon - \text{negl}(\lambda)$. \mathcal{A}_{CH} is described as follows. \mathcal{A}_{CH} runs $\text{RobustSim}^{\tilde{I}_2}(z)$, where \tilde{I}_2 behaves the same as \tilde{I}_1 except that, besides isolating the sWIAoK argument of Step bps_2 , also isolates CECom_{sb} of Step bps_1 of the t -th right session. While the CECom_{sb} commitment is forwarded to an external CECOM sender, \mathcal{A}_{CH} itself runs the honest verifier code of the isolated sWIAoK argument. If the sWIAoK verifier (run by \mathcal{A}_{CH}) accepts the sWIAoK argument, then \mathcal{A}_{CH} also runs the sWIAoK extractor on the isolated sWIAoK argument. Furthermore, \mathcal{A}_{CH} does not continue the interaction with \mathcal{M} after the sWIAoK argument. Crucially, note that the isolated sWIAoK argument (at Step bps_2) begins strictly after the completion of the isolated CECom_{sb} commitment (at Step bps_1) as they both belong to the same session. Furthermore, observe that the only other sub-protocols isolated by \tilde{I}_2 are the CECom_{sb} commitments at Step $\text{bps}_1^{\text{cfp}}$ of the *first* $i + 1$ commitments. Also recall that the ordering of the right sessions is by the order in which the CECom_{sb} commitments at Step bps_1 begin. Thus, since

in any given session Step $\text{bps}_1^{\text{cfp}}$ is completed well before Step bps_1 of that session, all the isolated CECOM commitments are completed before the beginning of the isolated CECom_{sb} commitments at Step bps_1 of the $i + 1$ -th right session. Hence, neither the sWIAoK rewindings nor the rewindings by the robust simulator interfere with the isolated CECom_{sb} commitment thus ensuring that the external CECOM sender will not be rewound.

Finally, note that the view of the MiM adversary \mathcal{M} when run by \mathcal{A}_{CH} is statistically close to the view of \mathcal{M} when run by P_1^* until the sWIAoK argument at Step bps_2 of the t -th right session, after which \mathcal{A}_{CH} aborts. We thus have that \mathcal{A}_{CH} breaks computational hiding of CECom_{sb} of Step bps_1 with probability $\epsilon - \text{negl}(\lambda)$.

Using knowledge-soundness of sWIAoK in Step bps_5 and CB of Com_{sh} at Step bps_2 . Now, we show that, for the $i + 1$ -th right session, one can extract the witness used in sWIAoK of Step bps_5 , and from its knowledge-soundness, we have that either we extract the value committed to in $\text{NMMXCom}_{\text{srs}}$ at Step bps_4 and a commitment information of CECom_{sh} at Step bps_{4+} such that both the openings are to a common value that is a valid sub-witness or we extract an opening of Com_{sh} at Step bps_2 to σ . Finally, we will see that CB of Com_{sh} at Step bps_2 implies that the extracted value is not an opening of Com_{sh} to σ . Putting it all together, we will have established soundness of the argument proved in the BPS phase. Finally, since $\text{NMMXCom}_{\text{srs}}$ is statistically binding, we will have that the value committed to in it is a valid sub-witness, rand with all but negligible probability.

Consider an adversarial prover P_2^* against knowledge-soundness of sWIAoK which behaves as follows. Recall that P_1^* executed $\text{RobustSim}^{\tilde{I}_1}(z)$, where \tilde{I}_1 behaved the same way as $I_{\text{extr}}^{(0,i:7)}$ except that it also isolated the sWIAoK argument in Step bps_2 of the t -th right session, but did not isolate any CECom_{sh} commitments at Step bps_{4+} of right sessions. P_2^* instead isolates the sWIAoK argument in Step bps_5 of the t -th right session. More formally, P_2^* runs $\text{RobustSim}^{\tilde{I}_3}(z)$, where \tilde{I}_3 behaves the same way as $I_{\text{extr}}^{(0,i:7)}$ except for the following modifications.

- Recall that the only sub-protocols isolated by $I_{\text{extr}}^{(0,i:7)}$ are the CECom_{sb} commitments at Step $\text{bps}_1^{\text{cfp}}$ and CECom_{sh} commitments at Step bps_{4+} of the first $i + 1$ right sessions. On the other hand, \tilde{I}_3 also isolates sWIAoK in Step bps_5 of the t -th right session and forwards it to an external sWIAoK verifier.

Upon completion of this sWIAoK protocol, if the sWIAoK verifier accepts, then it runs sWIAoK extractor on $\text{RobustSim}^{\tilde{I}_3}(z)$.

Recall that we proved that the view of the adversary \mathcal{M} when run by P_1^* is statistically close to the view of the adversary \mathcal{M} when run by $I_{\text{extr}}^{(0,i:7)}$. On precisely the same lines, it can be proven that the view of the adversary \mathcal{M} when run by P_2^* is statistically close to the view of the adversary \mathcal{M} when run by $I_{\text{extr}}^{(0,i:7)}$.

Moreover, since the sWIAoK argument is isolated and relayed to an honest external sWIAoK verifier, knowledge-soundness of sWIAoK implies that the sWIAoK extractor extracts a valid sWIAoK witness – either $\bar{y}_t, r_{nm}, r_{cecom}$ such that $(\bar{x}, \bar{y}_t) \in R_L$ with r_{nm} and r_{cecom} explaining the commitments $\text{NMMXCom}_{\text{srs}}$ of Step bps_4 and the CECOM commitment of Step bps_{4+} to \bar{y}_t , respectively, or an opening of Com_{sh} at Step bps_2 to σ . We shall shortly show that owing to CB of Com_{sh} of Step bps_2 , the extracted value is not a Com_{sh} opening to σ . Given this, we have that the extracted value contains a valid witness \bar{y}_t and a commitment information of the $\text{NMMXCom}_{\text{srs}}$ commitment of Step

bps_4 to \bar{y}_t together with a commitment information CECom_{sh} Step bps_{4+} to \bar{y}_t . Since $\text{NMMXCom}_{\text{srs}}$ is statistically binding, *the value committed to in it should be a valid witness itself* with all but negligible probability.

Now it remains to show that the extracted output of sWIAoK extractor above is not an opening of Com_{sh} to σ with all but negligible probability. Assume for contradiction that the extracted output is an opening of Com_{sh} to σ with some non-negligible probability ε . Then we construct an adversary \mathcal{A}'_{CB} that breaks CB of Com_{sh} also with probability $\varepsilon - \text{negl}(\lambda)$.

\mathcal{A}'_{CB} is described as follows. \mathcal{A}'_{CB} runs $\text{RobustSim}^{\tilde{I}_4}(z)$, where \tilde{I}_4 behaves the same as \tilde{I}_3 (run by P_2^*), except for the following modification.

- Unlike \tilde{I}_3 (or \tilde{I}_1) which isolates only one of the two sWIAoK protocols present in our protocol, \mathcal{A}'_{CB} isolates both the sWIAoK protocols, one at Step bps_2 and the other at Step bps_5 of the $i + 1$ -th right session.
- Furthermore, it also isolates Com_{sh} of Step bps_2 and forwards it to an external Com_{sh} receiver.

However, the sWIAoK verifiers' task for both the isolated sWIAoK arguments are run by \mathcal{A}'_{CB} itself. Furthermore, if both the isolated sWIAoK arguments are accepted, then \mathcal{A}'_{CB} also runs the sWIAoK extractor on $\text{RobustSim}^{\tilde{I}_4}(z)$ – once for each of the two isolated sWIAoK arguments.

Now observe that since the sWIAoK arguments are isolated, rewindings by the robust simulator do not interfere with these sWIAoK arguments. Knowledge-soundness of sWIAoK thus implies that the sWIAoK extractor extracts valid sWIAoK witnesses for both the sWIAoK arguments. Furthermore, since the Com_{sh} commitment is isolated and forwarded to an external receiver, it also holds that rewindings by the robust simulator do not interfere with this commitment. Thus, outputting two openings to two distinct values amounts to breaking binding of Com_{sh} .

As proven earlier, the extracted output of sWIAoK at Step bps_2 is an opening of Com_{sh} with all but negligible probability; as also proven earlier, the extracted value, however, is *not* σ with all but negligible probability. Thus, we have the extracted output obtained by \mathcal{A}'_{CB} out of this sWIAoK is $(\delta, \text{rand}_\delta)$, such that $\delta \neq \sigma$ and $(\delta, \text{rand}_\delta)$ is a valid opening of Com_{sh} . Furthermore, as also proven above, the extracted output of sWIAoK at Step bps_5 is *either* an opening of $\text{NMMXCom}_{\text{srs}}$ to a value rand such that rand explains CECom_{sh} at Step cfp_1 to be a commitment to r_V , *or* an opening of Com_{sh} at Step bps_2 to σ . If the extracted output is the latter, i.e., an opening of Com_{sh} to σ , then \mathcal{A}'_{CB} has openings of Com_{sh} (which is isolated) to two distinct values δ and σ . Thus, \mathcal{A}'_{CB} breaks computational binding of Com_{sh} with probability $\varepsilon - \text{negl}(\lambda)$, under the assumption that \mathcal{A}'_{CB} statistically simulates the view of \mathcal{M} while interacting with P_2^* . We shall argue validity of this assumption shortly. This proves that the value extracted from sWIAoK at Step bps_5 should be an opening of $\text{NMMXCom}_{\text{srs}}$ to rand . Since $\text{NMMXCom}_{\text{srs}}$ is SB, the value committed to in it should be a valid witness itself, with all but negligible probability, thus proving Sub-claim 2. Thus it only remains to prove statistical simulation by \mathcal{A}'_{CB} of the view of \mathcal{M} when incorporated by P_2^* .

Lastly, we will need to argue that the view of the adversary \mathcal{M} when run by \mathcal{A}'_{CB} is statistically close to the view of the adversary \mathcal{M} when run by P_2^* .

For this, we apply the Robust Extraction Lemma. Since P_2^* also isolates a sWIAoK protocol and \mathcal{A}'_{CB} two sWIAoK protocols and a Com_{sh} commitment, to apply the Robust Extraction Lemma, we will first create two intermediate hybrids $\text{hyb}_{0,i:7,C}^*$ and $\text{hyb}_{0,i:7,D}^*$, whose outputs are identical, and the output of the former is statistically close to the view of the MiM adversary \mathcal{M} when run

by P_2^* and the output of the latter is statistically close to the view of the MiM adversary \mathcal{M} when run by \mathcal{A}'_{CB} .

$\text{hyb}_{0,i:7,C}^*$ is described as follows. It simply outputs the view output by the online extractor, namely, $\text{REAL}_{\mathcal{E},\Pi}^{\tilde{I}_3}(\beta, z)$, where the external protocol Π here is the sWIAoK in Step bps_5 of the t -th right session. (Recall that \tilde{I}_3 is the interface incorporated by P_2^* . Now, we can invoke the Robust Extraction Lemma, for which we consider the following: the CECOM commitments, that are isolated and relayed to external CECOM receivers, have k_{cecom} -slots and the external protocol that the robust simulator is participating in, here, is the sWIAoK in Step bps_5 of the t -th right session. Thus, we have that $\ell_{\text{cecom}} = k_{\text{cecom}}$ and $\ell_{\text{external}} = k_{\text{swiaok}}$. Now by applying the Robust Extraction Lemma, we have that statistical distance between the view of the MiM adversary \mathcal{M} when run by P_2^* and the view $\text{REAL}_{\mathcal{E},\Pi}^{\tilde{I}_3}(\beta, z)$ is at most:

$$\Delta(\lambda) \leq 2^{-(k_{\text{cecom}} - k_{\text{swiaok}} \cdot \log(T(\lambda)))} \leq \text{negl}(\lambda),$$

since, $k_{\text{cecom}} \in \omega(\log(\lambda))$, k_{swiaok} is a constant, and T is at most a polynomial.

Next, we describe an intermediate hybrid $\text{hyb}_{0,i:7,D}^*$ whose output is identical to that of $\text{hyb}_{0,i:7,C}^*$. For this consider an interface, $I_{\text{extr}}^{(0,i:7,D)}$ behaves the same way as $I_{\text{extr}}^{(0,i:7)}$ except that it also isolates both the sWIAoK protocols, one at Step bps_2 and the other at Step bps_5 of the $i + 1$ -th right session; furthermore, it also isolates Com_{sh} of Step bps_2 . $\text{hyb}_{0,i:7,D}^*$ simply outputs the view output by the online extractor, namely, $\text{REAL}_{\mathcal{E},\Pi}^{I_{\text{extr}}^{(0,i:7,D)}}(\beta, z)$, where the external protocol Π here consists of the sWIAoK protocol at Step bps_2 , the sWIAoK protocol at Step bps_5 of the $i + 1$ -th right session, and Com_{sh} of Step bps_2 . Again, we can invoke the Robust Extraction Lemma, for which we consider the following: the CECOM commitments, that are isolated and relayed to external CECOM receivers, have k_{cecom} -slots and the external protocol that the robust simulator is participating in, here, Π here consists of two sWIAoK protocols and one Com_{sh} commitment. Thus, we have that $\ell_{\text{cecom}} = k_{\text{cecom}}$ and $\ell_{\text{external}} = k_{sh} + 2k_{\text{swiaok}}$. Now by applying the Robust Extraction Lemma, we have that statistical distance between the views $\text{REAL}_{\mathcal{E},\Pi}^{I_{\text{extr}}^{(0,i:7,D)}}(\beta, z)$ and the view of the \mathcal{M} during its interaction with \mathcal{A}'_{CB} is at most:

$$\Delta(\lambda) \leq 2^{-(k_{\text{cecom}} - k_{sh} + 2k_{\text{swiaok}} \cdot \log(T(\lambda)))} \leq \text{negl}(\lambda),$$

since, $k_{\text{cecom}} \in \omega(\log(\lambda))$, $k_{sh}, k_{\text{swiaok}}$ are constants, and T is at most a polynomial.

Thus, we have proven that the view of the MiM adversary \mathcal{M} during its interaction with \mathcal{A}'_{CB} is statistically close to the view during its interaction with P_2^* .

With this, we have proven that, $\text{view}_{\text{extr}}^{(1)}$, for any right session, if the verifier accepts the session then the value committed to in $\text{NMMXCom}_{\text{srs}}$ of every accepting right session should be a valid witness, with all but negligible probability, thus proving Sub-claim 15. □

Recall that our final objective in this proof is to prove that, in $\text{view}_{\text{extr}}^{(1)}$, for any $i \in [m_R]$, if the t -th verifier accepts the t -th right session, then the value committed to in $\text{NMMXCom}_{\text{srs}}$ at Step bps_4 of the t -th right session is a valid witness for the statement of that session. Towards this, we have completed the first step: namely, we have proven that in $\text{view}_{\text{extr}}^{(0,i:7)}$, if the $i + 1$ -th verifier accepts the $i + 1$ -th right session, then the value committed to in $\text{NMMXCom}_{\text{srs}}$ at Step bps_4 of the $i + 1$ -th right session is a valid witness for the statement of that session. Thus, to fulfill our

objective, it remains to show that as we move from $\mathcal{H}_{0,i:6}$ to $\mathcal{H}_{0,i:7}$, the modification of biasing the srs of the $i + 1$ -th right session from uniformly random to a random DDH tuple does not adversely affect the values committed to in $\text{NMMXCom}_{\text{srs}}$ at Step bps_4 of the j -th right session for every $j \leq i$. We shall establish this final step in the following.

Sub-Claim 16. $\forall i, \forall j \geq i, \beta_j^{((0,i:6))} \approx_c \beta_j^{((0,i:7))}$.

Proof. Fix any $i \in [m_R]$. We consider the views $\text{view}_{\text{extr}}^{(0,i:6)}$ and $\text{view}_{\text{extr}}^{(0,i:7)}$ generated by the hybrids $\mathcal{H}_{0,i:6}$ and $\mathcal{H}_{0,i:7}$, respectively. We shall prove that in moving from $\mathcal{H}_{0,i:6}$ and $\mathcal{H}_{0,i:7}$, where the only modification we introduce is that in the $i + 1$ -th right session the srs which used to be uniformly random in $\text{view}_{\text{extr}}^{(0,i:6)}$ is biased to a random DDH tuple, the values that were committed to in $\text{NMMXCom}_{\text{srs}}$ at Step bps_4 of j -th session for any $j \leq i$ does not adversely change. Before we proceed we provide some intuition to the proof.

Note that for all $j \leq i$, in $\text{view}_{\text{extr}}^{(0,i:7)}$, the srs is already biased to random DDH tuples. Also in $\text{view}_{\text{extr}}^{(0,i:6)}$, all the first i right sessions have their srs biased to random DDH tuples. Hence, in both the views $\text{view}_{\text{extr}}^{(0,i:6)}$ and $\text{view}_{\text{extr}}^{(0,i:7)}$, $\text{NMMXCom}_{\text{srs}}$ at Step bps_4 of the first i right sessions are SB and CH (see Proposition 3), (and thus, in both $\text{view}_{\text{extr}}^{(0,i:6)}$ and $\text{view}_{\text{extr}}^{(0,i:7)}$, for all $j \leq i$, $\beta_j^{((0,i:6))}$ and $\beta_j^{((0,i:7))}$ are well defined). Thus to now prove that the committed values in these $\text{NMMXCom}_{\text{srs}}$ commitments do not adversely change as we move from $\mathcal{H}_{0,i:6}$ to $\mathcal{H}_{0,i:7}$, we rely upon robust non-malleability of the Com_{nm} commitment scheme w.r.t. 1-round protocols (see Definition 15). The details follow.

Observe that the change introduced by $\mathcal{H}_{0,i:7}$ is only computationally indistinguishable, but not statistically so. Furthermore, as mentioned earlier, the value committed to in the SB NCom commitment $\text{NMMXCom}_{\text{srs}}$ at Step bps_{4+} of any right session is not revealed by \mathcal{M} in any part of the protocol. Hence, computational indistinguishability of DDH tuples from uniform strings is insufficient to argue that the value committed to by the adversary \mathcal{M} in the Com_{nm} commitment at Step $\text{bps}_{\text{cfp}_4}$ of any left session does not change adversely. For this, we shall rely upon robust non-malleability of the Com_{nm} commitment scheme w.r.t. 1-round protocols.

Assume for contradiction that there exists $j \in [m_R]$ such that $j \leq i$ and the values $\beta_j^{((0,i:6))}$ and $\beta_j^{((0,i:7))}$ committed to in $\text{NMMXCom}_{\text{srs}}$ at Step bps_{4+} at the j -th right session in $\text{view}_{\text{extr}}^{(0,i:6)}$ and $\text{view}_{\text{extr}}^{(0,i:7)}$, respectively, are computationally distinguishable by some PPT distinguisher with some non-negligible probability ϵ . Then, by relying on the DDH assumption, we construct an adversary $\mathcal{A}_{\text{rob-nmcom}}$ against robust non-malleability of the NCom commitment scheme w.r.t. 1-round protocols.

$\mathcal{A}_{\text{rob-nmcom}}$ behaves exactly the same way as $\mathcal{H}_{0,i:6}$ except that it also isolates the $\text{NMMXCom}_{\text{srs}}$ commitment at Step bps_{4+} of the j -th right session and the Step cfp_3 message of the $i + 1$ -th right session, and forwards them to an external NCom receiver and to an external DDH challenger, respectively. The DDH challenger chooses srs which is either a random DDH tuple or a uniform random string (and sets r_V accordingly).

Before we proceed to analyze the success probability of $\mathcal{A}_{\text{rob-nmcom}}$, we need to show that the view of the MiM adversary \mathcal{M} during its interaction with $\mathcal{A}_{\text{rob-nmcom}}$ in the case where the string received from the external DDH challenger is a uniform random string (in the isolated message of Step cfp_3 of the $i + 1$ -th right session) is statistically close to the view output by $\mathcal{H}_{0,i:6}$. For this, we apply the Robust Extraction Lemma. Since $\mathcal{A}_{\text{rob-nmcom}}$ also isolates the Com_{nm} commitment

at Step bps_{4+} of the j -th right session and the Step cfp_3 message of the $i + 1$ -th right session, and relays them to external parties, to apply the Robust Extraction Lemma, we will first create two intermediate hybrids $\text{hyb}_{0,i:6,C}^*$ and $\text{hyb}_{0,i:6,D}^*$, whose outputs are identical, and the output of the former is statistically close to the view of the MiM adversary \mathcal{M} when run by $\mathcal{H}_{0,i:6}$ and the output of the latter is statistically close to the view of the MiM adversary \mathcal{M} when run by $\mathcal{A}_{\text{rob-nmcom}}$ in the case where the string received from the external DDH challenger is a uniform random string.

$\text{hyb}_{0,i:6,C}^*$ is described as follows. It simply outputs the view output by the online extractor, namely, $\text{REAL}_{\mathcal{E},\Pi}^{I_{\text{extr}}^{(0,i:6)}}(\beta, z)$, where the external protocol Π here is the empty protocol. Now, we can invoke the Robust Extraction Lemma, for which we consider the following: the CECOM commitments, that are isolated and relayed to external CECOM receivers, have k_{cecom} -slots and the external protocol that the robust simulator is participating in, here, is the empty protocol. Thus, we have that $\ell_{\text{cecom}} = k_{\text{cecom}}$ and $\ell_{\text{external}} = 0$. Now by applying the Robust Extraction Lemma, we have that statistical distance between the views $\text{REAL}_{\mathcal{E},\Pi}^{I_{\text{extr}}^{(0,i:6)}}(\beta, z)$ and $\text{view}_{\text{extr}}^{(0,i:6)}$ is at most:

$$\Delta(\lambda) \leq 2^{-(k_{\text{cecom}} - 0 \cdot \log(T(\lambda)))} \leq \text{negl}(\lambda),$$

since, $k_{\text{cecom}} \in \omega(\log(\lambda))$ and T is at most a polynomial.

Next, we describe an intermediate hybrid $\text{hyb}_{0,i:6,D}^*$ whose output is identical to that of $\text{hyb}_{0,i:6,C}^*$. For this consider an interface, $I_{\text{extr}}^{(0,i:6,D)}$ behaves the same way as $I_{\text{extr}}^{(0,i:6)}$ except that it also isolates the Com_{nm} commitment at Step bps_{4+} of the j -th right session and the Step cfp_3 message of the $i + 1$ -th right session, and forwards them to an external NMCom receiver and a DDH challenger, respectively. $\text{hyb}_{0,i:6,D}^*$ simply outputs the view output by the online extractor, namely, $\text{REAL}_{\mathcal{E},\Pi}^{I_{\text{extr}}^{(0,i:6,D)}}(\beta, z)$, where the external protocol Π here consists of the Com_{nm} commitment at Step bps_{4+} of the j -th right session and the Step cfp_3 message of the $i + 1$ -th right session, and the role of the external party is to run the code of the NMCom receiver and the DDH challenger, respectively, for the isolated sub-protocols. Again, we can invoke the Robust Extraction Lemma, for which we consider the following: the CECOM commitments, that are isolated and relayed to external CECOM receivers, have k_{cecom} -slots and the external protocol that the robust simulator is participating in, here, Π here consists of the consists of two sub-protocols: namely, one Com_{nm} commitment of k_{nmcom} rounds and one 1-round protocol. Thus, we have that $\ell_{\text{cecom}} = k_{\text{cecom}}$ and $\ell_{\text{external}} = k_{\text{nmcom}} + 1$. Now by applying the Robust Extraction Lemma, we have that statistical distance between the views $\text{REAL}_{\mathcal{E},\Pi}^{I_{\text{extr}}^{(0,i:6,D)}}(\beta, z)$ and the view of the \mathcal{M} during its interaction with $\mathcal{A}_{\text{rob-nmcom}}$ in the case where the string received from the external DDH challenger is a uniform random string is at most:

$$\Delta(\lambda) \leq 2^{-(k_{\text{cecom}} - (k_{\text{nmcom}} + 1) \cdot \log(T(\lambda)))} \leq \text{negl}(\lambda),$$

since, $k_{\text{cecom}} \in \omega(\log(\lambda))$, $k_{\text{nmcom}} \in O(\log(\lambda))$, and T is at most a polynomial.

Thus, we have proven that the view of the MiM adversary \mathcal{M} during its interaction with $\mathcal{A}_{\text{rob-nmcom}}$ in the case where the string received from the external DDH challenger is a uniform random string is statistically close to the view output by $\mathcal{H}_{0,i:6}$.

Also, similarly, we have that the view of the MiM adversary \mathcal{M} during its interaction with $\mathcal{A}_{\text{rob-nmcom}}$ in the case where the string received from the external DDH challenger is a random DDH tuple is statistically close to the view output by $\mathcal{H}_{0,i:7}$.

We have thus proven that if the external DDH challenger sends a uniform random string, then the view of the adversary \mathcal{M} in its interaction with $\mathcal{A}_{rob-nmcom}$ is a statistical simulation of the view $\text{view}_{\text{extr}}^{(0,i:6)}$; on the other hand, if the external DDH challenger sends a random DDH tuple, then the view of the adversary \mathcal{M} in its interaction with $\mathcal{A}_{rob-nmcom}$ is a statistical simulation of the view $\text{view}_{\text{extr}}^{(0,i:7)}$. Thus, from the assumption that, for $\mathcal{H}_{0,i:7}$, for the j -th right session, \mathcal{P}_j accepts, but the value committed to by \mathcal{M} in $\text{NMMXCom}_{\text{srs}}$ commitment at Step bps_{4+} is not a valid witness, with some non-negligible probability ϵ , and from induction that this is not the case in $\mathcal{H}_{0,i:6}$, by relying on the DDH assumption, we have that $\mathcal{A}_{rob-nmcom}$ breaks robust non-malleability of the NMCCom commitment with probability $\epsilon - \text{negl}(\lambda)$. □

This completes the proof of Claim 12. □

Now we have established that for every accepting j -th right session for $j \leq i$, the value contained in $\text{NMMXCom}_{\text{srs}}$ at Step bps_4 is a valid witness. Next, we shall prove that the values $y_j^{(0,i:7)}$ extracted by $\mathcal{H}_{0,i:7}$ from CECom_{sh} at Step bps_4 of the j -th right sessions are valid witnesses.

Sub-Claim 17. $\forall i, \forall j \leq i$, in $\text{view}_{\text{extr}}^{(0,i:7)}$, $(b_j^{(0,i:7)} = 1) \implies (\bar{x}_j, y_j^{(0,i:7)}) \in R_L$.

Proof. We begin with a high-level sketch of the proof. Assume for contradiction that, with some non-negligible probability ϵ , there exists $j \in [m_R]$ such that $j \leq i$ and for hybrid $\mathcal{H}_{0,i:7}$ the j -th verifier accepts the j -th right session, but $(\bar{x}_j, y_j^{(0,i:7)}) \notin R_L$. Then we can construct an adversary \mathcal{A}_{CB}^* that breaks computational binding of CECom_{sh} at Step bps_4 . To pull through this reduction, we will along the way show that by knowledge-soundness of sWIAoK in Step bps_5 , the value extracted by $\mathcal{H}_{0,i:7}$ is in fact the value committed to in $\text{NMMXCom}_{\text{srs}}$ at Step bps_4 of that session, which we have already proven to be a valid witness. To in turn pull this argument, along the way, we will also invoke knowledge-soundness of sWIAoK in Step bps_2 and computational binding of Com_{sh} in Step bps_2 and computational hiding of CECom_{sb} at Step bps_1 .

Recall that in the proof of Sub-claim 15, we had proven that in a view that is statistically close to in $\text{view}_{\text{extr}}^{(0,i:7)}$ if the sWIAoK argument in Step bps_2 is isolated and forwarded to an external sWIAoK verifier then by running an sWIAoK extractor on it would give an opening of Com_{sh} commitment at Step bps_2 and by invoking computational hiding of the CECom_{sb} commitment at Step bps_1 , this extracted opening value is not the value committed to in the CECom_{sb} commitment at Step bps_1 . Also recall that in the proof of Sub-claim 15, we had proven that in a view that is statistically close to in $\text{view}_{\text{extr}}^{(0,i:7)}$ if the sWIAoK argument in Step bps_5 is isolated and forwarded to an external sWIAoK verifier then by running an sWIAoK extractor on it would give openings of $\text{NMMXCom}_{\text{srs}}$ at Step bps_4 and CECom_{sh} at Step bps_{4+} to the same value \bar{y}_j , which we have proven in Claim 12 to be a valid witness as $\text{NMMXCom}_{\text{srs}}$ of this session is SB.

With this, we are now ready to construct an adversary \mathcal{A}_{CB}^* that breaks computational binding of CECom_{sh} (at Step bps_4). This adversary behaves the same way as $\mathcal{H}_{0,i:7}$ except for a few modifications. Recall that $\mathcal{H}_{0,i:7}$ ran $\text{RobustSim}_{\text{extr}}^{I_{\text{extr}}^{(0,i:7)}}(z)$ where $I_{\text{extr}}^{(0,i:7)}$ isolated certain CECom commitments including CECom_{sh} commitments at Step bps_{4+} of the first $i+1$ right sessions. \mathcal{A}_{CB}^* runs $\text{RobustSim}_{\text{extr}}^{\tilde{I}_5}(z)$, where \tilde{I}_5 differs from $I_{\text{extr}}^{(0,i:7)}$ in the following sense.

- \tilde{I}_5 also isolates sWIAoK argument in Step bps_5 of the j -th right session. Let $(y_j^{(0,i:7,a)}, r_{cecom}^{robust})$ be the commitment information received from the outside for CECom_{sh} commitment at Step bps_{4+} of the j -th right session.

Upon completion of this sWIAoK protocol, if the sWIAoK verifier accepts, then it runs sWIAoK extractor on $\text{RobustSim}^{\tilde{I}_5}(z)$. If the sWIAoK extractor returns an opening of the CECom_{sh} commitment of Step bps_{4+} , then output this opening together with the opening $(y_j^{(0,i:7,a)}, r_{cecom}^{robust})$ to the same commitment.

We shall now analyze the success probability of \mathcal{A}_{CB}^* in breaking computational binding of CECom_{sh} commitment. Firstly, we recall that like we showed in the proof of Sub-claim 15, owing to knowledge-soundness of the sWIAoK argument at Step bps_5 and computational binding of Com_{sh} at Step bps_2 , the extracted value from the sWIAoK extractor is $\bar{y}_t, r_{nm}, r_{cecom}$ such that $(\bar{x}, \bar{y}_t) \in R_L$ with r_{nm} and r_{cecom} explaining the commitments $\text{NMMXCom}_{\text{srs}}$ of Step bps_4 and the CECom commitment of Step bps_{4+} to \bar{y}_t , respectively. Thus, \mathcal{A}_{CB}^* already has an opening (\bar{y}_t, r_{cecom}) of the isolated CECom_{sh} commitment to a valid witness \bar{y}_t . Furthermore, from the sWIAoK extractor, it also has an opening $(y_j^{(0,i:7,a)}, r_{cecom}^{robust})$ of the same commitment to $y_j^{(0,i:7,a)}$. Thus, if $y_j^{(0,i:7,a)} \neq \bar{y}_t$, then \mathcal{A}_{CB}^* breaks computational binding of CECom_{sh} commitment. In the following we shall show that the probability of the event that $y_j^{(0,i:7,a)} \neq \bar{y}_t$ for \mathcal{A}_{CB}^* is negligibly close to the probability that the extracted value by $\mathcal{H}_{0,i:7}$ for the j -th right session is not a valid witness.

Note that both $\mathcal{H}_{0,i:7}$ and \mathcal{A}_{CB}^* extract an opening of CECom_{sh} commitment at Step bps_{4+} of the j -th right session via robust simulator. Furthermore, both the algorithms run in polynomial time. Thus, we only need to show statistical indistinguishability in the outputs (views together with the values output) of these algorithms.

For this, we apply the Robust Extraction Lemma. Since \mathcal{A}_{CB}^* also isolates sWIAoK protocol and relays it to an external party, to apply the Robust Extraction Lemma, we will first create two intermediate hybrids $\text{hyb}_{0,i:7,A}^*$ and $\text{hyb}_{0,i:7,B}^*$, whose outputs are identical, and the output of the former is statistically close to the view of the MiM adversary \mathcal{M} when run by $\mathcal{H}_{0,i:7}$ and the output of the latter is statistically close to the view of the MiM adversary \mathcal{M} when run by \mathcal{A}_{CB}^* .

$\text{hyb}_{0,i:7,A}^*$ is described as follows. It simply outputs the view output by the online extractor, namely, $\text{REAL}_{\mathcal{E},\Pi}^{(0,i:7)}(\beta, z)$, where the external protocol Π here is the empty protocol. Now, we can invoke the Robust Extraction Lemma, for which we consider the following: the CECom commitments, that are isolated and relayed to external CECom receivers, have k_{cecom} -slots and the external protocol that the robust simulator is participating in, here, is the empty protocol. Thus, we have that $\ell_{cecom} = k_{cecom}$ and $\ell_{external} = 0$. Now by applying the Robust Extraction Lemma, we have that statistical distance between the views $\text{REAL}_{\mathcal{E},\Pi}^{(0,i:7)}(\beta, z)$ and $\text{view}_{\text{extr}}^{(0,i:7)}$ is at most:

$$\Delta(\lambda) \leq 2^{-(k_{cecom} - 0 \cdot \log(T(\lambda)))} \leq \text{negl}(\lambda),$$

since, $k_{cecom} \in \omega(\log(\lambda))$ and T is at most a polynomial.

Next, we describe an intermediate hybrid $\text{hyb}_{0,i:7,B}^*$ whose output is identical to that of $\text{hyb}_{0,i:7,A}^*$. For this consider an interface, $I_{\text{extr}}^{(0,i:7,B)}$ behaves the same way as $I_{\text{extr}}^{(0,i:7)}$ except that it also isolates the sWIAoK argument in Step bps_5 of the j -th right session. $\text{hyb}_{0,i:7,B}^*$ simply outputs the view output by the online extractor, namely, $\text{REAL}_{\mathcal{E},\Pi}^{(0,i:7,B)}(\beta, z)$, where the external protocol Π here is the sWIAoK argument in Step bps_5 of the j -th right session, and the external party running

the code of the sWIAoK verifier for the isolated sWIAoK argument. Again, we can invoke the Robust Extraction Lemma, for which we consider the following: the CECOM commitments, that are isolated and relayed to external CECOM receivers, have k_{cecom} -slots and the external protocol that the robust simulator is participating in, here, Π here is the sWIAoK protocol in Step bps_2 of the t -th right session, and the external party running the code of the sWIAoK verifier for the isolated sWIAoK argument. Thus, we have that $\ell_{\text{cecom}} = k_{\text{cecom}}$ and $\ell_{\text{external}} = k_{\text{swiaok}}$. Now by applying the Robust Extraction Lemma, we have that statistical distance between the views $\text{REAL}_{\mathcal{E}, \Pi}^{I_{\mathcal{E}, \Pi}^{\text{extr}}(0, i; 7, B)}(\beta, z)$ and the view of the \mathcal{M} during its interaction with \mathcal{A}_{CB}^* is at most:

$$\Delta(\lambda) \leq 2^{-(k_{\text{cecom}} - k_{\text{swiaok}} \cdot \log(T(\lambda)))} \leq \text{negl}(\lambda),$$

since, $k_{\text{cecom}} \in \omega(\log(\lambda))$, k_{swiaok} is a constant, and T is at most a polynomial.

Thus, we have that the adversary \mathcal{M} when run by \mathcal{A}_{CB}^* is statistically close to the view when run by $\mathcal{H}_{0, i; 7}$. Thus we have that \mathcal{A}_{CB}^* breaks binding of the CECOM_{sh} commitment with probability $\epsilon - \text{negl}(\lambda)$. □

\mathcal{H}_2 . \mathcal{H}_2 : \mathcal{H}_2 runs $\text{RobustSim}^{I_{\text{extr}}^{(2)}}(z)$, where we define $I_{\text{extr}}^{(2)}$ to be identical to $I_{\text{extr}}^{(1)}$, except for the following modification.

- $I_{\text{hyb}}^{(2)}$ also isolates CECOM_{sb} commitments of Step bps_1 of all left sessions and relays them to external CECOM receivers.

Claim 13.

$$\text{view}_{\text{extr}}^{(2)} \approx_s \text{view}_{\text{extr}}^{(1)}$$

Proof. Recall that the only sub-protocols isolated by $I_{\text{extr}}^{(1)}$ are certain CECOM commitments of the right sessions which are forwarded to external CECOM receivers. Also recall that the only difference we introduced as we moved from \mathcal{H}_1 to \mathcal{H}_2 is that $I_{\text{extr}}^{(2)}$ also isolated certain other CECOM commitments of the right sessions and forwarded them to external CECOM receivers.

With this, to argue statistical indistinguishability between the views output by these two hybrids, we invoke the Robust Extraction Lemma. To apply the Robust Extraction Lemma, we will first create an intermediate hybrid hyb_2^* . hyb_1^* simply outputs the view output by the online extractor, namely, $\text{REAL}_{\mathcal{E}, \Pi}^{I_{\mathcal{E}, \Pi}^{\text{real}}(2)}(\beta, z)$, where the external protocol Π here is the empty protocol. Since in hyb_1 , $I_{\text{real}}^{(1)}$ only isolates some CECOM commitments, we can invoke the Robust Extraction Lemma, for which we consider the following: the CECOM commitments, that are isolated and relayed to external CECOM receivers, have k_{cecom} -slots and the external protocol that the robust simulator is participating in, here, is the empty protocol. Thus, we have that $\ell_{\text{cecom}} = k_{\text{cecom}}$ and $\ell_{\text{external}} = 0$. Now by applying the Robust Extraction Lemma, we have that statistical distance between the views $\text{REAL}_{\mathcal{E}, \Pi}^{I_{\mathcal{E}, \Pi}^{\text{real}}(2)}(\beta, z)$ and $\text{view}^{(1)}$ output by \mathcal{H}_2^* and \mathcal{H}_1 , respectively, is at most:

$$\Delta(\lambda) \leq 2^{-(k_{\text{cecom}} - 0 \cdot \log(T(\lambda)))} \leq \text{negl}(\lambda),$$

since, $k_{\text{cecom}} \in \omega(\log(\lambda))$ and T is at most a polynomial. Thus,

$$\text{view}^{(2)} \approx_s \text{REAL}_{\mathcal{E}, \Pi}^{I_{\text{real}}^{(1)}}(\beta, z).$$

Furthermore, since $I_{\text{real}}^{(2)}$ also only isolates some CECOM commitments, we can invoke the Robust Extraction Lemma, just as before, to have that statistical distance between the views $\text{REAL}_{\mathcal{E}, \Pi}^{I_{\text{real}}^{(2)}}(\beta, z)$ and $\text{view}^{(2)}$ output by \mathcal{H}_2^* and \mathcal{H}_2 , respectively, is at most:

$$\Delta(\lambda) \leq 2^{-(k_{\text{cecom}} - 0 \cdot \log(T(\lambda)))} \leq \text{negl}(\lambda).$$

Thus,

$$\text{view}^{(2)} \approx_s \text{REAL}_{\mathcal{E}, \Pi}^{I_{\text{real}}^{(2)}}(\beta, z).$$

Thus, we have that $\text{view}_{\text{extr}}^{(2)} \approx_s \text{view}_{\text{extr}}^{(1)}$. □

\mathcal{H}_3 . \mathcal{H}_3 runs $\text{RobustSim}^{I_{\text{extr}}^{(3)}}(z)$, where we define $I_{\text{extr}}^{(3)}$ to be identical to $I_{\text{extr}}^{(2)}$, except for the following modification.

- Recall that $I_{\text{hyb}}^{(2)}$ (among other messages) isolated the CECom_{sb} commitments at Step bps_1 of the left session.
- Let value σ' be received from outside at the end of the CECom_{sb} commitment. Then commit to σ' using Com_{sh} at Step bps_2 ; also, use the same extracted value as the witness in proving sWIAoK at Step bps_2 .

Claim 14.

$$\text{view}_{\text{extr}}^{(3)} \approx_s \text{view}_{\text{extr}}^{(2)}.$$

Proof. Since Com_{sh} at Step $\text{bps}_2^{\text{cfp}}$ is a statistically hiding commitment scheme, and sWIAoK is statistically witness-indistinguishable, by applying the Robust Extraction Lemma exactly as before, the sub-claim follows. □

\mathcal{H}_4 . \mathcal{H}_4 runs $\text{RobustSim}^{I_{\text{extr}}^{(4)}}(z)$, where we define $I_{\text{extr}}^{(4)}$ to be identical to $I_{\text{extr}}^{(3)}$, except for the following modification.

- Recall that $I_{\text{hyb}}^{(3)}$ isolated the CECom_{sb} commitment at Step bps_1 of the left session. Let the extracted value received from the outside be σ' . Also, in Step bps_3 , let σ be the value that \mathcal{M} opens the CECom_{sb} commitment to. If $\sigma \neq \sigma'$, then abort.

Claim 15.

$$\text{view}_{\text{extr}}^{(4)} \approx_s \text{view}_{\text{extr}}^{(3)}.$$

Proof. Assume for contradiction that there exists $\ell \in [m_L]$ such that for the ℓ -th left session, $\sigma \neq \sigma'$, with some non-negligible probability ϵ . Then we construct an adversary that breaks CB of the CECom_{sb} commitment.

Recall that \mathcal{H}_4 isolates the CECom_{sb} commitment at Step $\text{bps}_{\text{cfp}_1}$ of the ℓ -th left session (among other messages). It thus receives an opening for it, say $(\sigma', \text{rand}_{\sigma'})$ from the outside. Furthermore, \mathcal{M} provides an opening to the same CECom_{sb} commitment at Step bps_3 ; call it $(\sigma, \text{rand}_\sigma)$. From the assumption in the proof that $\alpha \neq \alpha'$ with some non-negligible probability ϵ , we can construct an adversary that breaks CB of the CECom_{sb} commitment with the same probability. This clearly follows from the fact that the CECom_{sb} commitment in question is already isolated and forwarded to an external CECom receiver. Thus, \mathcal{H}_4 itself can be deemed our adversary against CB of the CECom_{sb} commitment, a contradiction. Thus, $\sigma = \sigma'$, with all but negligible probability, and the Claim follows. \square

\mathcal{H}_5 . \mathcal{H}_5 runs $\text{RobustSim}^{I_{\text{extr}}^{(5)}}(z)$, where we define $I_{\text{extr}}^{(5)}$ to be identical to $I_{\text{extr}}^{(4)}$, except for the following modification.

- Recall that, in the j -th left session in question, $I_{\text{hyb}}^{(4)}$ committed to a valid witness y_j in $\text{NMMXCom}_{\text{srs}}$ at Step bps_4 . Furthermore, it uses this committed value y_j as the witness in the CECom_{sh} commitment at Step bps_{4+} and also in proving sWIAoK at Step bps_5 . Here, the modification is that $I_{\text{hyb}}^{(5)}$ uses the commitment information from Com_{sh} at Step bps_2 where it committed to σ .

Claim 16.

$$\text{view}_{\text{extr}}^{(5)} \approx_s \text{view}_{\text{extr}}^{(4)}.$$

Proof. Note that we will no longer be able to directly apply the Robust Extraction Lemma as (one of) the external protocols in question here is a CECom commitment itself, while Robust Extraction Lemma can only be applied to external protocols of round complexity strictly less than that of the CECom commitments rewound by the robust simulator. However, we can still prove the claim by a simple hybrid argument: Firstly, the CECom_{sh} commitment is changed one sub-commitment at a time. We recall here that the the standard CECom_{sh} commitment from [PRS02] consists of multiple Com_{sh} sub-commitment each of which correspond to only three rounds. Since the external protocol in question now is just of three rounds, we apply the Robust Extraction Lemma exactly as before and get statistical indistinguishability. Secondly, sWIAoK at Step bps_5 is statistically witness-indistinguishable, again by applying the Robust Extraction Lemma, the sub-claim follows. \square

\mathcal{H}_6 . \mathcal{H}_6 runs $\text{RobustSim}^{I_{\text{extr}}^{(6)}}(z)$, where we define $I_{\text{extr}}^{(6)}$ to be identical to $I_{\text{extr}}^{(5)}$, except for the following modification.

- Recall that, in the j -th left session in question, $I_{\text{hyb}}^{(5)}$ committed to a valid witness y_j in $\text{NMMXCom}_{\text{srs}}$ at Step bps_4 . Here, the modification is that $I_{\text{hyb}}^{(6)}$ commits to 0^λ in $\text{NMMXCom}_{\text{srs}}$ at Step bps_4 .

Claim 17.

$$\text{view}_{\text{extr}}^{(6)} \approx_s \text{view}_{\text{extr}}^{(5)}.$$

Proof. By applying the Robust Extraction Lemma exactly as before, this immediately follows from the statistical hiding property of the NMCCom commitment $\text{NMMXCom}_{\text{srs}}$ at Step bps_4 of the left session. \square

Note that \mathcal{H}_6 does not use any witness for the left sessions and it outputs valid witness of the accepted right sessions. \square