

Round-Efficient Concurrently Composable Secure Computation via a Robust Extraction Lemma

Vipul Goyal¹, Huijia Lin², Omkant Pandey^{3,5,*,**}, Rafael Pass^{4,***},
and Amit Sahai^{5,*}

¹ Microsoft Research India, Bangalore, India
vipul@microsoft.com

² University of California Santa-Barbara
rachel.lin@cs.ucsb.edu

³ University of Illinois at Urbana Champaign
omkant@uiuc.edu

⁴ Cornell University
rafael@cs.cornell.edu

⁵ UCLA and Center for Encrypted Functionalities
sahai@cs.ucla.edu

Abstract. We consider the problem of constructing protocols for secure computation that achieve strong concurrent and composable notions of security in the plain model. Unfortunately UC-secure secure computation protocols are impossible in this setting, but the Angel-Based Composable Security notion offers a promising alternative. Until now, however, under standard (polynomial-time) assumptions, only protocols with polynomially many rounds were known to exist.

In this work, we give the first $\tilde{O}(\log n)$ -round secure computation protocol in the plain model that achieves angel-based composable security in the concurrent setting, under standard assumptions. We do so by constructing the first $\tilde{O}(\log n)$ -round CCA-secure commitment protocol. Our CCA-secure commitment protocol is secure based on the minimal assumption that one-way functions exist.

A central tool in obtaining our result is a new *robust concurrent extraction lemma* that we introduce and prove, based on the minimal

* Research supported in part from a DARPA/ONR PROCEED award, NSF Frontier Award 1413955, NSF grants 1228984, 1136174, 1118096, and 1065276, a Xerox Faculty Research Award, a Google Faculty Research Award, an equipment grant from Intel, and an Okawa Foundation Research Grant. This material is based upon work supported by the Defense Advanced Research Projects Agency through the U.S. Office of Naval Research under Contract N00014-11-1-0389. The views expressed are those of the author and do not reflect the official policy or position of the Department of Defense, the National Science Foundation, or the U.S. Government.

** Work done in part while visiting Microsoft Research India (Bangalore).

*** Pass is supported in part by an Alfred P. Sloan Fellowship, Microsoft New Faculty Fellowship, NSF CAREER Award CCF-0746990, NSF Award CCF-1214844, NSF Award CNS-1217821, AFOSR YIP Award FA9550-10-1-0093, and DARPA and AFRL under contract FA8750-11-2-0211.

assumptions that one-way functions exist. This robust concurrent extraction lemma shows how to build concurrent extraction procedures that work even in the context of an “external” protocol that cannot be rewound by the extractor. We believe this lemma can be used to simplify many existing works on concurrent security, and is of independent interest. In fact, our lemma when used in conjunction with the concurrent-simulation schedule of Pass and Venkatasubramanian (TCC’08), also yields a constant round construction based additionally on the existence of quasi-polynomial time (PQT) secure one-way functions.

1 Introduction

The notion of *secure multi-party computation* protocols is central to cryptography. Introduced in the seminal works of [Yao86, GMW87], secure multi-party computation allows a group of (mutually) distrustful parties P_1, \dots, P_n , with private inputs x_1, \dots, x_n , to jointly compute any functionality f in such a manner that the honest parties obtain correct outputs yet no group of malicious parties learn anything beyond their inputs and the prescribed outputs. These early results on secure computation [Yao86, GMW87], along with a rich body of followup works that further refined and developed the concept [GL90, GMW91, Bea91, MR91, Can00, PW01, Can01, Gol04], demonstrated that the delicate task of designing secure protocols can be captured by general secure computation.

Much of the early literature on secure computation only considered the *stand-alone* setting where security holds only if a single execution of the protocol takes place, in isolation with no other cryptographic activity in the system. We call this security *stand-alone security*. While stand-alone security may be sufficient for basic purposes, it does not suffice in today’s more complex networked environments where other cryptographic protocols might be running in the system simultaneously.

Concurrent Security. To deal with more complex systems, the last decade has seen a push towards obtaining protocols that have strong concurrent *composability* properties. For example, we could require concurrent self-composability: the protocol should remain secure even when there are multiple copies executing concurrently. The framework of *universal composability* (UC) was introduced by Canetti [Can01] to capture a more general security requirement for a protocol that may be executed concurrently with not only several copies of itself but also with other protocols in an arbitrary manner.

Unfortunately, strong impossibility results have been shown ruling out the existence of secure protocols in the concurrent setting. UC secure protocols for most functionalities of interest have been ruled out in [CF01, CKL03]. Protocols in even less demanding settings of concurrent security were ruled on in [Lin04, BPS06, AGJ⁺12, Goy12]. We stress that, in fact, the latest sequence of these impossibility results provide an *explicit attack* in the concurrent setting using which the adversary may even fully recover the input of an honest party (see, e.g., the chosen protocol attack in [BPS06]). Hence, designing secure protocols

in the concurrent setting is a question of great theoretical interest as well as practical motivation.

To overcome these impossibility results, UC secure protocols were proposed based on various “trusted setup assumptions” such as a common random string that is published by a *trusted party* [CF01, CLOS02, BCNP04, CPS07, Kat07, CGS08]. Nevertheless, a driving goal in cryptographic research is to eliminate the need to trust other parties. The main focus of this paper is to obtain concurrently-secure protocols in the *plain model*.

Relaxing the Security Notion. To address the problem of concurrent security for secure computation in the plain model, a few candidate definitions have been proposed, the most well studied one being that of *super-polynomial simulation* [Pas03, PS04, BS05]. The notion of security with *super-polynomial simulators* (SPS) [Pas03, PS04, BS05] is one where the adversary in the ideal world is allowed to run in (fixed) super-polynomial time. Very informally, SPS security guarantees that any polynomial-time attack in the real execution can also be mounted in the ideal world execution, albeit in super-polynomial time. This is directly applicable and meaningful in settings where ideal world security is guaranteed statistically or information-theoretically (which would be the case in most “end-user” functionalities that have been considered, from privacy-preserving data mining to electronic voting).

Angel-based UC security. To formalize the notion of SPS security in a way that allows modular analysis and provides composability, Prabhakaran and Sahai [PS04] put forward the notion of *angel-based composable security*. Very roughly, in the angel based security notion, the parties (including the simulator and the adversary) are all polynomial time but have access to an angel which will perform certain *specific* super-polynomial time computations. This angel-based definition is in contrast to the case where the simulator is given direct access to super-polynomial computation power: in this case, the resulting security notion is *not* closed under composition and thus does not permit a modular protocol design in the concurrent setting¹. A construction for concurrently secure computation in the angel based composable security model were given in [PS04, BS05], but only based on non-standard super-polynomial hardness assumptions.

Very recently, Canetti, Lin, and Pass [CLP10] obtained the first secure computation protocol that achieves angel-based composable security based on *standard polynomial-time assumptions*. Unfortunately, however, the improvement in terms of assumptions comes at the cost of the round complexity of the protocol. Specifically, the protocol of [CLP10] incurs *polynomial-round complexity*. A follow up work of Lin and Pass [LP12] considers the problem of designing *black-box* constructions for secure computation in the concurrent setting. They propose a protocol making only a black-box use of oblivious transfer satisfying the angel-based composable security notion. However the round complexity of their protocol continues to remain polynomial.

¹ However we note that according to this weaker SPS security notion, concurrently secure protocols in constant rounds are now known [GGJS12].

We note that the latency of sending messages back and forth has been shown to often be the dominating factor in the running time of cryptographic protocols [MNPS04, BNP08]. Indeed, round complexity has been the subject of a good deal of research in cryptography. For example, in the context of concurrent zero knowledge (ZK) proofs, round complexity was improved in a sequence of works [RK99, KP01, PRS02] from polynomial to slightly super-logarithmic (that nearly matches the lower bound w.r.t. black-box simulation [CKPR01]). The round complexity of non-malleable commitments in the stand-alone and concurrent settings has also been studied in several works [DDN91, Bar02, PR05b, PR05a, LP09, Wee10, Goy11, LP11], improving the round complexity from logarithmic rounds to constant rounds under minimal assumptions. We observe that for the setting of concurrently secure computation protocols with angel-based composable security, the situation is worse since the only known protocols that achieves angel-based composable security based on standard assumptions incurs polynomial-round complexity [CLP10, LP12]. This raises the following natural question:

“Do there exists round-efficient protocols in the concurrent setting satisfying the angel-based composable notion of security based on standard assumptions?”

Our Results. We answer the above question in the affirmative and provide a $\tilde{O}(\log n)$ round construction of concurrently secure computation in the plain model. Our construction satisfies the angel-based composable notion of security [PS04, CLP10]. To obtain our result, we construct a “CCA-secure commitment” protocol in $\tilde{O}(\log n)$ rounds, based only the assumption that one-way functions exist. CCA secure commitments were introduced in [CLP10]; roughly speaking, a commitment protocol is CCA-secure if it remains hiding even when the adversary is given an oracle that can open all commitment values (except the commitment given as a challenge to the adversary). In [CLP10], Canetti et al. show how to construct a protocol that securely realizes any functionality—under the angel-based composable notion of security—given an (appropriate) protocol for CCA secure commitments (see full version of [CLP10]). Prior to our work, the best known construction for CCA secure commitments required n^ϵ rounds [CLP10, LP12]. In contrast, the round complexity of our protocol matches that of the best known constructions for concurrent extractable commitment schemes [PRS02, MOSV06].

A Robust Concurrent Extraction Lemma. A key technical tool that we introduce is a lemma that allows *robust* extraction of secrets from an adversarial committer A^* in the *concurrent* setting. We call this lemma, the *robust concurrent extraction lemma*, which is of independent interest. Roughly speaking, the lemma is a strengthening of the concurrent extraction mechanism for the PRS preamble [PRS02] (we shall call this the PRS commitment), and states that concurrent extraction can be performed even in the presence of an external protocol which cannot be rewound by the “simulator.”

More precisely, consider an adversarial committer A^* who commits to multiple values in *concurrent* sessions of the PRS commitment to honest receivers; let us

label these sessions as the **right** sessions. Simultaneously, A^* participates in a **left** execution of an arbitrary k -round protocol, denoted $\Pi := \langle B, A \rangle$. Then, the robust concurrent-extraction lemma states that for every A^* there exists a simulator S which, *without rewinding the external party B in its execution of Π* , extracts the values committed to by A^* in every session of the PRS commitment. Furthermore, if ℓ is the round complexity of the PRS preamble, and T is the running time of A^* , then S only fails with probability that is exponentially small in $\ell - O(k \cdot \log T)$.

In order to capture correctness of extraction, we formulate our lemma by considering a “real-world” experiment in which A^* receives the values committed to in every valid PRS commitment from an exponentially powerful party \mathcal{E} , called the “online extractor.” The extraction of values are provided by \mathcal{E} as soon as a PRS session ends. Our lemma states, intuitively speaking, that both the adversary A^* and the exponentially powerful party \mathcal{E} can be replaced with a polynomial-time simulator, that still interacts with the external party B in the external protocol Π . We remark that formulating the lemma in a generic way, so that it can handle as general usage of the PRS commitment as we have seen in the literature, is a delicate task. Nevertheless, we show that it is possible to precisely capture the concurrent-extraction property of the PRS commitment in a generic way without referring to any specific protocol that uses it.

An immediate benefit of our formulation is that when the PRS commitment is used inside a larger protocol, the task of “concurrent extraction” can be *formally* isolated from other parts of the protocol. This allows one to design hybrid experiments without having to worry about the extraction. We provide two procedures for this purpose—a simulator S and an online extraction \mathcal{E} —and demonstrate their use in the security proof of our protocol.

We also wish to remark that the ability to extract without rewinding B turns out to be a very useful tool during concurrent security proofs, and we expect this will have significant applications elsewhere. This flexibility simplifies security proofs (of even previous works) to a great extent. For example, a situation similar to our lemma arises in previous works on concurrent non-malleable zero-knowledge [BPS06, LPTV10]. In these works, the problem is solved in an arguably ad-hoc fashion, which stops rewindings after a certain point in the simulation during certain hybrids. This, overall, leads to a rather delicate analysis, and the order of hybrid experiments becomes important.

By using our robust-extraction lemma, this problem can be avoided almost directly. We note however, that in our case, it is crucial that the rewindings not be stopped. This is because, in our situation, one needs to implement the super-polynomial angel from the beginning in each hybrid experiment. Hence, every session in every hybrid requires online extraction—which is possible either by using (all) rewindings or by using super-polynomial simulation.

Technical Overview. As mentioned before, the starting point of our construction is the robust extraction lemma. The basic problem encountered in proving the lemma is that given the entire transcript of interaction between A^* and the honest right parties, there are various “breakpoints” in the transcript (represent-

ing messages of the left protocol) which cannot be rewound. In addition, during rewindings (or look-ahead threads), the breakpoints can change their location and can even come earlier than expected (at which point the current thread must be discontinued and another one started). At a high level, we start by considering the necessary modification of the KP/PRS simulator so that when a breakpoint is encountered during the execution of a look-ahead thread, the look-ahead thread is stopped and abandoned.

Our proof shows that even with this modification, the simulation still succeeds. The key observation is that each breakpoint can “spoil” at most a $d = \log T$ number of “recursive blocks” used in the swapping argument of [PRS02]. Since there are k breakpoints, this incurs an additional loss of $k \cdot \log T$ blocks. However, to execute this proof strategy, it becomes crucial to use the information learned during a “look-ahead block” in certain special sibling blocks. This new feature of our analysis must be done carefully to maintain the correctness of our swapping argument.

We note that we strive to obtain the best possible bounds for the round complexity ℓ of PRS that is necessary to extract for a given value of k (rounds of the left protocol). For this reason, we choose to re-analyze the proof presented in [PRS02] (see also [PTV08])².

For the goal of constructing CCA secure commitments, a direct application of the robust extraction lemma will not be sufficient. This is because the number of rounds in the left and the right interaction will be the same, and the robust extraction lemma cannot work with respect to protocols with the same round complexity as the PRScommitment! To construct CCA secure commitments (which is our main technical goal), we instead build upon techniques from prior work on concurrent non-malleable zero-knowledge (CNMZK) [BPS06, GJO10, LPTV10]. At a high level, our protocol is simple: commit to the value using a regular commitment scheme, and then, prove the validity of the commitment using a concurrent zero-knowledge protocol that is also *simulation sound* [Sah99]. We note that we design our own protocol for this task since we strive to achieve a construction based on *one-way functions* only (for CCA secure commitments). Using techniques from [BPS06, LPTV10] is either not possible or results in stronger assumptions such as the existence of collision-resistant hash functions. Our protocol is presented in section 3.

A Constant Round Protocol from \mathcal{PQT} One-way Functions. Our techniques can be seen as a general method which reduce the task of concurrent-extraction to that of concurrent-simulation. The method requires only a constant-

² For example, one can consider the approach of applying a pigeonhole principle argument to argue that $\omega(\log n)$ slots must occur between some two breakpoints, and then trying to apply the PRS analysis simply to these slots. However, note that even if no breakpoints occur during these slots, *look-ahead threads* that are started during these slots can still encounter breakpoints, since the adversary can choose the scheduling adaptively. Dealing with this analytically would require further loss, and result in a worse asymptotic bound than ours for super-constant values of k . Our more direct approach shows how to amortize the gains made over all slots, even if only a few slots occur between some breakpoints.

factor blow up in the round-complexity (of a given concurrent-simulation method) and a polynomial-factor blow up in the running time of the (given) simulator. By applying our method to the concurrent-simulator of [PV08], we obtain a constant round protocol for CCA-secure commitments based only on the existence of one-way functions secure against adversaries running in time super quasi-polynomial time (\mathcal{PQT}) (see full version of this work [GLP⁺15]).

Assuming more complex, and somewhat non-standard assumptions—namely the existence of adaptive one-way functions [PPV08]—a constant round protocol for CCA secure commitments is already known [PPV08]. Recent progress on program obfuscation [GGH⁺13] has given rise to new non-black-box simulation techniques for fully concurrent zero-knowledge [PPS15, IPS15, CLP14]; it would be interesting to explore if these techniques can also lead to improved constructions of CCA secure commitments.

Related Works. The work of Garg, Goyal, Jain, and Sahai [GGJS12] is closely related to our work, who provide a constant round protocol under the non-composable SPS notion (instead of angel-based composable security). Their work requires the existence of statistically hiding commitment. Independently of [GGJS12], a recent work of Pass, Lin, and Venkatasubramanian [PLV12] also provides a constant round protocol achieving non-composable SPS security, using very different techniques.

Other security notions that deal with concurrent security were presented in [MPR06, GGJS12] who propose the notion of *input indistinguishable computation*, and in [GS09, GJO10] who considered a modified ideal world that allows the adversary to make more output queries to the ideal functionality (than just one) per session.

2 Robust Concurrent Extraction

In this section, we will prove the robust extraction lemma. We use standard notation. In particular, $A(x; r)$ denotes the process of evaluating (randomized) algorithm A on input x with random coins r , and $A(x)$ the process of sampling a uniform r and then evaluating $A(x; r)$. We define $A(x, y; r)$ and $A(x, y)$ analogously. The set of natural numbers is represented by \mathbb{N} . Unless specified otherwise, $n \in \mathbb{N}$ represents the security parameter available as an implicit input, when necessary. All inputs are assumed to be of length at most polynomial in n .

For two probability distributions D_1 and D_2 , we write $D_1 \stackrel{c}{\equiv} D_2$ to mean that D_1 and D_2 are computationally indistinguishable. For two interactive Turing machines (ITM) A and B , we write $\text{OUT}_B[A(1^n, x) \leftrightarrow B(1^n, y)]$, the output of B after an interaction with A where their inputs in the interaction are y and x respectively, and their random tapes are independent and uniform.

We assume familiarity with commitment schemes. Without loss of generality, we will be using commitment schemes with non-interactive reveal phase—i.e., the committer sends a single message (v, d) to decommit. For a commitment scheme

$\langle C, R \rangle$, we denote by $\text{open}_{\langle C, R \rangle}(c, v, d)$ the decommitment function. That is, the receiver accepts v as the value committed to in the commitment-transcript c if $\text{open}_{\langle C, R \rangle}(c, v, d)$ outputs 1, and rejects otherwise. For statistically binding commitments, v is uniquely determined given c with high probability.

2.1 The PRS Preamble

The robust extraction lemma deals with the commitment preamble of Prabhakaran, Rosen, and Sahai [PRS02]. This preamble has been used in many prior works, and is often referred to as the PRS preamble. The preamble uses an underlying commitment scheme Com . Roughly speaking, the committer first commits to many shares of the value v to be committed using Com . This is followed by a several rounds where in each round, the receiver sends a random challenge, and the committer responds with appropriate decommitments. Each round is called a *slot*.

Essentially, the PRS preamble is an interactive commitment scheme, which is statistically binding (resp. hiding) if the underlying scheme Com is statistically binding (resp., hiding). The formal description of the PRS preamble is given in figure 1. As before, we write $\text{open}_{\text{PRS}}(c, v, \rho) = 1$, to formally mean that there exists randomness ρ such that c is the transcript of the PRS preamble, executed between the honest committer with input v and randomness ρ and the honest receiver with randomness (equal to its challenges) appearing in c .

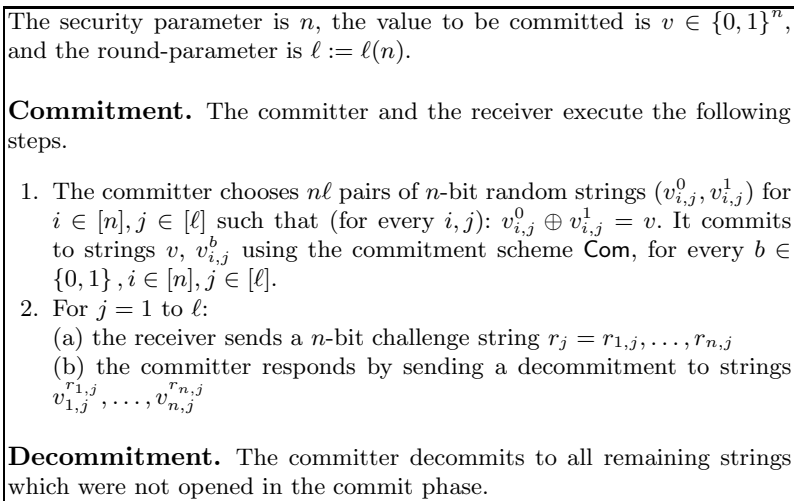


Fig. 1. The PRS Preamble based on Com

2.2 The Extraction Lemma

In this section, we present the robust extraction lemma. We will consider an adversary A^* who interacts in many sessions of the PRS preamble; simultaneously, A^* also participates in a single execution of a two party computation protocol Π . The running time of A^* is not necessarily polynomial in n . However, the lemma becomes trivial if the hiding/binding of the underlying commitment Com can be broken in $\text{poly}(n) \cdot T^2$ time, where $T = T(n)$ is the maximum number of PRS preambles A^* initiates.

SIMPLIFYING ASSUMPTION. We assume, for the clarity of presentation, that the commitment scheme Com underlying the PRS preamble is *statistically binding*. Later, we will present the general form which deals with both kinds of Com , as well as varying round complexity of the preamble.

PROTOCOL Π . Let $\Pi := \langle B, A \rangle$ be an arbitrary two-party computation protocol. We assume w.l.o.g. that both B and A receive a parameter $n \in \mathbb{N}$ as their first input. In addition, for a fixed $n \in \mathbb{N}$, let $\text{dom}_B(n)$ denote the domain of valid (second) input for algorithm B , and $k := k(n)$ denote the round complexity of Π .

The Robust-Concurrent Attack. Let A^* be an interactive Turing machine, called the adversary, $n \in \mathbb{N}$ the security parameter, and $x \in \text{dom}_B(n)$ an input. In the robust-concurrent attack, A^* interacts with a special, not necessarily polynomial time, party \mathcal{E} called the “online extractor.” Party \mathcal{E} simultaneously participates in one execution of the protocol Π , and several executions of the PRS preamble with A^* . Party \mathcal{E} follows the (honest) algorithm $B(1^n, x)$ in the execution of Π with A^* . Further, it follows the (honest) receiver algorithm in each execution of the PRS preamble. If A^* successfully completes a PRS preamble s , \mathcal{E} sends a string α_s to A^* , together with a special message END_s , to mark the completion of the preamble.

The scheduling of all messages in all sessions— Π as well as PRS preambles—is controlled by A^* including starting new sessions and finishing or aborting existing sessions. We adopt the following conventions [Ros04, PRS02]:

1. When A^* sends a round i message of session s , it immediately receives the next—i.e., $(i + 1)$ -st message of s ; this is without loss of generality,³ and holds for messages of Π as well.
2. If a session s has been aborted, A^* does not schedule any further messages of s .
3. If A^* starts a PRS preamble s , it also sends a special message, denoted START_s , immediately *after the last message* of step 1 of this preamble is completed (see figure 1). Message START_s indicates that the challenge-response phase is about to start.

³ This is because the next message can be stored and delivered whenever needed during the attack.

At some point, A^* halts. We say that A^* *launches* the robust-concurrent attack.

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

We are now ready to present the robust extraction lemma. Informally speaking, the lemma states that there exists an interactive Turing machine—a.k.a the robust simulator—whose output is statistically close to $\text{REAL}_{\mathcal{E}, \Pi}^{A^*}(n, x, z)$ *even if* the final response of \mathcal{E} at the end of a successful PRS session is actually the value A^* commits to in that session. Further, the robust simulator does not “rewind” B , and runs in time polynomial in total sessions opened by A^* .

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

- 1. Validity constraint.** *For every output ν of $\text{REAL}_{\mathcal{E}, \Pi}^{A^*}(n, x, z)$, for every PRS preamble s (appearing in ν) with transcript τ_s , if there exists a unique value $v \in \{0, 1\}^n$ and randomness ρ such that $\text{open}_{\text{PRS}}(\tau_s, v, \rho) = 1$, then:*

$$\alpha_s = v,$$

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

- 2. Statistical simulation.** *If $k = k(n)$ and $\ell = \ell(n)$ denote the round complexities of Π and the PRS preamble respectively, then the statistical distance between distributions $\text{REAL}_{\mathcal{E}, \Pi}^{A^*}(n, x, z)$ and $\text{OUT}_s [B(1^n, x) \leftrightarrow S^{A^*}(1^n, z)]$ is given by:*

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

where $T(n)$ is the maximum number of total PRS preambles between A^* and \mathcal{E} .⁴ Further, the running time of S is $\text{poly}(n) \cdot T(n)^2$.

We prove this lemma by presenting an explicit simulator S and a corresponding party \mathcal{E} . The explicit constructions appear in subsection 2.3, and the full proof of the lemma is given in the full version [GLP⁺15]. We now make some important remarks about the lemma.

⁴ The lemma allows for exponential $T(n)$ as well. However, if it is too large—e.g., $T(n) = 2^{2^n}$, the PRS preamble should be modified suitably. For example, the value v as well as the challenges in each slot, must be of length at least $n + 2 \log T(n)$.

Remarks

1. The special party \mathcal{E} is not completely defined by the lemma. In particular, when a PRS preamble s does not commit to a *unique and valid* value v , the value α_s sent by \mathcal{E} to A^* is not defined. This can happen, e.g., when not all shares committed to in step 1 XOR to v . In such situations, \mathcal{E} can choose whatever value α_s it wants. The only requirements on \mathcal{E} are that it uses honest algorithms during the robust-concurrent attack *and* that for each successfully completed PRS preamble it satisfies the validity constraint. Every \mathcal{E} satisfying these requirements is called a *valid* \mathcal{E} .
2. The PRS preamble is used in a variety of complex ways. For example, some protocols require opening the committed PRS-value (e.g., [PRS02, BPS06, OPV10]), whereas some others may never open this value (e.g., [LPTV10, GJO10]). To be able to capture such uses generically, we do not enforce any consistency requirements on the PRS preambles. The choice of not fully defining \mathcal{E} when PRS preamble is not valid provides sufficient flexibility to capture such generic uses of the preamble.
3. When the preamble is used in a larger protocol, a “main” simulator is used to prove the security of the larger protocol. Typically, the main simulator employs the rewinding strategy of [KP01, PRS02] to extract the PRS-values and simultaneously deals with other details of the protocol. Our lemma separates the task of extracting PRS-values from other necessary actions of the main simulator. This makes the overall proofs simpler. Party \mathcal{E} then only acts as mechanism to transfer the extracted PRS-values back to the main simulator. The main simulator takes upon the role of A^* to receive extracted values from \mathcal{E} , while only dealing with other details of the larger protocol.
4. A consequence of the above two remarks is that when a PRS preamble is not consistent, we do not know what value α_s actually gets extracted. Our choice of the order of quantifiers allows \mathcal{E} to depend on A^* as well as S . This essentially allows \mathcal{E} to extract and supply the same value α_s (by running S internally) that a typical “main” simulator would extract for inconsistent PRS preambles.
5. Requirements 1 and 2 of the lemma imply that if we sample an output of the simulator and consider a PRS preamble s with transcript τ_s which contains a *unique and valid* value v , and receives α_s as \mathcal{E} 's response in the end, then except with probability $\Delta(n)$, it holds that $\alpha_s = v$.

2.3 A Robust Simulator and an Online Extractor

In this section, we present an explicit construction of a robust simulator S , and the (online extractor) party \mathcal{E} for which (the robust extraction) lemma 1 holds. The simulator is a slight modification of [KP01, PRS02], to also deal with messages of Π , without rewinding them. We start by defining a few terms first.

The States of A^* . Recall that the scheduling of messages in the robust-concurrent attack is controlled by A^* , and when A^* sends the i -th message

of a session s (either PRS or Π), it immediately receives the next message of s , namely the $(i + 1)$ -st message. The *state* of A^* at any given point during the attack consists of its *view up to that point*: it includes all messages sent/received by A^* , its auxiliary input z and its randomness. The *starting* (or *original*) state of A^* —denoted throughout by st_0 —is its state before it receives the first message. If st denotes the state of A^* at some point during the robust-concurrent attack, the set of all PRS preambles which have not completed yet is denoted by $\text{LIVE}(\text{st})$.

The Robust Simulator S . The simulator receives as input an auxiliary string $z \in \{0, 1\}^*$, and the security parameter n . The simulator participates with an external party B of Π . Let $x \in \text{dom}_B(n)$ and γ denote the input and uniformly chosen randomness of B . The simulator incorporates the adversary A^* as a black-box; let $T = T(n)$ define the maximum number of PRS preambles that A^* can open during the robust-concurrent attack.

Simulator S starts by setting $(1^n, z)$ on A^* 's input tape, and a sufficiently long uniform string on its random tape. The simulator then initiates a helper procedure `recurse` as follows:

$$(\text{st}, \mathcal{T}) \leftarrow \text{recurse}(T, \text{st}_0, \emptyset, 1, \emptyset, 0).$$

Throughout its execution, messages of `recurse` are forwarded back and forth to $B(1^n, x; \gamma)$ and (the black-box) A^* as appropriate. Finally, the output of S is the first output of `recurse`, namely st ; the output st is also known as the *main thread*. Procedure `recurse` is given in figure 2, each execution of `recurse` is called a block, and has a unique name denoted by `id`. \square

The Online Extractor \mathcal{E} . Formally, an execution of \mathcal{E} begins during a robust-concurrent attack. Let (γ, ρ) denote the random tape of \mathcal{E} , and (n, x, z) denote its inputs. \mathcal{E} incorporates the program of A^* , and performs the following *internal steps* before it sends out its first message in the robust-concurrent attack,

1. \mathcal{E} proceeds identically to the robust simulator algorithm S , using ρ as its random tape and $(1^n, z)$ as its inputs. To successfully proceed in this step, \mathcal{E} uses (x, γ) to simulate the honest algorithm $B(1^n, x; \gamma)$, as well as black-box access to A^* .⁵ However, \mathcal{E} differs from S in its actions only when a PRS preamble s completes, in the manner described below.
2. Let s be a successfully completed PRS preamble; at this point S either extracts a value μ_s or reaches an `ExtractFail`. When this happens, \mathcal{E} neither sends μ_s nor aborts the simulation; instead it proceeds as follows. First, \mathcal{E} attempts to extract the *actual value* committed to in the preamble by inverting all instances of the underlying commitment `Com`. Then, it decides the value α_s , to be sent, as follows. If a valid and *unique* value v_s exists, set $\alpha_s = v_s$. Otherwise, it has following cases:

⁵ Observe that although \mathcal{E} depends on A^* here, it still uses A^* only as a black-box, as remarked earlier.

procedure $\text{recurse}(t, \text{st}, \mathcal{T}, f, \text{aux}, \text{id})$:

1. If $t = 1$, **repeat**:
 - (a) If the next message is **START**, start a new session s .
 - send $r \leftarrow \{0, 1\}^n$ as the challenge of the first slot of s .
 - add entry $(s : 1, r, _)$ to \mathcal{T} .
 - (b) If the next message is the slot- i challenge of an existing session s .
 - send $r \leftarrow \{0, 1\}^n$ as the slot- i challenge of s .
 - add entry $(s : i, r, _)$ to \mathcal{T} .
 - (c) If the next message is the slot- i *response*, say β , of an existing session s .
 - If β is a valid message:
 - update entry $(s : i, r_i, _)$ to $(s : i, r_i, \beta)$.
 - if $i = \ell$, i.e., it is the last slot, send $(\text{END}, \text{extract}(s, \text{id}, \mathcal{T}, \text{aux}))$.
 - Otherwise, if $\beta = \perp$, abort session j , and add $(s : \perp, \perp, \perp)$ to \mathcal{T} .
 - Update st to be the current state of A^*
 - **return** (st, \mathcal{T}) .
 - (d) If the next message is a *response* from A^* for the external protocol Π .
 - If $f = 0$, i.e., it is a look-ahead block, then **return** (st, \mathcal{T}) ;
 - If $f = 1$, i.e., it is the *main thread*), do the following:
 - send A^* 's message to the external party of Π , return the response to A^* .
 - Update st to be the current state of A^*
 - For every live session $s \in \text{LIVE}(\text{st})$, do the following:
 - $\times_{s, \text{id}} = \text{true}$,
 - for every block id' that contains the block id , set: $\times_{s, \text{id}'} = \text{true}$.
2. If $t > 1$,
 - # Rewind the first half twice:
 - (a) $(\text{st}_1, \mathcal{T}_1) \leftarrow \text{recurse}(t/2, \text{st}, \mathcal{T}, 0, \text{aux}, \text{id} \circ 1)$ [look-ahead block C']
 - (b) Let $\text{aux}_2 = (\text{aux}, \mathcal{T}_1 \setminus \mathcal{T})$,
 $(\text{st}_2, \mathcal{T}_2) \leftarrow \text{recurse}(t/2, \text{st}, \mathcal{T}, f, \text{aux}_2, \text{id} \circ 2)$ [main block C]
 - # Rewind the second half twice:
 - (c) $(\text{st}_3, \mathcal{T}_3) \leftarrow \text{recurse}(t/2, \text{st}, \mathcal{T}^*, 0, \text{aux}, \text{id} \circ 3)$ [look-ahead block D']
 - (d) Let $\mathcal{T}^* = \mathcal{T}_1 \cup \mathcal{T}_2$ and $\text{aux}_4 = (\text{aux}, \mathcal{T}_3 \setminus \mathcal{T}^*)$,
 $(\text{st}_4, \mathcal{T}_4) \leftarrow \text{recurse}(t/2, \text{st}, \mathcal{T}^*, f, \text{aux}_4, \text{id} \circ 4)$ [main block D]
 - (e) **return** $(\text{st}_4, \mathcal{T}_3 \cup \mathcal{T}_4)$.

procedure `extract`($s, \text{id}, \mathcal{T}, \text{aux}$):

1. Attempt to extract a value for s from \mathcal{T} .
2. If extraction fails, consider every block id_1 for which $\times_{s, \text{id}_1} = \text{true}$.
 - Let id'_1 be the sibling of id_1 , with input/output tables $\mathcal{T}_{\text{in}}, \mathcal{T}_{\text{out}}$ respectively.
 - Attempt to extract from $\text{aux}_{\text{id}'_1} := \mathcal{T}_{\text{out}} \setminus \mathcal{T}_{\text{in}}$; (included in aux).
3. If all attempts fail, abort the simulation and **return** `ExtractFail`.
Otherwise **return** the extracted value.

Fig. 2. Procedures `recurse` and `extract` used by the robust simulator S

- (a) If there are more than one valid v_s , proceed as follows: if μ_s equals to any of them, set $\alpha_s = \mu_s$, otherwise, set α_s to be one of them chosen at random.
 - (b) If no valid v_s exists, then proceed as follows: if $\mu_s = \text{ExtractFail}$, set α_s to be a random value; otherwise, set $\alpha_s = \mu_s$.
3. After reaching the end of the simulation, \mathcal{E} internally stores the randomness ρ_s and the values α_s for every PRS preamble s appearing on the *main thread*.

Having completed the steps above,⁶ \mathcal{E} is now ready to interact with the (outside) A^* launching the robust-concurrent attack, and proceeds as follows.

- If A^* sends a message intended for the (only) session of Π , \mathcal{E} interacts with A^* by following actions of $B(1^n, x; \gamma)$. Likewise, if A^* sends a messages for a PRS preamble s , A^* follows the honest receiver algorithm of PRS preamble with randomness ρ_s already computed internally.
- If A^* successfully completes a PRS preamble s , \mathcal{E} sends the already stored value α_s to A^* . □

3 CCA Secure Commitments in $\tilde{O}(\log n)$ Rounds

In this section we apply our robust extraction lemma to construct a $\tilde{O}(\log n)$ -round protocol for CCA secure commitments. We will need the generalized version of the lemma which allows for *statistically hiding* PRS preamble as well; the general version appears in 4. We will also use a non-malleable commitment scheme, denoted `NMCom`, that is *robust* w.r.t. constant round protocols [LP09]. Constant round constructions for such protocols are now known (see the non-malleable commitments section in the full version [GLP⁺15]).

⁶ We insist that all the steps above are *internal* to \mathcal{E} and that as of now it has not sent any external message in the robust-concurrent attack. Further, it can successfully complete these steps since it has all the required inputs.

The protocol has a very intuitive “commit-and-prove” structure where the committer commits to the value v using a PRS preamble, and then proves its consistency using what resembles a “concurrent simulation-sound” protocol. We start by recalling the notion of CCA secure commitments from [CLP10].

3.1 CCA Secure Commitments

Let $\langle C, R \rangle$ denote a statistically binding and computationally hiding commitment scheme. We assume w.l.o.g. that $\langle C, R \rangle$ has a non-interactive reveal phase—i.e., the committer simply sends (v, d) . The decommitment is verified using a function $\text{open}(c, v, d)$; that is, the receiver accepts v as the value committed in the commitment-transcript c if $\text{open}(c, v, d)$ outputs 1, and rejects otherwise. A tag-based commitment scheme with $l(n)$ -bit identities [PR05b, DDN91] is a commitment-scheme where in addition to 1^n , C and R also receive a “tag” (or *identity*) of length $l(n)$ as common input. We will consider schemes which are *efficiently checkable*: meaning that if R accepts in the interaction (with transcript c) then there exists a decommitment pair (v, d) such that $\text{open}(c, v, d) = 1$.

In CCA-secure commitments, we consider an adversarial receiver A , who has access to an oracle \mathcal{O} , called the “decommitment oracle.” The oracle participates with A in many *concurrent* sessions of (the commit phase of) $\langle C, R \rangle$, using tags of length $l(n)$, chosen adaptively by A . At the end of each session, if the session is accepting, the oracle returns the (unique) value committed by A in that session; otherwise it returns \perp . (In case there is more than one possible decommitment, \mathcal{O} returns any one of them.)⁷

Roughly speaking, we say that a tag-based scheme $\langle C, R \rangle$ is CCA-secure if there exists a decommitment oracle \mathcal{O} for $\langle C, R \rangle$, such that the hiding property of the scheme holds even for adversaries A with access to \mathcal{O} . Formally, let $\text{IND}_b(\langle C, R \rangle, \mathcal{O}, A, n, z)$ denote the output of the following probabilistic experiment: on common input 1^n and auxiliary input z , the PPT adversary $A^\mathcal{O}$ (adaptively) chooses a pair of challenge values $(v_0, v_1) \in \{0, 1\}^n$ and a tag $\text{id} \in \{0, 1\}^{l(n)}$, and receives a commitment to v_b using the tag id (note that A interacts with \mathcal{O} throughout the experiment as described before); finally, when $A^\mathcal{O}$ halts, the experiment returns the output y of $A^\mathcal{O}$; y is replaced by \perp if during the experiment, A sends \mathcal{O} any commitment using the tag id .

Definition 1 (CCA-secure Commitments, [CLP10]). *Let $\langle C, R \rangle$ be a tag-based commitment scheme with tag-length $l(n)$, and \mathcal{O} be a decommitment oracle for it. We say that $\langle C, R \rangle$ is CCA-secure w.r.t. \mathcal{O} , if for every PPT ITM A , every $z \in \{0, 1\}^*$, and every sufficiently large n , it holds that:*

$$\text{IND}_0(\langle C, R \rangle, \mathcal{O}, A, n, z) \stackrel{c}{\equiv} \text{IND}_1(\langle C, R \rangle, \mathcal{O}, A, n, z).$$

⁷ Note that since $\langle C, R \rangle$ is efficiently checkable, and the session is accepting, such a valid decommitment always exists. In addition, note that since we only have statistical binding, this value is unique except with negligible probability.

We say that $\langle C, R \rangle$ is CCA-secure if there exists a decommitment oracle \mathcal{O}' such that $\langle C, R \rangle$ is CCA-secure w.r.t. \mathcal{O}' .

An analogous version of the definition which considers many concurrent executions on left (instead of just one), is known to be equivalent to the current definition (via a simple hybrid argument). Also note that, for this reason, this definition implies concurrent non-malleable security for commitments [PR05a].

3.2 Our Protocol for CCA Secure Commitments

We are now ready to present our CCA secure commitment protocol, denoted CCA-Com. The protocol employs the PRS preamble in *both* directions (similar to [GJO10, LPTV10]). However, our protocol is much simpler, and admits an easier security proof.

We now provide a quick overview of our protocol. Our protocol will also use a constant-round non-malleable commitment scheme—NCom—that is robust w.r.t. 4 rounds. The formal description of our protocol appears in figure 3.

Let $v \in \{0, 1\}^n$ be the value to be committed. Our commitment protocol, CCA-Com, consists of five phases. In the first phase, the committer C commits to v using a statistically-binding PRS preamble. We denote this instance of the preamble by PRS_1 . In phase 2, the receiver R uses a statistically-hiding PRS preamble to commit to a random value $\sigma \in \{0, 1\}^n$; we denote this instance of the PRS preamble by PRS_2 .

In phase 3, C commits to 0^n using the *robust* NCom scheme. In phase 4, R decommits to the value σ (of phase PRS_2). Finally, in phase 5, C uses a witness-indistinguishable (WI) *proof* system to prove that: “either there exists a valid value v in phase 1 or the value committed in NCom is σ .” It can use, for example, Blum’s 3-round protocol repeated in parallel n times [Blu87]. We remark that it is necessary to use a proof system which ensures soundness against unbounded provers.

To decommit, C sends decommitments corresponding to the first phase, namely PRS_1 . The round-complexity of both PRS preambles is $\ell \in \omega(\log n)$. We have the following theorem, whose proof is given in the full version [GLP⁺15].

Theorem 1. *Assuming the existence of collision-resistant hash functions, protocol CCA-Com presented in figure 3 is a $\tilde{O}(\log n)$ -round CCA secure commitment scheme for identities of length n .*

3.3 Construction Based on One-Way Functions

Protocol CCA-Com of the previous section requires the use of a statistically-hiding phase PRS_2 . If we change the protocol so that PRS_2 is statistically-binding, we will not be able to prove that the right sessions of CCA-Com remain statistically binding. This is because of the presence of the super-polynomial time oracle \mathcal{O} . Indeed, without the oracle we can make PRS_2 statistically-binding and

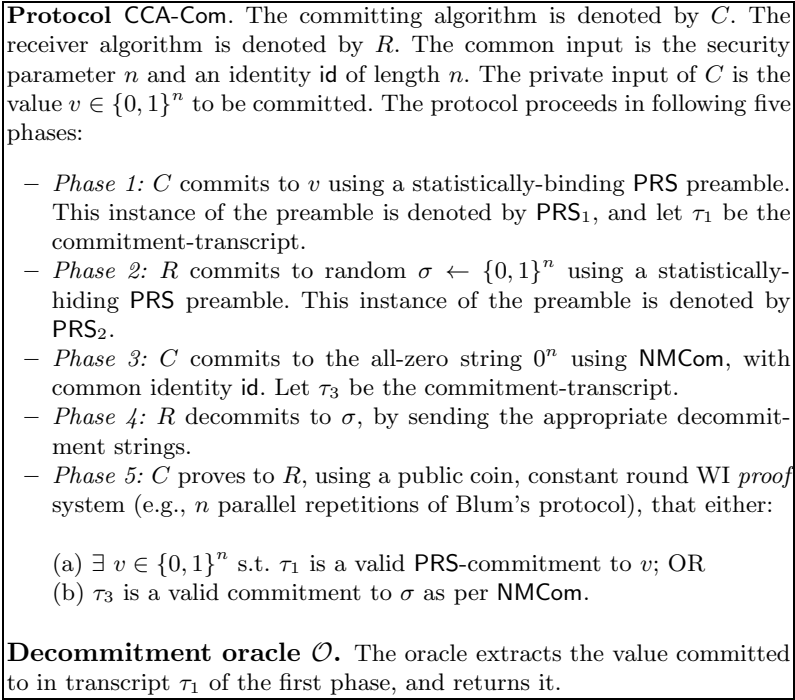


Fig. 3. Protocol CCA-Com

have a protocol based on one-way functions; but in the presence of oracle, it may happen that A^* is able to ensure that for some right session s , $\tilde{u}_s = \tilde{\sigma}_s$.

To avoid this problem, we must somehow remove PRS_2 yet still be able to later modify the left PRS slot-by-slot. The key observation is that in case of commitments, there is only one left session. Therefore, we can use a different simulation strategy which guarantees “soundness” w.r.t. unbounded committers as well.

Based on this observation, our one-way functions based protocol does not use PRS_2 . Instead it uses a CoinFlip protocol in which no unbounded prover can succeed in “setting up a trapdoor” but a rewinding party can. Our new protocol, denoted CCA-Com^* , appears in figure 4.

Theorem 2. *Assuming the existence of one way functions, protocol CCA-Com^* presented in figure 4 is a $\tilde{O}(\log n)$ -round CCA secure commitment scheme for identities of length n .*

For the proof of this theorem see appendix ??. Although the proof is very similar to that of theorem 1, one crucial point is that we need to re-use the idea of robust-simulation, and be careful about how we apply it.

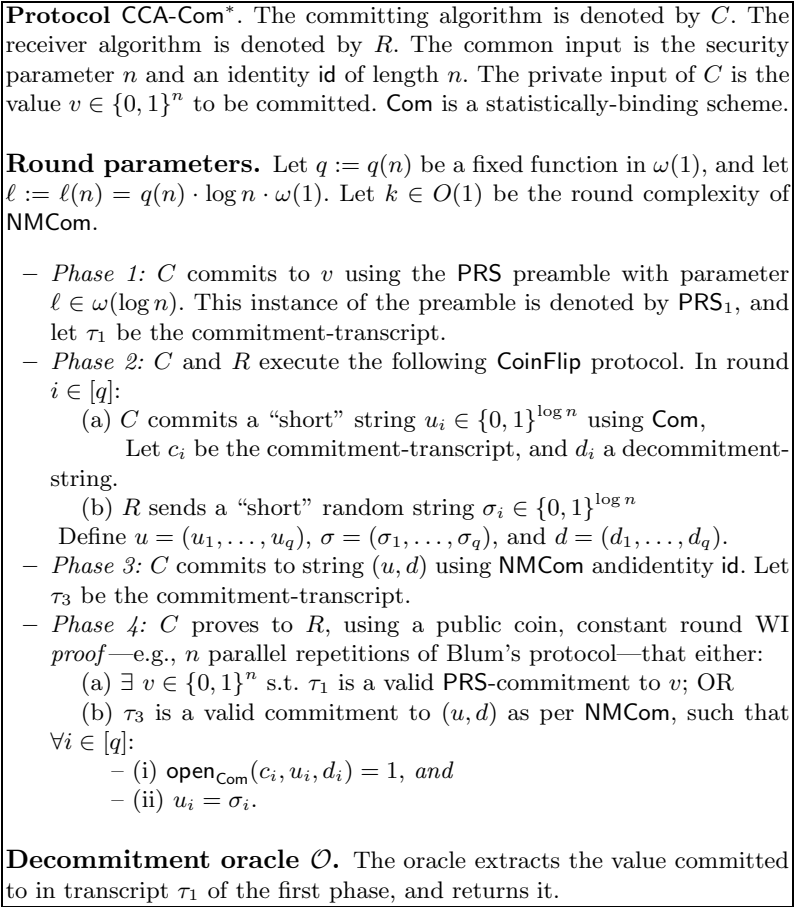


Fig. 4. Protocol CCA-Com* based on one-way functions

4 Generalized Version of the Robust Extraction Lemma

In the generalized version of the lemma, we allow A^* to open sessions of the PRS preambles that are *statistically hiding*. At the start of the preamble, it is already understood whether it is statistically-binding or statistically-hiding. Since statistically-hiding preambles are only computationally-binding, we require that A^* is a PPT machine. Furthermore, we make following new adjustments to the robust-concurrent attack:

At the successful completion of a *statistically-hiding* PRS preamble s , when A^* receives the string α_s from \mathcal{E} , it can choose to respond with an opening of the committed value v_s . It does so by sending appropriate decommitment strings (v_s, d_s) . Furthermore, the scheduling of this message is decided by A^* .

Other than this, the attack remains unchanged. We now present the generalized version of the lemma. The essence of the lemma is still the same as before. We

only need to add conditions to deal with the statistically-hiding preambles. For such preambles, the validity constraint requires that α_s be equal to the opened value v_s (if any), except with negligible probability.

Lemma 2 (Robust Extraction: General Version). *There exists an interactive Turing machine S (“robust simulator”), such that for every PPT A^* , for every $\Pi := \langle B, A \rangle$, there exists a party \mathcal{E} (“online extractor”), such that for every $n \in \mathbb{N}$, for every $x \in \text{dom}_B(n)$, and every $z \in \{0, 1\}^*$, the following conditions hold:*

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

(a) *for every statistically-binding preamble s (appearing in ν) with transcript τ_s , if there exists a unique value $v \in \{0, 1\}^n$ in the commitment-transcript τ_s , then $\alpha_s = v$,*

(b) *for every statistically-hiding preamble s (appearing in ν) with transcript τ_s , if there exists a valid opening (v_s, d_s) in the view ν , then $\alpha_s = v_s$, where α_s is the value \mathcal{E} sends at the completion of s .*

2. Statistical simulation. *If $k = k(n)$ and $\ell = \ell(n)$ denote the round complexities of Π and the PRS preamble respectively, then the statistical distance between distributions $\text{REAL}_{\mathcal{E}, \Pi}^{A^*}(n, x, z)$ and $\text{OUT}_S[B(1^n, x) \leftrightarrow S^{A^*}(1^n, z)]$ is given by:*

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

where $T(n)$ is the maximum number of total PRS preambles between A^* and \mathcal{E} .⁸ Further, the running time of S is $\text{poly}(n) \cdot T(n)^2$.

The proof of this general version of the lemma is identical to the proof of the original lemma and appears in the full version [GLP⁺15]. We only need to show that condition 1(b) also holds. We show that if it does not then we can violate the computational binding of the statistically-hiding PRS. Suppose that Com_{sh} is the underlying commitment scheme of the statistically hiding PRS. Then, if α_s is not equal to v_s , look at the two opened challenges in the execution of `recurse` whose XOR results in α_s ; let them belong to slot i of this session. Both of these strings must have been decommitted to correctly (for the scheme Com_{sh}). Further, since opening v_s requires opening of all slots of PRS such that pairs in each slot XOR to the same v_s , we have that all pairs of slot i must have been correctly opened to strings that are different from what `recurse` learned. Therefore, we must have an instance of Com_{sh} with correct openings to two different values. The details are standard and omitted. We note that the value of $\Delta(n)$ in the second condition does not change since both \mathcal{E} and S extract value α_s identically for the statistically-hiding PRS.

⁸ The lemma allows for exponential $T(n)$ as well. However, if it is too large—e.g., $T(n) = 2^{2^n}$, the PRS preamble should be modified suitably. For example, the value v as well as the challenges in each slot, must be of length at least $n + 2 \log T(n)$.

5 Proof of Security for CCA-Com

We start by noting that the scheme is statistically-binding even against an unbounded cheating committer algorithm C^* . This is because of the following. PRS_2 is statistically-hiding for σ , and therefore except with negligible probability, the value committed to in NMCom (which is statistically-binding) does not equal σ . Then, from the soundness of the WI *proof*, it follows that C^* can succeed only when the PRS_1 is consistent. It will be important to have statistical-binding w.r.t. unbounded C^* since we will be dealing with the super-polynomial time oracle \mathcal{O} in various hybrid experiments.

We now proceed to demonstrate the CCA security of our scheme. To do so, we will directly construct a PPT simulator SIM , which simulates the view of any CCA adversary $A^{*\mathcal{O}}$. SIM will not have access to the decommitment oracle \mathcal{O} , and uses A^* only as a black-box. Recall that \mathcal{O} extracts the value from the transcript of PRS_1 . SIM only receives 1^n and auxiliary-input z for A^* , as its own inputs.

Lemma 3. *There exists a strict polynomial time machine SIM such that for every PPT ITM machine A^* , every $z \in \{0, 1\}^*$, every sufficiently large n , and every $b \in \{0, 1\}$, it holds that:*

$$\text{SIM}^A(1^n, z) \stackrel{c}{\equiv} \text{IND}_b(\langle C, R \rangle, \mathcal{O}, A, n, z).$$

Proof. Recall that during its execution A^* opens one session of $\text{CCA-Com} := \langle C, R \rangle$ on left, while simultaneously interacting in multiple concurrent sessions of CCA-Com with \mathcal{O} , called the right sessions. Let $T = T(n)$ be a polynomial upper-bounding the number of right sessions of A^* .

Algorithm SIM will use the robust simulator S guaranteed by the (general version of) the robust-extraction lemma (see sections 2 and 4), allowing it to extract the values committed to by A^* in all PRS preambles. It then uses these values to simulate the answers of \mathcal{O} , as well as to succeed in WI proof (by committing the extracted value in NMCom of the left session).

SIM uses two helper procedures: the robust simulator S and an “interface” algorithm I to be described shortly. Procedure I essentially “decouples” the PRS preambles from the rest of the CCA-Com protocol. It incorporates A^* as a black-box, and handles all messages of all sessions of CCA-Com internally and honestly, *except* for all the PRS preambles in which A^* acts as the *committer*. All these preambles are forwarded to outside PRS -receivers. That is, I participates in a robust-concurrent attack, interacting with the party \mathcal{E} . It executes PRS preambles with \mathcal{E} , and at the end of each PRS preamble s , I expects to receive a value α_s from \mathcal{E} . This value will be used internally.

SIM is a polynomial time machine without access to any super-polynomial time helper. Therefore, to run I , it runs the robust-simulator S providing it black-box access to the “adversary” procedure I . SIM outputs whatever S outputs. Formal descriptions follow.

Algorithm $\text{SIM}(1^n, z)$. Return the output of $S^I(1^n, z)$, where procedure I —which has black-box access to adversary A^* —is described below.

Procedure $I(1^n, z)$. Procedure I launches the robust-concurrent attack, by committing in several PRS sessions to external receivers denoted R_1, \dots, R_T . At the end of each preamble, it expects to receive a string α_s . I incorporates the CCA adversary A^* internally, as a black-box. I initiates an execution of A^* , simulating various sessions of CCA-Com that A^* opens as follows.

1. If A^* starts a new session s of CCA-Com on *right*, I starts by initiating a new session of the statistically-binding PRS₁ with an external receiver R_s . Then, messages of phase-1 of s are relayed between A^* and R_s . I simulates all other phases of s internally by following the honest algorithm for each phase.
2. If A^* starts a new session s of CCA-Com on *left*, I initiates a new session of the statistically-*hiding* preamble to be used as PRS₂ of s with an external receiver R'_s . I completes various phases of s as enumerated below.
 - (a) *Phase 1:* I commits to an all zero string to A^* .
 - (b) *Phase 2:* I simply relays messages between A^* and R'_s .
 - (c) *Phase 3:* I commits to value α_s using NMCom (instead of 0^n). Value α_s was received from outside at the end of PRS₂ of sessions s .
 - (d) *Phase 4:* If A^* correctly opens the value in PRS₂ of session s , I checks that the opened value is equal to the “fake witness” α_s . If not, it outputs a special symbol ExtractFail, and halts. Otherwise it continues the execution.
 - (e) *Phase 5:* I uses α_s and the randomness used in phase 3 (NMCom) to complete the WI proof in phase 5.⁹
3. *Oracle answers:* If A^* successfully finishes a session s of CCA-Com on right, I sends the (already extracted) value α_s to A^* .

When A^* halts, I outputs the view of A^* , and halts. □

Proving Indistinguishability. We are now ready to prove our lemma. We will prove this by using a series of hybrid experiments. Our hybrid experiments will be designed by making step-by-step changes to how I communicates with A^* internally in various phases of the protocol. For $i \in [7]$, we denote by ν^i the output of hybrid H_i .

Hybrid H_0 : This hybrid is identical the experiment $\text{IND}_b(\langle C, R \rangle, \mathcal{O}, A, n, z)$. Recall that in this experiment, $A^{*\mathcal{O}}$ receives a commitment to value $v_b \in \{0, 1\}^n$ from the honest committer C , while interacting with the oracle \mathcal{O} . The output of the experiment consists of the view of A^* , which is also the output of H_0 .

⁹ Note that, technically, I can use the valid witness corresponding to the PRS₁ phase s . This is since it committed to a valid value, namely 0^n . However, we choose to use α_s so that SIM will in fact be a valid simulator even for our concurrent non-malleable zero-knowledge protocols.

Hybrid H_1 : This hybrid is identical to H_0 , except for the following differences:

- (a) H_1 does not forward the right sessions to \mathcal{O} . Instead, it executes all right sessions on its own, by playing the honest receiver strategy R . Denote by R_s the instance of R executed in session s .
- (b) When A^* successfully completes the session R_s , H_1 queries a different oracle \mathcal{O}' , which acts as follows. The query to \mathcal{O}' consists of only part of the commitment-transcript PRS_1 that belongs for the value to be committed. All other messages (e.g., commitments of shares and slot-messages) are not part of the query. In particular, if Com is non-interactive, the query will be commitment of the value \tilde{v}_s . If there is a unique value defined by the query, \mathcal{O}' extracts that value, and returns it as the answer. In all other cases, it behaves exactly as \mathcal{O} .

Hybrid H_2 : This hybrid is identical to H_1 ; we use it to set up some notation. Define a procedure I_2 , which is identical to procedure I (defined above), except that it it executes (all internal phases of) the left session honestly and that it uses the oracle \mathcal{O}' as in H_1 . That is, it only forwards the PRS preambles received from A^* outside, but executes all other phases internally and honestly. Formally, I_2 is identical to I except:

- (a) I_2 receives the value v_b to be committed in left sessions of CCA-Com , as an input. Furthermore, I_2 commits to value v_b in PRS_1 of the left session (instead of 0^n , step 2(a)).
- (b) When a session s of PRS_1 ends, I_2 expects to receive a value α_s from outside. If s is a statistically-binding sessions, then I_2 gives α_s to A^* . If s is statistically-hiding, it *ignores* the value s —there is only one such session corresponding to the left session.
- (c) I_2 uses the valid witness, v_b and the randomness of PRS_1 to complete its WI proof in the left session (step 2(e)). I_2 does not check the validity of the “fake witness” (step 2(d)).

Observe that I_2 is essentially launching a robust-concurrent attack. It does not perform any rewindings. Hybrid H_2 simply runs the procedure $I_2^{A^*}$, and simulates PRS-receivers for it exactly as H_1 does; furthermore, when a statistically-binding session s finishes, H_2 forwards its commitment to \mathcal{O}' and returns the oracles answer, denoted α_s , to I . On the other hand, if the statistically-hiding session of PRS ends, H_2 sends a random string α_s (which, by construction, is ignored by I_2 , see point (b)). The output of H_2 is the view of I_2 (which in turn is the view of A^*).

Hybrid H_3 : For procedure I_2 , let \mathcal{E} be our online extractor as defined in section 2.3 (with addendum from section 4 to deal with statistically-hiding PRS as well). Recall that \mathcal{E} is a super-polynomial time machine, which facilitates the robust-concurrent attack. In particular, it acts honest receivers in all PRS preambles, as well as provides extractions for each one of them when they finish.

On input $(1^n, v_b, z)$, hybrid H_3 starts an execution of $I_2(1^n, v_b, z)$, making it interact with party \mathcal{E} . The output of H_3 is the output of the robust-concurrent

attack, which in turn consists of the view of I_2 (and hence A^*). Therefore, H_3 differs significantly from H_2 : it does not run PRS receivers and does not have access to \mathcal{O}' , since these are automatically done by \mathcal{E} for H_3 .

From hereon, we will only be making changes to the interface procedure I_2 . All future hybrids, except for the last one, differ from H_2 in only that they use a modified version of I_2 . Furthermore, changes are made to the phases of the *left session* only, of which there is only one.

Hybrid H_4 : This hybrid is identical to H_3 except that instead of using I_2 , it uses procedure I_4 . Let α_s be the value that I_2 receives from \mathcal{E} at the end of statistically-hiding PRS₂ of the *left session*. Furthermore, let σ_s be the value opened by A^* which appears internally, in the *left session* of CCA-Com (simulated for A^* by I_2).

Procedure I_4 is identical to I_2 except that when A^* sends σ_s along with valid openings, I_4 tests that $\alpha_s = \sigma_s$; it aborts the entire simulation if this test fails, and outputs `BindingFail`. Recall that I_2 simply ignores α_s .

Hybrid H_5 : This hybrid is identical to H_4 except that instead of using I_4 , it uses a modified procedure I_5 . Procedure I_5 is identical to I_4 except that instead of committing to 0^n , it commits to α_s in protocol NMCCom of the *left session* s .

Hybrid H_6 : Identical to H_5 except that instead of using I_5 , it uses a modified procedure I_6 . Procedure I_6 is identical to I_5 except that it uses the “fake witness” in the WI proof of the *left session* s . Recall that the “fake witness” consists of the value α_s and the randomness used in NMCCom.

Hybrid H_7 : Identical to H_6 except that instead of using I_6 , it uses a modified procedure I_7 . Procedure I_7 is identical to I_6 except it does not receive the input v_b , and commits to 0^n in the first phase PRS₁ of the *left session*. Observe that I_7 is in fact identical to I .

Hybrid H_8 : This hybrid differs from H_7 in two crucial places. First, it does not receive the value v_b as input. Therefore, its only inputs are $(1^n, z)$. In addition, it does not use the party \mathcal{E} . Instead, H_8 simply runs the robust simulator S with inputs $(1^n, z)$ and black-box I_7 . It outputs whatever S outputs. Observe that H_8 is in fact our original simulator SIM, presented earlier. \square

Notation. Let ν^0 denote the output of H_0 . Let u^0 be the variable denoting the value committed by C in NMCCom in the *left session*, and let σ^0 be the variable denoting PRS-value opened by A^* in (phase-4 of) the *left session*. Since there is only one *left session*, we will not use any subscripts. Analogously, define variables \tilde{u}_s^0 and $\tilde{\sigma}_s^0$ for the s -th *right session* in hybrid H_0 . Finally, we denote by \tilde{v}_s^0 the value A^* commits in PRS₁ of the s -th *right session* in H_0 . For $i \in \{0, \dots, 8\}$ define values $\nu^i, \tilde{u}_s^i, \tilde{\sigma}_s^i$ and \tilde{v}_s^i w.r.t. the hybrid H_i analogously. Identity of the *left session* is referred to by `id`, and that of the s -th *right sessions* by `ids` without

mention of the hybrid. Further, unless specified otherwise, index s is in $[T]$, and $\widetilde{\text{id}}_s$ is not equal to id , and the probability is taken over the randomness of the hybrid in consideration.

We start by noting that ν^0 is identical to $\text{IND}_b(\langle C, R \rangle, \mathcal{O}, A, n, z)$. Next, we claim the following.

Claim. For every hybrid $i \in \{0, \dots, 4\}$ and for every right session $s \in [T]$,

$$\nu^0 \stackrel{s}{\equiv} \nu^1 \equiv \nu^2 \stackrel{s}{\equiv} \nu^3 \stackrel{s}{\equiv} \nu^4 \quad (1)$$

$$\Pr [\widetilde{a}_s^i = \widetilde{\sigma}_s^i] \leq \text{negl}(n) \quad (2)$$

PROOF. It is seen by way of construction of hybrids H_1 and H_2 , that $\nu^0 \stackrel{s}{\equiv} \nu^1 \equiv \nu^2$.¹⁰ In case of hybrid H_3 party \mathcal{E} , which also simulates PRS sessions exactly as H_2 does (since \mathcal{E} is valid). Furthermore, \mathcal{E} extracts the values α_s in the (first message) of the PRS preamble exactly as \mathcal{O}' does in H_2 . If a unique and valid α_s exists, the value α_s returned to I_2 is the same in both hybrids. However, when the value is not unique, \mathcal{E} uses a different decision procedure to decide the value of α_s . Nevertheless, statistical-binding of PRS ensures that this happens with only negligible probability. It follows that the distribution of answers α_s in both hybrids is statistically close, and hence $\nu^2 \stackrel{s}{\equiv} \nu^3$.

Finally, in H_4 the only difference is that I_4 verifies that $\alpha_s = \sigma_s$ (i.e., the fake witness is correct). From the validity constraint 1(b) on \mathcal{E} (see lemma 2, section 4) this condition fails with only negligible probability. Therefore, $\nu^3 \stackrel{s}{\equiv} \nu^4$.

We first prove the second equation for hybrid H_0 . Fix any right-session s of H_0 . Observe that PRS_2 of s is statistically-hiding. Therefore, except with negligible probability, we have that value $\widetilde{\sigma}_s^0$ is not defined until after the completion of NMCom of session s . Since NMCom is statistically-binding, and there are exponentially many possible values for $\widetilde{\sigma}_s^0$, the claim follows (for hybrid H_0). Now, observe that the same argument applies for hybrids H_1 and H_2 as well. In case of H_3 and H_4 , since \mathcal{E} simulates PRS receivers honestly, the same argument applies to these hybrids as well. \square

A corollary of the second equation is that in each of these hybrids, for every right session s that is accepting, there exists a *unique* and *valid*¹¹ value \widetilde{v}_s^i to which A^* is committed to (except with negligible probability). This is because if the second equation holds, from the soundness of WI proof against unbounded adversaries,¹² phase PRS_1 must be consistent defining a unique and valid value.

¹⁰ First two hybrids are not identical since the oracle changes: \mathcal{O} extracts from the full PRS transcript (and therefore always checks for consistency), whereas \mathcal{O}' simply extracts from the (first) committing message of PRS. Statistical-binding ensures that this difference happens only in negligible cases.

¹¹ Recall that v is valid if all shares XOR to v ; formally, there exists randomness ρ such that $\text{open}_{\text{PRS}}(\tau_1, v_s, \rho) = 1$ where τ_1 is the commitment-transcript of PRS_1 of s .

¹² We need this since A^* does have access to super-polynomial computations via \mathcal{O} , and we do not know what he might be learning.

Therefore, in all future hybrids, we will continue to maintain the second equation as an invariant.

Claim. We have that, $\nu^4 \stackrel{c}{=} \nu^5$; and $\forall s : \Pr [\tilde{u}_s^5 = \tilde{\sigma}_s^5] \leq \text{negl}(n)$.

PROOF. Both hybrids H_4 and H_5 use party \mathcal{E} which is super-polynomial time. Therefore, to prove the claim, we need to first need to eliminate the use of \mathcal{E} . Observe that this can be done by employing the robust simulator S . However, since we aim to reduce the claim to the non-malleability of NMCom , we would like one left execution and one right execution of NMCom that the robust-simulator S does not rewind.

Recall that the only difference between H_4 and H_5 is that they use different procedures: I_4 and I_5 respectively. The only difference between I_4 and I_5 is that the first one commits to $x = 0^n$ whereas the second one commits to $x = \alpha_s$ in phase NMCom of the left session. Recall that α_s is the value sent by \mathcal{E} at the conclusion of PRS_2 of the left session.

We design an intermediate procedure I_* , which is identical to I_4 except for the following differences:

1. When the NMCom phase of the left session is about to begin, I_* sends $(0^n, \alpha_s)$ to an external committer.
2. I_* does not execute the NMCom of the left session internally. Instead, it expects to receive it from an external honest committer, denoted C_{NM} , who either commits to 0^n or α_s .
3. For a randomly chosen right session j , I_* does not execute the the NMCom of session s internally. Instead, it forwards the messages of this NMCom to an external honest receiver, denoted R_{NM}^j .

Therefore, I_* is executed several PRS preambles, and at the same time acting as a man-in-the-middle for protocol NMCom by receiving a commitment and making a commitment at the same time. Let B denote the party who runs algorithm C_{NM} as well as R_{NM}^j for I_* , both *honestly*. Observe that the number of rounds of interaction between B and I_* are $2k$ where k is the round-complexity of NMCom . The input x to B consists of the value to be committed by C_{NM} .

Viewed this way, I_* is an adversary who launches the robust-concurrent attack with respect to party B , and the party \mathcal{E} . Furthermore, if B commits to $x = 0^n$ then the execution is identical to that of H_4 with I_4 ; on the other hand, if $x = \alpha_s$, the execution is identical to that of H_5 with procedure I_5 . That is,

$$\begin{aligned} \nu^4 &\equiv \text{REAL}_{\mathcal{E}, \Pi}^{I_*}(n, 0^n, z), \\ \nu^5 &\equiv \text{REAL}_{\mathcal{E}, \Pi}^{I_*}(n, \alpha_s, z). \end{aligned}$$

where protocol $\Pi = (B, P_2)$, and P_2 is “converse of B —that is P_2 acts as a receiver in one NMCom in which B is acting as a committer and vice-versa in the other.

Now, we can remove \mathcal{E} and instead use the robust-simulator S to sample statistically close views. That is, suppose that we run S with I_* and consider

the output $\text{OUT}_s(x) := \text{OUT}_s [B(1^n, 0^n) \leftrightarrow S^{I^*}(1^n, z)]$. By applying the robust concurrent extraction lemma, we have that statistical distance between $\text{OUT}_s(x)$ and $\text{REAL}_{\mathcal{E}, \Pi}^{I^*}(n, x, z)$ is at most:

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

for every $x \in \{0^n, \alpha_s\}$ since $\ell \in \omega(\log n)$, k is a constant, and T is at most a polynomial. Using this with the equations above, we get:

$$\nu^4 \stackrel{s}{\equiv} \text{OUT}_s(0^n) \tag{3}$$

$$\nu^5 \stackrel{s}{\equiv} \text{OUT}_s(\alpha_s) \tag{4}$$

By construction, $\text{OUT}_s(x)$ is the output of S^{I^*} in an interaction with B on input x . Algorithm S^{I^*} is a PPT man-in-the-middle adversary for NMCom who receives a commitment to x from C_{NM} , and commits a value, say $\tilde{u}_j(x)$, to R_{NM}^j . From non-malleability of NMCom w.r.t. to itself, we have that:

$$(\tilde{u}_j(0^n), \text{OUT}_s(0^n)) \stackrel{c}{\equiv} (\tilde{u}_j(\alpha_s), \text{OUT}_s(\alpha_s)) \tag{5}$$

It immediately follows from (3), (4), and (5), that $\nu^4 \stackrel{c}{\equiv} \nu^5$. Furthermore, suppose that the other part of the claim is false, so that for some session $s \in [T]$ value $\tilde{u}_s^5 = \tilde{\sigma}_s^5$ with noticeable probability p . Define this to be event bad_5 for hybrid H_5 , and analogously define bad_4 for H_4 .

Now observe that since j was chosen uniformly from T sessions, $j = s$ with probability $1/T$. Therefore, value \tilde{u}_s^5 appears as the variable $\tilde{u}_j(\alpha_s)$ with probability $1/T$. In addition, value $\tilde{\sigma}_j(\alpha_s)$ is a part of $\text{OUT}_s(\alpha_s)$, and therefore, event bad_5 is efficiently observable given both $(\tilde{u}_j(\alpha_s), \text{OUT}_s(\alpha_s))$; and it occurs with probability p/T which is noticeable. By an analogous argument, event bad_4 is also efficiently observable; furthermore bad_4 must occur with noticeable probability as well due to equation (5). This contradicts equation (2). Hence the claim. \square

Claim. We have that, $\nu^5 \stackrel{c}{\equiv} \nu^6$; and $\forall s : \Pr [\tilde{u}_s^6 = \tilde{\sigma}_s^6] \leq \text{negl}(n)$.

PROOF. The proof of this claim is almost identical to the proof of claim 5. The only difference is that instead of using non-malleability w.r.t. itself property of NMCom , we use the fact that NMCom is robust w.r.t. every *interesting* 3-round protocol (i.e., one that “hides” its input, see the non-malleable commitments definition in the full version [GLP⁺15]). Since the only difference between H_5 and H_6 is in the WI part, the proof follows; we omit the details. \square

Claim. We have that, $\nu^6 \stackrel{c}{\equiv} \nu^7$; and $\forall s : \Pr [\tilde{u}_s^7 = \tilde{\sigma}_s^7] \leq \text{negl}(n)$.

PROOF. Observe that in H_6 , since the “fake witness” is being used in the WI part, the PRS_1 phase of the left session (which commits to input value v_b) does not have to be consistent. Therefore, we proceed as follows:

1. Design $\ell + 1$ intermediate hybrids $H_{6:i}$ for $i = \{0, \dots, \ell - 1\}$ where $H_{6:0} = H_6$, $H_{6:\ell} = H_7$.
2. Hybrid $H_{6:i}$ is the same as $H_{6:i+1}$ except that in slot- i of PRS_1 phase of the left session, $H_{6:i+1}$ commits to shares of an all-zero string.
3. Now, following the proof of claim 5 and by using the robustness of NMCom w.r.t. the 3-round protocols, we conclude that $\nu^{6:i} \stackrel{c}{\equiv} \nu^{6:i+1}$ and that $\forall s : \Pr [\tilde{u}_s^{6:i} = \tilde{\sigma}_s^{6:i+1}] \leq \text{negl}(n)$, where variables are analogously defined. The details of this part are repetitive, and omitted.

The claim now follows. □

Completing the Proof. Note that the output of H_8 , is statistically-close to ν^7 due to the robust-concurrent-extraction lemma. Therefore, combining all the equations we have that $\nu^0 \stackrel{c}{\equiv} \nu^8$. Since H_8 is identical to the simulator, the lemma follows. □

References

- [AGJ⁺12] Agrawal, S., Goyal, V., Jain, A., Prabhakaran, M., Sahai, A.: New impossibility results for concurrent composition and a non-interactive completeness theorem for secure computation. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 443–460. Springer, Heidelberg (2012)
- [Bar02] Barak, B.: Constant-round coin-tossing with a man in the middle or realizing the shared random string model. In: FOCS (2002)
- [BCNP04] Barak, B., Canetti, R., Nielsen, J.B., Pass, R.: Universally composable protocols with relaxed set-up assumptions. In: Proc.45th FOCS, pp. 186–195 (2004)
- [Bea91] Beaver, D.: Foundations of secure interactive computing. In: Feigenbaum, J. (ed.) Advances in Cryptology - CRYPTO 1991. LNCS, vol. 576, pp. 377–391. Springer, Heidelberg (1992)
- [Blu87] Blum, M.: How to prove a theorem so no one else can claim it. In: Proceedings of the International Congress of Mathematicians, pp. 1444–1451 (1987)
- [BNP08] Ben-David, A., Nisan, N., Pinkas, B.: Fairplaymp: A system for secure multi-party computation. In: ACM Conference on Computer and Communications Security, pp. 257–266 (2008)
- [BPS06] Barak, B., Prabhakaran, M., Sahai, A.: Concurrent non-malleable zero knowledge. In: FOCS, pp. 345–354 (2006), Full version available on eprint archive
- [BS05] Barak, B., Sahai, A.: How to play almost any mental game over the net - concurrent composition using super-polynomial simulation. In: Proc. 46th FOCS (2005)
- [Can00] Canetti, R.: Security and composition of multiparty cryptographic protocols. Journal of Cryptology: The Journal of the International Association for Cryptologic Research 13(1), 143–202 (2000)

- [Can01] Canetti, R.: Universally composable security: A new paradigm for cryptographic protocols. In: Werner, B. (ed.) Proc. 42nd FOCS, pp. 136–147 (2000), Preliminary full version available as Cryptology ePrint Archive Report 2000/067
- [CF01] Canetti, R., Fischlin, M.: Universally composable commitments. Report 2001/055, Cryptology ePrint Archive (July 2001), Extended abstract appeared in CRYPTO 2001
- [CGS08] Chandran, N., Goyal, V., Sahai, A.: New constructions for UC secure computation using tamper-proof hardware. In: Smart, N.P. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 545–562. Springer, Heidelberg (2008)
- [CKL03] Canetti, R., Kushilevitz, E., Lindell, Y.: On the limitations of universally composable two-party computation without set-up assumptions. In: Biham, E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, Springer, Heidelberg (2003)
- [CKPR01] Canetti, R., Kilian, J., Petrank, E., Rosen, A.: Black-box concurrent zero-knowledge requires (almost) logarithmically many rounds. In: STOC (2001)
- [CLOS02] Canetti, R., Lindell, Y., Ostrovsky, R., Sahai, A.: Universally composable two-party computation. In: Proc. 34th STOC, pp. 494–503 (2002)
- [CLP10] Canetti, R., Lin, H., Pass, R.: Adaptive hardness and composable security in the plain model from standard assumptions. In: FOCS, pp. 541–550 (2010), Full version:
<http://www.cs.cornell.edu/~rafael/papers/ccaccommit.pdf>
- [CLP14] Chung, K.-M., Lin, H., Pass, R.: Constant-round concurrent zero-knowledge from indistinguishability obfuscation. Cryptology ePrint Archive, Report 2014/991 (2014), <http://eprint.iacr.org/>
- [CPS07] Canetti, R., Pass, R., Shelat, A.: Cryptography from sunspots: How to use an imperfect reference string. In: FOCS, pp. 249–259 (2007)
- [DDN91] Dolev, D., Dwork, C., Naor, M.: Non-malleable cryptography (extended abstract). In: STOC, pp. 542–552 (1991)
- [GGH⁺13] Garg, S., Gentry, C., Halevi, S., Raykova, M., Sahai, A., Waters, B.: Candidate indistinguishability obfuscation and functional encryption for all circuits. In: FOCS, pp. 40–49 (2013)
- [GGJS12] Garg, S., Goyal, V., Jain, A., Sahai, A.: Concurrently secure computation in constant rounds. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 99–116. Springer, Heidelberg (2012)
- [GJO10] Goyal, V., Jain, A., Ostrovsky, R.: Password-authenticated session-key generation on the internet in the plain model. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 277–294. Springer, Heidelberg (2010), <http://research.microsoft.com/en-us/um/people/vipul/pke.pdf>
- [GL90] Goldwasser, S., Levin, L.A.: Fair computation of general functions in presence of immoral majority. In: Menezes, A., Vanstone, S.A. (eds.) Advances in Cryptology - CRYPTO 1990. LNCS, vol. 537, pp. 77–93. Springer, Heidelberg (1991)
- [GLP⁺15] Goyal, V., Lin, H., Pandey, O., Pass, R., Sahai, A.: Round-efficient concurrently composable secure computation via a robust extraction lemma. In: Dodis, Y., Nielsen, J.B. (eds.) TCC 2015, Part I. LNCS, vol. 9014, pp. 260–289. Springer, Heidelberg (2015); Full version of this work available as IACR Eprint Report 2012/652
- [GMW87] Goldreich, O., Micali, S., Wigderson, A.: How to play ANY mental game. In: ACM (ed.) Proc. 19th STOC, pp. 218–229 (1987), See [Gol04, ch. 7] for more details

- [GMW91] Goldreich, O., Micali, S., Wigderson, A.: Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. *Journal of the ACM* 38(3), 691–729 (1991), Preliminary version in FOCS 1986
- [Gol04] Goldreich, O.: *Foundations of Cryptography: Basic Applications*. Cambridge University Press (2004)
- [Goy11] Goyal, V.: Constant round non-malleable protocols using one way functions. In: STOC, pp. 695–704 (2011), See full version available on <http://research.microsoft.com/en-us/people/vipul/nmcom.pdf>
- [Goy12] Goyal, V.: Positive results for concurrently secure computation in the plain model. In: FOCS, pp. 695–704 (2012)
- [GS09] Goyal, V., Sahai, A.: Resettable secure computation. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 54–71. Springer, Heidelberg (2009)
- [IPS15] Ishai, Y., Pandey, O., Sahai, A.: Public-coin differing-inputs obfuscation and its applications. In: Dodis, Y., Nielsen, J.B. (eds.) TCC 2015, Part II. LNCS, vol. 9015, pp. 668–697. Springer, Heidelberg (2015)
- [Kat07] Katz, J.: Universally composable multi-party computation using tamper-proof hardware. In: Naor, M. (ed.) EUROCRYPT 2007. LNCS, vol. 4515, pp. 115–128. Springer, Heidelberg (2007)
- [KP01] Kilian, J., Petrank, E.: Concurrent and resettable zero-knowledge in poly-logarithm rounds. In: Proc. 33th STOC, pp. 560–569 (2000), Preliminary full version published as cryptology ePrint report 2000/013
- [Lin04] Lindell, Y.: Lower bounds for concurrent self composition. In: Naor, M. (ed.) TCC 2004. LNCS, vol. 2951, pp. 203–222. Springer, Heidelberg (2004)
- [LP09] Lin, H., Pass, R.: Non-malleability amplification. In: STOC, pp. 189–198 (2009)
- [LP11] Lin, H., Pass, R.: Constant-round non-malleable commitments from any one-way function. In: STOC (2011)
- [LP12] Lin, H., Pass, R.: Black-box constructions of composable protocols without set-up. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 461–478. Springer, Heidelberg (2012)
- [LPTV10] Lin, H., Pass, R., Tseng, W.-L.D., Venkatasubramanian, M.: Concurrent non-malleable zero knowledge proofs. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 429–446. Springer, Heidelberg (2010)
- [MNPS04] Malkhi, D., Nisan, N., Pinkas, B., Sella, Y.: Fairplay - secure two-party computation system. In: USENIX Security Symposium, pp. 287–302 (2004)
- [MOSV06] Micciancio, D., Ong, S.J., Sahai, A., Vadhan, S.P.: Concurrent zero knowledge without complexity assumptions. In: Halevi, S., Rabin, T. (eds.) TCC 2006. LNCS, vol. 3876, pp. 1–20. Springer, Heidelberg (2006)
- [MPR06] Micali, S., Pass, R., Rosen, A.: Input-indistinguishable computation. In: FOCS, pp. 367–378 (2006)
- [MR91] Micali, S., Rogaway, P.: Secure computation. In: Feigenbaum, J. (ed.) CRYPTO 1991. LNCS, vol. 576, pp. 392–404. Springer, Heidelberg (1992)
- [OPV10] Ostrovsky, R., Pandey, O., Visconti, I.: Efficiency preserving transformations for concurrent non-malleable zero knowledge. In: Micciancio, D. (ed.) TCC 2010. LNCS, vol. 5978, pp. 535–552. Springer, Heidelberg (2010)
- [Pas03] Pass, R.: Simulation in quasi-polynomial time, and its application to protocol composition. In: Biham, E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, pp. 160–176. Springer, Heidelberg (2003)

- [PLV12] Pass, R., Lin, H., Venkatasubramanian, M.: A unified framework for UC from only OT. In: Wang, X., Sako, K. (eds.) ASIACRYPT 2012. LNCS, vol. 7658, pp. 699–717. Springer, Heidelberg (2012)
- [PPS15] Pandey, O., Prabhakaran, M., Sahai, A.: Obfuscation-Based Non-black-box Simulation and Four Message Concurrent Zero Knowledge for NP. In: Dodis, Y., Nielsen, J.B. (eds.) TCC 2015, Part II. LNCS, vol. 9015, pp. 638–667. Springer, Heidelberg (2015)
- [PPV08] Pandey, O., Pass, R., Vaikuntanathan, V.: Adaptive one-way functions and applications. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 57–74. Springer, Heidelberg (2008)
- [PR05a] Pass, R., Rosen, A.: Concurrent non-malleable commitments. In: FOCS (2005)
- [PR05b] Pass, R., Rosen, A.: New and improved constructions of non-malleable cryptographic protocols. In: STOC (2005)
- [PRS02] Prabhakaran, M., Rosen, A., Sahai, A.: Concurrent zero knowledge with logarithmic round-complexity. In: FOCS (2002)
- [PS04] Prabhakaran, M., Sahai, A.: New notions of security: achieving universal composability without trusted setup. In: STOC, pp. 242–251 (2004)
- [PTV08] Pass, R., Tseng, W.-L.D., Venkatasubramanian, M.: Concurrent zero knowledge: Simplifications and generalizations (2008) (manuscript), <http://hdl.handle.net/1813/10772>
- [PV08] Pass, R., Venkatasubramanian, M.: On constant-round concurrent zero-knowledge. In: Canetti, R. (ed.) TCC 2008. LNCS, vol. 4948, pp. 553–570. Springer, Heidelberg (2008)
- [PW01] Pfitzmann, B., Waidner, M.: A model for asynchronous reactive systems and its application to secure message transmission. In: IEEE Symposium on Security and Privacy (2001)
- [RK99] Richardson, R., Kilian, J.: On the concurrent composition of zero-knowledge proofs. In: Stern, J. (ed.) EUROCRYPT 1999. LNCS, vol. 1592, pp. 415–432. Eurocrypt 1999, Heidelberg (1999)
- [Ros04] Rosen, A.: The Round-Complexity of Black-Box Concurrent Zero-Knowledge. PhD thesis, Department of Computer Science and Applied Mathematics, Weizmann Institute of Science, Rehovot, Israel (2004)
- [Sah99] Sahai, A.: Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In: Proc. 40th FOCS, pp. 543–553 (1999)
- [Wee10] Wee, H.: Black-box, round-efficient secure computation via non-malleability amplification. In: FOCS (2010)
- [Yao86] Yao, A.C.-C.: How to generate and exchange secrets. In: Proc. 27th FOCS, pp. 162–167 (1986)