

Statistical Concurrent Non-malleable Zero-knowledge from One-way Functions

Susumu Kiyoshima

NTT Secure Platform Laboratories, Japan
kiyoshima.susumu@lab.ntt.co.jp

June 29, 2015

Abstract

Concurrent non-malleable zero-knowledge (CNMZK) protocols are zero-knowledge protocols that are secure even against adversaries that interact with multiple provers and verifiers simultaneously. Recently, the first *statistical* CNMZK argument for \mathcal{NP} was constructed under the DDH assumption (Orlandi et al., TCC'14).

In this paper, we construct a statistical CNMZK argument for \mathcal{NP} assuming only the existence of one-way functions. The security is proven via black-box simulation, and the round complexity is $\text{poly}(n)$. Under the existence of collision-resistant hash functions, the round complexity can be reduced to $\omega(\log n)$, which is known to be essentially optimal for black-box concurrent zero-knowledge protocols.

1 Introduction

Zero-knowledge (ZK) proofs and arguments are protocols that enable the prover to convince the verifier of the correctness of a mathematical statement while providing *zero additional knowledge*. This “zero additional knowledge” property is formalized by using the *simulation paradigm*: An interactive proof or argument is said to be zero-knowledge if for any adversarial verifier there exists a *simulator* that can output a simulated view of the adversary. In the original definition of the ZK property, the adversary interacts with a single prover at a time. Thus, the original definition guarantees the ZK property in the stand-alone setting.

Non-malleable zero-knowledge (NMZK) [DDN00] and *concurrent zero-knowledge (CZK)* [DNS04] are security notions that guarantee the ZK property in the concurrent setting. Specifically, NMZK guarantees the ZK property in the setting where the adversary concurrently interacts with a dishonest prover in the *left session* and a honest verifier in the *right session*, and CZK guarantees the ZK property in the setting where the adversary concurrently interacts with unbounded number of honest provers.

As a security notion that implies both NMZK and CZK, Barak et al. [BPS06] proposed *concurrent non-malleable zero-knowledge (CNMZK)*. CNMZK guarantees the ZK property in the setting where the adversary concurrently interacts with many provers in the left sessions and many verifiers in the right sessions. In particular, it guarantees that receiving proofs in the left session does not help the adversary to give proofs in the right sessions—that is, it guarantees that if the adversary can prove some statements in the right sessions while receiving proofs in the left sessions, the adversary could prove the same statements even without receiving proofs in the left sessions. In the definition of CNMZK, this guarantee is formalized as the existence of a *simulator-extractor* that can simulate the adversary’s view in the left and right sessions while extracting witnesses from the adversary in the simulated right sessions.

The first CNMZK argument was constructed by Barak et al. [BPS06]. Subsequently, a computationally efficient construction was shown by Ostrovsky et al. [OPV10]. The first CNMZK *proof* was constructed by Lin et al. [LPTV10], and a variant of their protocol was shown to be secure with adaptively chosen inputs by Lin and Pass [LP11]. Additionally, a CNMZK argument that is secure with “fully” adaptively chosen inputs was recently constructed by Venkatasubramanian [Ven14].

Very recently, Orlandi et al. [OOR⁺14] constructed the first *statistical* CNMZK argument—that is, a CNMZK argument such that the view simulated by the simulator-extractor is statistically indistinguishable from the adversary’s view. Statistical CNMZK is clearly of great interest since it guarantees quite strong security in the concurrent setting. However, statistical CNMZK is hard to achieve, and the existing techniques of computational CNMZK protocols seem to be insufficient for constructing statistical CNMZK protocols (see Section 2.1).

An important open question on statistical CNMZK protocols is what hardness assumption is needed for constructing them. The statistical CNMZK argument of Orlandi et al. [OOR⁺14] was constructed under the DDH assumption (or the existence of dense cryptosystems); hence, we already know that statistical CNMZK protocols can be constructed under standard assumptions. However, since the existence of one-way functions is known to be sufficient for constructing statistical ZK protocols and computational CNMZK protocols [HNO⁺09, BPS06], it is important to study the following question.

Can we construct statistical concurrent non-malleable zero-knowledge protocols by assuming only the existence of one-way functions?

1.1 Our Result

In this paper, we answer the above question affirmatively.

Theorem 1. *Assume the existence of one-way functions. Then, there exists a statistical concurrent non-malleable zero-knowledge argument for \mathcal{NP} with round complexity $\text{poly}(n)$. Furthermore, if there exists a family of collision-resistant hash functions, the round complexity can be reduced to $\omega(\log n)$.*

The round complexity of our statistical CNMZK argument— $\text{poly}(n)$ rounds when only the existence of one-way functions is assumed and $\omega(\log n)$ rounds when the existence of a family of collision-resistant hash functions is assumed—is the same as the round complexity of the known statistical CZK arguments [GMOS07]. Thus, our result closes the gap between statistical CNMZK arguments and statistical CZK arguments. Furthermore, since the security of our statistical CNMZK protocol is proven via black-box simulation, the logarithmic round complexity of our hash-function-based protocol is essentially tight due to the lower bound on black-box CZK protocols [CKPR02].

2 Techniques

2.1 Previous Techniques

Before explaining our technique, we explain the difficulty of constructing statistical CNMZK protocols with the techniques of existing computational CNMZK protocols [BPS06, LPTV10].

We first recall the protocols of [BPS06, LPTV10]. The definition of CNMZK requires the existence of a simulator-extractor that simulates the adversary’s view while extracting the witnesses for the statements proven by the adversary in the simulated view. To satisfy this definition, protocols need to satisfy the following properties: (i) the proofs in the left sessions can be simulated for the adversary, and (ii) even when the adversary receives simulated proofs in the left sessions, the witnesses can be extracted from the adversary in the right sessions. In the protocol of [BPS06, LPTV10], the simulatability of the left sessions is guaranteed by requiring the verifier to commit to a random trapdoor by using a *concurrently extractable commitment scheme* CCom [MOSV06]. Since the committed values of CCom can be extracted by a rewinding extractor even in the concurrent setting, the proofs in the left sessions can be simulated by extracting the trapdoors from CCom. On the other hand, the witness-extractability of the right sessions is guaranteed by requiring the prover to commit to the witness with a non-malleable commitment scheme NCom [DDN00] and additionally designing the protocols so that the following hold.

1. When the adversary receives honest proofs in the left sessions, the committed value of the NCom commitment is indeed a valid witness in every accepted right session.
2. When the proofs in the left sessions are switched to the simulated ones, the committed values of the NCom commitments do not change in the right sessions due to the non-malleability of NCom.

It follows from these that even when the adversary receives simulated proofs in the left sessions, the committed value of the NCom commitment is a witness for the statement in every accepted right session. Therefore, the witnesses can be extracted in the right sessions by extracting the committed values of the NCom commitments.

As mentioned above, the techniques of [BPS06, LPTV10] alone seem to be insufficient for constructing statistical CNMZK protocols. The main problem is that the techniques of [BPS06, LPTV10]

requires the prover to commit to the witness by using NMCom , which is only computationally hiding.¹ Since in the simulation the committed values of NMCom need to be switched to another values (e.g., 0^n) in the left sessions, the simulated view can be only computational indistinguishable from the real view.

Recently, Orlandi et al. [OOR⁺14] constructed a statistical CNMZK protocol by modifying the CNMZK protocol of [BPS06] with *mixed non-malleable commitment scheme* MXNMCom . MXNMCom is parametrized by a string and is either statistically hiding or non-malleable depending on the string.² Very roughly speaking, Orlandi et al. circumvent the above problem by carefully switching the parameter string of MXNMCom in the security proof—when proving the statistical indistinguishability of the simulation, the string is set so that MXNMCom is statistically hiding, and when proving the non-malleability, the string is set so that MXNMCom is non-malleable. The use of MXNMCom , however, requires assumptions that are stronger than the existence of one-way functions (such as the DDH assumption or the existence of dense cryptosystems). Thus, the technique of Orlandi et al. cannot be used to construct statistical CNMZK protocols from one-way functions.

2.2 Our Technique

Since the reason why the techniques of [BPS06, LPTV10] cannot be used for statistical CNMZK protocols is that the committed values of NMCom need to be switched during the simulation, one potential strategy for constructing statistical CNMZK is to construct a protocol such that the adversary’s view can be simulated without switching the committed value of NMCom (and of any other computationally hiding commitment). However, when the simulator commits to the same value in NMCom as a honest prover, it is not clear how non-malleability of NMCom can be used in the security proof. Below, we show that the CNMZK property can be shown even in this case if we use a stronger variant of NMCom .

A key technical tool in our technique is *CCA-secure commitment schemes* [CLP10], which is a stronger variant of (concurrent) non-malleable commitment schemes. Roughly speaking, CCA security guarantees that the scheme is hiding even against adversaries that have access to the *committed-value oracle*, which receives concurrent commitments from the adversary and returns their committed values to the adversary. (In non-malleability, the oracle receives only parallel commitments from the adversary and returns the committed values only after the adversary finishes the interaction with the committer.) Several CCA-secure commitment schemes were constructed from one-way functions [CLP10, LP12, Kiy14, GLP⁺15]; furthermore, although CCA security itself does not provide any extractability, all of these schemes satisfy concurrent extractability as well.

Using CCA-secure commitment schemes, we consider the following protocol as a starting point.

Stage 1. (V commits to trapdoor)

1. The verifier V chooses random $r_V \in \{0, 1\}^n$ and commits to r_V by using a statistically binding commitment scheme Com , which can be constructed from one-way functions [Nao91, HILL99]. Let (r_V, d) be the decommitment.
2. V commits to (r_V, d) by using CCA-CECom , where CCA-CECom is a CCA-secure commitment scheme that is also concurrent extractable [CLP10, LP12, Kiy14, GLP⁺15].

¹ NMCom need to be *non-malleable w.r.t. commitment* [DDN00], which roughly says that the committed value of the commitment that the man-in-the-middle adversary gives is independent of the committed value of the commitment that adversary receives. Since the definition of non-malleability w.r.t. commitment is meaningless when the committed values cannot be uniquely determined, NMCom cannot be statistically hiding.

²Specifically, Orlandi et al. [OOR⁺14] used the scheme such that (i) when the string is sampled from a uniform distribution, the scheme is statistically hiding and (ii) when the string is taken from another (computationally indistinguishable) distribution, the scheme is non-malleable.

Stage 2. (P proves $x \in L$ or knowledge of trapdoor) The prover P proves that it knows a witness for $x \in L$ or a valid decommitment (r_V, d) of the Com commitment that V gives in Stage 1. P proves this statement by using a statistical witness-indistinguishable argument of knowledge sWIAOK, which can be constructed from one-way functions by instantiating Blum’s Hamiltonian-cycle protocol with the statistically hiding commitment scheme of [HNO⁺09].

In this protocol, the verifier’s view can be statistically simulated by a simulator that extracts (r_V, d) from CCA-CECom and uses it as a witness in sWIAOK. (Note that this simulator executes Stage 1 honestly; thus, even if computationally hiding commitment schemes are used as building blocks in CCA-CECom, the simulator commits to the same values as a honest prover in these schemes.) Also, intuitively this protocol seems to be CNMZK from the following reasons.

- The CCA security of CCA-CECom guarantees that the trapdoors of the right sessions are hidden from the adversary even when the trapdoors of the left sessions are extracted and returned to the adversary.
- Then, since the simulated proofs are generated in the left sessions by extracting the trapdoors, the trapdoors in the right sessions are hidden from the adversary even when the adversary receives simulated proofs in the left sessions.
- Thus, even when the adversary receives the simulated proofs in the left sessions, the adversary cannot “cheat” in the right sessions, and therefore witnesses for the statements must be extractable from sWIAOK in the right sessions.

Of course, to formally show the statistical CNMZK property, we need to show a simulator-extractor that statistically simulates the adversary’s view and also extracts witnesses for the statements in the right sessions.

As the simulator-extractor, we consider the following \mathcal{SE} .

1. First, \mathcal{SE} simulates the view of the adversary \mathcal{A} by executing the following simulator \mathcal{S} : Simulator \mathcal{S} internally invokes \mathcal{A} and interacts with it in the left and right sessions honestly except that in each left session, \mathcal{S} extracts (r_V, d) by using the concurrent extractor of CCA-CECom and uses it as a witness in sWIAOK.
2. After simulating the view of \mathcal{A} as above, \mathcal{SE} extracts witnesses from the right sessions by doing the following for each right session. First, \mathcal{SE} rewinds \mathcal{S} until the point just before \mathcal{S} sends the challenge message of sWIAOK to \mathcal{A} .³ Then, \mathcal{SE} repeatedly executes \mathcal{S} from this point with fresh randomness until it obtains another accepted transcript of sWIAOK. After obtaining another accepted transcript, \mathcal{SE} extracts a witness by using the argument-of-knowledge property of sWIAOK.

It is easy to see that \mathcal{SE} statistically simulates the real view of \mathcal{A} . Thus, it remains to show that \mathcal{SE} extracts witnesses for the statements in the right sessions.

To show the witness extractability of \mathcal{SE} , a natural approach is to follow the above-mentioned approach of [BPS06, LPTV10] and show the following.

1. When \mathcal{A} receives honest proofs in the left sessions, a witness for the statement is extracted from the sWIAOK proof in every accepted right session.
2. When the honest proofs in the left sessions are switched to the simulated ones, the value extracted from sWIAOK does not change in every accepted right session.

³Since \mathcal{S} rewinds \mathcal{A} during the concurrent extraction of CCA-CECom, \mathcal{S} may send the challenge message of sWIAOK of a right session to \mathcal{A} multiple times. Here, \mathcal{SE} rewinds \mathcal{S} until the point just before \mathcal{S} sends it to \mathcal{A} on the “main thread.”

Note that here we argue about the extracted values instead of the committed values. At first sight, it seems that this is not a big difference and it seems that the above can be shown by using an argument similar to the one used in [BPS06, LPTV10].

However, this approach does not work. In particular, we cannot show the second part—that is, we cannot show that the extracted values remain to be the same when the honest proofs in the left sessions are switched to the simulated ones. To see this, observe the following. Since the witnesses used in sWIAOK are switched in the simulated proofs, we need to use the witness indistinguishability of sWIAOK of the left sessions to show the indistinguishability of the extracted values. However, since \mathcal{A} is rewound during the witness extraction of the sWIAOK proofs of the right sessions, if the left and the right sessions are scheduled so that the sWIAOK proofs of the left sessions are executed in parallel with the sWIAOK proofs of the right sessions, the sWIAOK proofs of the left sessions are also rewound, and thus we cannot use their witness indistinguishability.⁴

Thus, we instead use the following approach. Informally, the above approach does not work because the honest proofs and the simulated proofs are “too different.” We thus introduce a hybrid experiment in which \mathcal{A} receives *hybrid proofs* in the left sessions, where a hybrid proof is generated by extracting (r_V, d) by brute force and using it as a witness in sWIAOK . (Notice that the only difference between the hybrid proofs and the simulated proofs is how the trapdoors are extracted.) We then show that (i) witnesses for the statements are extracted in the right sessions when \mathcal{A} receives hybrid proofs in the left sessions, and (ii) when hybrid proofs are switched to the simulated ones, the extracted values do not change. More precisely, our analysis proceeds as follows.

- First, we show the second part, i.e., we show that the values extracted in the right sessions do not change when the proofs in the left sessions are switched from the hybrid proofs to the simulated ones. Since the only difference between the hybrid proofs and the simulated ones is how the committed values of the CCA-CECom commitments are extracted (by brute-force or by the concurrent extractability), we can show this by using the concurrent extractability of CCA-CECom . We note however that there is a subtlety since CCA-CECom in the left sessions can be rewound not only by the concurrent extractor of CCA-CECom but also by the extractor of sWIAOK . Nonetheless, by carefully using a standard technique (the “good prefix” argument), we can show that the concurrent extractor of CCA-CECom works even in this case.
- Next, we show that in the hybrid experiment, witnesses for the statements are extracted from the right sessions. Since the simulated proofs can be efficiently generated given access to the committed-value oracle of CCA-CECom , at first sight it seems that this follows directly from the CCA security of CCA-CECom and argument-of-knowledge property of sWIAOK —if a witness for the statement is not extracted, (r_V, d) must be extracted, and thus we can break the CCA security of CCA-CECom . However, there are two problems.
 1. Since CCA-CECom in the left sessions can be rewound during the witness extraction of sWIAOK of the right sessions, the hybrid experiment cannot be emulated even given access to the committed-value oracle of CCA-CECom . Hence, the CCA -secure commitments in the right sessions may not be hiding in the hybrid experiment.
 2. Since the adversary obtains hybrid proofs, which are generated in super-polynomial time, the argument-of-knowledge property of sWIAOK may not hold in the hybrid experiment.

⁴If we use the robust extraction technique [GLP⁺15], for each left session there exists a rewinding strategy that allows us to extract witnesses from the right sessions without rewinding sWIAOK of this left session. However, since what we want to show is that the values extracted in the right sessions *by the rewinding strategy that \mathcal{SE} uses* are unchanged, the robust extraction technique cannot be used here (unless there exists a rewinding strategy that allows us to extract witnesses from the right sessions without rewinding the sWIAOK proof of *every* left session).

We remark that although existing CCA-secure commitment schemes provides *robustness*, which guarantees that arbitrary “small”-round protocol remains secure even when adversaries have access to the committed-value oracle, we cannot use robustness here since CCA-CECom in the left sessions can be rewound during the witness extraction of sWIAOK of the right sessions and therefore the hybrid experiment cannot be emulated even given access to the committed-value oracle.

Because of these problems, we cannot use the security of CCA-CECom directly in the analysis. Thus, instead of using existing CCA-secure commitment schemes in a modular way, we directly use their building blocks in the protocol and directly use their proof technique in the analysis. (In particular, we use the robust concurrent extraction technique of [GLP⁺15] and a one-one CCA-secure commitment scheme of [KMO14].) The proof techniques of existing CCA-secure commitment schemes are strong enough to solve the above problems, and thus we can show that witnesses for the statements are extracted in the hybrid experiment.

From the above two, it follows that even when \mathcal{A} receives simulated proofs in the left session, valid witnesses are extracted in right sessions. This completes the overview of our technique.

3 Preliminaries

3.1 Notations

We use n to denote the security parameter. For any $k \in \mathbb{N}$, let $[k] \stackrel{\text{def}}{=} \{1, \dots, k\}$.

3.2 Commitment Schemes

Recall that commitment schemes are two-party protocols between the committer C and the receiver R . A transcript of the commit phase is *valid* if there exists a valid decommitment of this transcript.

It is known that a two-round statistically binding commitment scheme Com_{SB} can be constructed from one-way functions [Nao91, HILL99]. It is also known that a $\text{poly}(n)$ -round statistically hiding commitment scheme can be constructed from one-way functions [HNO⁺09], and a constant-round one can be constructed from a family of collision-resistant hash functions [NY89, DPP98].

3.3 Concurrently Extractable Commitment Schemes

Roughly speaking, a commitment scheme is *concurrently extractable* if there exists a polynomial-time extractor such that for any adversarial committer that concurrently commits to many values by using the scheme, the extractor can extract the committed value from the adversarial committer in every valid commitment.

Micciancio et al. [MOSV06] proposed a $\omega(\log n)$ -round concurrently extractable commitment CECom (Figure 1), which is an abstraction of the preamble stage of the concurrent zero-knowledge protocol of [PRS02] and can be constructed from one-way functions. The extractor of CECom performs the extraction by rewinding the adversarial committer according to the rewinding strategy of [PRS02, PTV12]. Specifically, while interacting with the adversarial committer C^* on the “main thread” as honest receivers, the extractor rewinds the main thread and generates many “look-ahead threads” on which it interacts with C^* again as honest receivers with fresh randomness; at the end of each commitment on each thread, the extractor extracts the committed values by using the information collected on the other threads.

CECom can be seen as a concurrent execution of the extractable commitment scheme ExtCom in Figure 2, which consists of three stages—`commit`, `challenge`, and `reply`—and can be constructed from one-way functions.

Commit Phase

The committer C and the receiver R receive common input 1^n and parameter ℓ . (In [MOSV06], $\ell = \omega(\log n)$.) To commit to $v \in \{0, 1\}^n$, the committer C commits to v concurrently ℓ times by using ExtCom as follows.

1. C and R execute `commit` stage of ExtCom ℓ times in parallel.
2. For each $j \in [\ell]$ in sequence, C and R do the following.
 - (a) R sends the `challenge` message of ExtCom for the j -th session.
 - (b) C sends the `reply` message of ExtCom for the j -th session.

Decommit Phase

C sends v to R and decommits all the ExtCom commitments.

Figure 1: Concurrently extractable commitment CECom [MOSV06].

Commit Phase

The committer C and the receiver R receive common inputs 1^n . To commit to $v \in \{0, 1\}^n$, the committer C does the following with the receiver R .

`commit` stage.

For each $i \in [n]$, the committer C chooses a pair of random n -bit strings (a_i^0, a_i^1) such that $a_i^0 \oplus a_i^1 = v$. Then, for each $i \in [n]$ in parallel, C commits to a_i^0 and a_i^1 by using Com_{SB} . For each $i \in [n]$ and $b \in \{0, 1\}$, let c_i^b be the commitment to a_i^b .

`challenge` stage.

R sends random n -bit string $e = (e_1, \dots, e_n)$ to C .

`reply` stage.

For each $i \in [n]$, C decommits $c_i^{e_i}$ to $a_i^{e_i}$.

Decommit Phase

C sends v to R and decommits c_i^b to a_i^b for all $i \in [n]$ and $b \in \{0, 1\}$. R checks whether $a_1^0 \oplus a_1^1 = \dots = a_n^0 \oplus a_n^1 = v$.

Figure 2: Extractable commitment ExtCom [PW09], which is used as a building block in the concurrently extractable commitment scheme CECom of [MOSV06].

Robust Concurrent Extraction.

On the concurrently extractable commitment scheme CECOM of [MOSV06], Goyal et al. [GLP⁺15] showed a very useful lemma called the *robust concurrent extraction lemma*. Roughly speaking, this lemma states that even when the adversarial committer additionally participates in an external protocol, the committed values can be extracted from the adversarial committer *without rewinding the external protocol* as long as the round complexity of the external protocol is “small.” More precisely, consider any PPT adversarial committer \mathcal{A} that commits to multiple values in concurrent sessions of CECOM—these sessions are denoted as the *right sessions*—and simultaneously participates in an execution of an arbitrary protocol $\Pi := \langle B, A \rangle$ with a honest B —this session is denoted as the *left session*. The robust concurrent extraction lemma states that for every \mathcal{A} , there exists an extractor E that extracts the committed values from \mathcal{A} in every valid right session without rewinding the external party B in the left session. The extractor E fails with probability that is exponentially small in $\ell - O(k \log n)$, where ℓ is the parameter of CECOM and k is the round complexity of Π ; hence, E fails only with negligible probability if we set $\ell := \omega(k \log n)$. The formal statement of the robust concurrent extraction lemma is given in Appendix A. We remark that the extractor E shown by [GLP⁺15] performs the extraction by generating the main thread and the look-ahead threads as in the rewinding strategies of [PRS02, PTV12].

3.4 (One-one) CCA-secure Commitment Schemes

We recall the definition of (one-one) CCA security and κ -robustness of commitment schemes [CLP10, LP12, KMO14].

(One-one) CCA security. Roughly speaking, a tag-based commitment scheme $\langle C, R \rangle$ (i.e., a commitment scheme that takes an n -bit string—a *tag*—as an additional input) is *CCA-secure* if it is hiding even against adversary \mathcal{A} that interacts with the following *committed-value oracle*: The committed-value oracle \mathcal{O} interacts with \mathcal{A} as an honest receiver in many concurrent sessions of the commit phase of $\langle C, R \rangle$ using tags chosen adaptively by \mathcal{A} ; at the end of each session, if the commitment of this session is invalid or has multiple committed values, \mathcal{O} returns \perp to \mathcal{A} ; otherwise, \mathcal{O} returns the unique committed value to \mathcal{A} .

More precisely, CCA-secure commitment scheme is defined as follows. Consider the following probabilistic experiment $\text{IND}_b(\langle C, R \rangle, \mathcal{A}, n, z)$ for each $b \in \{0, 1\}$. On input 1^n and auxiliary input z , adversary $\mathcal{A}^{\mathcal{O}}$ adaptively chooses a pair of challenge values $v_0, v_1 \in \{0, 1\}^n$ and an n -bit tag $\text{id} \in \{0, 1\}^n$. Then, $\mathcal{A}^{\mathcal{O}}$ interacts with the challenger and obtains a commitment to v_b with tag id . Let y be the output of \mathcal{A} . The output of the experiment is \perp if during the experiment, \mathcal{A} sends \mathcal{O} any commitment using tag id . Otherwise, the output of the experiment is y . Let $\text{IND}_b(\langle C, R \rangle, \mathcal{A}, n, z)$ denote the output of experiment $\text{IND}_b(\langle C, R \rangle, \mathcal{A}, n, z)$.

Definition 1. Let $\langle C, R \rangle$ be a tag-based commitment scheme and \mathcal{O} be the committed-value oracle of $\langle C, R \rangle$. Then, $\langle C, R \rangle$ is **CCA-secure (w.r.t the committed-value oracle)** if for any PPT adversary \mathcal{A} , the following are computationally indistinguishable:

- $\{\text{IND}_0(\langle C, R \rangle, \mathcal{A}, n, z)\}_{n \in \mathbb{N}, z \in \{0, 1\}^*}$
- $\{\text{IND}_1(\langle C, R \rangle, \mathcal{A}, n, z)\}_{n \in \mathbb{N}, z \in \{0, 1\}^*}$ ◇

If $\langle C, R \rangle$ is CCA secure only against adversaries that start a single session with \mathcal{O} , we say that $\langle C, R \rangle$ is *one-one CCA secure*. That is, one-one CCA security is defined as follows. Let *one-session committed-value oracle* be an oracle that is the same as the committed-value oracle except that it interacts with the adversary only in a single session of the commit phase of $\langle C, R \rangle$. Then, one-one

CCA security is defined by replacing the committed-value oracle in the definition of CCA security with the one-session committed-value oracle.

Robustness. Roughly speaking, a tag-based commitment scheme is κ -robust if for any adversary \mathcal{A} and any ITM B , the joint output of a κ -round interaction between \mathcal{A}^O and B can be simulated without O by a PPT simulator.

Definition 2. Let $\langle C, R \rangle$ be a tag-based commitment scheme and O be the committed-value oracle of $\langle C, R \rangle$. For any constant $\kappa \in \mathbb{N}$, we say that $\langle C, R \rangle$ is κ -robust (w.r.t. the committed-value oracle) if there exists a PPT oracle machine (called **simulator**) \mathcal{S} such that for any PPT adversary \mathcal{A} and any κ -round PPT ITM B , the following are computationally indistinguishable:

- $\left\{ \text{out}_{B, \mathcal{A}^O} \left\langle B(1^n, x, y) \leftrightarrow \mathcal{A}^O(1^n, x, z) \right\rangle \right\}_{n \in \mathbb{N}, x, y, z \in \{0, 1\}^n}$
- $\left\{ \text{out}_{B, \mathcal{S}^{\mathcal{A}}} \left\langle B(1^n, x, y) \leftrightarrow \mathcal{S}^{\mathcal{A}}(1^n, x, z) \right\rangle \right\}_{n \in \mathbb{N}, x, y, z \in \{0, 1\}^n}$

Here, for any ITM A and B , we use $\text{out}_{A, B} \langle A(x, y) \leftrightarrow B(x, z) \rangle$ to denote the joint output of A and B in an interaction between them on inputs x, y to A and x, z to B respectively. \diamond

Intuitively, the κ -robustness guarantees that the security of any κ -round protocol (say, the hiding property of a κ -round commitment scheme) holds even against the adversary that interacts with O . In fact, it is easy to see that the following proposition holds.

Proposition 1. Let $\langle C, R \rangle$ be a κ -robust commitment scheme for a constant $\kappa \in \mathbb{N}$, and let B be any κ -round PPT ITM. Then, for every two sequences $\{y_n^1\}_{n \in \mathbb{N}}$ and $\{y_n^2\}_{n \in \mathbb{N}}$ such that for every PPT adversary \mathcal{A}' it holds that

- $\left\{ \text{out}_{B, \mathcal{A}'} \left\langle B(1^n, x, y_n^1) \leftrightarrow \mathcal{A}'(1^n, x, z) \right\rangle \right\}_{n \in \mathbb{N}, x, y, z \in \{0, 1\}^n}$ and
- $\left\{ \text{out}_{B, \mathcal{A}'} \left\langle B(1^n, x, y_n^2) \leftrightarrow \mathcal{A}'(1^n, x, z) \right\rangle \right\}_{n \in \mathbb{N}, x, y, z \in \{0, 1\}^n}$

are computationally indistinguishable, then it also holds that for every PPT adversary \mathcal{A} ,

- $\left\{ \text{out}_{B, \mathcal{A}^O} \left\langle B(1^n, x, y_n^1) \leftrightarrow \mathcal{A}^O(1^n, x, z) \right\rangle \right\}_{n \in \mathbb{N}, x, y, z \in \{0, 1\}^n}$ and
- $\left\{ \text{out}_{B, \mathcal{A}^O} \left\langle B(1^n, x, y_n^2) \leftrightarrow \mathcal{A}^O(1^n, x, z) \right\rangle \right\}_{n \in \mathbb{N}, x, y, z \in \{0, 1\}^n}$

are computationally indistinguishable.

The scheme we use. From a result shown in [GLP⁺15], we can obtain a constant-round κ -robust one-one CCA-secure commitment scheme for every constant $\kappa \in \mathbb{N}$ from one-way functions. In [GLP⁺15], Goyal et al. constructed a $\omega(\log n)$ -round CCA-secure commitment scheme from one-way functions. This scheme has $\omega(\log n)$ rounds because CECOM with parameter $\ell = \omega(\log n)$ is used as a building block. The reason why ℓ is set to be $\omega(\log n)$ is that in the security analysis, the committed values of CECOM need to be extracted when polynomially many CECOM commitments are concurrently executed. In the setting of *one-one* CCA security, however, the security analysis works even if the committed values of CECOM are extractable only when a single CECOM commitment is executed. Hence, by setting $\ell := O(1)$, we can obtain a constant-round one-one CCA-secure commitment scheme. For completeness, we give the protocol and the proof of one-one CCA security in Appendix B.

3.5 Witness Indistinguishable Proofs and Arguments, and Special Soundness

Definition 3. An interactive proof (or argument) system $\langle P, V \rangle$ for an \mathcal{NP} language L with witness relation \mathbf{R}_L is said to be **witness indistinguishable** if for every probabilistic polynomial-time adversarial verifier V^* and for every two sequences $\{w_x^1\}_{x \in L}$ and $\{w_x^2\}_{x \in L}$ such that $w_x^1, w_x^2 \in \mathbf{R}_L(x)$, the following ensembles are computationally indistinguishable.

- $\left\{ \text{view}_{V^*} \left\langle P(x, w_x^1) \leftrightarrow V^*(x) \right\rangle \right\}_{x \in L}$
- $\left\{ \text{view}_{V^*} \left\langle P(x, w_x^2) \leftrightarrow V^*(x) \right\rangle \right\}_{x \in L}$

Here, $\text{view}_{V^*} \langle P(x, w) \leftrightarrow V^*(x) \rangle$ denotes the view of $V^*(x)$ in the interaction with $P(x, w)$.

If the above ensembles are statistically indistinguishable, $\langle P, V \rangle$ is said to be **statistically witness indistinguishable**. \diamond

We can obtain a four-round witness-indistinguishable proof system WIPOK from one-way functions by executing Blum’s Hamiltonian-cycle protocol in parallel. Recall that WIPOK consists of three stages—commit, challenge, and response—and has a useful property called *special soundness*: Given two accepting transcripts $\langle \alpha_1, \alpha'_1, \beta_1, \gamma_1 \rangle$ and $\langle \alpha_2, \alpha'_2, \beta_2, \gamma_2 \rangle$ of statement $x \in L$ such that $\alpha_1 = \alpha_2$, $\alpha'_1 = \alpha'_2$, and $\beta_1 \neq \beta_2$, we can compute a valid witness for $x \in L$.

We can obtain a statistical witness-indistinguishable argument system sWIAOK from any statistically hiding commitment scheme by instantiating Blum’s Hamiltonian-cycle protocol with the statistically hiding commitment scheme. It is easy to see that this argument system satisfies special soundness in the following sense: Let us say that two accepting transcripts $\langle \vec{\alpha}_1, \beta_1, \gamma_1 \rangle$ and $\langle \vec{\alpha}_2, \beta_2, \gamma_2 \rangle$ are *admissible* if $\vec{\alpha}_1 = \vec{\alpha}_2$ and $\beta_1 \neq \beta_2$; then, given admissible transcripts that are generated in polynomial time, we can compute a valid witness. In particular, given admissible transcripts $\langle \vec{\alpha}, \beta_1, \gamma_1 \rangle$ and $\langle \vec{\alpha}, \beta_2, \gamma_2 \rangle$, we can compute a valid witness or we can decommit a commitment given in $\vec{\alpha}$ to two different values.

3.6 Statistical Concurrent Non-malleable Zero-knowledge Arguments

We recall the definition of (statistical) concurrent non-malleable zero-knowledge from [BPS06, OOR⁺14], which is closely related to the definition of simulation extractability of [PR05]. Let $\langle P, V \rangle$ be an interactive argument system for a language $L \in \mathcal{NP}$ with witness relation \mathbf{R}_L . For any man-in-the-middle adversary \mathcal{A} , let us consider a probabilistic experiment in which \mathcal{A} participates in the following left and right interactions (see Figure 3). In the left interaction, \mathcal{A} interacts with a honest prover P of $\langle P, V \rangle$ and verifies the validity of statements x_1, \dots, x_m using identities $\text{id}_1, \dots, \text{id}_m$. In the right interaction, \mathcal{A} interacts with a honest verifier V of $\langle P, V \rangle$ and proves the validity of statements $\tilde{x}_1, \dots, \tilde{x}_m$ using identities $\tilde{\text{id}}_1, \dots, \tilde{\text{id}}_m$. The statements proven in the left interaction, x_1, \dots, x_m , are given to P and \mathcal{A} prior to the experiment. In contrast, the statements proven in the right interaction, $\tilde{x}_1, \dots, \tilde{x}_m$, and the identities used in the left and the right interactions, $\text{id}_1, \dots, \text{id}_m$ and $\tilde{\text{id}}_1, \dots, \tilde{\text{id}}_m$, are chosen by \mathcal{A} during the experiment. Then, roughly speaking, $\langle P, V \rangle$ is *statistical concurrent non-malleable zero-knowledge* (statistical CNMZK) if for any adversary \mathcal{A} , there exists a ppt machine called the *simulator-extractor* that can statistically simulate the view of \mathcal{A} in the above experiment while extracting witnesses for the statements proven by \mathcal{A} in the accepted right interactions that use different identities from the left interactions. The formal definition is given below.

Definition 4. An interactive proof $\langle P, V \rangle$ for language L with witness relation \mathbf{R}_L is said to be **statistical concurrent non-malleable zero-knowledge** if for every polynomial $m(\cdot)$ and every probabilistic polynomial-time man-in-the-middle adversary \mathcal{A} that participates in at most $m = m(n)$ concurrent executions, there exists a probabilistic polynomial-time machine \mathcal{SE} called **simulator-extractor** such that the following hold.

1. Let $\text{sim-view}(n, x_1, \dots, x_m, z)$ be a random variable representing the first output of $\mathcal{SE}(n, x_1, \dots, x_m, z)$. Then, the following ensembles are statistically indistinguishable.
 - $\{\text{view}_{\mathcal{A}}(n, x_1, \dots, x_m, z)\}_{n \in \mathbb{N}, x_1, \dots, x_m \in L \cap \{0,1\}^n, z \in \{0,1\}^*}$
 - $\{\text{sim-view}(n, x_1, \dots, x_m, z)\}_{n \in \mathbb{N}, x_1, \dots, x_m \in L \cap \{0,1\}^n, z \in \{0,1\}^*}$
2. For every $n \in \mathbb{N}$, $x_1, \dots, x_m \in L \cap \{0,1\}^n$, and $z \in \{0,1\}^*$, the following holds. Let $(\text{view}, \{\widetilde{w}_i\}_{i \in [m]})$ denote the output of $\mathcal{SE}(n, x_1, \dots, x_m, z)$. Let $\widetilde{x}_1, \dots, \widetilde{x}_m$ be the statements of the right interaction in view , and let $\text{id}_1, \dots, \text{id}_m$ and $\widetilde{\text{id}}_1, \dots, \widetilde{\text{id}}_m$ be the identities of the left and the right interactions in view , respectively. Then, except with negligible probability, we have $(\widetilde{x}_i, \widetilde{w}_i) \in \mathbf{R}_L$ for every $i \in [m]$ if the i -th right interaction is accepting and $\text{id}_i \neq \text{id}_j$ for every $j \in [m]$. \diamond

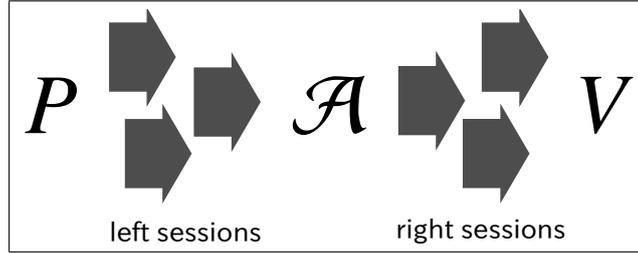


Figure 3: left sessions and right sessions.

4 Our Statistical Concurrent Non-malleable ZK Argument

We show that a statistical concurrent non-malleable zero-knowledge argument can be constructed from any statistically hiding commitment scheme.

Theorem 2. *Assume the existence of statistically hiding commitment schemes with round complexity $R_{\text{SH}}(n)$. Then, there exists an $\omega(R_{\text{SH}}(n) \log n)$ -round statistical concurrent non-malleable zero-knowledge argument sCNMZK .*

Since $\text{poly}(n)$ -round statistically hiding commitment schemes can be constructed from one-way functions [HNO⁺09] and constant-round ones can be constructed from a family of collision-resistant hash functions [NY89, DPP98], our main theorem (Theorem 1) follows from Theorem 2.

Proof of Theorem 2. In sCNMZK , we use the following building blocks.

- Two-round statistically binding commitment scheme Com_{SB} .
- Constant-round 4-robust one-one CCA-secure commitment scheme $\text{CCACom}^{1:1}$.
- Four-round witness-indistinguishable proof of knowledge WIPOK , which is a parallel version of Blum’s Hamiltonian-cycle protocol.
- $(R_{\text{SH}}(n)+2)$ -round statistical witness-indistinguishable argument of knowledge sWIAOK , which is a parallel version of Blum’s Hamiltonian-cycle protocol that is instantiated with a $R_{\text{SH}}(n)$ -round statistically hiding commitment scheme Com_{SH} .
- $\omega(R_{\text{SH}}(n) \log n)$ -round concurrently extractable commitment scheme CECom , which is the scheme of [MOSV06] with parameter $\ell = \omega(R_{\text{SH}}(n) \log n)$. From the robust concurrent extraction lemma [GLP⁺15], we can extract the committed values from any adversarial committer even when it additionally participates in any $O(R_{\text{SH}}(n))$ -round external protocol.

As explained in Section 3, all of the above building blocks can be constructed from $R_{\text{SH}}(n)$ -round statistically hiding commitment schemes (or from one-way functions, which can be obtained from statistically hiding commitment schemes).

Protocol sCNMZK is shown in Figure 4. We prove its soundness in Section 4.1 and prove its statistical CNMZK property in Section 4.2.

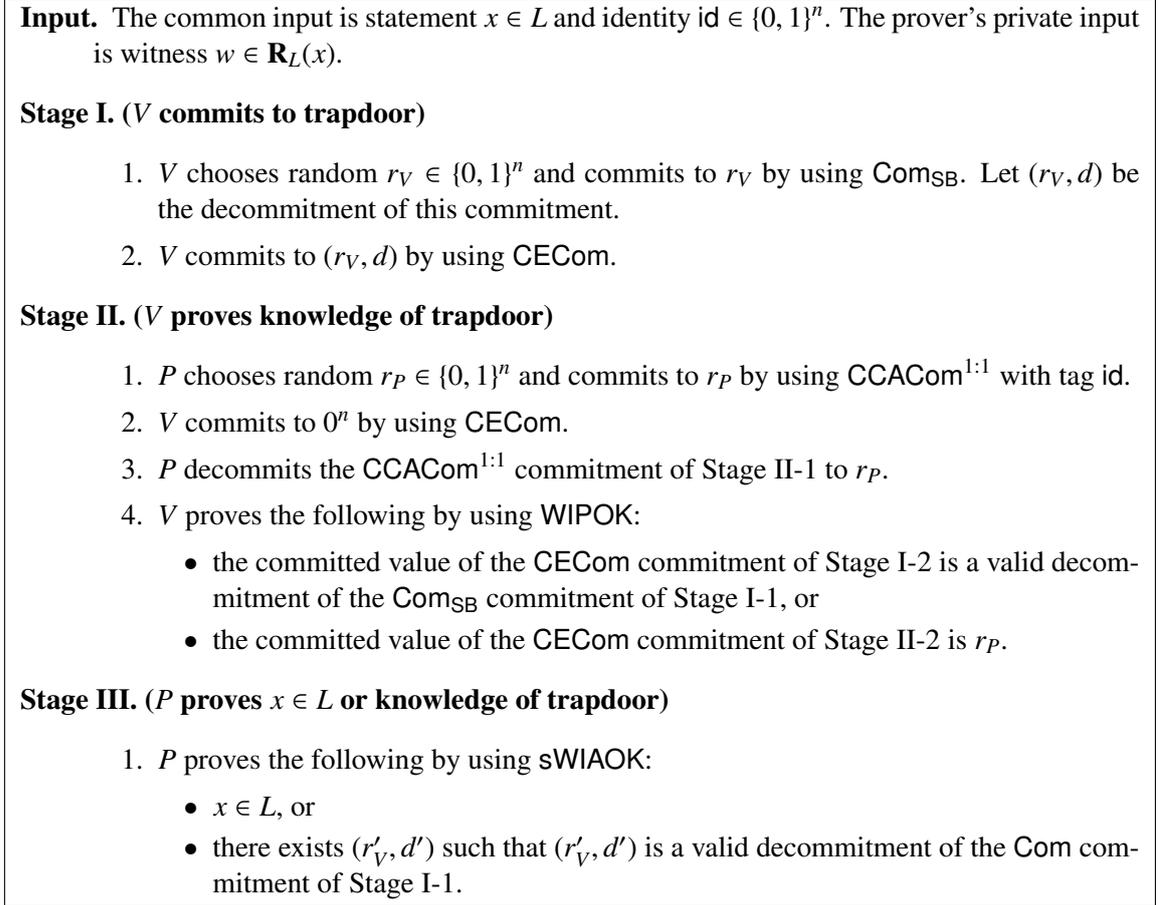


Figure 4: Statistical concurrent non-malleable zero-knowledge argument sCNMZK .

4.1 Proof of Soundness

Lemma 1. *Protocol sCNMZK is sound.*

Proof. Assume for contradiction that there exists an adversarial prover P^* that breaks the soundness of sCNMZK . It follows from the argument-of-knowledge property of sWIAOK that we can extract r_V from P^* in Stage III with non-negligible probability, where r_V is the value committed to by the verifier in Stage I-1. In the following, we consider a sequence of hybrid experiments in which the verifier is gradually modified so that P^* receives no information about r_V in the last hybrid, and then we derive a contradiction by showing that r_V is still extracted with non-negligible probability in the last hybrid.

Hybrid H_0 is an experiment in which a honest verifier interacts with P^* and then a witness is extracted in Stage III by the knowledge extractor of sWIAOK . The output of H_0 is the witness

extracted in Stage III. From the above observation, the output of H_0 is r_V with non-negligible probability.

Hybrid H_1 is the same as H_0 except that (i) the committed value r_P of the $\text{CCACom}^{1:1}$ commitment of Stage II-1 is extracted by the one-session committed-value oracle O of $\text{CCACom}^{1:1}$ and (ii) the committed value of the CECom commitment of Stage II-2 is switched from 0^n to r_P .

Note that, basically, the only difference between H_0 and H_1 is the value committed to with CECom in Stage II-2. However, since the execution of H_1 involves a super-polynomial-time computation (i.e., the extraction of r_P), we cannot directly use the hiding property of CECom to argue that the output of H_1 is indistinguishable from that of H_0 . Nevertheless, since H_1 can be executed in polynomial-time given access to the one-session committed-value oracle of $\text{CCACom}^{1:1}$, we can show the indistinguishability between the output of H_1 and that of H_0 by combining the hiding property of CECom with the robustness of $\text{CCACom}^{1:1}$ (cf. Proposition 1 in Section 3.4).

COMMENT: Formally, since $\text{CCACom}^{1:1}$ is robust only w.r.t. 4-round protocols, we need to consider a sequence of intermediate hybrids in which the CECom commitment is gradually modified by switching the committed values of the ExtCom commitments one by one in CECom . Since ExtCom has only four rounds, the 4-robustness of $\text{CCACom}^{1:1}$ guarantees that the outputs of these intermediate hybrids are indistinguishable.

Hybrid H_2 is the same as H_1 except that the WIPOK proof in Stage II-4 is computed by using a witness for the fact that the committed value of the CECom commitment of Stage II-2 is r_P .

Similar to the above, the indistinguishability between the output of H_2 and that of H_1 follows from the witness-indistinguishability of WIPOK and the robustness of $\text{CCACom}^{1:1}$.

Hybrid H_3 is the same as H_2 except that in Stage I-2, the committed value of the CECom commitment is switched from (r_V, d) to $(0^{|r_V|}, 0^{|d|})$.

The indistinguishability between the output of H_3 and H_2 follows from the hiding property of CECom (or, more precisely, the hiding property of ExtCom used in CECom) and the robustness of $\text{CCACom}^{1:1}$.

Hybrid H_4 is the same as H_3 except that in Stage I-1, the committed value of the Com_{SB} commitment is switched from r_V to 0^n .

The indistinguishability between the output of H_4 and that of H_3 follows from the hiding property of Com_{SB} and the robustness of $\text{CCACom}^{1:1}$.

From the above, the probability that the output of H_4 is r_V is non-negligible. However, since P^* receives no information about r_V in H_4 , this probability must be negligible. Thus, we reach a contradiction. \square

4.2 Proof of Statistical CNMZK Property

Simulator-extractor \mathcal{SE} .

Recall that to prove the statistical CNMZK property, we need to show a simulator-extractor that statistically simulates the view of the adversary \mathcal{A} and also extracts a witness in every accepted right session. We construct our simulator-extractor step by step. First, we construct a super-polynomial-time simulator $\hat{\mathcal{S}}$ that simulates the view of \mathcal{A} but does not extract witnesses in the right sessions. Next, we construct a super-polynomial-time simulator-extractor $\hat{\mathcal{SE}}$ that simulates the view of \mathcal{A} by

executing $\hat{\mathcal{S}}$ and then extracts the witnesses by rewinding $\hat{\mathcal{S}}$. Finally, we construct a polynomial-time simulator-extractor \mathcal{SE} that emulates the execution of $\hat{\mathcal{SE}}$ in polynomial time.

Remark 1. In the following, we use the hat symbol in the names of simulators and simulator-extractors if they run in super-polynomial time (e.g., $\hat{\mathcal{S}}$ and $\hat{\mathcal{SE}}$).

Remark 2. In the following, we use the tilde symbol in the names of the messages of sCNMZK if they are the messages of the right sessions (e.g., \tilde{r}_V and \tilde{r}_P). If necessary, we use subscript to denote the index of the session.

Super-polynomial-time simulator $\hat{\mathcal{S}}$. First, we show the simulator $\hat{\mathcal{S}}$, which simulates the view of \mathcal{A} in super-polynomial time as follows. $\hat{\mathcal{S}}$ internally invokes \mathcal{A} and interacts with \mathcal{A} as provers and verifiers in the following way.

- In each left session, $\hat{\mathcal{S}}$ interacts with \mathcal{A} in the same way as a honest prover except for the following. In Stage I-2, $\hat{\mathcal{S}}$ extracts the committed value (r_V, d) of the CECOM commitment by brute force. (If the committed value is not uniquely determined, (r_V, d) is defined to be (\perp, \perp) .) In Stage III, $\hat{\mathcal{S}}$ checks whether (r_V, d) is a valid decommitment of the $\text{Com}_{\mathbb{S}\mathbb{B}}$ commitment of Stage I-1; if so, $\hat{\mathcal{S}}$ gives a sWIAOK proof by using (r_V, d) as a witness; otherwise, $\hat{\mathcal{S}}$ terminates with output fail.
- In each right session, $\hat{\mathcal{S}}$ interacts with \mathcal{A} in the same way as a honest verifier.

Finally, $\hat{\mathcal{S}}$ outputs the view of internal \mathcal{A} . Notice that $\hat{\mathcal{S}}$ does not rewind \mathcal{A} .

Super-polynomial-time simulator-extractor $\hat{\mathcal{SE}}$. Next, we show the simulator-extractor $\hat{\mathcal{SE}}$, which simulates the view of \mathcal{A} in super-polynomial time and also extracts witnesses in every accepted right session as follows. First, $\hat{\mathcal{SE}}$ simulates the view of \mathcal{A} by executing $\hat{\mathcal{S}}$. We call this execution of $\hat{\mathcal{S}}$ the *wi-main thread*. Next, for each $i \in [m]$, if the i -th right session is accepted on the wi-main thread and uses a different identity from every left session, $\hat{\mathcal{SE}}$ extracts a witness from this session as follows.

- $\hat{\mathcal{SE}}$ rewinds the wi-main thread until the point just before the challenge message of sWIAOK of the i -th right session is sent. Then, from this point, $\hat{\mathcal{SE}}$ executes $\hat{\mathcal{S}}$ again with fresh randomness (i.e., interacts with \mathcal{A} as $\hat{\mathcal{S}}$ does with fresh randomness). $\hat{\mathcal{SE}}$ repeats this rewinding until it obtains another accepting transcript of the i -th right session. We call each execution of $\hat{\mathcal{S}}$ in this step a *wi-auxiliary thread*.
- After obtaining two accepting transcripts of the i -th right session (one is on the wi-main thread and the other is on an wi-auxiliary thread), $\hat{\mathcal{SE}}$ extracts a witness from sWIAOK by using the witness extractability of sWIAOK. If $\hat{\mathcal{SE}}$ fails to extract a witness for $\tilde{x}_i \in L$ (the statement proven in the i -th right session), $\hat{\mathcal{SE}}$ terminates with output fail_{WI} . Otherwise, let \tilde{w}_i be the extracted witness.

If the i -th right session is not accepted or uses the same identity as a left session, define $\tilde{w}_i \stackrel{\text{def}}{=} \perp$. The output of $\hat{\mathcal{SE}}$ is $(\text{view}, \{\tilde{w}_i\}_{i \in [m]})$, where *view* is the view of \mathcal{A} on the wi-main thread.

Polynomial-time simulator-extractor \mathcal{SE} . Finally, we show the simulator-extractor \mathcal{SE} , which emulates the execution of $\hat{\mathcal{SE}}$ in polynomial time as follows. First, \mathcal{SE} emulates the wi-main thread in polynomial time as follows.

- \mathcal{SE} internally invokes \mathcal{A} and interacts with \mathcal{A} as $\hat{\mathcal{S}}$ does except that in each left session, \mathcal{SE} extracts (r_V, d) by using the concurrent extractability of CECom instead of by brute force. Recall that a concurrent extraction of CECom involves the generation of a main thread and many look-ahead threads. We call the main thread generated during the concurrent extraction of CECom the *cec-main thread*, and call the look-ahead threads generated during the concurrent extraction of CECom the *cec-auxiliary threads*.⁵ (See Figure 5.)

Next, for each $i \in [m]$, if the i -th right session is accepted on the emulated w_i -main thread and uses a different identity from every left session, \mathcal{SE} emulates w_i -auxiliary threads as follows.

- \mathcal{SE} rewinds the emulation of the w_i -main thread until the point just before the challenge message of sWIAOK of the i -th right session is sent on the cec-main thread. Then, from this point, \mathcal{SE} emulates the w_i -main thread again with fresh randomness (i.e., generates the rest of cec-main thread and cec-auxiliary threads with fresh randomness). \mathcal{SE} repeats this rewinding until it obtains another accepted transcript of the i -th right session on an emulated w_i -auxiliary thread.

Let $(\text{view}, \{\tilde{w}_i\}_{i \in [m]})$ be the output of the emulated $\hat{\mathcal{S}}$. Then, \mathcal{SE} outputs $(\text{view}, \{\tilde{w}_i\}_{i \in [m]})$.

Analysis of poly-time simulator-extractor \mathcal{SE} .

To prove the statistical CNMZK property, we show that \mathcal{SE} statistically simulates the view of \mathcal{A} and also extracts witnesses for the statements in the right sessions.

Lemma 2. *The view of \mathcal{A} simulated by \mathcal{SE} is statistically indistinguishable from the view of \mathcal{A} in the real experiment. Furthermore, except with negligible probability, \mathcal{SE} outputs witnesses for the statements proven by \mathcal{A} in the accepted right sessions that use different identities from the left sessions.*

Proof. In this proof, we use the following claim, which states that the super-polynomial-time simulator-extractor $\hat{\mathcal{S}}$ statistically simulates the view of \mathcal{A} and also extracts the witnesses from the right sessions.

Claim 1. *The view of \mathcal{A} simulated by $\hat{\mathcal{S}}$ is statistically indistinguishable from the view of \mathcal{A} in the real experiment. Furthermore, except with negligible probability, $\hat{\mathcal{S}}$ outputs witnesses for the statements proven by \mathcal{A} in the accepted right sessions that use different identities from the left sessions.*

Before proving this claim, we finish the proof of Lemma 2. Given Claim 1, we can prove Lemma 2 by showing that the output of \mathcal{SE} is statistically indistinguishable from that of $\hat{\mathcal{S}}$. Roughly speaking, this indistinguishability can be shown by observing the following.

- In \mathcal{SE} , the emulation of $\hat{\mathcal{S}}$ is perfect if in every left session that reaches Stage III, the value extracted by the concurrent extractability of CECom is equal to the value that would be extracted by brute force.
- In every such left session, the value extracted by the concurrent extractability of CECom is indeed equal to the value that would be extracted by brute force. This is because the CECom commitment in Stage I-2 is valid in every such left session except with negligible probability, which in turn is because of the soundness of WIPOK and the hiding property of $\text{CCACom}^{1:1}$.

⁵Note that the w_i -main thread is also a cec-main thread.

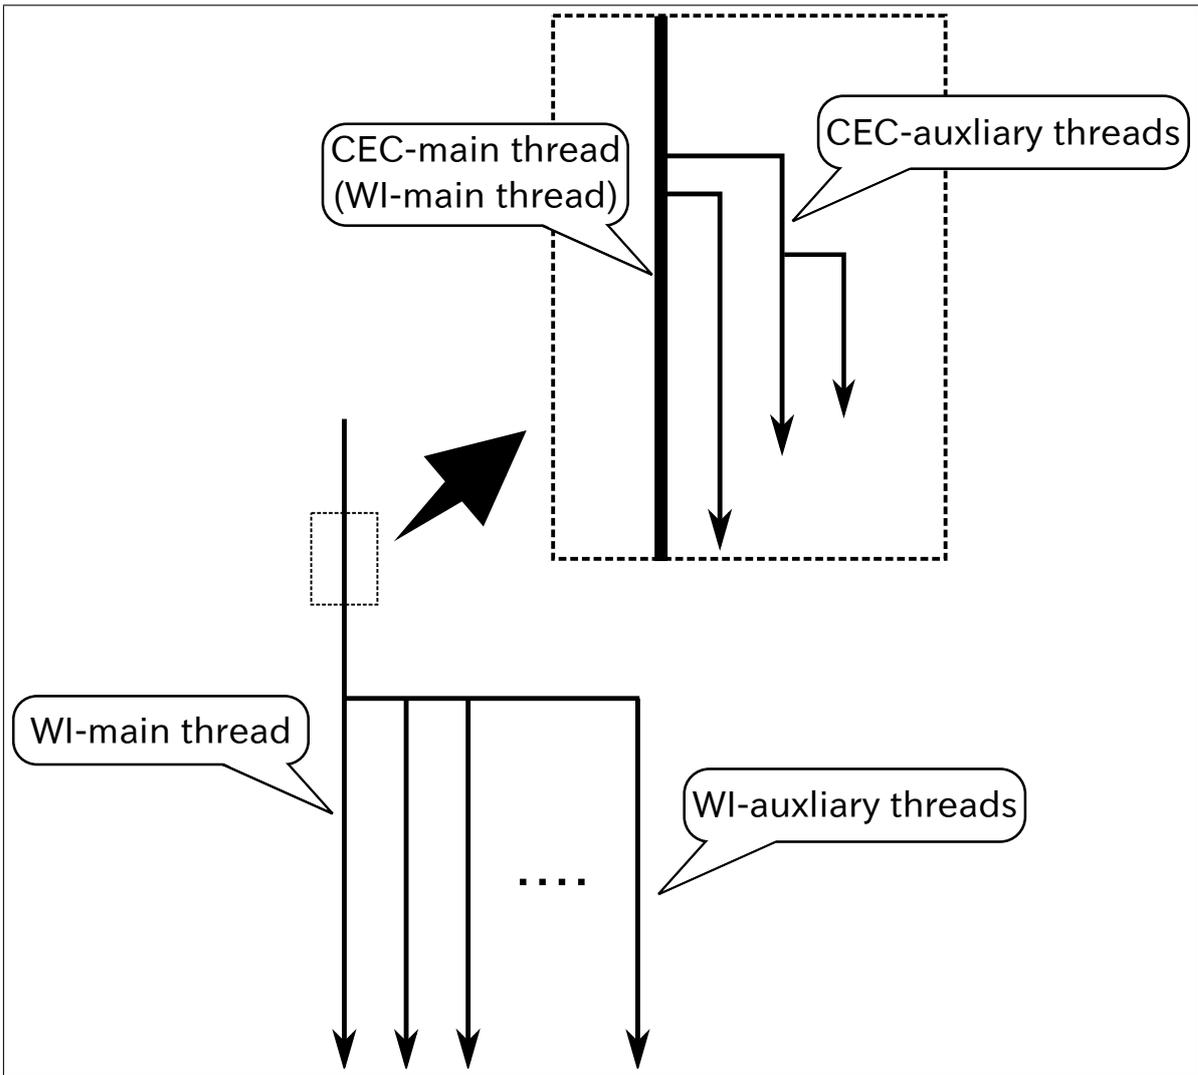


Figure 5: wi-main thread, wi-auxiliary thread, cec-main thread, and cec-auxiliary thread.

We note that there is a subtlety since the concurrent extraction of CECOM itself is rewound in \mathcal{SE} when the witnesses are extracted from the right sessions. A formal argument is given below.

Let CEC-BAD be the event that during the execution of \mathcal{SE} , in a left session that reaches Stage III, the value extracted from the CECOM commitment in Stage I-2 is different from the value that would be extracted by the brute-force extraction. Let ϵ be the probability that CEC-BAD occurs during the emulation of the w_1 -main thread (i.e., during the emulation of the execution of $\hat{\mathcal{S}}$). From the concurrent extractability of CECOM, if the CECOM commitment in Stage I-2 is valid except with negligible probability, CEC-BAD occurs only with negligible probability. Hence, we obtain $\epsilon = \text{negl}(n)$ from the following claim.

Claim 2. *In $\hat{\mathcal{S}}$, the following holds except with negligible probability: In every left session that reaches Stage III, the CECOM commitment in Stage I-2 of this session is valid and its committed value is a valid decommitment of the COM_{SB} commitment of Stage I-1.*

The proof of Claim 2 is given after this proof.

To show the indistinguishability between the output of $\hat{\mathcal{SE}}$ and that of \mathcal{SE} , we consider the following hybrid simulator-extractor $\hat{\mathcal{SE}}'$ and \mathcal{SE}' .

- $\hat{\mathcal{SE}}'$ (resp., \mathcal{SE}') is the same as $\hat{\mathcal{SE}}$ (resp., \mathcal{SE}) except that for each $i \in [m]$, it terminates with output time-out if it does not obtain another accepting transcript of the i -th right session after rewinding the w_1 -main thread (resp., the emulation of the w_1 -main thread) $1/\epsilon^{1/4}$ times.

First, we show that the output of $\hat{\mathcal{SE}}$ and that of $\hat{\mathcal{SE}}'$ are statistically indistinguishable. From the definition of $\hat{\mathcal{SE}}'$, this indistinguishability holds if $\hat{\mathcal{SE}}'$ outputs time-out with at most negligible probability. Thus, to show this indistinguishability, it suffices to show that the probability that $\hat{\mathcal{SE}}$ rewinds w_1 -main thread more than $1/\epsilon^{1/4}$ times during the witness extraction of a right session is negligible. To show that this probability is negligible, we do the following. For each $i \in [m]$, let T_i be the random variable representing the number of rewinding during the witness extraction of the i -th right session in $\hat{\mathcal{SE}}$. From a standard argument, we can show that we have $\mathbb{E}[T_i] = 1$ for every $i \in [m]$.⁶ Thus, from Marcov's inequality, we have

$$\Pr[T_i > 1/\epsilon^{1/4}] \leq \epsilon^{1/4} = \text{negl}(n)$$

for every $i \in [m]$. Thus, from union bound, we have

$$\Pr[\exists i \in [m] \text{ s.t. } T_i > 1/\epsilon^{1/4}] \leq m \cdot \text{negl}(n) = \text{negl}(n) .$$

As noted above, this implies that the output of $\hat{\mathcal{SE}}$ and that of $\hat{\mathcal{SE}}'$ are statistically indistinguishable.

Next, we show that the output of $\hat{\mathcal{SE}}'$ and that of \mathcal{SE}' are statistically indistinguishable. Since the only difference between $\hat{\mathcal{SE}}'$ and \mathcal{SE}' is how the committed values are extracted from CECOM, this indistinguishability holds if CEC-BAD occurs in \mathcal{SE}' with at most negligible probability. For $\ell \in \mathbb{N}$, let ST_ℓ be the random variable representing the internal state of \mathcal{SE}' at the time that \mathcal{A} has sent the ℓ -th messages on the CEC-main thread during the emulation of w_1 -main thread. Let $\text{CEC-BAD}_{\text{main}}$ be the event that CEC-BAD occurs during the emulation of the w_1 -main thread. We say that an internal state st of \mathcal{SE}' is *good w.r.t ℓ* if we have $\Pr[\text{CEC-BAD}_{\text{main}} \mid ST_\ell = \text{st}] \leq \epsilon^{1/2}$. Let GOOD_ℓ be the event that $ST_\ell = \text{st}$ holds for an internal state st that is good w.r.t. ℓ . Then, for any ℓ , we have

$$\Pr[\text{CEC-BAD}_{\text{main}}] \geq \Pr[\text{CEC-BAD}_{\text{main}} \mid \neg\text{GOOD}_\ell] \Pr[\neg\text{GOOD}_\ell] \geq \epsilon^{1/2} \cdot \Pr[\neg\text{GOOD}_\ell] .$$

⁶ For any prefix ρ of the transcript immediately before the challenge message of swIAOK of the i -th right session, let p be the probability that the i -th right session is accepted when the prefix of the transcript is ρ . Then, we have $\mathbb{E}[T_i \mid \text{prefix}_\rho] = p \cdot 1/p = 1$, where prefix_ρ is the event that the prefix of the transcript is ρ . Thus, we have $\mathbb{E}[T_i] = \sum_\rho \mathbb{E}[T_i \mid \text{prefix}_\rho] \Pr[\text{prefix}_\rho] = 1$.

Then, since we have $\Pr [\text{CEC-BAD}_{\text{main}}] = \epsilon$ from the definition of ϵ , we have

$$\Pr [\neg \text{GOOD}_\ell] \leq \epsilon^{1/2} = \text{negl}(n)$$

for every ℓ . Thus, from union bound, we have

$$\Pr \left[\bigvee_{\ell} \neg \text{GOOD}_\ell \right] \leq \text{negl}(n) . \quad (1)$$

Let GOOD be the event that GOOD_ℓ occurs for every ℓ . Then, from Equation (1), we have $\Pr [\text{GOOD}] \geq 1 - \text{negl}(n)$. For $i \in [m]$, let CEC-BAD_i be the event that CEC-BAD occurs during the witness extraction of the i -th right session. Then, since the emulation of each wi-auxiliary thread proceeds identically with that of the wi-main thread, and since for each $i \in [m]$ there are at most $1/\epsilon^{1/4}$ wi-auxiliary threads during the witness extraction of the i -th right session, we have

$$\Pr [\text{CEC-BAD}_i \mid \text{GOOD}] \leq \frac{1}{\epsilon^{1/4}} \cdot \epsilon^{1/2} = \epsilon^{1/4} .$$

Thus, we have

$$\begin{aligned} \Pr [\text{CEC-BAD}] &= \Pr [\text{CEC-BAD}_{\text{main}}] + \sum_{i=1}^m \Pr [\text{CEC-BAD}_i] \\ &\leq \Pr [\text{CEC-BAD}_{\text{main}}] + \sum_{i=1}^m (\Pr [\neg \text{GOOD}] + \Pr [\text{CEC-BAD}_i \mid \text{GOOD}]) \\ &\leq \epsilon + \sum_{i=1}^m (\text{negl}(n) + \epsilon^{1/4}) \\ &= \text{negl}(n) . \end{aligned}$$

As noted above, this implies that the output of $\hat{\mathcal{SE}}'$ and that of \mathcal{SE}' are statistically indistinguishable.

Finally, we show that the output of \mathcal{SE}' and that of \mathcal{SE} are statistically indistinguishable. Since the output of $\hat{\mathcal{SE}}'$ and that of \mathcal{SE}' are statistically indistinguishable and since $\hat{\mathcal{SE}}'$ outputs time-out with at most negligible probability, \mathcal{SE}' outputs time-out with at most negligible probability. Then, since \mathcal{SE}' is identical to \mathcal{SE} unless \mathcal{SE}' outputs time-out, the output of \mathcal{SE}' and that of \mathcal{SE} are statistically indistinguishable. \square

Proof of Claim 2. Recall that Claim 2 states that during the execution of $\hat{\mathcal{S}}$, in every left session that reaches Stage III, the CECom commitment in Stage I-2 is valid and its committed value is a valid decommitment of the Com_{SB} commitment of Stage I-1.

Let us say that a left session is *bad* if it reaches Stage III and either the CECom commitment in Stage I-2 is invalid or its committed value is not a valid decommitment of the Com_{SB} commitment in Stage I-1; a left session is *good* if it is not bad. What we need to prove is that every left session is good except with negligible probability.

Roughly speaking, the proof proceeds as follows. From the soundness of the WIPOK proof of Stage II-4, in a bad left session, the committed value of the CECom commitment of Stage II-2 is r_P , which is the committed value of the $\text{CCACom}^{1:1}$ commitment of Stage II-1; thus, before r_P is decommitted to in Stage II-3, we can obtain r_P by extracting the committed value from CECom in Stage II-2. This itself does not contradict to the hiding property of $\text{CCACom}^{1:1}$ since $\hat{\mathcal{S}}$ runs in super-polynomial time in the brute-force extraction of CECom . Thus, we consider a hybrid simulator-extractor in which the brute-force extraction of CECom is replaced with the concurrent extraction of

CECom. Here, since we want to use the hiding property of $\text{CCACom}^{1:1}$, we use the robust concurrent extraction of CECom so that the $\text{CCACom}^{1:1}$ commitment in a left session is not rewound. For details, see below.

Assume for contradiction that there exists $i \in [m]$ such that the i -th left session is bad with non-negligible probability. (Here, the indices of the left sessions are defined by the order in which Stage III begins; the reason why we define the indices in this way will become clear later.) Then, there exists $i^* \in [m]$ such that the first $(i^* - 1)$ left sessions are good except with negligible probability but the i^* -th left session is bad with non-negligible probability. Note that from the soundness of WIPOK, when the i^* -th left session is bad, then the CECom commitment of Stage II-2 of the i^* -th left session is valid and its committed value is r_P except with negligible probability, where r_P is the value committed to in Stage II-1 of the i^* -th left session. In the following, we use BAD to denote the event that the i^* -th left session is bad, and use CHEAT to denote the event that the committed value of the CECom commitment of Stage II-2 is r_P in the i^* -th left session. Then, let us consider the following hybrids.

Hybrid $\hat{\mathcal{S}}_0$ is the same as $\hat{\mathcal{S}}$. From our assumption, BAD occurs in $\hat{\mathcal{S}}_0$ with non-negligible probability. Thus, from the above argument, CHEAT occurs in $\hat{\mathcal{S}}_0$ with non-negligible probability.

Hybrid $\hat{\mathcal{S}}_1$ is the same as $\hat{\mathcal{S}}_0$ except that $\hat{\mathcal{S}}_1$ terminates just before Stage III of the i^* -th left session begins. Clearly, BAD and CHEAT occur in $\hat{\mathcal{S}}_1$ with non-negligible probability.

Hybrid \mathcal{S}_1 emulates $\hat{\mathcal{S}}_1$ in polynomial time as follows.

- At the beginning, a random left session s is chosen. (Here, we guess that session s is the i^* -th left session.)
- In every left session, in Stage I-2, the committed value (r_V, d) is extracted by the robust concurrent extractor of CECom in such a way that the $\text{CCACom}^{1:1}$ commitment of session s is not rewound. In addition, in session s , the committed value is also extracted in Stage II-2.

Note that in every left session in which Stage III is executed, the CECom commitment in Stage I-2 is valid except with negligible probability (since such a session is one of the first $(i^* - 1)$ left sessions and therefore it is good except with negligible probability). Thus, the values extracted from the concurrent extractor are equal to the values that would be extracted by the brute-force extraction except with negligible probability; therefore, \mathcal{S}_1 statistically emulates $\hat{\mathcal{S}}_1$, and BAD and CHEAT occur in \mathcal{S}_1 with non-negligible probability.

Note that session s is the i^* -th left session with non-negligible probability. Then, since CHEAT occurs in \mathcal{S}_1 with non-negligible probability, the value extracted from the CECom commitment in Stage II-2 of session s is r_P with non-negligible probability, where r_P is the value committed to in Stage II-1 of session s . Then, since the $\text{CCACom}^{1:1}$ commitment in Stage II-1 of session s is not rewound in \mathcal{S}_1 , we can break the hiding property of $\text{CCACom}^{1:1}$. Thus, we reach a contradiction. \square

Analysis of super-poly-time simulator-extractor $\hat{\mathcal{S}}\mathcal{E}$.

It remains to prove Claim 1, which states that (i) super-polynomial-time simulator-extractor $\hat{\mathcal{S}}\mathcal{E}$ statistically simulates the real view of \mathcal{A} and (ii) $\hat{\mathcal{S}}\mathcal{E}$ also extracts a valid witness from every accepted right session in the simulated view.

Proof of Claim 1. First, we observe that the output of $\hat{\mathcal{S}}$ is statistically indistinguishable from the real view of \mathcal{A} . Since $\hat{\mathcal{S}}\mathcal{E}$ simulates the view of \mathcal{A} by executing $\hat{\mathcal{S}}$, this implies that $\hat{\mathcal{S}}\mathcal{E}$ statistically simulates the real view of \mathcal{A} . Recall that in $\hat{\mathcal{S}}$, each left session is simulated by extracting (r_V, d) from

the CECOM commitment in Stage I-2 and giving a sWIAOK proof in Stage III with witness (r_V, d) . From Claim 2 (which states that the CECOM commitment in Stage I-2 is a valid commitment to a valid decommitment of the COM_{SB} commitment in Stage I-1 in every session that reaches Stage III), the value (r_V, d) that is extracted from the CECOM commitment in Stage I-2 is a valid decommitment of the COM_{SB} commitment of Stage I-1 in each left session that reaches Stage III. Thus, from the statistical witness indistinguishability of sWIAOK, the output of $\hat{\mathcal{S}}$ is statistically indistinguishable from the real view of \mathcal{A} .

Next, we show that $\hat{\mathcal{S}}\mathcal{E}$ outputs fail_{WI} with at most negligible probability. Since $\hat{\mathcal{S}}\mathcal{E}$ outputs fail_{WI} when it fails to extract a witness in an accepted right session, this implies that $\hat{\mathcal{S}}\mathcal{E}$ extracts a valid witness from every accepted right session except with negligible probability. Assume for contradiction that there exists $\tilde{i}^* \in [m]$ such that $\hat{\mathcal{S}}\mathcal{E}$ outputs fail_{WI} during the witness extraction of the \tilde{i}^* -th right session with non-negligible probability. Then, let us consider the following hybrid simulator-extractor $\hat{\mathcal{S}}\mathcal{E}_{\tilde{i}^*}$.

- $\hat{\mathcal{S}}\mathcal{E}_{\tilde{i}^*}$ is the same as $\hat{\mathcal{S}}\mathcal{E}$ except that $\hat{\mathcal{S}}\mathcal{E}_{\tilde{i}^*}$ tries to extract a witness only from the \tilde{i}^* -th right session (and therefore rewinds the WI-main thread only from the challenge message of sWIAOK of the \tilde{i}^* -th right session).

Clearly, $\hat{\mathcal{S}}\mathcal{E}_{\tilde{i}^*}$ outputs fail_{WI} with non-negligible probability. Then, we reach a contradiction roughly as follows.

Step 1. First, we show that in $\hat{\mathcal{S}}\mathcal{E}_{\tilde{i}^*}$, the probability that \tilde{r}_V is extracted as a witness during the witness extraction of the \tilde{i}^* -th right session is non-negligible, where \tilde{r}_V is the value chosen by the verifier in Stage I-1 of the \tilde{i}^* -th right session.

Step 2. Next, we define a sequence of hybrid simulator-extractors, where the first hybrid is the same as $\hat{\mathcal{S}}\mathcal{E}_{\tilde{i}^*}$, and we gradually modify the \tilde{i}^* -th right session so that it is independent of \tilde{r}_V in the last hybrid.

Step 3. Finally, we show that even in the last hybrid, the probability that \tilde{r}_V is extracted during the witness extraction of the \tilde{i}^* -th right session is non-negligible. Since the \tilde{i}^* -th right session is independent of \tilde{r}_V in the last hybrid, we reach a contradiction.

Details are given below.

Step 1. Prove that $\hat{\mathcal{S}}\mathcal{E}_{\tilde{i}^*}$ extracts \tilde{r}_V . We first prove the following claim.

Claim 3. *Let \tilde{r}_V be the value chosen by the verifier in Stage I-1 of the \tilde{i}^* -th right session. If $\hat{\mathcal{S}}\mathcal{E}_{\tilde{i}^*}$ outputs fail_{WI} with non-negligible probability, then in $\hat{\mathcal{S}}\mathcal{E}_{\tilde{i}^*}$ the probability that \tilde{r}_V is extracted during the witness extraction of the \tilde{i}^* -th right session is non-negligible.*

Proof. Assume for contradiction that \tilde{r}_V is extracted during the witness extraction of the \tilde{i}^* -th right session with at most negligible probability. Then, since we assume that $\hat{\mathcal{S}}\mathcal{E}_{\tilde{i}^*}$ outputs fail_{WI} with non-negligible probability, the following occurs in $\hat{\mathcal{S}}\mathcal{E}_{\tilde{i}^*}$ with non-negligible probability:

- $\hat{\mathcal{S}}\mathcal{E}_{\tilde{i}^*}$ obtains two accepting transcript of the \tilde{i}^* -th right session (and therefore that of sWIAOK) such that the commit-messages of sWIAOK are the same,⁷ but
- from these two transcript, $\hat{\mathcal{S}}\mathcal{E}_{\tilde{i}^*}$ fails to extract any witness from sWIAOK (i.e., a witness for $\tilde{x}_{\tilde{i}^*} \in L$ or a valid decommitment of the Stage I-1 commitment).

⁷Recall that WIPOK consists of three stages: commit, challenge, and response.

We first show that when the above occurs, the two accepting sWIAOK transcripts are admissible except with negligible probability. (Recall that a pair of accepted transcripts of sWIAOK are admissible if their commit-messages are the same but their challenge-messages are different.) Toward this end, it suffices to show that $\hat{\mathcal{S}}\mathcal{E}_{\tilde{i}^*}$ chooses the same challenge-message of sWIAOK on two wi-auxiliary threads with at most negligible probability. This can be shown as follows.

- From a standard argument, we can show that the expected number of rewinding of the wi-main thread is 1 in $\hat{\mathcal{S}}\mathcal{E}_{\tilde{i}^*}$.⁸ Thus, the probability that $\hat{\mathcal{S}}\mathcal{E}_{\tilde{i}^*}$ rewinds the wi-main thread more than $2^{n/2}$ times is at most $2^{-n/2}$. Furthermore, under the condition that $\hat{\mathcal{S}}\mathcal{E}_{\tilde{i}^*}$ rewinds the wi-main thread at most $2^{n/2}$ times, the probability that $\hat{\mathcal{S}}\mathcal{E}_{\tilde{i}^*}$ chooses the same challenge-message on two wi-auxiliary threads is at most $2^{n/2} \cdot 2^{-n} = 2^{-n/2}$. Thus, the probability that $\hat{\mathcal{S}}\mathcal{E}_{\tilde{i}^*}$ chooses the same challenge-message in two wi-auxiliary thread is at most $2^{-n/2} + 2^{-n/2} = \text{negl}(n)$.

Thus, with non-negligible probability $\hat{\mathcal{S}}\mathcal{E}_{\tilde{i}^*}$ obtains two admissible transcripts of sWIAOK from which no witness can be computed.

We then reach a contradiction as follows. Since sWIAOK is a parallel version of Blum's Hamiltonian-cycle protocol, if no witness is extracted from two admissible transcripts of sWIAOK, a Com_{SH} commitment in the commit-messages is decommitted to two different values in the transcripts. Thus, we derive a contradiction by breaking the binding property of Com_{SH} using $\hat{\mathcal{S}}\mathcal{E}_{\tilde{i}^*}$. A problem is that since $\hat{\mathcal{S}}\mathcal{E}_{\tilde{i}^*}$ runs in super-polynomial time, the *computational* hiding property of Com_{SH} may not hold in $\hat{\mathcal{S}}\mathcal{E}_{\tilde{i}^*}$. To overcome this problem, we consider hybrid simulator-extractor $\mathcal{S}\mathcal{E}_{\tilde{i}^*}$ that emulates the execution of $\hat{\mathcal{S}}\mathcal{E}_{\tilde{i}^*}$ in polynomial time. Specifically, $\mathcal{S}\mathcal{E}_{\tilde{i}^*}$ emulates $\hat{\mathcal{S}}\mathcal{E}_{\tilde{i}^*}$ in the same way as $\mathcal{S}\mathcal{E}$ emulates $\hat{\mathcal{S}}\mathcal{E}$ (i.e., by using the concurrent extractability of CECom instead of the brute-force extraction) except for the following.

- During the emulation of the wi-main thread, the value (r_V, d) is extracted in Stage I-2 of each left session by using the *robust* concurrent extractability of CECom so that the commit-message of sWIAOK of the i^* -th right session is not rewound.

As in the proof of Lemma 2, we can show that $\mathcal{S}\mathcal{E}_{\tilde{i}^*}$ statistically emulates the execution of $\hat{\mathcal{S}}\mathcal{E}_{\tilde{i}^*}$. Thus, with non-negligible probability, $\mathcal{S}\mathcal{E}_{\tilde{i}^*}$ obtains two valid decommitments of a Com_{SH} commitment (in the commit-messages of sWIAOK of the \tilde{i}^* -th right session) such that decommitted values are different. Then, since $\mathcal{S}\mathcal{E}_{\tilde{i}^*}$ runs in polynomial time and since the commit-messages of sWIAOK (and therefore the Com_{SH} commitment) of the \tilde{i}^* -th right session is not rewound in $\mathcal{S}\mathcal{E}_{\tilde{i}^*}$,⁹ we can break the binding property of Com_{SH} . Thus, we reach a contradiction. \square

Step 2. Introduce hybrid simulator-extractor. Next, we introduce hybrid simulator-extractors. To clarify the exposition, we first define a sequence of hybrid simulators by gradually modifying $\hat{\mathcal{S}}$ and then define the hybrid simulator-extractors by using them. Below, when we refer to a particular stage of sCNMZK, we always means the corresponding stage of sCNMZK in the \tilde{i}^* -th right session.

Hybrid simulator $h\text{-}\hat{\mathcal{S}}_0$ is identical with $\hat{\mathcal{S}}$.

Hybrid simulator $h\text{-}\hat{\mathcal{S}}_1$ is the same as $h\text{-}\hat{\mathcal{S}}_0$ except that \tilde{r}_P is extracted by brute force in Stage II-1 and the committed value of the CECom commitment in Stage II-2 is switched from 0^n to \tilde{r}_P .

Hybrid simulator $h\text{-}\hat{\mathcal{S}}_2$ is the same as $h\text{-}\hat{\mathcal{S}}_1$ except that in Stage II-4, the WIPOK proof is computed by using a witness for the fact that the committed value of the CECom commitment of Stage II-2 is \tilde{r}_P .

⁸See Footnote 6.

⁹Note that the commit-messages of sWIAOK of the \tilde{i}^* -th right session appear only on the wi-main thread.

Hybrid simulator $h\text{-}\hat{S}_3$ is the same as $h\text{-}\hat{S}_2$ except that in Stage I-2, the committed value of the CECOM commitment is switched from (\tilde{r}_V, \tilde{d}) to $(0^{|\tilde{r}_V|}, 0^{|\tilde{d}|})$.

Hybrid simulator $h\text{-}\hat{S}_4$ is the same as $h\text{-}\hat{S}_3$ except that in Stage I-1, the committed value of the Com_{SB} commitment is switched from \tilde{r}_V to 0^n .

Then, for each $k \in \{0, \dots, 4\}$, hybrid simulator-extractor $h\text{-}\hat{S}\mathcal{E}_k$ is defined as follows.

Hybrid simulator-extractor $h\text{-}\hat{S}\mathcal{E}_k$ is the same as $\hat{S}\mathcal{E}_{\tilde{i}^*}$ except that the execution of \hat{S} is replaced with that of $h\text{-}\hat{S}_k$. The output of $h\text{-}\hat{S}\mathcal{E}_k$ is the value extracted during the witness extraction of the \tilde{i}^* -th right session.

Note that the value \tilde{r}_V is not used anywhere in $h\text{-}\hat{S}\mathcal{E}_4$.

Step 3. Prove that \tilde{r}_V is extracted in every hybrid. Finally, we show that \tilde{r}_V is extracted with non-negligible probability in each hybrid. First, we consider $h\text{-}\hat{S}\mathcal{E}_1$.

Claim 4. *Let \tilde{r}_V be the value chosen by the verifier in Stage I-1 of the \tilde{i}^* -th right session. If $\hat{S}\mathcal{E}_{\tilde{i}^*}$ outputs fail_{W} with non-negligible probability, then in $h\text{-}\hat{S}\mathcal{E}_1$ the probability that \tilde{r}_V is extracted during the witness extraction of the \tilde{i}^* -th right session is non-negligible.*

Proof. In this proof, we use intermediate hybrid simulator-extractors in which the CECOM commitment in Stage II-2 of the \tilde{i}^* -th right session is gradually modified. Again, we first introduce hybrid simulators. Recall that a CECOM commitment consists of $\ell = \omega(R_{\text{SH}}(n) \log n)$ ExtCOM commitments. Then, the intermediate hybrid simulators $h\text{-}\hat{S}_{0:0}, \dots, h\text{-}\hat{S}_{0:\ell}$ are defined as follows.

Hybrid simulator $h\text{-}\hat{S}_{0:0}$ is the same as $h\text{-}\hat{S}_0$ except that \tilde{r}_P is extracted by brute force in Stage II-1 of the \tilde{i}^* -th right session.

Hybrid simulator $h\text{-}\hat{S}_{0:k}$ ($k \in [\ell]$) is the same as $h\text{-}\hat{S}_{0:k-1}$ except that the committed value of the k -th ExtCOM commitment in the CECOM commitment of Stage II-2 is switched from 0^n to \tilde{r}_P in the \tilde{i}^* -th right session.

Then, for each $k \in \{0, \dots, \ell\}$, hybrid simulator-extractor $h\text{-}\hat{S}\mathcal{E}_{0:k}$ is defined as follows.

Hybrid simulator-extractor $h\text{-}\hat{S}\mathcal{E}_{0:k}$ is the same as $h\text{-}\hat{S}\mathcal{E}_0$ except that the execution of $h\text{-}\hat{S}_0$ is replaced with that of $h\text{-}\hat{S}_{0:k}$.

Note that $h\text{-}\hat{S}\mathcal{E}_{0:\ell}$ is identical with $h\text{-}\hat{S}\mathcal{E}_1$.

Below, we show that for every $k \in [\ell]$, the output of $h\text{-}\hat{S}\mathcal{E}_{0:k-1}$ and that of $h\text{-}\hat{S}\mathcal{E}_{0:k}$ are indistinguishable. (Recall that the outputs of $h\text{-}\hat{S}\mathcal{E}_{0:k-1}$ and $h\text{-}\hat{S}\mathcal{E}_{0:k}$ are the value extracted in the \tilde{i}^* -th right session.) Since the probability that \tilde{r}_V is extracted in $h\text{-}\hat{S}\mathcal{E}_{0:0}$ is non-negligible from Claim 3, this suffices to prove Claim 4.

Roughly speaking, we show this indistinguishability as follows. Since $h\text{-}\hat{S}\mathcal{E}_{0:k-1}$ and $h\text{-}\hat{S}\mathcal{E}_{0:k}$ differ only in the committed values of a ExtCOM commitment, we use the hiding property of the ExtCOM commitment to show the indistinguishability. A problem is that we cannot use it directly since $h\text{-}\hat{S}\mathcal{E}_{0:k-1}$ and $h\text{-}\hat{S}\mathcal{E}_{0:k}$ run in super-polynomial time. To overcome this problem, we observe that the only super-polynomial computations in $h\text{-}\hat{S}\mathcal{E}_{0:k-1}$ and $h\text{-}\hat{S}\mathcal{E}_{0:k}$ are the brute-force extraction of $\text{CCACOM}^{1:1}$ (in the \tilde{i}^* -th right session) and those of CECOM (in each left session). Based on this observation, we first show that the execution of $h\text{-}\hat{S}\mathcal{E}_{0:k-1}$ and $h\text{-}\hat{S}\mathcal{E}_{0:k}$ can be emulated in polynomial-time by using the one-session committed-value oracle \mathcal{O} of $\text{CCACOM}^{1:1}$ and the concurrent extractability of CECOM. We then combine the 4-robustness of $\text{CCACOM}^{1:1}$ with the hiding

property of ExtCom (which has only four rounds) to argue that the output of $h\text{-}\hat{\mathcal{S}}\mathcal{E}_{0:k-1}$ and that of $h\text{-}\hat{\mathcal{S}}\mathcal{E}_{0:k}$ are indistinguishable. To formally implement this idea, we need to make sure that the ExtCom commitment and the $\text{CCACom}^{1:1}$ commitment are not rewound during the concurrent extraction of GECOM. Details are given below.

First, we introduce hybrid simulator-extractors $h\text{-}\mathcal{SE}_{0:k-1}^{\mathcal{O}}$ and $h\text{-}\mathcal{SE}_{0:k}^{\mathcal{O}}$, where \mathcal{O} is the one-session committed-value oracle of $\text{CCACom}^{1:1}$. Hybrid $h\text{-}\mathcal{SE}_{0:k}^{\mathcal{O}}$ (resp., $h\text{-}\mathcal{SE}_{0:k-1}^{\mathcal{O}}$) emulates $h\text{-}\hat{\mathcal{S}}\mathcal{E}_{0:k}$ (resp., $h\text{-}\hat{\mathcal{S}}\mathcal{E}_{0:k-1}$) in the same way as \mathcal{SE} emulates $\hat{\mathcal{S}}\mathcal{E}$ except for the following.

- During the emulation of the wi-main thread, the value (r_V, d) is extracted in Stage I-2 of each left session by using the robust concurrent extractability so that the $\text{CCACom}^{1:1}$ commitment of Stage II-1 and the k -th ExtCom commitment of the CECOM commitment of Stage II-2 are not rewound in the \tilde{i}^* -th right session. In addition, in the \tilde{i}^* -th right session, the committed value of $\text{CCACom}^{1:1}$ is extracted by forwarding the commitment to \mathcal{O} . Note that the $\text{CCACom}^{1:1}$ commitment in the \tilde{i}^* -th right session is not rewound and therefore it can be forwarded to \mathcal{O} .

Next, we show that for each $h \in \{k-1, k\}$, the output of $h\text{-}\hat{\mathcal{S}}\mathcal{E}_{0:h}$ and that of $h\text{-}\mathcal{SE}_{0:h}^{\mathcal{O}}$ are indistinguishable. This can be proven in a similar way to Lemma 2. In particular, we can use the same argument if we use the following claim instead of Claim 2.

Claim 5. *In $h\text{-}\hat{\mathcal{S}}_{0:h}$ for each $h \in \{k-1, k\}$, the following holds except with negligible probability: In every left session that reaches Stage III, the CECOM commitment in Stage I-2 of this session is valid and its committed value is a valid decommitment of the Com_{SB} commitment of Stage I-1.*

Claim 5 can be proven in a similar way to Claim 2. For completeness, we give the proof below. (Many texts are taken verbatim from the proof of Claim 2)

Proof of Claim 5. Let us say that a left session is *bad* if it reaches Stage III and either the CECOM commitment in Stage I-2 is invalid or its committed value is not a valid decommitment of the Com_{SB} commitment in Stage I-1; a left session is *good* if it is not bad. What we want to prove is that every left session is good except with negligible probability.

Roughly speaking, the proof proceeds as follows. From the soundness of WIPOK, if a left session is bad, then in Stage II-2 of this left session, the committed value of the CECOM commitment is r_P , which is the committed value of the $\text{CCACom}^{1:1}$ commitment of Stage II-1; thus, before r_P is decommitted to in Stage II-3, we can obtain r_P by extracting the committed value from CECOM in Stage II-2. This itself does not contradict to the hiding property of $\text{CCACom}^{1:1}$ since $h\text{-}\hat{\mathcal{S}}_{0:h}$ runs in super-polynomial time in the brute-force extraction of CECOM and $\text{CCACom}^{1:1}$. Thus, we again replace the brute-force extraction with the concurrent extraction of CECOM and an oracle access to the one-session committed-value oracle \mathcal{O} of $\text{CCACom}^{1:1}$, and use the one-one CCA-security of $\text{CCACom}^{1:1}$ instead of its hiding property. Here, since we want to use the one-one CCA-security of $\text{CCACom}^{1:1}$, we perform the concurrent extraction of CECOM so that the $\text{CCACom}^{1:1}$ commitment in a left session and the $\text{CCACom}^{1:1}$ in the \tilde{i}^* -th right session are not rewound. Details are given below.

Assume for contradiction that there exists $h \in \{k-1, k\}$ such that in $h\text{-}\hat{\mathcal{S}}_{0:h}$, a left session is bad with non-negligible probability. (Here, the indices of the left sessions are determined by the order in which Stage III begins; the reason why we define the indices in this way will become clear later.) Then, there exists $i^* \in [m]$ such that in $h\text{-}\hat{\mathcal{S}}_{0:h}$, the first $(i^* - 1)$ left sessions are good except with negligible probability but the i^* -th left session is bad with non-negligible probability. Note that from the soundness of WIPOK, when the i^* -th left session is bad, the committed value of the CECOM commitment in Stage II-2 is r_P in the i^* -th left session except with negligible probability, where r_P is the value committed to in Stage II-1 of the i^* -th left session. In the following, we use BAD to denote

the event that the i^* -th left session is bad, and use CHEAT to denote the event that the committed value of the CECOM commitment in Stage II-2 is r_P in the i^* -th left session. Then, let us consider the following hybrids.

Hybrid simulator $h\text{-}\hat{\mathcal{S}}_{0:h:0}$ is the same as $h\text{-}\hat{\mathcal{S}}_{0:h}$. From our assumption, BAD occurs in $h\text{-}\hat{\mathcal{S}}_{0:h:0}$ with non-negligible probability. Thus, from the above argument, CHEAT occurs in $h\text{-}\hat{\mathcal{S}}_{0:h:0}$ with non-negligible probability.

Hybrid simulator $h\text{-}\hat{\mathcal{S}}_{0:h:1}$ is the same as $h\text{-}\hat{\mathcal{S}}_{0:h:0}$ except that $h\text{-}\hat{\mathcal{S}}_{0:h:1}$ terminates just before Stage III of the i^* -th left session begins. Clearly, BAD and CHEAT also occur in $h\text{-}\hat{\mathcal{S}}_{0:h:1}$ with non-negligible probability.

Hybrid simulator $h\text{-}\mathcal{S}_{0:h:1}^O$ emulates $h\text{-}\hat{\mathcal{S}}_{0:h:1}$ in polynomial time as follows.

- At the beginning, a random left session s is chosen. (Here, we guess that session s is the i^* -th left session.)
- In every left session, in Stage I-2, the committed value (r_V, d) is extracted by the robust concurrent extractor of CECOM in such a way that the $\text{CCACOM}^{1:1}$ commitment of left session s and the $\widetilde{\text{CCACOM}}^{1:1}$ commitment of the $\widetilde{i^*}$ -th right session are not rewind. In addition, in the $\widetilde{i^*}$ -th right session, the committed value of $\text{CCACOM}^{1:1}$ is extracted by forwarding the commitment to \mathcal{O} .
- In left session s , the committed value is also extracted in Stage II-2 by the robust concurrent extractor of CECOM without rewinding the $\text{CCACOM}^{1:1}$ commitment of the $\widetilde{i^*}$ -th right session.

Note that when Stage III of a left session is executed, the CECOM commitment in Stage I-2 of that session is valid except with negligible probability (since that session is one of the first $(i^* - 1)$ left sessions and therefore it is good except with negligible probability). Thus, the values extracted from the concurrent extractor are equal to the values that would be extracted by brute force except with negligible probability; therefore, $h\text{-}\mathcal{S}_{0:h:1}^O$ statistically emulates $h\text{-}\hat{\mathcal{S}}_{0:h:1}$, and BAD and CHEAT occur in $h\text{-}\mathcal{S}_{0:h:1}^O$ with non-negligible probability.

Note that session s is the i^* -th left session with non-negligible probability. Then, since CHEAT occurs in $h\text{-}\mathcal{S}_{0:h:1}^O$ with non-negligible probability, r_P is extracted from the CECOM commitment in Stage II-2 of session s with non-negligible probability, where r_P is the value committed to in Stage II-1 of session s . Then, since the $\text{CCACOM}^{1:1}$ commitment of session s is not rewind in $h\text{-}\mathcal{S}_{0:h:1}^O$, we can break the one-one CCA security of $\text{CCACOM}^{1:1}$. Thus, we reach a contradiction. \square

As argued above, Claim 5 implies that for each $h \in \{k - 1, k\}$, the outputs of $h\text{-}\hat{\mathcal{S}}\mathcal{E}_{0:h}$ and $h\text{-}\mathcal{S}\mathcal{E}_{0:h}^O$ are indistinguishable.

To show that the outputs of $h\text{-}\hat{\mathcal{S}}\mathcal{E}_{0:k-1}$ and $h\text{-}\hat{\mathcal{S}}\mathcal{E}_{0:k}$ are indistinguishable, it remains to prove that the outputs of $h\text{-}\mathcal{S}\mathcal{E}_{0:k-1}^O$ and $h\text{-}\mathcal{S}\mathcal{E}_{0:k}^O$ are indistinguishable. This can be shown as follows. Observe that $h\text{-}\mathcal{S}\mathcal{E}_{0:k-1}^O$ and $h\text{-}\mathcal{S}\mathcal{E}_{0:k}^O$ differ only in the k -th ExtCOM commitment of the CECOM commitment of the $\widetilde{i^*}$ -th right session, and this ExtCOM commitment is not rewind in $h\text{-}\mathcal{S}\mathcal{E}_{0:k-1}^O$ and $h\text{-}\mathcal{S}\mathcal{E}_{0:k}^O$. In addition, $h\text{-}\mathcal{S}\mathcal{E}_{0:k-1}^O$ and $h\text{-}\mathcal{S}\mathcal{E}_{0:k}^O$ run in polynomial time given oracle access to the one-session committed-value oracle \mathcal{O} of $\text{CCACOM}^{1:1}$. Thus, from the hiding property of ExtCOM and the 4-robustness of $\text{CCACOM}^{1:1}$, the output of $\mathcal{S}\mathcal{E}_{0:k-1}^O$ and that of $h\text{-}\mathcal{S}\mathcal{E}_{0:k}^O$ are indistinguishable.

Thus, we conclude that the probability that \widetilde{r}_V is extracted in $h\text{-}\hat{\mathcal{S}}\mathcal{E}_1$ is non-negligible. This concludes the proof of Claim 4. \square

By using essentially the same argument as in the proof of Claim 4, we can show that \tilde{r}_V is extracted with non-negligible probability also in $h\text{-}\hat{\mathcal{S}}\mathcal{E}_2$, $h\text{-}\hat{\mathcal{S}}\mathcal{E}_3$, and $h\text{-}\hat{\mathcal{S}}\mathcal{E}_4$. For example, let us consider $h\text{-}\hat{\mathcal{S}}\mathcal{E}_2$. Recall that $h\text{-}\hat{\mathcal{S}}\mathcal{E}_2$ differs from $h\text{-}\hat{\mathcal{S}}\mathcal{E}_1$ only in that the different witness is used in WIPOK of the \tilde{i}^* -th right session. Then, in the same way as in the proof of Claim 4, we can define hybrid simulator-extractors $h\text{-}\mathcal{SE}_1^O$ and $h\text{-}\mathcal{SE}_2^O$ such that the following hold.

- Given oracle access to the one-session committed-value oracle O of $\text{CCACom}^{1:1}$, both $h\text{-}\mathcal{SE}_1^O$ and $h\text{-}\mathcal{SE}_2^O$ run in polynomial-time.
- For each $k \in \{1, 2\}$, the probability that \tilde{r}_V is extracted in $h\text{-}\mathcal{SE}_k^O$ is statistically close to the probability in $h\text{-}\hat{\mathcal{S}}\mathcal{E}_k$.
- $h\text{-}\mathcal{SE}_1^O$ and $h\text{-}\mathcal{SE}_2^O$ differ only in the witness used in WIPOK of the \tilde{i}^* -th right session, and this WIPOK is not rewound in both $h\text{-}\mathcal{SE}_1^O$ and $h\text{-}\mathcal{SE}_2^O$.

Assume for contradiction that \tilde{r}_V is extracted in $h\text{-}\hat{\mathcal{S}}\mathcal{E}_2$ only with negligible probability. Then, since \tilde{r}_V is extracted in $h\text{-}\hat{\mathcal{S}}\mathcal{E}_1$ with non-negligible probability, we can break witness indistinguishability of WIPOK and the 4-robustness of $\text{CCACom}^{1:1}$ by using $h\text{-}\mathcal{SE}_1^O$ and $h\text{-}\mathcal{SE}_2^O$. Thus, \tilde{r}_V is extracted in $h\text{-}\hat{\mathcal{S}}\mathcal{E}_2$ with non-negligible probability. In this way, we can show that \tilde{r}_V is extracted also in $h\text{-}\hat{\mathcal{S}}\mathcal{E}_3$ and $h\text{-}\hat{\mathcal{S}}\mathcal{E}_4$ with non-negligible probability.

Concluding the proof of Claim 1. In $h\text{-}\hat{\mathcal{S}}\mathcal{E}_4$, the \tilde{i}^* -th right session is independent of \tilde{r}_V , and therefore the probability that \tilde{r}_V is extracted is negligible. However, we show above that this probability is non-negligible. Thus, we reach a contradiction. \square

This concludes the proof of Theorem 2. \square

References

- [BPS06] Boaz Barak, Manoj Prabhakaran, and Amit Sahai. Concurrent non-malleable zero knowledge. In *FOCS*, pages 345–354, 2006.
- [CKPR02] Ran Canetti, Joe Kilian, Erez Petrank, and Alon Rosen. Black-box concurrent zero-knowledge requires (almost) logarithmically many rounds. *SIAM J. Comput.*, 32(1):1–47, 2002.
- [CLP10] Ran Canetti, Huijia Lin, and Rafael Pass. Adaptive hardness and composable security in the plain model from standard assumptions. In *FOCS*, pages 541–550, 2010.
- [DDN00] Danny Dolev, Cynthia Dwork, and Moni Naor. Nonmalleable cryptography. *SIAM J. Comput.*, 30(2):391–437, 2000.
- [DNS04] Cynthia Dwork, Moni Naor, and Amit Sahai. Concurrent zero-knowledge. *J. ACM*, 51(6):851–898, 2004.
- [DPP98] Ivan Damgård, Torben P. Pedersen, and Birgit Pfitzmann. Statistical secrecy and multibit commitments. *IEEE Transactions on Information Theory*, 44(3):1143–1151, 1998.
- [GK96] Oded Goldreich and Ariel Kahan. How to construct constant-round zero-knowledge proof systems for NP. *J. Cryptology*, 9(3):167–190, 1996.

- [GLP⁺15] Vipul Goyal, Huijia Lin, Omkant Pandey, Rafael Pass, and Amit Sahai. Round-efficient concurrently composable secure computation via a robust extraction lemma. In *TCC*, 2015.
- [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.
- [HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28(4):1364–1396, 1999.
- [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.
- [Kiy14] Susumu Kiyoshima. Round-efficient black-box construction of composable multi-party computation. In *CRYPTO*, pages 351–368, 2014.
- [KMO14] Susumu Kiyoshima, Yoshifumi Manabe, and Tatsuaki Okamoto. Constant-round black-box construction of composable multi-party computation protocol. In *TCC*, pages 343–367, 2014.
- [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 *TCC*, pages 274–292, 2011.
- [LP12] Huijia Lin and Rafael Pass. Black-box constructions of composable protocols without set-up. In *CRYPTO*, pages 461–478, 2012.
- [LPTV10] Huijia Lin, Rafael Pass, Wei-Lung Dustin Tseng, and Muthuramakrishnan Venkatasubramanian. Concurrent non-malleable zero knowledge proofs. In *CRYPTO*, 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.
- [Nao91] Moni Naor. Bit commitment using pseudorandomness. *J. Cryptology*, 4(2):151–158, 1991.
- [NY89] Moni Naor and Moti Yung. Universal one-way hash functions and their cryptographic applications. In *STOC*, pages 33–43, 1989.
- [OOR⁺14] Claudio Orlandi, Rafail Ostrovsky, Vanishree Rao, Amit Sahai, and Ivan Visconti. Statistical concurrent non-malleable zero knowledge. In *TCC*, pages 167–191, 2014.
- [OPV10] Rafail Ostrovsky, Omkant Pandey, and Ivan Visconti. Efficiency preserving transformations for concurrent non-malleable zero knowledge. In *TCC*, pages 535–552, 2010.
- [PR05] Rafael Pass and Alon Rosen. New and improved constructions of non-malleable cryptographic protocols. In *STOC*, pages 533–542, 2005.

- [PRS02] Manoj Prabhakaran, Alon Rosen, and Amit Sahai. Concurrent zero knowledge with logarithmic round-complexity. In *FOCS*, pages 366–375, 2002.
- [PTV12] Rafael Pass, Wei-Lung Dustin Tseng, and Muthuramakrishnan Venkatasubramanian. Concurrent zero knowledge, revisited. *J. Cryptology*, pages 1–22, 2012.
- [PW09] Rafael Pass and Hoeteck Wee. Black-box constructions of two-party protocols from one-way functions. In *TCC*, pages 403–418, 2009.
- [PW10] Rafael Pass and Hoeteck Wee. Constant-round non-malleable commitments from sub-exponential one-way functions. In *EUROCRYPT*, pages 638–655, 2010.
- [Ven14] Muthuramakrishnan Venkatasubramanian. On adaptively secure protocols. In *SCN*, pages 455–475, 2014.

A Robust Concurrent Extraction Lemma [GLP⁺15]

Below, we give a formal description of the robust concurrent extraction lemma [GLP⁺15].

The external protocol Π . Let $\Pi := \langle B, A \rangle$ be an arbitrary two-party computation protocol. For security parameter n , let $\text{dom}_B(n)$ denote the domain of the input for B and $k := k(n)$ denote the round complexity of Π .

The robust-concurrent attack. Let $x \in \text{dom}_B(n)$. In the *robust-concurrent attack*, the adversary \mathcal{A} interacts with a special (possibly super-polynomial-time) party \mathcal{E} called the *online extractor*. Online extractor \mathcal{E} simultaneously participates in one execution of Π and several executions of CECom , where \mathcal{E} interacts with \mathcal{A} as honest $B(1^n, x)$ in the execution of Π and interacts with \mathcal{A} as a honest receiver in each execution of CECom . The scheduling of all messages in all sessions— Π as well as CECom —is controlled by \mathcal{A} . When \mathcal{A} successfully completes a CECom commitment s , online extractor \mathcal{E} sends a string α_s to \mathcal{A} .

For $n \in \mathbb{N}, x \in \text{dom}_B(n), z \in \{0, 1\}^*$, let $\text{REAL}_{\mathcal{E}, \Pi}^{\mathcal{A}}(n, x, z)$ denote the output of the following probabilistic experiment. On input 1^n and auxiliary input z , the experiment starts an execution of \mathcal{A} , which launches the robust-concurrent attack by interacting with \mathcal{E} . The output of the experiment is the view of \mathcal{A} .

The robust concurrent extraction lemma. Informally speaking, the lemma states that there exists an interactive Turing machine—called *robust simulator*—whose output is statistically close to $\text{REAL}_{\mathcal{E}, \Pi}^{\mathcal{A}}(n, x, z)$ even if the value that the online extractor \mathcal{E} returns to \mathcal{A} at the end of each successful CECom commitment is the committed value of this commitment. Furthermore, the robust simulator does not “rewind” B and runs in time polynomial in total sessions opened by \mathcal{A} .

Lemma 3 (Robust Concurrent Extraction Lemma [GLP⁺15]). *There exists an interactive Turing machine \mathcal{S} (called **robust simulator**) such that for every adversary \mathcal{A} and every two-party protocol $\Pi := \langle B, A \rangle$, there exists a party \mathcal{E} (called **online extractor**) such that for every $n \in \mathbb{N}, x \in \text{dom}_B(n)$, and $z \in \{0, 1\}^*$, the following conditions hold:*

1. **Validity constraint.** *For every output ρ of $\text{REAL}_{\mathcal{E}, \Pi}^{\mathcal{A}}(n, x, z)$ and for every CECom commitment s appearing in ρ , if there exists a unique value $v \in \{0, 1\}^n$ to which the commitment s can be decommitted, then:*

$$\alpha_s = v,$$

where α_s is the value \mathcal{E} sends to \mathcal{A} at the end of s .

2. **Statistical simulation.** Let $k = k(n)$ be the round complexity of Π . Then the statistical distance between $\text{REAL}_{\mathcal{E}, \Pi}^{\mathcal{A}}(n, x, z)$ and $\text{out}_{\mathcal{S}}(B(1^n, x) \leftrightarrow \mathcal{S}^{\mathcal{A}}(1^n, z))$ is given by

$$\Delta(n) \leq 2^{-\Omega(\ell - k \cdot \log T(n))},$$

where $\ell := \ell(n)$ is the parameter of CECom and $T(n)$ is the number of the CECom commitments between \mathcal{A} and \mathcal{E} . Furthermore, the running time of \mathcal{S} is $\text{poly}(n) \cdot T(n)^2$.

B Constant-round One-one CCA-secure Commitment Scheme from OWF

In this section, we observe that from a result by Goyal et al. [GLP⁺15], it follows almost immediately that we can obtain a constant-round one-one CCA-secure commitment scheme from one-way functions.

Theorem 3. *Assume the existence of one-way functions. Then, for any constant $\kappa \in \mathbb{N}$, there exists a constant-round κ -robust one-one CCA-secure commitment scheme $\text{CCACom}^{1:1}$.*

We use the following building blocks, where all of them can be constructed from one-way functions.

- Constant-round commitment scheme NMCom that is non-malleable w.r.t. itself and any 4-round protocol. Specifically, we use the scheme by Lin and Pass [PW10]. We remark that, like many other non-malleable commitment schemes, the scheme by [PW10] also satisfies extractability.¹⁰ (For the definitions of non-malleability and extractability, see Appendix C.)
- Four-round witness-indistinguishable proof of knowledge WIPOK .
- Constant-round zero-knowledge argument ZKArg [GK96].
- Concurrently extractable commitment scheme CECom of Micciancio et al. [MOSV06] with parameter $\ell = \max(\kappa, r_{\text{NM}}, 4) + 1$, where r_{NM} is the round complexity of NMCom (cf. Figure 1.) When $\ell = \max(\kappa, r_{\text{NM}}, 4) + 1 = O(1)$, CECom does not guarantee concurrent extractability. It is easy to see, however, that it guarantees the following “robust extractability” property: For any adversarial committer C^* that commits to a value in a *single* session of CECom and simultaneously participates an arbitrary $\max(\kappa, r_{\text{NM}}, 4)$ -round protocol Π , the extractor can extract the committed value from C^* without rewinding Π . (For details, see Appendix D.)

$\text{CCACom}^{1:1}$ is shown in Figure 6. We remark that $\text{CCACom}^{1:1}$ is almost identical to the CCA-secure commitment scheme of Goyal et al. [GLP⁺15]; essentially, the only difference is the parameter ℓ of CECom . We prove its one-one CCA security in Section B.1 and prove its robustness in Section B.2.

B.1 Proof of One-one CCA Security

For any adversary \mathcal{A} that interacts with the committed-value oracle only in a single session, we show that the following ensembles are computationally indistinguishable.

- $\{\text{IND}_0(\langle C, R \rangle, \mathcal{A}, n, z)\}_{n \in \mathbb{N}, z \in \{0,1\}^*}$

¹⁰In the scheme of [PW10], the committer proves by a witness-indistinguishable proof of knowledge system that it knows either the committed value or trapdoor information. Since the scheme is designed so that the trapdoor is hidden from the committer, the committed value can be extracted by extracting the witness from the witness-indistinguishable proof.

Below, CECOM is the scheme of Micciancio et al. [MOSV06] with parameter $\ell = \max(\kappa, r_{\text{NM}}, 4) + 1$, where r_{NM} is the round complexity of NMCom.

Commit Phase

To commit to $v \in \{0, 1\}^n$, the committer C does the following with the receiver R .

Stage 1. C commits to v by using CECOM.

Stage 2. R chooses random $r \in \{0, 1\}^n$ and commits to r by using CECOM. R then prove the validity of this CECOM commitment by using ZKArg.

Stage 3. C commits to 0^n by using NMCom.

Stage 4. R decommits the CECOM commitment in Stage 2 to r .

Stage 5. C prove the following by using WIPOK:

- the CECOM commitment in Stage 1 is valid, or
- the committed value of the NMCom commitment in Stage 3 is r .

Decommit Phase

C decommits the CECOM commitment in Stage 1 to v .

Figure 6: Constant-round one-one CCA-secure commitment scheme CCACOM^{1:1}.

- $\{\text{IND}_1(\langle C, R \rangle, \mathcal{A}, n, z)\}_{n \in \mathbb{N}, z \in \{0, 1\}^*}$

Toward this end, we consider a sequence of hybrid experiments in which the left session of $\text{IND}_b(\langle C, R \rangle, \mathcal{A}, n, z)$ is gradually modified so that \mathcal{A} receives no information about v_b in the last hybrid.

Hybrid $H_0^b(n, z)$ is the same as $\text{IND}_b(\langle C, R \rangle, \mathcal{A}, n, z)$.

Hybrid $H_1^b(n, z)$ is the same as $H_0^b(n, z)$ except for the following.

- In Stage 2 of the left session, the committed value r of the CECOM commitment is extracted by brute force. If the CECOM commitment is invalid or has more then one committed value, r is defined to be \perp .
- In Stage 3 of the left session, the committed value of NMCom is switched from 0^n to r .

Hybrid $H_2^b(n, z)$ is the same as $H_1^b(n, z)$ except that in Stage 5 of the left session, the WIPOK proof is computed by using the witness for the fact that the committed value of the NMCom commitment in Stage 3 is r .

Hybrid $H_3^b(n, z)$ is the same as $H_2^b(n, z)$ except that in Stage 1 of the left session, the committed value of CECOM is switched from v_b to 0^n .

For each $i \in \{0, 1, 2, 3\}$ and $b \in \{0, 1\}$, let $H_i^b(n, z)$ be the random variable representing the output of $H_i^b(n, z)$. From the construction, \mathcal{A} receives no information about v_b in $H_3^0(n, z)$ and $H_3^1(n, z)$, and hence $H_3^0(n, z)$ and $H_3^1(n, z)$ are identically distributed. Therefore, to show the indistinguishability

between the above two ensembles, it suffices to prove that the outputs of each neighboring hybrids are computationally indistinguishable.

Our strategy for proving the indistinguishability of each neighboring hybrids is to reduce their indistinguishability to the security of NMCom, WIPOK, and CECOM. The problem in this strategy is the existence of the committed-value oracle: Since the oracle runs in super-polynomial time, the security of NMCom, WIPOK, and CECOM might not hold against the adversaries that interact with the oracle. We overcome this problem by showing that the oracle can be emulated efficiently without “disturbing” the security of NMCom, WIPOK, and CECOM. Specifically, we show that the oracle can be emulated by extracting the committed value of the CECOM commitment in Stage 1 of the right session using the extractability of CECOM; since CECOM provides a *robust* extractability property, the extraction from CECOM does not disturb the security of NMCom, WIPOK, and CECOM. We remark that in the formal argument given below, we first show that the probability that \mathcal{A} “cheats” in the hybrids is negligible, meaning that in the right session, the committed value of the NMCom commitment in Stage 3 is equal to the committed value of the CECOM commitment in Stage 2 only with negligible probability. Showing that \mathcal{A} cheats only with negligible probability is crucial to showing that the oracle can be efficiently emulated. In particular, once we show that \mathcal{A} cheats only with negligible probability, we can use the soundness of WIPOK to argue that the CECOM commitment in Stage 1 is valid in the accepted right session except with negligible probability, and thus we can conclude that the extracted value is equal to the committed value when the right session is accepted. The formal argument is given below.

Let us say that \mathcal{A} *cheats* if the committed value of NMCom in Stage 3 is equal to the committed value \tilde{r} of CECOM in Stage 2 in the accepted right session. First, we show that \mathcal{A} cheats in $H_0^b(n, z)$ only with negligible probability.

Claim 6. *The probability that \mathcal{A} cheats in $H_0^b(n, z)$ is negligible for each $b \in \{0, 1\}$.*

Proof. Roughly speaking, this claim follows from the hiding property of CECOM—when the adversary cheats, we can obtain \tilde{r} by extracting the committed value from NMCom, and thus we can obtain the committed value of a CECOM commitment before it is decommitted to. To formally implement this idea, it is important that no super-polynomial-time computation is performed during the execution of CECOM in Stage 2 of the right session. Fortunately, in $H_0^b(n, z)$ no super-polynomial-time computation is indeed performed during CECOM of the right session, as super-polynomial-time computation is performed only at the end of the right session. (Recall that in the setting of one-one CCA security, \mathcal{A} interacts with the oracle only in a single session.) The formal argument is given below.

Assume for contradiction that there exists $b \in \{0, 1\}$ such that \mathcal{A} cheats in $H_0^b(n, z)$ with non-negligible probability. Fix any such b . To derive a contradiction, we consider the following hybrid experiments.

Hybrid $H_{0:1}^b(n, z)$ is the same as $H_0^b(n, z)$ except that in Stage 3 of the right session, the committed value of the NMCom commitment is extracted by using the extractability of NMCom.

Clearly, the probability that \mathcal{A} cheats is still non-negligible in $H_{0:1}^b(n, z)$. Furthermore, from the extractability of NMCom, the extracted value is equal to \tilde{r} with non-negligible probability.

Hybrid $H_{0:2}^b(n, z)$ is the same as $H_{0:1}^b(n, z)$ except that in Stage 2 of the right session, the ZKArg proof is generated by using the simulator of ZKArg.

From the zero-knowledge property of ZKArg, the probability that \tilde{r} is extracted from NMCom is still non-negligible in $H_{0:2}^b(n, z)$.

We derive a contradiction by constructing an adversary \mathcal{B} that breaks the hiding property of CECOM. Externally, \mathcal{B} interacts with a committer of CECOM: It sends random $\tilde{r}_0, \tilde{r}_1 \in \{0, 1\}^n$ to the committer

and receives a CECOM commitment in which either \tilde{r}_0 or \tilde{r}_1 is committed to. Internally, \mathcal{B} invokes \mathcal{A} and emulates $H_{0:2}^b(n, z)$ for \mathcal{A} honestly except that in Stage 2 of the right session, \mathcal{B} forwards the CECOM commitment from the external committer to internal \mathcal{A} . Finally, if the value extracted from NMCom is \tilde{r}_1 in internally emulated $H_{0:2}^b(n, z)$, \mathcal{B} outputs 1, and otherwise, it outputs 0. If \mathcal{B} receives a commitment to \tilde{r}_1 , it outputs 1 with non-negligible probability from the above argument. On the other hand, if \mathcal{B} receives a commitment to \tilde{r}_0 , it outputs 1 only with negligible probability since internal \mathcal{A} receives no information about \tilde{r}_1 . Hence, \mathcal{B} breaks the hiding property of CECOM. \square

Next, we show that \mathcal{A} cheats only with negligible probability in $H_1^b(n, z)$, and we use it to prove that $H_0^b(n, z)$ and $H_1^b(n, z)$ are indistinguishable.

Claim 7. *For each $b \in \{0, 1\}$, the following hold.*

- *The probability that \mathcal{A} cheats in $H_1^b(n, z)$ is negligible.*
- *$\{H_0^b(n, z)\}_{n \in \mathbb{N}, z \in \{0,1\}^*}$ and $\{H_1^b(n, z)\}_{n \in \mathbb{N}, z \in \{0,1\}^*}$ are computationally indistinguishable.*

Proof. First, we show that \mathcal{A} cheats in $H_1^b(n, z)$ with negligible for each $b \in \{0, 1\}$. Roughly speaking, this follows from the non-malleability of NMCom: Since $H_1^b(n, z)$ differs from $H_0^b(n, z)$ only in the value committed to in NMCom in the left session, the value that \mathcal{A} commits to by using NMCom in the right session of $H_1^b(n, z)$ is indistinguishable from the value that \mathcal{A} commits to by using NMCom in the right session of $H_0^b(n, z)$; hence, from Claim 6, the probability that \mathcal{A} cheats in $H_1^b(n, z)$ is negligible. We remark that since the left session in $H_1^b(n, z)$ involves the brute-force extraction of CECOM in Stage 2, in the formal argument given below we consider a hybrid experiment in which brute-force extraction is replaced with the rewinding extraction. Since we want to use the non-malleability of NMCom, this extraction is performed in such a way that NMCom in the right session is not rewind. The formal argument is given below.

Assume for contradiction that there exists $b \in \{0, 1\}$ such that \mathcal{A} cheats in $H_1^b(n, z)$ with non-negligible probability. Fix any such b . To derive a contradiction, we consider the following hybrid experiment for $i \in \{0, 1\}$.

Hybrid $G_i^b(n, z)$ is the same as $H_i^b(n, z)$ except for the following.

- In Stage 1 of the left session, the committed value r of the CECOM commitment is extracted by using the extractability of CECOM instead of by brute force. Furthermore, this extraction is performed in such a way that the NMCom commitment in the right session is not rewind.
- $G_i^b(n, z)$ terminates immediately after NMCom ends in Stage 3 of the right session.

From the soundness of ZKArg, the CECOM commitment in Stage 2 of the left session is valid when the ZKArg proof in Stage 2 of the left session is accepted. Hence, when Stage 3 is executed in the left session, the value extracted from the CECOM commitment in Stage 2 is equal to its (unique) committed value. Since the only difference from $G_i^b(n, z)$ and $H_i^b(n, z)$ is how r is extracted, the view of \mathcal{A} in $G_i^b(n, z)$ is statistically close to that in $H_i^b(n, z)$. Therefore, \mathcal{A} cheats in $G_0^b(n, z)$ with negligible probability from Claim 6, and \mathcal{A} cheats in $G_1^b(n, z)$ with non-negligible probability from our hypothesis.

We then derive a contradiction by constructing an adversary \mathcal{M} that breaks the non-malleability of NMCom. Externally, \mathcal{M} interacts with a committer and a receiver of NMCom: It sends 0^n and $r \in \{0, 1\}^n$ to the committer and receives a NMCom commitment in which either 0^n or r is committed to; at the same time, it sends a CECOM commitment to the receiver. Internally, \mathcal{M} invokes \mathcal{A} and emulates $G_0^b(n, z)$ for \mathcal{A} honestly except for the following.

- After r is extracted in Stage 2 of the left session, \mathcal{M} sends 0^n and r to the external committer.
- In Stage 3 of the left session, \mathcal{M} forwards the NMCCom commitment from the external committer to internal \mathcal{A} .
- In Stage 3 of the right session, \mathcal{M} forwards the NMCCom commitment from the internal \mathcal{A} to the external receiver.

From the construction, \mathcal{M} perfectly emulates $G_0^b(n, z)$ when it receives a NMCCom commitment to 0^n , and it perfectly emulates $G_1^b(n, z)$ when it receives a NMCCom commitment to r . Hence, when \mathcal{M} receives a NMCCom commitment to 0^n , internal \mathcal{A} cheats with negligible probability, and when \mathcal{M} receives a NMCCom commitment to r , internal \mathcal{A} cheats with non-negligible probability. Then, since the cheating of \mathcal{A} is efficiently recognizable given the view of \mathcal{M} and the committed value of the NMCCom commitment in the right session, \mathcal{M} breaks the non-malleability of NMCCom.

Next, we show that $\{H_0^b(n, z)\}_{n \in \mathbb{N}, z \in \{0,1\}^*}$ and $\{H_1^b(n, z)\}_{n \in \mathbb{N}, z \in \{0,1\}^*}$ are computationally indistinguishable. Roughly speaking, this indistinguishability follows from the hiding of NMCCom: Since \mathcal{A} cheats only with negligible probability both in $H_0^b(n, z)$ and in $H_1^b(n, z)$, the CECCom commitment in Stage 1 is valid in the accepted right session in both hybrids; hence the committed-value oracle can be efficiently emulated by extracting the committed value of the CECCom commitment in Stage 1, and thus the indistinguishability follows from the hiding property of NMCCom. Here, since we want to use the hiding property of NMCCom, the extraction from CECCom is performed in such a way that NMCCom in the left session is not rewound. The formal argument is given below.

Assume for contradiction that there exists $b \in \{0, 1\}$ such that $\{H_0^b(n, z)\}_{n \in \mathbb{N}, z \in \{0,1\}^*}$ and $\{H_1^b(n, z)\}_{n \in \mathbb{N}, z \in \{0,1\}^*}$ are distinguishable. Fix any such b . From Claim 6 and what is shown above, \mathcal{A} cheats only with negligible probability both in $H_0^b(n, z)$ and in $H_1^b(n, z)$. Hence, from the soundness of WIPOK, the CECCom commitment in Stage 1 is invalid in the accepted right session only with negligible probability. Therefore, for infinitely many n , there exists $z \in \{0, 1\}^*$ such that (i) $H_0^b(n, z)$ and $H_1^b(n, z)$ are distinguishable with advantage $1/\text{poly}(n)$ and (ii) the CECCom commitment in Stage 1 of the left session is invalid in the accepted right session with probability at most $1/2\text{poly}(n)$ in both $H_0^b(n, z)$ and $H_1^b(n, z)$. Fix any such n and z . From an average argument, there exists a partial transcript ρ of $H_0^b(n, z)$ up until the end of Stage 2 of the left session such that under the condition that a prefix of the transcript is ρ , both of the above (i) and (ii) hold. Let r be the value that is committed to in Stage 2 of the left session in ρ . We consider the following two cases.

Case 1. Stage 1 of the right session has already started in ρ . Since the committed value of a CECCom commitment is determined by the first message, ρ uniquely determined the committed value \bar{v} of the CECCom commitment in Stage 1 of the right session. Notice that given ρ , r , and \bar{v} as auxiliary input, $H_0^b(n, z)$ and $H_1^b(n, z)$ can be executed from ρ in polynomial time. Hence, we can derive a contradiction by considering an adversary that breaks the hiding property of NMCCom by internally emulating $H_0^b(n, z)$ from ρ and forwarding a NMCCom commitment from the external committer (who commits to either 0^n or r) to internally emulated \mathcal{A} .

Case 2. Stage 1 of the right session starts after ρ . We consider the following hybrid experiment.

In Hybrid $F_i^b(n, z)$, $H_i^b(n, z)$ is executed from ρ honestly except for the following.

- In the left session, brute-force extraction of r is not performed, and hardwired r is used.
- In Stage 1 of the right session, the committed value \bar{v} of the CECCom commitment is extracted by using the extractability of CECCom in such a way that NMCCom in Stage 3 is not rewound in the left session.

- At the end of the right session, the extracted value \widetilde{v} is returned to \mathcal{A} as the committed value of the right session.

From the definition of ρ , the CECOM commitment in Stage 1 of the right session is invalid in the accepted right session with probability at most $1/2\text{poly}(n)$. Since the output of $F_i^b(n, z)$ differs from that of $H_i^b(n, z)$ only when the correct committed value is not extracted in the accepted right session (which occurs with probability at most $1/2\text{poly}(n)$ from the above), from our hypothesis, the outputs of $F_0^b(n, z)$ and $F_1^b(n, z)$ are distinguishable with advantage $1/2\text{poly}(n)$. Then, since $F_0^b(n, z)$ and $F_1^b(n, z)$ differ only in the value committed to in NMCom and since both experiments run in polynomial time, we can derive a contradiction by considering an adversary that internally emulates $F_0^b(n, z)$ and forwards a NMCom commitment from the external committer to internally emulated \mathcal{A} . \square

In the same way above, we can prove that the outputs of the other neighboring hybrids are also indistinguishable.

Claim 8. *For each $b \in \{0, 1\}$, the following hold.*

- *The probability that \mathcal{A} cheats in $H_2^b(n, z)$ is negligible.*
- *$\{H_1^b(n, z)\}_{n \in \mathbb{N}, z \in \{0, 1\}^*}$ and $\{H_2^b(n, z)\}_{n \in \mathbb{N}, z \in \{0, 1\}^*}$ are computationally indistinguishable.*

Proof. This claim can be proven in essentially the same way as Claim 7. First, we can show that \mathcal{A} cheats in $H_2^b(n, z)$ only with negligible probability by using the same argument except that we use the non-malleability w.r.t. $\max(\kappa, r_{\text{NM}}, 4)$ -round protocols of NMCom instead of the non-malleability w.r.t. itself. (Recall that $H_2^b(n, z)$ differs from $H_1^b(n, z)$ only in the witness used in WIPOK, which has four rounds.) Next, we can show the indistinguishability between $\{H_1^b(n, z)\}_{n \in \mathbb{N}, z \in \{0, 1\}^*}$ and $\{H_2^b(n, z)\}_{n \in \mathbb{N}, z \in \{0, 1\}^*}$ by using the same argument except that we use the witness indistinguishability of WIPOK instead of the hiding property of NMCom. We omit the formal proof. \square

Claim 9. *For each $b \in \{0, 1\}$, the following hold.*

- *The probability that \mathcal{A} cheats in $H_3^b(n, z)$ is negligible.*
- *$\{H_2^b(n, z)\}_{n \in \mathbb{N}, z \in \{0, 1\}^*}$ and $\{H_3^b(n, z)\}_{n \in \mathbb{N}, z \in \{0, 1\}^*}$ are computationally indistinguishable.*

Proof. Like Claim 8, this claim can be proven in essentially the same way as Claim 7. We remark however that since the round complexity of CECOM is much more than four, we need to consider a sequence of intermediate hybrid experiments in which the committed value of ExtCom in CECOM are switched one by one. We omit the formal proof. \square

This concludes the proof of one-one CCA security.

B.2 Proof of κ -robustness

We show that there exists a simulator \mathcal{S} such that for any adversary \mathcal{A} that interacts with the committed-value oracle only in a single session, and for any κ -round PPT ITM B , the following are computationally indistinguishable:

- $\left\{ \text{out}_{B, \mathcal{A}^O} \left\langle B(1^n, x, y) \leftrightarrow \mathcal{A}^O(1^n, x, z) \right\rangle \right\}_{n \in \mathbb{N}, x, y, z \in \{0, 1\}^n}$
- $\left\{ \text{out}_{B, \mathcal{S}^{\mathcal{A}}} \left\langle B(1^n, x, y) \leftrightarrow \mathcal{S}^{\mathcal{A}}(1^n, x, z) \right\rangle \right\}_{n \in \mathbb{N}, x, y, z \in \{0, 1\}^n}$

This can be shown easily by using the argument we used in the proof of one-one CCA security. Roughly, we consider a simulator that emulates O for \mathcal{A} efficiently by extracting the committed value of the CECOM commitment in Stage 1 using the robust extractability of CECOM in such a way that the interaction with B is not rewound. (Since we set $\ell = \max(\kappa, r_{\text{NM}}, 4) + 1$, such extraction is possible.) To show that this simulator indeed emulates the oracle for \mathcal{A} , we need to show that the CECOM commitment in Stage 1 is invalid in the accepted right session only with negligible probability. This can be shown by using the argument in the proof of Claim 6. Hence, by using this simulator, we can prove the κ -robustness. The formal proof is omitted.

C Additional Definitions

In this section, we give the definitions that are used in Appendix B.

C.1 Non-malleable Commitment Schemes

Non-malleability w.r.t. Itself

We recall the definition of non-malleable commitment schemes from [LPV08]. For convenience, we use a slightly different presentation (based on indistinguishability rather than simulation), which is used in [LP09, LP11]. Let $\langle C, R \rangle$ be a tag-based commitment scheme. For any man-in-the-middle adversary \mathcal{M} , consider the following experiment. On input security parameter $n \in \mathbb{N}$ and auxiliary input $z \in \{0, 1\}^*$, \mathcal{M} participates in one left and one right interactions simultaneously. In the left interaction, \mathcal{M} interacts with the committer of $\langle C, R \rangle$ and receives a commitment to value v using identity $\text{id} \in \{0, 1\}^n$ of its choice. In the right interaction, \mathcal{M} interacts with the receiver of $\langle C, R \rangle$ and gives a commitment using identity $\widetilde{\text{id}}$ of its choice. Let \widetilde{v} be the value that \mathcal{M} commits to on the right. If the right commitment is invalid or undefined, \widetilde{v} is defined to be \perp . If $\text{id} = \widetilde{\text{id}}$, value \widetilde{v} is also defined to be \perp . Let $\text{mim}(\langle C, R \rangle, \mathcal{M}, v, z)$ denote a random variable representing \widetilde{v} and the view of \mathcal{M} in the above experiment.

Definition 5. A commitment scheme $\langle C, R \rangle$ is **non-malleable** if for any PPT man-in-the-middle adversary \mathcal{M} , the following are computationally indistinguishable.

- $\{\text{mim}(\langle C, R \rangle, \mathcal{M}, v, z)\}_{n \in \mathbb{N}, v \in \{0, 1\}^n, v' \in \{0, 1\}^n, z \in \{0, 1\}^*}$
- $\{\text{mim}(\langle C, R \rangle, \mathcal{M}, v', z)\}_{n \in \mathbb{N}, v \in \{0, 1\}^n, v' \in \{0, 1\}^n, z \in \{0, 1\}^*}$

Non-malleability w.r.t. κ -round Protocols

We recall the definition of non-malleability w.r.t. κ -round protocols from [LP09]. In [LP09], this property is also referred to as κ -robustness. We refer to this property as non-malleability w.r.t. κ -round protocols to distinguish it from the κ -robustness for CCA secure commitment schemes, which is also used in this work.

Consider a man-in-the-middle adversary \mathcal{M} that participates in a left interaction—communicating with a machine B —and a right interaction—acting as a committer by using the commitment scheme $\langle C, R \rangle$. As in the standard definition of non-malleability, \mathcal{M} can choose the identity in the right interaction. We denote by $\text{mim}(\langle C, R \rangle, B, \mathcal{M}, y, z)$ the random variable consisting of the view of $\mathcal{M}(z)$ in a man-in-the-middle execution when communicating with $B(y)$ on the left and a honest receiver on the right, combined with the values $\mathcal{M}(z)$ commits to on the right. Intuitively, we say that $\langle C, R \rangle$ is non-malleable w.r.t. B if $\text{mim}(\langle C, R \rangle, B, \mathcal{M}, y_1, z)$ and $\text{mim}(\langle C, R \rangle, B, \mathcal{M}, y_2, z)$ are indistinguishable whenever interactions with $B(y_1)$ and $B(y_2)$ cannot be distinguished.

Definition 6. Let $\langle C, R \rangle$ be a commitment scheme and B be a PPT ITM. We say that the commitment scheme $\langle C, R \rangle$ is **non-malleable w.r.t. B** if for every two sequences $\{y_n^1\}_{n \in \mathbb{N}}$ and $\{y_n^2\}_{n \in \mathbb{N}}$, $y_n^1, y_n^2 \in \{0, 1\}^n$ such that for all PPT ITM \mathcal{A} it holds that

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

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

$$\{\text{mim}(\langle C, R \rangle, B, \mathcal{M}, y_n^1, z)\}_{n \in \mathbb{N}, z \in \{0, 1\}^*} \approx \{\text{mim}(\langle C, R \rangle, B, \mathcal{M}, y_n^2, z)\}_{n \in \mathbb{N}, z \in \{0, 1\}^*} .$$

We say that $\langle C, R \rangle$ is **non-malleable w.r.t. κ -round protocols** if $\langle C, R \rangle$ is non-malleable w.r.t. any machine B that interacts with the man-in-the-middle adversary in κ rounds.

C.2 Extractable Commitment Scheme

We recall the definition of *extractable commitment schemes* from [PW09]. Roughly speaking, a commitment scheme is *extractable* if there exists an expected polynomial-time oracle machine (called an *extractor*) E such that for any committer C^* that generates a commitment, E^{C^*} extracts the committed value when the commitment is valid. We note that when the commitment is invalid, E can output an arbitrary garbage value.

Formally, extractable commitment schemes are defined as follows. A commitment scheme $\langle C, R \rangle$ is *extractable* if there exists an expected polynomial-time probabilistic extractor E such that for any PPT committer C^* , extractor E^{C^*} outputs a pair (τ, σ) such that

- τ is identically distributed with the view of C^* interacting with honest receiver R in the commit phase.
- If τ is accepted, then $\sigma \neq \perp$ except with negligible probability.
- If $\sigma \neq \perp$, then it is statistically impossible to decommit τ to any value other than σ .

D On the Robust Extractability of CCom

In this section, we observe that for any constant $\kappa \in \mathbb{N}$, CCom with parameter $\ell = \kappa + 1$ satisfies the following robust extractability property: For any adversarial committer C^* that commits to a value in a single session of CCom and simultaneously participates an arbitrary κ -round protocol Π , the extractor can extract the committed value from C^* without rewinding Π . This property is used to obtain constant-round one-one CCA-secure commitment scheme in Appendix B.

Recall that in CCom, the extractable commitment scheme of [PW09] is executed ℓ times in the following schedule:

1. First, the commit-stage messages of all sessions (of ExtCom) are exchanged in parallel.
2. Subsequently, the challenge-stage message and the reply-stage message of the i -th session are exchanged for each $i \in [\ell]$ in sequence.

Let us call the pair of the challenge-stage message and the reply-stage message of a ExtCom commitment a *slot*. Since the committed value of a ExtCom commitment can be extracted by rewinding the slot and obtaining a new pair of the challenge-stage message and the reply-stage message (cf. Figure 2 in Section 3.3), the committed value of a CCom commitment can be extracted by rewinding any of the ℓ sequential slots.

Consider the following extractor E . Externally, E participates in a κ -round protocol Π . Internally, E invokes the cheating committer C^* and forwards all messages of Π from the external party to internal C^* and vice versa; additionally, E interacts with C^* in a session of CECOM as a honest receiver. When the session of CECOM ends, E extracts the committed value of the session by rewinding C^* in a slot that does not “interleave” any message of Π (i.e, a slot such that C^* does not exchange any message of Π after receiving the challenge message of the slot until it sends the reply-message of the slot; notice that such a slot always exists because there are $\ell = \kappa + 1$ sequential slots). Specifically, E continues to rewind such a slot until it obtains a new pair of the challenge-stage message and the reply-stage message. If C^* requires a message of Π after being rewound, E cuts off the execution of C^* immediately and rewinds C^* again. After obtaining a new pair of the challenge-stage message and the reply-stage message, it extracts the committed value by using them.

From the construction, E perfectly emulates the view of C^* and does not rewind the external protocol Π . Also, from the extractability of ExtCom, the extraction fails only with negligible probability. Hence, it remains to show that E runs in (expected) polynomial time. This can be shown easily by using the standard “ $p \times 1/p$ ” argument as follows. For any $i \in [\ell]$ and any partial view ρ_i of C^* from which the i -th slot can start, let prefix_{ρ_i} be the event that in the execution of E , the view of internal C^* up until the beginning of the i -th slot is ρ_i . Let T_i be the random variable representing the number of rewinding in the i -th slot in E , and let p_{ρ_i} be the probability that under the condition that prefix_{ρ_i} occurs, the accepting i -th slot starts immediately after ρ and it does not interleave any message of Π . We then have

$$\mathbb{E}[T_i \mid \text{prefix}_{\rho_i}] \leq p_{\rho_i} \cdot 1/p_{\rho_i} = 1$$

for any ρ_i . Thus, we have

$$\mathbb{E}[T_i] = \sum_{\rho_i} \mathbb{E}[T_i \mid \text{prefix}_{\rho_i}] \Pr[\text{prefix}_{\rho_i}] \leq \sum_{\rho_i} \Pr[\text{prefix}_{\rho_i}] \leq 1 .$$

Hence, from the linearity of expectation, the expected number of rewinding of C^* in the execution of E is at most ℓ , and thus the expected running time of E can be bounded by a polynomial.