Constant-Round Concurrent Non-Malleable Statistically Binding Commitments and Decommitments

Zhenfu Cao[1], Ivan Visconti[2,⋆], and Zongyang Zhang[1]

[1] Department of Computer Science and Engineering,
Shanghai Jiao Tong University, P.R. China
{zfcao,zongyangzhang}@sjtu.edu.cn

[2] Dipartimento di Informatica ed Applicazioni,
University of Salerno, Italy
visconti@dia.unisa.it

Abstract. When commitment schemes are used in complex environments, e.g., the Internet, the issue of malleability appears, i.e., a concurrent man-in-the-middle adversary might generate commitments to values related to ones committed to by honest players. In the plain model, the current best solution towards resolving this problem in a constant number of rounds is the work of Ostrovsky, Persiano and Visconti (TCC' 09). They constructed a constant-round commitment scheme that is concurrent non-malleable with respect to both commitment and decommitment. However, the scheme is only computationally binding. For application scenarios where the security of receivers is of a great concern, computational binding may not suffice.

In this work, we follow the line of their work and give a construction of statistically binding commitment scheme which is concurrent non-malleable with respect to both commitment and decommitment. Our work can be seen as a complement of the work of Ostrovsky et al. in the plain model. Our construction relies on the existence of a family of pairs of claw-free permutations and only needs a constant number of communication rounds in the plain model. Our proof of security uses non-black-box techniques and satisfies the (most powerful) simulation-based definitions of non-malleability.

Keywords: commitment schemes, statistically binding, non-malleability.

1 Introduction

A commitment scheme is a two-phase interactive protocol between two parties, the committer, who holds a value, and the receiver. It enables the committer to commit itself a value while keeping it secret from the receiver. Two basic properties of a commitment scheme are the hiding property (the receiver can not learn the committed value before the decommitment phase) and the binding

⋆ Work done while Ivan Visconti was visiting UCLA, USA.

P.Q. Nguyen and D. Pointcheval (Eds.): PKC 2010, LNCS 6056, pp. 193–208, 2010.

property (the committer is bounded to one value after the commitment phase). In the literature, two fundamental types of commitment schemes, statistical hiding and statistical binding, are considered.

It is well known that the basic properties of commitment schemes can not prevent "malleability" attacks mounted by a probabilistic polynomial-time (PPT) man-in-the-middle (MIM) adversary who has full control of the communication channel between the committer and the receiver. The concept of non-malleability was first introduced by Dolev et al. [1] to capture security concerns in such settings. Loosely speaking, a commitment scheme is non-malleable if one can not transform the commitment of a value into a commitment of a related value. This kind of non-malleability is called *non-malleability with respect to commitment* (NMc for short) [1]. This definition is based on the independence of the committed messages played by the MIM adversary with respect to the ones played by the committer. The notion of non-malleability used by Di Crescenzo et al. [2] is called *non-malleability with respect to decommitment or opening* (NMd for short), i.e., the adversary can not construct a commitment from a given one, such that after having seen the opening of the original commitment, the adversary is able to correctly open his commitment with a related value. This definition requires that the success probability of a MIM adversary is maintained by a stand-alone simulator. Subsequent NMc definitions are modified in a similar way [3,4,5,6,7,8]. Simulation-based definitions are much more useful when a commitment scheme is used as a building block in a larger protocol since the existence of a simulator heavily simplifies the task of proving the security of the larger protocol.

Intuitively, it seems that NMc is stronger than NMd. However, this depends on the subtleties of the definitions. Indeed this does not necessarily always hold at least with respect to non-malleability definitions in [2,1,3]. In a journal version of [3], the authors [9] presented a stringent definition of non-malleability w.r.t commitment in order to imply the notion of non-malleability w.r.t opening.

Several previous results focused on designing statistically hiding commitment schemes which are NMd. Based on number-theoretic assumptions, NMd commitment schemes were designed in [10,3] assuming the existence of a common reference string (CRS) that is shared by the two players before the protocol execution. Thus, their schemes do not work in the plain model (i.e., without setup assumptions). Recently, Pass and Rosen [4,5] presented a slightly different definition of NMd.[1] They then constructed a commitment scheme under their NMd definition based on a family of collision-resistant hash functions in the plain model. Their scheme is round-efficient and needs only constant-round communication. More recently, based on the work of [11,12], Zhang et al. [13] presented a non-malleable commitment scheme under the weakest assumption, i.e., the existence of one-way functions.

[1] More precisely, the NMd definitions in [2,10] do not take into account possible a priori information the adversary might have about the commitment received in the left interaction, while the definitions in [3,4,5] do. The definitions in [2,10,3] do not provide the stand-alone simulator the value committed in the left interaction after the commitment phase is finished, while the definitions in [4,5] do.

Before the work of [8], it was commonly believed that NMd (compared with NMc) is the only notion that makes sense in a computationally binding commitment scheme [1]. However, Ostrovsky et al. [8] argued that by slightly relaxing the NMc definition,[2] NMc can also be achieved for computationally binding commitment schemes. They considered concurrent MIM attacks where the adversary can simultaneously participate in any polynomial number of executions as a receiver and as a committer. Based on the work of [6,7], and using some techniques already introduced in [14,15] they gave a computationally hiding and computationally binding commitment scheme which is both concurrent NMc and concurrent NMd. In a full version of [8], they [16] further gave a construction of a constant-round statistically hiding commitment scheme which is concurrent NMd and that actually consists of a simplified protocol with respect to the one presented in [8]. The above schemes assume the existence of a family of pairs of claw-free permutations, require constant number of communication rounds only and assume that commitment phase and decommitment phase do not overlap in time.

For statistically binding commitment schemes, the first NMc one was designed by Dolev et al. [1] assuming the existence of one-way functions. However, the scheme requires $O(\log n)$ rounds, where n is the security parameter. In the CRS model, Di Crescenzo et al. [10] constructed very efficient NMc commitment schemes based on any public-key cryptosystem that is non-malleable under chosen plaintext attacks in addition to any shared-key cryptosystem that enjoys indistinguishability under plaintext oracle CCA-post attack. In the plain model, Pass and Rosen [4,5] first constructed a *constant-round* NMc commitment scheme assuming the existence of collision resistant hash functions. Pass and Rosen [6,7] then showed the NMc scheme of [4,5] is actually a concurrent NMc one under a stronger simulation-based definition.[3] The security proofs of [4,5,6,7] requires a non-black-box use of the code of the adversary and moreover the one of [6,7] assumes that commitment phase and decommitment phase do not overlap in time. Lin et al. [12] reconsidered the scheme of [1] and presented a concurrent NMc commitment scheme using only black-box techniques. Their scheme requires a polynomial number of communication rounds and is based on the minimal assumption, i.e., existence of one-way functions. In addition to the above results focusing on NMc, the only one that *explicitly* claimed NMd commitment schemes was designed in [9] (see Sec. 3) in the CRS model.

Before the clarification of [8], another folklore belief about a statistically binding commitment scheme is that if it is NMc then it is NMd. However, at least

[2] The values committed to by the adversary in a MIM execution are uniquely defined for all algorithms in the NMc definition [1,3], but only for PPT algorithms in the relaxed definition. More recently, the NMc definition formulated in [9] can also be applied to computationally binding commitment scheme.

[3] The NMd definition in [6,7] is stronger than that in [2]. The former is a indistinguishability-based definition, i.e., there exists a PPT stand-alone simulator that commits to a value which is computationally indistinguishable from the value committed to by the MIM adversary. The latter is a relation-based definition, i.e., the stand-alone simulator is less likely to commit to a value satisfying any polynomial-time computable relation than the value committed to by the MIM adversary.

this can not be deduced just from the simulation-based definitions in [4,5,6,7] in the plain model.[4] The main problem is that the success probability of the stand-alone simulator is required to be only negligible close to the success probability of the MIM adversary [8]. Recall in the NMc proof [4,5,6,7], a stand-alone simulator will internally simulate the left interaction for the MIM adversary by committing to a bogus value 0^n. It seems that this simulator can not handle the NMd proof, because after receiving a committed value m, the simulator is stuck to open the bogus commitment to m.

Therefore, achieving simultaneously concurrent NMc and NMd in a constant number of rounds and under the simulation-based notions, the work of [8] achieves the strongest security for commitment schemes in the plain model. However, the scheme is only computationally binding. When the security of receivers is of a great concern in some application scenarios, it may not be sufficient. Thus, there remains an open problem as to whether or not constant-round statistically binding commitment scheme that is both concurrent NMc and concurrent NMd exists in the plain model, under the stronger simulation-based definition [6,5,6,7,8].

1.1 Our Contribution

We solve the above problem by presenting a round-efficient protocol for concurrent non-malleable statistically binding commitment scheme. We show the following theorem.

Theorem 1. *Suppose that there exists a family of pairs of claw-free permutations. Then there exists a constant-round statistically binding commitment scheme that is both concurrent* NMc *and concurrent* NMd.

On a high level view, the commitment phase of our scheme is almost identical with that in [8]. The technique used in this phase is also the same. More precisely, in addition to the technique used by [4,5,6,7], the two-witness technique of Feige [17] is also employed. Our contribution lies in the modification of the open phase in order to simultaneously achieve concurrent NMd and statistical binding property. We borrow the idea of [18] in designing concurrent zero-knowledge *proofs*, i.e., we let the committer guess the private values committed to by the receiver in the commitment phase, and then use a witness-indistinguishable *proof* system to prove a carefully designed statement. In this way, the scheme is guaranteed to prevent any unbounded adversary from opening the commitment in two different ways.

Our work can be viewed as a complement of the work of [8]. Both of the work resolve the non-malleability issues against concurrent man-in-the-middle attacks and achieve the same-level of security in the plain model. The main difference between the two results lies in that the work of [8] focuses mainly on computationally binding commitment schemes, whereas our work considers statistically binding ones. Compared with the work of [6,7], our work also achieves both concurrent NMc and concurrent NMd, whereas they only achieve concurrent NMc.

[4] There is no problem in the CRS model. The reader is refereed to [8] for more details.

We emphasize here that our scheme inherits the limitation from [6,7,8], i.e., the non-malleability proof heavily relies on the assumption that commitment phase and the decommitment phase do not overlap in time.

2 Preliminaries

We assume the reader is familiar with witness-indistinguishable protocols, zero-knowledge protocols and commitment schemes. For more details, the reader is refereed to [19] for references.

2.1 Concurrent Non-Malleable Commitments and Decommitments

Next, we formulate the definitions of concurrent NMc and concurrent NMd. As stated in [6,7] we formalize the notion of non-malleability by a comparison between a *man-in-the-middle* execution and a *simulated* execution. Let $\langle C, R\rangle$ be a commitment scheme. Let $n \in \mathbb{N}$ be a security parameter.

The man-in-the-middle execution. In the MIM execution, the adversary $\mathcal{A}$ is simultaneously participating in $m(n) = \mathsf{poly}(n)$ left and $m(n)$ right interactions (WLOG, the number of commitments is the same in the left and right execution). In the i^{th} left interaction, $\mathcal{A}$ interacts with the committer C to receive a commitment to a value v_i. In the i^{th} right interaction, $\mathcal{A}$ interacts with the receiver R and tries to commit to a value $\tilde{v}_i$ of its choice. After the execution of the commitments in all interactions, $\mathcal{A}$ executes the decommitments with C and the decommitments with R. Prior to the interaction, the value vector $\mathbb{V} = (v_1, \ldots, v_m)$ is given to C as local inputs. $\mathcal{A}$ also receives an auxiliary input z, which might contain a priori information about $\mathbb{V}$.

Let the random variable $\mathsf{mim}^{\mathcal{A}}_{\mathsf{com}}(\mathbb{V}, z)$ denote the values $\tilde{v}_1, \ldots, \tilde{v}_m$ to which the adversary has committed in the right interactions. If the i^{th} right commitment fails, or its transcript (commitment phase) equals to the transcript of any left interaction, the value $\tilde{v}_i$ is set to $\perp$.

Similarly, we let the random variable $\mathsf{mim}^{\mathcal{A}}_{\mathsf{open}}(\mathbb{V}, z)$ denote the values $\tilde{v}_1, \ldots, \tilde{v}_m$ to which the adversary has opened in the right interactions. If the i^{th} right commitment or decommitment fails, or its transcript (both commitment phase and decommitment phase) equals to the transcript of any left interaction, the value $\tilde{v}_i$ is set to $\perp$.

The simulated execution. In the simulated execution, a simulator S directly interacts with an honest receiver R in $m(n)$ interactions. As in the MIM execution, the value vector $\mathbb{V} = (v_1, \ldots, v_m)$ is chosen prior to the interaction, and S receives some a prior information about $\mathbb{V}$ as part of its auxiliary input z. S first executes the commitment phases with R. Once all the commitment phases have been completed, S receives the value vector $\mathbb{V}$ and attempts to decommit to values $\tilde{v}_1, \ldots, \tilde{v}_m$.

Let the random variable $\mathsf{sim}^{S}_{\mathsf{com}}(\mathbb{V}, z)$ denote the values $\tilde{v}_1, \ldots, \tilde{v}_m$ committed to by S. The value $\tilde{v}_i$ is set to $\perp$ if S fails in the i^{th} commitment phase. Let the random variable $\mathsf{sim}^S_{\mathsf{open}}(\mathbb{V}, z)$ denote the values $\tilde{v}_1, \ldots, \tilde{v}_m$ opened by S. The value $\tilde{v}_i$ is set to $\perp$ if S fails in the i^{th} commitment phase or decommitment phase.

Definition 1 (Concurrent Non-Malleable Commitment w.r.t Commitment [6,7]). *A commitment scheme $\langle C, R\rangle$ is said to be* concurrent non-malleable with respect to commitment *if for every PPT man-in-the-middle adversary $\mathcal{A}$ that participates in at most $m(n)$ left and $m(n)$ right interactions, there exists a PPT simulator S such that the following two ensembles are computationally indistinguishable:*

- $\{\mathsf{mim}^{\mathcal{A}}_{\mathsf{com}}(\mathbb{V}, z)\}_{\mathbb{V}=(v_1,\ldots,v_m)\in\{0,1\}^{n*m}, z\in\{0,1\}^*}$
- $\{\mathsf{sim}^{S}_{\mathsf{com}}(\mathbb{V}, z)\}_{\mathbb{V}=(v_1,\ldots,v_m)\in\{0,1\}^{n*m}, z\in\{0,1\}^*}$

Definition 2 (Concurrent Non-Malleable Commitment w.r.t Decommitment). *A commitment scheme $\langle C, R\rangle$ is said to be* concurrent non-malleable with respect to decommitment *if for every PPT man-in-the-middle adversary $\mathcal{A}$ that participates in at most $m(n)$ left and $m(n)$ right interactions, there exists an expected PPT simulator S such that the following two ensembles are computationally indistinguishable:*

- $\{\mathsf{mim}^{\mathcal{A}}_{\mathsf{open}}(\mathbb{V}, z)\}_{\mathbb{V}=(v_1,\ldots,v_m)\in\{0,1\}^{n*m}, z\in\{0,1\}^*}$
- $\{\mathsf{sim}^{S}_{\mathsf{open}}(\mathbb{V}, z)\}_{\mathbb{V}=(v_1,\ldots,v_m)\in\{0,1\}^{n*m}, z\in\{0,1\}^*}$

A commitment scheme that is non-malleable according to Definition 2 is liberal non-malleable rather than strict non-malleable [1,3]. Note we follow [4,5,8] in that non-malleability is guaranteed only if the commitment phase and the decommitment phase do not overlap in time.

Strong signature schemes. A signature scheme $\mathsf{SS} = (\mathsf{Sgen}, \mathsf{Ssig}, \mathsf{Sver})$ is said to be *strongly unforgeable* under adaptive chosen-message attack if no efficient adversary, with access to signature oracle with respect to the verification key VK, can output a valid message/signature pair (m, σ) with non-negligible probability. Here "valid" means that $\mathsf{Sver}(\mathsf{VK}, \mathsf{m}, \sigma) = 1$ and (m, σ) does not correspond to any message/signature pair that was output by the signature oracle. A *strong signature scheme* is a signature scheme that is strongly unforgeable.

3 Constant-Round Statistically Binding Concurrent NMc and Concurrent NMd

In this section, we present a constant-round statistically binding commitment scheme that is concurrent NMc and concurrent NMd. Denote by SBCom the statistically binding commitment scheme from any one-way function [20]. Denote by SHCom the statistically hiding commitment scheme from any collection of claw-free permutation with an efficiently-recognizable index set [21]. Denote by $\{\langle \mathcal{P}_{\text{tag}}, \mathcal{V}_{\text{tag}}\rangle\}_{\text{tag}}$ the constant-round tag-based perfect non-malleable zero-knowledge argument of knowledge ($\mathsf{NMZKAOK}$) for NP [4,5]. Denote by $\langle \mathrm{swi}\mathcal{P}, \mathrm{swi}\mathcal{V}\rangle$ the constant-round statistically witness-indistinguishable argument of knowledge (WIAOK) for NP [22,23].[5] Let $\langle \mathrm{cwi}\mathcal{P}, \mathrm{cwi}\mathcal{V}\rangle$ be a constant-round computationally witness-indistinguishable proof of knowledge (WIPOK) for NP. Let $\mathsf{SS} = (\mathsf{Sgen}, \mathsf{Ssig}, \mathsf{Sver})$ be a strong signature scheme. The commitment scheme is shown in Fig. 1. Note that all the tools used above can be achieved assuming the existence of a family of pairs of claw-free permutations.

Our commitment scheme is a statistically binding variant of the one in [8]. The commitment phase is almost identical with that of the commitment scheme in [8] with the following exception: in Stage 2, the receiver R uses a statistically hiding commitment scheme SHCom instead of a statistically binding one. It also invokes the statistical WIAOK $\langle \mathrm{swi}\mathcal{P}, \mathrm{swi}\mathcal{V}\rangle$ instead of a computational WIPOK. Roughly, in Stage 1, the committer generates a commitment c to v and proves knowledge of opening of c. In Stage 2, the receiver generates two commitments c_0, c_1 to two secretes v_0, v_1 respectively and proves knowledge of either secret. In Stage 3, the committer generates a signature to the transcripts up to now and the receiver then verifies the correctness of the signature.

The decommitment phase is more involved and needs more careful design. The main difficulty lies in simultaneously achieving concurrent NMd and statistical binding properties. We are inspired by the work of [18] on concurrent zero-knowledge *proofs*. We modify the scheme in [8] by letting the committer guess the private values committed to in the commitment phase and then use a WIPOK to prove a carefully designed OR statements. The construction employs the two-witness technique by Feige [17] and the well known FLS-technique [24]. Roughly, in Stage $1'$, the committer first generates a commitment c' to a dummy value 0^n. After receiving c', the receiver then opens the values v_0, v_1 committed to in the commitment phase and proves knowledge of opening of either commitment c_0 or c_1. In Stage $2'$, the committer sends the committed value v and runs a computational WIPOK to prove the statements that either c is a commitment to v, or c' is a commitment to v_0 or v_1. In Stage $3'$, the committer proves that c is a commitment to v, or it knows opening of c_{b^*} to v_{b^*} for some $b^* \in \{0,1\}$. In Stage $4'$, the committer generates a signature to the transcripts up to now and the receiver then verifies the correctness of the signature. Note that the FLS-technique is used both in Stage $2'$ and Stage $3'$.

Theorem 2. *Suppose that* SBCom *is a statistically binding commitment scheme,* SHCom *is a statistically hiding commitment scheme and* $\mathsf{SS} = (\mathsf{Sgen}, \mathsf{Ssig}, \mathsf{Sver})$ *is a strong signature scheme. Suppose that* $\{\langle \mathcal{P}_{\mathrm{tag}}, \mathcal{V}_{\mathrm{tag}}\rangle\}_{\mathrm{tag}}$ *is an one-many concurrent perfect* $\mathsf{NMZKAOK}$ *for* NP*,* $\langle \mathrm{swi}\mathcal{P}, \mathrm{swi}\mathcal{V}\rangle$ *is a statistical* WIAOK *for* NP *and* $\langle \mathrm{cwi}\mathcal{P}, \mathrm{cwi}\mathcal{V}\rangle$ *is a computational* WIPOK *for* NP*. Then* $\langle C, R\rangle$ *is a statistically binding commitment scheme that is both concurrent* NMc *and concurrent* NMd*.*

[5] Blum's basic protocol for Hamiltonicity [22] is only computational zero-knowledge with soundness error $\frac{1}{2}$. Moreover, the protocol includes three rounds of interaction. By running the basic protocol polynomial times in parallel, we get a computational WIPOK for Hamiltonicity. If the prover uses a statistically hiding commitment scheme [21] in the first round, then we get a statistical WIAOK for Hamiltonicity.

Protocol $\langle C, R\rangle$

Security Parameter: 1^n

String to be committed: $v \in \{0,1\}^n$

Commitment Phase:

Stage 1:

$C \rightarrow R$: Let $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Sgen}(1^{\mathsf{n}})$. Pick uniformly $r \in \{0,1\}^n$ and compute $c \leftarrow \mathsf{SBCom}(v; r)$. Send pk, c.

$C \Leftrightarrow R$: C uses witness (v, r) and proves using $\langle \mathcal{P}_{\mathsf{pk}}, \mathcal{V}_{\mathsf{pk}}\rangle$ (with tag pk) the statement that there exist values $v, r \in \{0,1\}^n$ such that $c = \mathsf{SBCom}(v; r)$.

R : Abort if the above proof fails.

Stage 2:

$R \rightarrow C$: Pick uniformly $v_0, r_0, v_1, r_1 \in \{0,1\}^n$ and compute $c_0 = \mathsf{SHCom}(v_0; r_0), c_1 = \mathsf{SHCom}(v_1; r_1)$. Send c_0, c_1.

$R \rightarrow C$: Pick a random bit $b \in \{0,1\}$. R uses witness (v_b, r_b) and proves using $\langle \mathrm{swi}\mathcal{P}, \mathrm{swi}\mathcal{V}\rangle$ that there exist values $v^*, r^* \in \{0,1\}^n$ such that $c_0 = \mathsf{SHCom}(v^*; r^*)$ or $c_1 = \mathsf{SHCom}(v^*; r^*)$.

C : Abort if the above proof fails.

Stage 3:

$C \rightarrow R$: Let tr_0 be the transcript of the above interaction. Compute $\sigma_0 \leftarrow \mathsf{Ssig}(\mathsf{sk}, \mathsf{tr}_0)$ and send σ_0.

R : Verify that $\mathsf{Sver}(\mathsf{pk}, \mathsf{tr}_0, \sigma_0) = 1$.

Decommitment Phase:

Stage $1'$:

$C \rightarrow R$: Pick uniformly $r' \in \{0,1\}^n$. Compute $c' = \mathsf{SBCom}(0^n; r')$ and send c'.

$R \rightarrow C$: Send v_0, v_1.

$R \Leftrightarrow C$: R uses witness r_b and proves using $\langle \mathrm{swi}\mathcal{P}, \mathrm{swi}\mathcal{V}\rangle$ (with tag pk) the statement that there exists a value $r^* \in \{0,1\}^n$ such that $c_0 = \mathsf{SHCom}(v_0; r^*)$ or $c_1 = \mathsf{SHCom}(v_1; r^*)$.

C : Abort if the above proof fails.

Stage $2'$:

$C \rightarrow R$: Send v.

$C \Leftrightarrow R$: C uses witness r and proves using $\langle \mathrm{cwi}\mathcal{P}, \mathrm{cwi}\mathcal{V}\rangle$ the OR of the following statements

1. $\exists\, r \in \{0,1\}^n$ s.t $c = \mathsf{SBCom}(v; r)$,
2. $\exists\, b^* \in \{0,1\}, r^* \in \{0,1\}^n$ s.t $c' = \mathsf{SBCom}(v_{b^*}; r^*)$.

R**:** Abort if the above proof fails.

Stage $3'$:

$C \Leftrightarrow R$: C uses witness r and proves using $\langle \mathcal{P}_{\mathsf{pk}}, \mathcal{V}_{\mathsf{pk}}\rangle$ (with tag pk) the statement that either there exists $r \in \{0,1\}^n$ such that $c = \mathsf{SBCom}(v; r)$, or there exist $b^* \in \{0,1\}, r^* \in \{0,1\}^n$ such that $c_{b^*} = \mathsf{SHCom}(v_{b^*}; r^*)$.

R : Abort if the above proof fails.

Stage $4'$:

$C \rightarrow R$: Let tr_1 be the transcript of the above interaction. Compute $\sigma_1 \leftarrow \mathsf{Ssig}(\mathsf{sk}, \mathsf{tr}_1)$ and send σ_1.

R : Verify that $\mathsf{Sver}(\mathsf{pk}, \mathsf{tr}_1, \sigma_1) = 1$.

Fig. 1. Concurrent non-malleable statistically binding commitment scheme $\langle C, R\rangle$

Proof. We need to prove the scheme satisfies the following three properties: computational hiding, statistical binding, and concurrent NMc and concurrent NMd.

Computational hiding. Intuitively, the hiding property follows from the hiding property of SBCom and perfect zero-knowledge property of $\langle \mathcal{P}_{\text{tag}}, \mathcal{V}_{\text{tag}} \rangle$. Suppose, on the contrary, there exists an adversary R^* that violates the hiding property of $\langle C, R \rangle$. Then we design an efficient adversary R' that breaks the hiding property of SBCom. R' proceeds as follows. On input a challenge com (i.e., a commitment to m_0 or m_1) from the committer of SBCom, R' internally incorporates R^* and forwards the external commitment com to R^* in $\mathsf{Stage\ 1}$. All other executions are emulated by R' by following the honest committer strategy except that R' runs the simulator for $\langle \mathcal{P}_{\text{tag}}, \mathcal{V}_{\text{tag}} \rangle$ in $\mathsf{Stage\ 1}$. Finally, R' outputs whatever R^* outputs. From the perfect zero-knowledge property of $\langle \mathcal{P}_{\text{tag}}, \mathcal{V}_{\text{tag}} \rangle$, if R^* distinguishes the commitment made using $\langle C, R \rangle$, then R' distinguishes the commitment made using SBCom.

Statistical binding. The proof of binding property is more subtle. We show that any malicious adversary C^* can not violate the binding property of $\langle C, R \rangle$. Intuitively, if C^* can open the commitment in two different ways, then due to the soundness property of $\langle \text{cwi}\mathcal{P}, \text{cwi}\mathcal{V} \rangle$ and the statistical binding property of SBCom, C^* must use a fake witness in the execution of $\langle \text{cwi}\mathcal{P}, \text{cwi}\mathcal{V} \rangle$ in Stage $2'$, i.e., it knows the witness to the statement that c' is a commitment to v_0 or v_1. Note that the only place that C^* might learn v_0 or v_1 before $\mathsf{Stage}\ 1'$ is in $\mathsf{Stage\ 2}$ of the commitment phase. Since both the commitments c_0, c_1 are statistically hiding and $\langle \text{swi}\mathcal{P}, \text{swi}\mathcal{V} \rangle$ is statistical WI, C^* learns v_0 or v_1 only with negligible probability in $\mathsf{Stage\ 2}$. Thus, C^* makes a commitment c' to the value v_0 or v_1 only with negligible probability in $\mathsf{Stage}\ 1'$. Moreover, C^* commits using SBCom. So the second statement proved in $\mathsf{Stage}\ 2'$ is a false statement (even to an unbounded machine). According to the property of $\langle \text{cwi}\mathcal{P}, \text{cwi}\mathcal{V} \rangle$, even an unbounded C^* can not successfully execute the proof with non-negligible probability in $\mathsf{Stage}\ 2'$. This reaches a contradiction.

More in details, assume for contradiction that there exists some adversary (not necessarily PPT) $\mathcal{A}$ that is able to violate the binding property of $\langle C, R \rangle$. We show how to construct an algorithm (not necessarily PPT) A' that violates the binding property of SBCom or the hiding property of SHCom or the WI property of $\langle \text{swi}\mathcal{P}, \text{swi}\mathcal{V} \rangle$. A' interacts with $\mathcal{A}$ and follows honest receiver strategy. Once the decommitment phase is finished, A' runs the extractor of $\langle \text{cwi}\mathcal{P}, \text{cwi}\mathcal{V} \rangle$ in $\mathsf{Stage}\ 2'$. According to the property of the extractor of $\langle \text{cwi}\mathcal{P}, \text{cwi}\mathcal{V} \rangle$, with overwhelming probability, A' gets a witness w. Then it must be the case that (1) $w = r$ s.t $c = \mathsf{SBCom}(v; r)$. (2) $w = r^*$ s.t $c' = \mathsf{SBCom}(v_b; r^*)$. (3) $w = r^*$ s.t $c' = \mathsf{SBCom}(v_{1-b}; r^*)$.

We now show however that case 1 happens with negligible probability only. Assume by contradiction, that it can happen with non-negligible probability. Then we can design an algorithm $\mathcal{B}$ that breaks the binding property of SBCom. $\mathcal{B}$ proceeds exactly as A'. $\mathcal{B}$ then succeeds extracting r such that $c = \mathsf{SBCom}(v; r)$. Next, we let $\mathcal{B}$ keep rewinding $\mathcal{A}$ to the beginning of the decommitment phase until case 1 happens again. $\mathcal{B}$ again extracts a witness and we denote by r^* the extracted witness. Let v^* be the opened value by $\mathcal{A}$. Now we get $c = \mathsf{SBCom}(v^*; r^*)$. According to the assumption of $\mathcal{A}$, $v \neq v^*$ with non-negligible probability. Now we find a commitment c that can be opened in two different ways. Thus, we break the binding property of SBCom.

Next we show case 2 happens with negligible probability. Suppose on the contrary, with some non-negligible probability it happens that $c' = \mathsf{SBCom}(v_b; r^*)$. We can design an algorithm $\mathcal{B}$ that breaks the WI property of $\langle \mathrm{swi}\mathcal{P}, \mathrm{swi}\mathcal{V}\rangle$. $\mathcal{B}$ then internally executes all the interactions with $\mathcal{A}$ and proceeds exactly as A' with the only exception that the proof of $\langle \mathrm{swi}\mathcal{P}, \mathrm{swi}\mathcal{V}\rangle$ in $\mathsf{Stage\ 2}$ is generated by relaying all the messages with an external prover ($\mathcal{B}$ submits opening information of c_0 and c_1 to the external prover. The external prover then proves using a witness for c_{b^*} for some $b^* \in \{0,1\}$). We emphasize here that $\mathcal{B}$ generates the proof of $\mathsf{Stage}\ 2'$ itself. Then $\mathcal{B}$ successfully simulates the interactions with $\mathcal{A}$ in the decommitment, and $\mathcal{B}$ runs the extractor of $\langle \mathrm{cwi}\mathcal{P}, \mathrm{cwi}\mathcal{V}\rangle$. By looking at the extracted witness, $\mathcal{B}$ will guess the witness used by the external prover.

Finally, we show case 3 happens with negligible probability. Suppose on the contrary, with non-negligible probability it happens that $c' = \mathsf{SBCom}(v_{1-b}; r^*)$. We then design an algorithm $\mathcal{B}$ that breaks the hiding property of SHCom. On input a challenge commitment c^* (to value $\hat{v}_0, \hat{v}_1$), $\mathcal{B}$ has to decide which value corresponds to c^*. $\mathcal{B}$ proceeds exactly as A' with the following two exceptions. The first exception lies in the handling of interaction in $\mathsf{Stage\ 2}$ of the commitment. Here $\mathcal{B}$ first picks a random bit $b \in \{0,1\}$, a random string $v_b \in \{0,1\}^n$ and a uniform random string $r_b \in \{0,1\}^n$. $\mathcal{B}$ then computes $c_b = \mathsf{SHCom}(v_b; r_b)$ and sets $c_{1-b} = c^*$. Next $\mathcal{B}$ continues the execution of $\mathsf{Stage\ 2}$ of commitment by following the honest prover strategy of $\langle \mathrm{swi}\mathcal{P}, \mathrm{swi}\mathcal{V}\rangle$ using (v_b, r_b) as witness. The second exception lies in the handling of interaction in $\mathsf{Stage}\ 1'$ of decommitment. Here $\mathcal{B}$ randomly chooses a bit $b^* \in \{0,1\}$, sends $v_b, v_{1-b} = \hat{v}_{b^*}$ to $\mathcal{A}$ and then uses witness r_b to complete the proof $\langle \mathrm{swi}\mathcal{P}, \mathrm{swi}\mathcal{V}\rangle$. Finally, if the witness r^* extracted satisfies $c' = \mathsf{SBCom}(v_{1-b}; r^*)$, $\mathcal{B}$ then outputs v_{1-b}; otherwise, $\mathcal{B}$ outputs $\hat{v}_{b'}$ for randomly chosen $b' \in \{0,1\}$. Therefore the probability that $\mathcal{B}$ breaks the hiding property of SHCom is also non-negligible.

Concurrent non-malleability. We need to show that the scheme is concurrent NMc and concurrent NMd. The proof of concurrent NMc is almost identical with that of the proof in [8]. Note NMc only concerns the commitment phase and as we discussed previously, the commitment phase of our scheme only deviates from that of [8] when R sends commitments and plays the WIAOK in $\mathsf{Stage\ 2}$. Indeed, in our scheme we use statistical versions of these tools while the protocol of [8] only needs the computational versions. The proof however goes through precisely as the one of [8]. We omit the details here and defer the proof in the full version.

Next, we show it is concurrent NMd. We show that for every PPT man-in-the-middle adversary $\mathcal{A}$ that participates in $m(n)$ left commitments and $m(n)$ right commitments, there exists an expected PPT simulator S such that for every PPT distinguisher D and every negligible function μ, for every value vector $\mathbb{V} = (v_1, \ldots, v_m)$ where $v_i \in \{0,1\}^n$ and every $z \in \{0,1\}^*$, it holds that

$$\left|\Pr[D(\mathsf{mim}^{\mathcal{A}}_{\mathsf{open}}(\mathbb{V}, z)) = 1] - \Pr[D(\mathsf{sim}^{S}_{\mathsf{open}}(\mathbb{V}, z)) = 1]\right| \le \mu(n). \qquad (1)$$

Denote by $\mathcal{A}_{\mathsf{dec}}$ the state of $\mathcal{A}$ after the commitment phase, i.e., $\mathcal{A}_{\mathsf{dec}}$ contains $\mathcal{A}$'s description along with its configuration at that time just before the decommitment phase starts.

We proceed by giving the description of the simulator S. S on input z and security parameter 1^n interacts with external honest receivers and runs the adversary $\mathcal{A}$ internally. During the commitment phases, on a high level, S internally incorporates $\mathcal{A}$ and emulates the commitment phases of all left interactions for adversary $\mathcal{A}$ by honestly committing to 0^n, while internally emulating the right interactions as honest receivers. After all the commitment phases end, S invokes the extractors for all the proofs provided by $\mathcal{A}$ in the left and right commitments to extract all the corresponding witnesses. More precisely, for each right commitment, S runs the extractor of $\langle \mathcal{P}_{\text{tag}}, \mathcal{V}_{\text{tag}} \rangle$ and we denote by $(\tilde{v}_i, \tilde{r}_i)$ the witness extracted in the i^{th} right commitment. For each left commitment, S runs the extractor of $\langle \text{swi}\mathcal{P}, \text{swi}\mathcal{V} \rangle$ to get witness $(v_{b_i,i}, r_{b_i,i})$ $(b_i \in \{0,1\})$. Next, S plays the commitment phases with external receivers. S follows the honest committer strategy and commits to $\tilde{v}_i$ in the i^{th} commitment phase.

Once all the commitment phases are finished, S receives a value vector $\mathbb{V} = (v_1, \ldots, v_m)$ and has to perform the decommitment phases internally with $\mathcal{A}_{\mathsf{dec}}$. S follows the honest receiver strategy in all right decommitments. The simulation of the i^{th} left decommitment is as follows. In $\mathsf{Stage}\ 1'$, S acts identically as an honest committer with the exception that S commits to $v_{b_i,i}$ instead of 0^n (using randomness $r^*_{b_i,i}$). In $\mathsf{Stage}\ 2'$, S follows the honest committer strategy with the exception that it uses the "fake" witness $r^*_{b_i,i}$ to open the commitment to v_i. In $\mathsf{Stage}\ 3'$, S uses the fake witness $r_{b_i,i}$ to complete the proof. S follows the honest committer strategy in $\mathsf{Stage}\ 4'$. Finally, for each i, if $\mathcal{A}_{\mathsf{dec}}$ has successfully completed the i^{th} right decommitment, then S completes the decommitment phase of the external execution with honest receivers by opening the commitment to $\tilde{v}_i$.

Running time of S. From the construction of S, we know that S performs at most $2m$ extraction procedures in the commitment phases. Note that the running time of the extractions in both $\langle \mathcal{P}_{\text{tag}}, \mathcal{V}_{\text{tag}} \rangle$ and $\langle \text{swi}\mathcal{P}, \text{swi}\mathcal{V} \rangle$ are all expected PPT. Since the extractions are executed sequentially, the running time of all extractions is also expected PPT. Furthermore, the MIM adversary $\mathcal{A}$ is a PPT algorithm and therefore invoking a copy of $\mathcal{A}$ also requires PPT. Thus, S runs in expected PPT in the commitment phases. In the decommitment phases, since S runs in a straight-line manner and no rewinding is involved, the running time of S is strict PPT. Finally, we conclude that the overall running time of S is expected PPT.

Next, we prove that the distribution of the messages opened by $\mathcal{A}$ when interacting with honest committers and honest receivers is indistinguishable from the distribution of the messages opened by $\mathcal{A}$ when interacting with S.

Indistinguishability of the simulation. We first consider the case when there is only one left commitment and $m(n)$ right commitments. Towards of showing

Equation (1) (note $\mathbb{V}$ contains only a value v), we define a sequence of hybrid experiments $\{\mathsf{HYB}_i(v,z)\}_{1\leq i\leq 7}$ that receive v and z as auxiliary inputs. The output of each experiment is the output of a PPT distinguisher D on input a value v and a vector of values $\tilde{V}$ whose i^{th} element is defined as follows. If the i^{th} right decommitment completes successfully and its transcript is different from the left interaction, then $\tilde{v}_i$ is the value opened in the i^{th} right interaction. Otherwise, $\tilde{v}_i$ is set to $\perp$. Let $p_i = \Pr[\mathsf{HYB}_i(v,z) = 1]$.

$\mathsf{HYB}_1(v,z)$ proceeds exactly as S except that in $\mathsf{Stage\ 1}$ of the left commitment phase, it runs the simulator of $\langle \mathcal{P}_{\text{tag}}, \mathcal{V}_{\text{tag}} \rangle$. Since the simulation is perfect we conclude that $p_1 = \Pr[D(\mathsf{sim}^S_{\mathsf{open}}(v,z)) = 1]$.

$\mathsf{HYB}_2(v,z)$ proceeds exactly as HYB_1 except that in the left commitment phase, instead of feeding $\mathcal{A}$ a commitment to 0^n in $\mathsf{Stage\ 1}$, HYB_2 feeds $\mathcal{A}$ a commitment to v using SBCom. Since both HYB_1 and HYB_2 are efficiently computable, that $|p_1 - p_2|$ is negligible follows directly from the computational hiding property of SBCom.

$\mathsf{HYB}_3(v,z)$ proceeds exactly as HYB_2 except that it runs the simulator of $\langle \mathcal{P}_{\text{tag}}, \mathcal{V}_{\text{tag}} \rangle$ in $\mathsf{Stage}\ 3'$ of the left decommitment. It follows from the perfect zero-knowledge property of $\langle \mathcal{P}_{\text{tag}}, \mathcal{V}_{\text{tag}} \rangle$ that $p_3 = p_2$.

$\mathsf{HYB}_4(v,z)$ differs from HYB_3 in that it uses the real witness (i.e., decommitment of c) to complete the proof in $\mathsf{Stage}\ 2'$ of the left decommitment. It follows from the computational WI property of $\langle \text{cwi}\mathcal{P}, \text{cwi}\mathcal{V} \rangle$ that $|p_4 - p_3|$ is negligible.

$\mathsf{HYB}_5(v,z)$ differs from HYB_4 in that it commits to 0^n in $\mathsf{Stage}\ 1'$ of the left decommitment. It follows from the computational hiding property of SBCom that $|p_5 - p_4|$ is negligible.

$\mathsf{HYB}_6(v,z)$ proceeds exactly as HYB_5 except that it uses the real witness (i.e., decommitment of c) to complete the proof in $\mathsf{Stage}\ 3'$ of the left decommitment. It follows from the perfect zero-knowledge property of $\langle \mathcal{P}_{\text{tag}}, \mathcal{V}_{\text{tag}} \rangle$ that $p_6 = p_5$.

$\mathsf{HYB}_7(v,z)$ proceeds exactly as HYB_6 except that it does not need to run the extractor of $\langle \text{swi}\mathcal{P}, \text{swi}\mathcal{V} \rangle$ in $\mathsf{Stage\ 2}$ of left commitment phase. Since the extraction fails with negligible probability we have that $|p_7 - p_6|$ is negligible.

Note that HYB_7 differs from the real game in that it runs the simulator-extractor of $\langle \mathcal{P}_{\text{tag}}, \mathcal{V}_{\text{tag}} \rangle$ in $\mathsf{Stage\ 1}$ of the commitment phase. Following the description of HYB_7 we know that it opens to external receivers the values it extracts from $\mathcal{A}$ at the end of the commitment phase in the simulated game. Moreover, in the real game the adversary $\mathcal{A}$ can not open its commitments in a different way (cf. Claim 1). Thus, $\mathcal{A}$ opens to external receivers the values it commits to in the commitment phase. It follows from the simulation-extractability property of $\langle \mathcal{P}_{\text{tag}}, \mathcal{V}_{\text{tag}} \rangle$ that the simulation is perfect and the extraction fails with negligible probability in each right commitment where the tag is different from that in the left decommitment (when tags are the same, the security of signature scheme is violated). Therefore, except with negligible probability, we have that HYB_7 opens to external receivers the same values opened by $\mathcal{A}$ in the real game, i.e., we have that $\big|p_7 - \Pr[D(\mathsf{mim}^{\mathcal{A}}_{\mathsf{open}}(v,z)) = 1]\big|$ is negligible.

Finally we conclude Equation (1). This completes the proof of Theorem 2.

Claim 1. *In the real game $\mathcal{A}$ can not open in a different way.*

Proof. Assume that, with some probability p, there exists $i \in [m]$ such that the value $\tilde{v}'$ opened in the i^{th} right decommitment is different from the value $\tilde{v}$ committed in the i^{th} right commitment.[6] We denote by c_0, c_1 the commitments made by R in $\mathsf{Stage\ 2}$ of the i^{th} right commitment. Denote by b the bit such that R uses the decommitment information corresponding to the commitment c_b to complete the proof of $\langle \text{swi}\mathcal{P}, \text{swi}\mathcal{V} \rangle$ in $\mathsf{Stage\ 2}$ and $\mathsf{Stage}\ 1'$ of i^{th} right interaction. If $\mathcal{A}$ successfully completes the i^{th} right decommitment, then we consider the following experiment B. B on inputs i, v and z interacts with $\mathcal{A}$ and works as follows. We let B commit to v in the left commitment phase. Furthermore, B follows the honest committer strategy in the left interaction and the honest receiver strategy in all right interactions, except in $\mathsf{Stage}\ 1'$ of i^{th} decommitment B sends v_b and a randomly chosen v^*_{1-b} (With overwhelming probability v^*_{1-b} will be different from the v_{1-b} chosen in the commitment phase. This is important for the proof of Case 3 below.). Once the i^{th} right decommitment is over, B runs the extractor of $\langle \mathcal{P}_{\text{tag}}, \mathcal{V}_{\text{tag}} \rangle$ in $\mathsf{Stage}\ 3'$. Note that the views of $\mathcal{A}$ in a real execution and an execution of B are identical. So the probability that $\mathcal{A}$ opens in a different way in the execution of B is also p. According to the property of the extractor of $\langle \mathcal{P}_{\text{tag}}, \mathcal{V}_{\text{tag}} \rangle$, with probability $p' = p - \epsilon(n)$ where ϵ is a negligible function, B gets a witness $\tilde{w}$, and one of the following three cases must happen:

1. $\tilde{w} = \tilde{r}$ s.t $\tilde{c} = \mathsf{SBCom}(\tilde{v}'; \tilde{r})$.($\tilde{c}$ is the commitment generated by $\mathcal{A}$ in the i^{th} right commitment.)
2. $\tilde{w} = r^*$ s.t $c_b = \mathsf{SHCom}(v_b; r^*)$. $(v_b, v^*_{1-b}$ are the values opened by B in Stage $1'$ of i^{th} right decommitment.)
3. $\tilde{w} = r^*$ s.t $c_{1-b} = \mathsf{SHCom}(v^*_{1-b}; r^*)$.

Let $p' = p_1 + p_2 + p_3$, where p_i is the probability that Case i happens (for $i = 1, 2, 3$). By the statistical binding property of SBCom, we have that $\tilde{v}'$ must correspond to $\tilde{v}$ and thus in this case $\mathcal{A}$ does not open in a different way, therefore p_1 must be negligible. Recall in the i^{th} right commitment, B already generates commitments c_0 and c_1 to two different values v_0 and v_1 respectively and B uses knowledge of a decommitment of v_b for a random bit b in both $\langle \text{swi}\mathcal{P}, \text{swi}\mathcal{V} \rangle$ of $\mathsf{Stage\ 2}$ and $\langle \text{swi}\mathcal{P}, \text{swi}\mathcal{V} \rangle$ of $\mathsf{Stage}\ 1'$. By the statistical hiding property of SHCom and statistical WI property of $\langle \text{swi}\mathcal{P}, \text{swi}\mathcal{V} \rangle$, we have that p_2 is essentially identical to p_3, and therefore both p_2 and p_3 roughly correspond to $p'/2 - \epsilon(n)$ where ϵ is a negligible function.

We can now conclude the proof showing that p_3 must be negligible, and thus summing up p is negligible as well, therefore $\mathcal{A}$ can not open in a different way with non-negligible probability. Indeed, notice that when Case 3 happens, we have that $\mathcal{A}$ committed to some value v^*_{1-b} in c_{1-b} without having never used any opening of c_{1-b} in the two executions of $\langle \text{swi}\mathcal{P}, \text{swi}\mathcal{V} \rangle$. Now in addition to the opening in the commitment phase (note B generates the commitment itself), we

[6] The committed value is the one uniquely specified by the statistically binding commitment scheme SBCom.

get two openings for SHCom. Therefore being SHCom computationally binding, this happens with negligible probability.

Extending to many-many concurrent NMd*.* Next, we present the proof sketch for the many-many concurrent case. We show that the two ensembles $\{\mathsf{mim}^{\mathcal{A}}_{\mathsf{open}}(\mathbb{V}, z)\}$ and $\{\mathsf{sim}^{S}_{\mathsf{open}}(\mathbb{V}, z)\}$ are computationally indistinguishable. Suppose, for contradiction, this is not the case. That is, there exists a PPT distinguisher D and a polynomial $p(n)$ such that for infinitely many $n \in \mathbb{N}$, there exists a value vector $\mathbb{V} = (v_1, \ldots, v_m), z \in \{0,1\}^*$ such that D distinguishes $\mathsf{mim}^{\mathcal{A}}_{\mathsf{open}}(\mathbb{V}, z)$ and $\mathsf{sim}^{S}_{\mathsf{open}}(\mathbb{V}, z)$ with probability at least $\frac{1}{p(n)}$. For a generic n for which this happens. We design a sequence of hybrid experiments $\{\mathsf{HYB}_i(\mathbb{V}, z)\}_{0\leq i\leq m}$ where $\mathsf{HYB}_i(\mathbb{V}, z)$ is defined as follows. HYB_i proceeds as S except that it emulates the i^{th} left commitment phase by committing to v_i, if $j \leq i$, and 0^n otherwise. Moreover, it emulates the i^{th} left decommitment by using a legal witness to v_i, if $j \leq i$, and a false witness otherwise. It directly follows that $\mathsf{sim}^{\mathsf{HYB}_m}_{\mathsf{open}}(\mathbb{V}, z) = \mathsf{mim}^{\mathcal{A}}_{\mathsf{open}}(\mathbb{V}, z)$ and $\mathsf{sim}^{\mathsf{HYB}_0}_{\mathsf{open}}(\mathbb{V}, z) = \mathsf{sim}^{S}_{\mathsf{open}}(\mathbb{V}, z)$. By a standard hybrid argument there exists an $i \in [m]$ such that

$$\left|\Pr[D(\mathsf{sim}^{\mathsf{HYB}_{i-1}}_{\mathsf{open}}(\mathbb{V}, z)) = 1] - \Pr[D(\mathsf{sim}^{\mathsf{HYB}_i}_{\mathsf{open}}(\mathbb{V}, z)) = 1]\right| > \frac{1}{m \cdot p(n)} \quad (2)$$

Note that the only difference between experiment $\mathsf{HYB}_{i-1}(\mathbb{V}, z)$ and $\mathsf{HYB}_i(\mathbb{V}, z)$ is that in the former $\mathcal{A}$ receives a commitment to v_i and its corresponding decommitment generated using a valid witness in the i^{th} interaction, whereas in the latter it receives a commitment to 0^n and its corresponding decommitment generated using a false witness.

Then we design an efficient MIM adversary $\tilde{A}$ that breaks the one-many concurrent non-malleability of $\langle C, R\rangle$. $\tilde{A}$ on auxiliary inputs $z' = (n, i, \mathbb{V}, z)$ proceeds as follows. $\tilde{A}$ internally incorporates $A(z)$ and emulates the left and right interactions for $\mathcal{A}$. $\tilde{A}$ relays the messages in all right interactions between $\mathcal{A}$ and external receivers. In the i^{th} left interaction, $\tilde{A}$ relays either messages between an external committer and $\mathcal{A}$, or messages between the simulator of the one-many concurrent case and $\mathcal{A}$. For $j \in [m]$ and $j \neq i$, $\tilde{A}$ internally emulates the i^{th} left commitment phase for A by committing to v_i, if $j \leq i$, and 0^n otherwise. Moreover, $\tilde{A}$ emulates the i^{th} left decommitment for $\mathcal{A}$ by using a valid witness, if $j \leq i$, and a false witness otherwise. By construction, it follows that $\mathsf{sim}^{S}_{\mathsf{open}}(\mathbb{V}, z) = \mathsf{sim}^{\mathsf{HYB}_{i-1}}_{\mathsf{open}}(z')$ and $\mathsf{mim}^{\mathcal{A}}_{\mathsf{open}}(\mathbb{V}, z) = \mathsf{sim}^{\mathsf{HYB}_i}_{\mathsf{open}}(z')$.

Therefore, $\tilde{A}$ breaks the one-many concurrent non-malleability of $\langle C, R\rangle$.

4 Concluding Remarks

Our result on top of previous work shows that there exist constant-round commitment schemes that are secure also against very powerful adversaries, as long as there is a barrier in time between commitment and decommitment phase. An interesting open question concerns the possibility of achieving commitment schemes that remain secure even without such a barrier. The question is interesting even without requiring a constant round complexity.[7]

Acknowledgments. We thank the anonymous reviewers for their helpful suggestions. The work of the first and third authors are supported in part by the National Natural Science Foundation of China under grant No.60773086, No.60972034 and No.60970110. The work of the second author has been supported in part by the European Commission through the EU ICT program under Contract ICT-2007-216646 ECRYPT II, and in part by a research grant received from Provincia di Salerno.

References

1. Dolev, D., Dwork, C., Naor, M.: Nonmalleable cryptography. SIAM J. Comput. 30(2), 391–437 (2000)
2. Crescenzo, G.D., Ishai, Y., Ostrovsky, R.: Non-interactive and non-malleable commitment. In: STOC 1998: Proceedings of the thirtieth annual ACM symposium on Theory of computing, pp. 141–150. ACM, New York (1998)
3. Fischlin, M., Fischlin, R.: Efficient non-malleable commitment schemes. In: Bellare, M. (ed.) CRYPTO 2000. LNCS, vol. 1880, pp. 413–431. Springer, Heidelberg (2000)
4. Pass, R., Rosen, A.: New and improved constructions of non-malleable cryptographic protocols. In: Gabow, H.N., Fagin, R. (eds.) STOC, pp. 533–542. ACM, New York (2005)
5. Pass, R., Rosen, A.: New and improved constructions of nonmalleable cryptographic protocols. SIAM J. Comput. 38(2), 702–752 (2008)
6. Pass, R., Rosen, A.: Concurrent non-malleable commitments. In: FOCS, pp. 563–572. IEEE Computer Society, Los Alamitos (2005)
7. Pass, R., Rosen, A.: Concurrent nonmalleable commitments. SIAM J. Comput. 37(6), 1891–1925 (2008)
8. Ostrovsky, R., Persiano, G., Visconti, I.: Simulation-based concurrent non-malleable commitments and decommitments. In: Reingold, O. (ed.) Theory of Cryptography. LNCS, vol. 5444, pp. 91–108. Springer, Heidelberg (2009)
9. Fischlin, M., Fischlin, R.: Efficient non-malleable commitment schemes. J. Cryptology 22(4), 530–571 (2009)
10. Di Crescenzo, G., Katz, J., Ostrovsky, R., Smith, A.: Efficient and non-interactive non-malleable commitment. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, pp. 40–59. Springer, Heidelberg (2001)
11. Haitner, I., Reingold, O.: Statistically-hiding commitment from any one-way function. In: Johnson, D.S., Feige, U. (eds.) STOC, pp. 1–10. ACM, New York (2007)
12. Lin, H., Pass, R., Venkitasubramaniam, M.: Concurrent non-malleable commitments from any one-way function. In: Canetti, R. (ed.) TCC 2008. LNCS, vol. 4948, pp. 571–588. Springer, Heidelberg (2008)
13. Zhang, Z., Cao, Z., Ding, N., Ma, R.: Non-malleable statistically hiding commitment from any one-way function. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 303–318. Springer, Heidelberg (2009)

[7] Indeed, in this case one could also use the concurrent non-malleable zero knowledge argument of [25], and the one of [26] when some efficiency is also required.

14. Ostrovsky, R., Persiano, G., Visconti, I.: Concurrent non-malleable witness indistinguishability and its applications. Electronic Colloquium on Computational Complexity (ECCC) 13(95) (2006)
15. Ostrovsky, R., Persiano, G., Visconti, I.: Constant-round concurrent non-malleable zero knowledge in the bare public-key model. In: Aceto, L., Damgård, I., Goldberg, L.A., Halldórsson, M.M., Ingólfsdóttir, A., Walukiewicz, I. (eds.) ICALP 2008, Part II. LNCS, vol. 5126, pp. 548–559. Springer, Heidelberg (2008)
16. Ostrovsky, R., Persiano, G., Visconti, I.: Concurrent non-malleable commitments and decommitments. Full version, unpublished manuscript (2009)
17. Feige, U.: Alternative Models for Zero Knowledge Interactive Proofs. PhD thesis, The Weizmann Institute of Science, Rehovot, Israel (1990)
18. Richardson, R., Kilian, J.: On the concurrent composition of zero-knowledge proofs. In: Stern, J. (ed.) EUROCRYPT 1999. LNCS, vol. 1592, pp. 415–431. Springer, Heidelberg (1999)
19. Goldreich, O.: Foundations of Cryptography Volume II Basic Applications. Cambridge University Press, Cambridge (2004)
20. Naor, M.: Bit commitment using pseudorandomness. J. Cryptology 4(2), 151–158 (1991)
21. Goldreich, O., Kahan, A.: How to construct constant-round zero-knowledge proof systems for NP. J. Cryptology 9(3), 167–190 (1996)
22. 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 (1986)
23. Feige, U., Shamir, A.: Witness indistinguishable and witness hiding protocols. In: STOC, pp. 416–426. ACM, New York (1990)
24. Feige, U., Lapidot, D., Shamir, A.: Multiple noninteractive zero knowledge proofs under general assumptions. SIAM J. Comput. 29(1), 1–28 (1999)
25. Barak, B., Prabhakaran, M., Sahai, A.: Concurrent non-malleable zero knowledge. In: FOCS, pp. 345–354. IEEE Computer Society, Los Alamitos (2006)
26. 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)