Constant-Round Non-Malleable Commitments from Sub-Exponential One-Way Functions

Rafael Pass* and Hoeteck Wee**

[1] Cornell University
[2] Queens College, CUNY

Abstract. We present a constant-round non-malleable commitment scheme based on the existence of sub-exponential one-way functions and using a black-box proof of security. As far as we know, this is the first construction of a constant-round non-malleable protocol based on only one-wayness, or to admit a black-box proof of security under any standard-type assumption.

Keywords. commitment schemes, non-malleability, complexity leveraging

1 Introduction

We consider the execution of two-party protocols in the presence of an adversary that has full control of the communication channel between the parties. The adversary may omit, insert, or modify messages at will. The honest parties are not necessarily aware of the existence of the adversary, and not use any kind of trusted set-up (such as a common reference string). The above kind of attack is often referred to as a *man-in-the-middle* attack. Protocols that are secure against such attacks are said to be *non-malleable*, and were first studied in the seminal work of Dolev, Dwork and Naor [6]. Due to the hostile environment in which they operate, the design and analysis of non-malleable protocols is a notoriously difficult task. The task becomes even more challenging if the honest parties are not allowed to use any kind of trusted set-up. Indeed, only a handful of such protocols have been constructed so far.

In their paper, Dolev et al. presented non-malleable protocols for the tasks of commitment and zero-knowledge. The protocols rely on the existence of one-way functions, and require $O(\log n)$ rounds of interaction, where n is a security parameter. More recently, Barak [2] presented the first constant-round non-malleable protocols for commitment and zero-knowledge whose security relies on the existence of trapdoor permutations and collision-resistant hash functions with sub-exponential hardness. The result was subsequently improved by Pass and Rosen [24], who obtained constant-round

* Supported in part by a Microsoft New Faculty Fellowship, NSF CAREER Award CCF-0746990, AFOSR Award FA9550-08-1-0197, and BSF Grant 2006317

** Most of this work was done while a post-doc at Columbia University; supported in part by NSF Grants CNS-0716245, SBE-0245014, NSF CAREER Award CNS-0953626 and PSC-CUNY Award # 6014939 40

protocols assuming only collision-resistant hash functions with standard hardness. There has been a series of follow-up work on non-malleable commitments [25, 16, 21, 20, 15], but none of which reduces the assumptions in [24] for constant-round non-malleable commitments. This raises the following natural question:

> *What are the minimal assumptions under which we can construct constant-round non-malleable commitment schemes? Specifically, is one-wayness alone sufficient to construct constant-round non-malleable commitment schemes?*

1.1 Our results

In this work, we address the above question. Our main result is that one-wayness alone—with sub-exponential hardness—suffices for constructing constant-round non-malleable commitments.

> **Main Theorem (informal):** Suppose there exists one-way functions secure against sub-exponential size circuits. Then, there exists a constant-round non-malleable commitment scheme.

We note that while all known candidates of one-way functions are conceivably also secure against sub-exponential size circuits, there are several natural candidates which do not appear to yield collision-resistant hash functions. Our result should be compared with the very recent work of Lin and Pass [15], which gave a $O(1)^{\log^* n}$-round non-malleable commitment schemes under the minimal assumption of one-way functions with standard (super-polynomial) hardness. Comparing the two, our result may be viewed as offering a new trade-off between round complexity and quantitative hardness assumptions. As with [15, 25, 16, 21], our commitment scheme achieves a very strong notion of non-malleability—that of concurrent non-malleability—which guarantees independence of the committed values even when multiple executions of the commitment schemes are executed at the same time. Before providing further details about our construction, we provide some additional context and applications.

On black-box proofs of security. While the original [6] construction only relies on "elementary" techniques and has a black-box proof of security, basically all constant-round non-malleable commitment schemes rely on *non-black-box* simulation techniques [1] and inherit the sophisticated machinery (e.g. the PCP theorem) associated with them, along with the need for qualitatively stronger assumptions (that of collision-resistant hash functions). As such, the problem of reducing the cryptographic assumptions for constant-round non-malleable commitment schemes appears to be intimately related to the question of whether non-black-box techniques are necessary for constructing

constant-round non-malleable protocols, without resorting to non-standard assumptions (c.f. [21][3]).

Understanding the power and limitations of black-box techniques has been an important goal in the foundations of cryptography, starting from [13]. For the usage of a primitive in cryptographic constructions, a recent line of work has narrowed the gap between what can be achieved using black-box and non-black-box techniques. On the other hand, for usage of the adversary's code in the proof of security, we do know for a fact that non-black-box techniques are inherently more powerful, as evidenced by the works on constant-round public-coin zero-knowledge protocols [1, 9]. A natural question is whether such a separation extends beyond the realm of zero knowledge. Given the state-of-the-art for non-malleability, it is tempting to conjecture that such a separation extends also to constant-round non-malleable commitment schemes. Our construction refutes such a conjecture since it admits a black-box proof of security (which is to be expected since we do not require collision-resistant hash functions).

On constant-round secure multi-party computation. The early work of Goldreich, Micali and Wigderson [10] showed that we may realize secure multi-party computation in the presence of a dishonest majority assuming the existence of enhanced trapdoor permutations, where the round complexity of the protocol grows linearly with the number of parties.[4] Subsequent improvements by Katz, Ostrovsky and Smith [14] (relying on [2]) and Pass [22] culminated a constant-round protocol, assuming in addition the existence of collision-resistant hash functions. As with previous constant-round non-malleable protocols, both of these constructions exploit non-black-box techniques in the proof of security.

More recently, Lin, Pass and Venkitasubramaniam [17] showed that constant-round protocols for secure multi-party computation may be based on enhanced trapdoor permutations and any "natural" constant-round non-malleable commitment scheme. Combining their construction with our commitment scheme yields the following corollary:

> **Corollary (informal):** Suppose there exists one-way functions secure against sub-exponential size circuits and standard enhanced trapdoor permutations. Then, there exists a constant-round protocol that secure computes any multi-party functionality against a malicious adversary corrupting any number of parties.

[3] Pandey, Pass and Vaikuntanathan constructed non-interactive non-malleable commitment schemes assuming the existence of, so called, adaptive one-way permutations – namely permutations which remain one-way even when the adversary has access to an inversion oracle. Note that this assumption has a strong non-malleability flavor and as such provide limited insight into realizing non-malleability "from scratch".

[4] In the protocol, each player takes turns to sequentially commit to its input (along with a "proof of knowledge"); any non-trivial improvement in round complexity will require interweaving these input commitments, which could potentially allow an adversary to violate input independence via a man-in-the-middle attack.

As with our non-malleable commitment scheme, the ensuing protocol for secure multi-party computation admits a black-box proof of security.

Perspective. Prior to this work, the trade-offs between computational assumptions and round complexity for non-malleable commitments and secure computation looked fairly similar to those for (computational) zero-knowledge *proofs* for NP (c.f. [11, 8]): we have constant-round protocols based on collision-resistant hash functions whereas those based on the minimal assumption of one-way functions require at least a super-constant number of rounds (for secure computation, we also require oblivious transfer). An interesting open problem is whether we can also base constant-round zero-knowledge proofs for NP on one-way functions with sub-exponential hardness.

1.2 Our techniques

Our construction of the non-malleable commitment scheme proceeds in two steps:

Step 1: Short identities from sub-exponential hardness. First, we construct a constant-round concurrent non-malleable commitment scheme for identities of length $\log\log\log n + O(1)$ (again, n here refers to the security parameter). Our main technical contribution lies in this step. The starting point of this construction is "two-slot message length" technique from [22] underlying the recent constructions of constant-round non-malleable protocols in [24, 25].[5] The basic (and very much simplified) idea is to let the receiver *sequentially* send two challenges—one "long" and one "short"; the length of the challenges are determined by the identity of the sender. Intuitively, the protocol is designed to have the property that the response to a shorter challenge does not help an adversary to provide a response to a longer challenge. If done appropriately, this guarantees that the man-in-the-middle adversary needs to act independently. Our key conceptual insight is to rely on the complexity leveraging technique from [4] to construct these challenges.[6] More precisely, for one-way functions with sub-exponential hardness, an oracle for inverting challenges of length $n^{o(1)}$ (the "short" challenge) does not help invert invert random challenges of length n (the "long" challenge), since we may simulate such an oracle by brute force in time $2^{n^{o(1)}}$.

Step 2: Non-malleability amplification. Next, we transform the initial construction into a constant-round concurrent non-malleable commitment scheme for identities of length $\mathrm{poly}(n)$. This relies on *non-malleability amplification* techniques of Lin and Pass [15]. This is a transformation of so-called "natural" commitment schemes that are non-malleable for identities of length t into ones for identities of length $\Omega(2^t)$ while incurring only a constant multiplicative blow-up in round complexity.

[5] Our protocol, like that in [24, 25], also has a "commit and prove" structure.

[6] This appears to be the first work to exploit complexity leveraging with a super-constant levels of challenges.

Primitive	Hardness	Rounds	Black-box?	Reference
one-way functions	standard	$O(\log n)$	yes	[6]
one-way functions	standard	$O(1)^{\log^* n}$	yes	[15]
one-way functions	sub-exp	$O(1)$	yes	this work
collision-resistance	standard	$O(1)$	no	[24]
collision-resistance, TDP	sub-exp	$O(1)$	no	[2, 5]
adaptive OWP	standard	1	yes	[21]

Fig. 1. Summary of non-malleable statistically binding commitments.

We modify our initial construction to satisfy naturality by using the "multiple slots" approach from [22] (introduced in the context of handling longer identities) to boost the number of rewinding slots. Applying the [15] transformation to the modified construction a constant number of times yields the final construction.[7]

Our final protocol has a conceptually simple and "elementary" proof of security. This is a welcome respite from the technical subtleties and complexity and/or heavy technical machinery that arise in much of the previous literature on non-malleability. We also point out that complexity leveraging has been previously used in [18, 23] - as in Step 1 - to achieve similar but weaker notions of "independence". The constructions therein use a single challenge slot and achieve only "uni-directional" independence as they require that the challenge in the left interaction be shorter than that on the right. This appears a prori to be an inherent limitation of the complexity leveraging approach[8], because with two challenge slots, the long challenge in the left interaction may be longer than both challenges on the right, so that solving the challenge on the left via brute force violates soundness for both challenges on the right. We show precisely how to overcome this difficulty in our construction and in the analysis.

Organization. We present our construction for short identities in Section 3. For simplicity, we first present the construction assuming one-way *permutations* secure against circuits of size 2^{n^δ} for some constant $\delta < 1$. In Section 4, we apply non-malleability amplification to handle identities of length $\mathrm{poly}(n)$. In Section 5, we modify our constructions to work with general one-way *functions* (as opposed to permutations).

[7] In [15], the transformation is applied to the [6] protocol for constant-length identities (for which the protocol is constant-round) a total of $O(\log^* n)$ times.

[8] and indeed, [23] —the pre-cursor to [22]— handles the "opposite direction" via non-black-box techniques

2 Concurrent non-malleable commitments

We recall the definition of concurrent non-malleability from [16], which builds upon those in [6, 25]. Let $(\mathcal{C}, \mathcal{R})$ be a commitment scheme with identities, and 1^n be the security parameter.

The man-in-the-middle execution. In the man-in-the-middle execution, the adversary $\mathcal{A}$ is participating m left interactions and m right interactions. In the left interactions, $\mathcal{A}$ interacts with $\mathcal{C}$ receiving a commitment to m values $v_1, \ldots, v_m$, using identities $\mathsf{id}_1, \ldots, \mathsf{id}_m$ of its choice. In the right interactions, $\mathcal{A}$ interacts with $\mathcal{R}$ attempting to commit to a sequence of m related values $\tilde{v}_1, \ldots, \tilde{v}_m$, again using identities $\tilde{\mathsf{id}}_1, \ldots, \tilde{\mathsf{id}}_m$ of its choice. $\mathcal{A}$ also receives an auxiliary z. If any of the right commitments (as determined by the transcript) are invalid or undefined, its value is set to $\perp$. For any i such that $\tilde{\mathsf{id}}_i = \tilde{\mathsf{id}}_j$ for some j, the value $\tilde{v}_i$ is also set to $\perp$ (that is, any commitment where adversary uses the same identity as that in one of the left interactions is considered invalid). We write $\mathsf{mim}^{\mathcal{A}}(v_1, \ldots, v_m, z)$ to denote a random variable comprising the view of $\mathcal{A}$ along with the m-tuple of values $(\tilde{v}_1, \ldots, \tilde{v}_m)$.

The simulated execution. In the simulated execution, a simulator $\mathcal{S}$ receives the auxiliary input z and interacts directly with $\mathcal{R}$ in m right interactions. We write $\mathsf{sta}^{\mathcal{S}}(1^n, z)$ to denote a random variable comprising the output of $\mathcal{S}$ along with the m-tuple of values $(\tilde{v}_1, \ldots, \tilde{v}_m)$ that the simulator has committed to as determined by the transcript.

Definition 1 ([16, 6, 25]). *A commitment scheme $(\mathcal{C}, \mathcal{R})$ is* concurrent non-malleable *if for every* PPT *$\mathcal{A}$ and every polynomial $m = m(n)$, there exists a* PPT *$\mathcal{S}$ such that*

$$\left\{\mathsf{mim}^{\mathcal{A}}(v_1, \ldots, v_m, z)\right\}_{v_1,\ldots,v_m \in \{0,1\}^n, z \in \{0,1\}^*} \text{ and}$$
$$\left\{\mathsf{sta}^{\mathcal{S}}(1^n, z)\right\}_{v_1,\ldots,v_m \in \{0,1\}^n, z \in \{0,1\}^*, \mathsf{id} \in \{0,1\}^m}$$

are computationally indistinguishable.

We will also consider a restricted notion of concurrent non-malleability where in the left and right interactions, the adversary $\mathcal{A}$ may only use identities of length at most d. In addition, we will refer to relaxed notions of concurrent non-malleability: one-many and one-one non-malleability. In the former, the adversary participates in one interaction on the left and m interactions on the right, and in the latter, the adversary participates in one interaction on the left and one interaction on the right. As shown in [16], any commitment scheme that is one-many non-malleable is also concurrent non-malleable.

3 Short identities from sub-exponential hardness

3.1 Overview of our construction

We construct a family of $d = \Theta(\log \log n)$ protocols (corresponding to d different identities) as follows. Let $n^{\omega(1)} = T_0 \ll T_1 \ll \cdots \ll T_{d-1}$ be a hierarchy of running times. The ith protocol in the family, $i = 0, 1, \ldots, d-1$ is as follows: to commit to a string v (with identity i),

- Commit to v using a statistically binding commitment Com that is hiding against adversaries of size T_d.
- Slot 1: prove knowledge of v using a zero-knowledge argument of knowledge that is computationally sound against adversaries of size T_i and can be (straight-line) simulated in time $o(T_{i+1})$.
- Slot 2: prove knowledge of v using a zero-knowledge argument of knowledge that is computationally sound against adversaries of size T_{d-1-i} and can be (straight-line) simulated in time $o(T_{d-i})$.

The intuition is that for one of the two slots, the man-in-middle adversary must prove knowledge of the string $\tilde{v}$ committed to in the right interaction without getting much help from the left interaction. Roughly speaking, we will then "extract" from that slot on the right (by rewinding) while simulating on the left (c.f. [6]). To guarantee that the extraction succeeds we need to ensure that the simulation does not violate the soundness of the right interaction; this property is often called *simulation soundness* [26].

For concreteness, consider a synchronizing adversary participating in the ith protocol on the left and the jth protocol on the right. If $i < j$, we may extract the string committed to on the right as follows: run the knowledge extractor for first slot on the right while simulating the first slot on the left. This works because we may simulate on the left in time $o(T_{i+1}) \ll T_j$ without rewinding, without knowing the string committed to on the left, and without violating soundness for the first slot on the right. Similarly, if $i > j$, we can extract the string committed to on the right by running the knowledge extractor for the second slot in the right while simulating the second slot on the left in time $o(T_{d-i+1}) \ll T_{d-j}$. In either case, we may achieve strict polynomial-time simulation by running the man-in-the-middle adversary and committing to 0^n on the left (cf. [25, 16]).

We point out several technical difficulties that arise in turning the above intuition into a proof (indeed, the actual analysis is quite different from that suggested by the above line of reasoning).

- *Simulation may violate soundness.* Consider the case $i > j$, where we need to extract from the second slot on the right. To reach the second slot, we will still need to simulate the first slot on the left, and simply running the straight-line simulator may violate soundness for the second slot on the right. Roughly speaking, we get around this specific problem by using non-uniformity.

– *Which slot should we extract from?* In the analysis, we need to know which slot to simulate and which one to extract from. This is problematic because we allow the identity on the right to be adaptively chosen, and because we do not know the message schedule in advance. To make things worse, the messages may be adaptively and dynamically scheduled.

The key insight in the analysis is to decouple the issue of extraction and the issue of simulation-soundness (this is similar to the approach in [21]). Specifically, we will always simulate both slots on the left and extract from both slots on the right, no matter what the scheduling is. We will then carefully argue that extraction succeeds in at least one slot even though we may be violating soundness while simulating on the left. This is where we reason about the scheduling of messages. For technical reasons, we will also require that we can break Com via brute force in less time than it takes to break the zero-knowledge property.

3.2 Handling identities of length $\log\log\log n + O(1)$

Let π denote a one-way permutation secure against circuits of size 2^{n^δ} (where $\delta < 1$) and let Com be a statistically binding commitment scheme. In addition, let $\langle \mathcal{P}_{\text{WIPOK}}, \mathcal{V}_{\text{WIPOK}} \rangle$ denote the 3-round public-coin witness-indistinguishable proof of knowledge based on the Feige-Shamir protocol from [7], which satisfies the following properties:

– The first two messages depend only on the length of the instance and the security parameter and can be computed efficiently without knowing the instance or the witness.
– The third message can be computed efficiently given the instance, the witness, and the randomness used to generate the first message.
– The protocol is *special-sound*—namely, given any two accepting proofs of x, $(\alpha, \beta, \gamma), (\alpha, \beta', \gamma')$ such that $\beta \neq \beta'$. a witness to x can be efficiently recovered.

We consider a hierarchy of security levels for the one-wayness of the d permutations $\pi_0, \ldots, \pi_{d-1}$ and the hiding properties of Com and $\langle \mathcal{P}_{\text{WIPOK}}, \mathcal{V}_{\text{WIPOK}} \rangle$, that is given by:

$$\pi_0 \ll \pi_1 \ll \cdots \ll \pi_{d-1} \ll \mathsf{Com} \ll \langle \mathcal{P}_{\text{WIPOK}}, \mathcal{V}_{\text{WIPOK}} \rangle$$

– For each $i = 0, 1, \ldots, d-1$: π_i is T_i-one-way but can be broken in time $T_{i+1}^{1/2}$.
– Com is T_d-hiding but can be broken in time $T_{d+1}^{1/2}$.
– $\langle \mathcal{P}_{\text{WIPOK}}, \mathcal{V}_{\text{WIPOK}} \rangle$ is T_{d+1}-witness-indistinguishable (by using a T_{d+1}-hiding commitment). We denote the messages of the protocol by (α, β, γ).

Specifically, we pick π_i to be π restricted to $\{0,1\}^{\ell_i}$, where $\ell_i = (\log n)^{(4/\delta)^{i+1}}$ so that $\text{poly}(n) \cdot 2^{\ell_i} \ll 2^{\ell_{i+1}^\delta}$. Taking $\ell_{d-1} = \text{poly}(n)$, yields $(4/\delta)^d = \Theta(\log n / \log\log n)$ and thus $d = \Theta(\log\log n)$. We will instantiate Com from π on $(\log n)^{(4/\delta)^{d+2}}$ bits and $\langle \mathcal{P}_{\text{WIPOK}}, \mathcal{V}_{\text{WIPOK}} \rangle$ from π on $(\log n)^{(4/\delta)^{d+3}}$ bits. We present the protocol in Fig 2.

Common input: security parameter 1^n and an identity $\mathsf{id} \in \{0, 1, \ldots, d-1\}$.

Sender's input: a value $v \in \{0,1\}^n$.

Commit Phase:

Stage 0:

$\mathcal{C} \rightarrow \mathcal{R}$: Pick uniformly $r \in \{0,1\}^{\mathrm{poly}(n)}$ and send $c = \mathsf{Com}(v; r)$.

Stage 1 (Slot 1):

$\mathcal{R} \rightarrow \mathcal{C}$: Pick uniformly $\sigma_1 \in \{0,1\}^{\ell_{\mathsf{id}}}$.

$\mathcal{C} \Leftrightarrow \mathcal{R}$: Prove statement (c, σ_1) using $\langle \mathcal{P}_{\text{WIPOK}}, \mathcal{V}_{\text{WIPOK}} \rangle$ and witness $(v, r, \perp)$ w.r.t. the relation

$$\Lambda_{\mathsf{Com}} = \{((c, \sigma), (v, r, s)), |s| = |\sigma| \mid c = \mathsf{Com}(v; r) \text{ OR } \pi(s) = \sigma\}$$

Stage 2 (Slot 2):

$\mathcal{R} \rightarrow \mathcal{C}$: Pick uniformly $\sigma_2 \in \{0,1\}^{\ell_{d-1-\mathsf{id}}}$.

$\mathcal{C} \Leftrightarrow \mathcal{R}$: Prove statement (c, σ_2) using $\langle \mathcal{P}_{\text{WIPOK}}, \mathcal{V}_{\text{WIPOK}} \rangle$ and witness $(v, r, \perp)$ w.r.t. Λ_{Com}.

Reveal Phase:

$\mathcal{C} \rightarrow \mathcal{R}$: Send v, r.

$\mathcal{R}$: Verify that $c = \mathsf{Com}(v; r)$.

Fig. 2. The commitment scheme $\mathsf{nmCom} = (\mathcal{C}, \mathcal{R})$ for short identities. We denote the 4 messages exchanged in stage b by $\sigma_b, \alpha_b, \beta_b, \gamma_b$, for $b = 1, 2$. The values $\ell_0, \ldots, \ell_d$ are specified in Section 3.2.

Lemma 1. *The protocol* nmCom *is a statistically binding commitment scheme.*

Proof. The binding property follows readily from the fact that Com is itself statistically binding. To establish hiding, we construct a simulator $\mathcal{C}'$ that plays the role of the sender in nmCom. $\mathcal{C}'$ on input a commitment c to a string under Com and an identity id interacts with $\mathcal{R}$ as follows:

Stage 0: Sends c.

Stage 1: Computes $s_1 = \pi^{-1}(\sigma_1)$ and proves the statement (c, σ_1) using the witness $(\perp, \perp, s_1)$.

Stage 2: Computes $s_2 = \pi^{-1}(\sigma_2)$ and proves the statement (c, σ_2) using the witness $(\perp, \perp, s_2)$.

We allow $\mathcal{C}'$ to run in time $o(T_d)$ so that it can invert π on σ_1, σ_2. Then, witness indistinguishability of $\langle \mathcal{P}_{\text{WIPOK}}, \mathcal{V}_{\text{WIPOK}} \rangle$ implies that for all v:

$$\mathsf{view}_{\mathcal{R}^*}\langle \mathcal{C}(v), \mathcal{R}^* \rangle \cong_c \mathsf{view}_{\mathcal{R}^*}\langle \mathcal{C}'(\mathsf{Com}(v)), \mathcal{R}^* \rangle$$

On the other hand, Com is T_d-hiding and $\mathcal{C}'$ runs in time $o(T_d)$, so we have

$$\mathsf{view}_{\mathcal{R}^*}\langle \mathcal{C}'(\mathsf{Com}(v)), \mathcal{R}^* \rangle \cong_c \mathsf{view}_{\mathcal{R}^*}\langle \mathcal{C}'(\mathsf{Com}(0^n)), \mathcal{R}^* \rangle$$

Combining, we obtain $\mathsf{view}_{\mathcal{R}^*}\langle \mathcal{C}(v), \mathcal{R}^* \rangle \cong_c \mathsf{view}_{\mathcal{R}^*}\langle \mathcal{C}(0^n), \mathcal{R}^* \rangle$, from which hiding follows. □

Lemma 2. *The protocol* nmCom *is one-one non-malleable for identities of length* $\log\log\log n + O(1)$.

Proof. Consider a man-in-the-middle adversary $\mathcal{A}$. We assume WLOG that $\mathcal{A}$ is deterministic. Following [25, 16], the stand-alone adversary $\mathcal{S}$ uses $\mathcal{A}$ as a black box and emulates the left interaction by honestly committing to the string 0^n. Messages from the right interaction are forwarded externally. As such, it suffices to show that for all v:

$$\mathsf{mim}^{\mathcal{A}}_{\mathsf{nmCom}}(v) \cong_c \mathsf{mim}^{\mathcal{A}}_{\mathsf{nmCom}}(0^n) \qquad (*)$$

On a high level, the proof consists of a series of hybrid arguments:

STEP 1: *Simulate the left interaction using $\mathcal{C}'$ instead of $\mathcal{C}$.*

Specifically, let $\mathcal{S}'$ denote the stand-alone adversary that like $\mathcal{S}$, uses $\mathcal{A}$ as a black box and forwards message from the right interaction externally; the difference is that it emulates the left interaction by running $\mathcal{C}'$ (on input $\mathsf{Com}(v)$) instead of $\mathcal{C}$. We denote the output of this experiment by $\mathsf{sta}^{\mathcal{S}'}(\mathsf{Com}(v))$. By T_{d+1}-witness-indistinguishability, the transcripts of the right interaction when we use $\mathcal{C}$ and when we use $\mathcal{C}'$ on the left will be T_{d+1}-indistinguishable; in particular, the commitments in Stage 0 on the right are T_{d+1}-indistinguishable. Recall that we can extract the values in these Stage 0 commitments in time $o(T_{d+1})$. This implies:

$$\mathsf{mim}^{\mathcal{A}}_{\mathsf{nmCom}}(v) \cong_c \mathsf{sta}^{\mathcal{S}'}(\mathsf{Com}(v))$$

STEP 2: *Extract $\tilde{v}$ on the right.*

Using the knowledge extractor for $\langle \mathcal{P}_{\text{WIPOK}}, \mathcal{V}_{\text{WIPOK}} \rangle$ on the right, we may extract the witnesses for both slots on the right in the experiment $\mathsf{sta}^{\mathcal{S}'}(\mathsf{Com}(v))$, the idea being one of the two witnesses should contain the witness $(\tilde{v}, \tilde{r})$ for the commitment on the right. More precisely, let $\mathsf{ext\text{-}sta}^{\mathcal{S}'}(c)$ denote the output of the following experiment (a pictorial representation is provided in Fig 4):

1. Fix a random tape for $\mathcal{S}'(c)$ by fixing one for $\mathcal{C}'(c)$. This allows us to treat $\mathcal{S}'(c)$ as a single deterministic entity.

2. Fix a random tape for $\mathcal{R}$ and compute $\tau = \langle \mathcal{S}'(c), \mathcal{R} \rangle$. Let $\tilde{\mathsf{id}}$ denote the tag of the right interaction and τ denote the transcript $(\tilde{c}, \tilde{\sigma}_1, \tilde{\alpha}_1, \tilde{\beta}_1, \tilde{\gamma}_1, \tilde{\sigma}_2, \tilde{\alpha}_2, \tilde{\beta}_2, \tilde{\gamma}_2)$. If $\tilde{\mathsf{id}} = \mathsf{id}$ or $\mathcal{S}'$ aborts or if $\mathcal{R}$ rejects, output the view of $\mathcal{A}$ and $\perp$ and halt.
3. Rewind and attempt to extract witnesses $\tilde{w}_1, \tilde{w}_2$ for the respective statements $(\tilde{c}, \tilde{\sigma}_1)$ and $(\tilde{c}, \tilde{\sigma}_2)$ w.r.t. Λ_{Com}, relying on the special-soundness of $\langle \mathcal{P}_{\text{WIPOK}}, \mathcal{V}_{\text{WIPOK}} \rangle$. This is done as usual, by sending new random messages $\tilde{\beta}'_1$, $\tilde{\beta}'_2$, but with the following important exception: if $\mathcal{A}$ schedules messages in a different way than in τ (or if $\mathcal{R}$ rejects), the rewinding is aborted, and restarted. Let Γ denote the set of all possible scheduling; clearly, $|\Gamma| = O(1)$ since the protocol is constant-round. We will show that the expected number of rewindings for Slot 1 is given by $|\Gamma| = O(1)$; the same argument applies to Slot 2. Let τ_1 denote the prefix of τ up to Slot 1. For each schedule $\rho \in \Gamma$, let $\Pr[\rho \mid \tau_1]$ denote the probability that $\mathcal{R}$ accepts (i.e. obtaining convincing proof both slots) using the scheduling ρ conditioned on the prefix being τ_1. For a fixed τ_1, ρ, the expected number of rewindings is given by $\frac{1}{\Pr[\rho|\tau_1]}$. Therefore, the total expected number of rewindings for Slot 1 is given by:

$$\sum_{\tau_1} \Pr[\tau_1] \sum_{\rho \in \Gamma} \Pr[\rho \mid \tau_1] \cdot \frac{1}{\Pr[\rho \mid \tau_1]} = \sum_{\tau_1} \Pr[\tau_1] \cdot |\Gamma| = |\Gamma|$$

By linearity of expectations, the total expected number of rewindings for both slots is also $O(1)$. We now only need to make sure that we indeed extracted a valid opening: if either $\tilde{w}_1$ or $\tilde{w}_2$ is a valid opening $(\tilde{v}, \tilde{r})$ for $\tilde{c}$, output $(\tau, \tilde{v})$, else output `fail`.

We know that whenever $\mathsf{ext\text{-}sta}^{\mathcal{S}'}(\mathsf{Com}(v))$ does not output `fail`, its output contains the correct value $\tilde{v}$ and therefore the distributions

$$\mathsf{sta}^{\mathcal{S}'}(\mathsf{Com}(v)) \quad \text{and} \quad \mathsf{ext\text{-}sta}^{\mathcal{S}'}(\mathsf{Com}(v))$$

are identical. In the next subsection, we will establish the following claim:

Claim (simulation-soundness). For all v, $\Pr[\mathsf{ext\text{-}sta}^{\mathcal{S}'}(\mathsf{Com}(v)) = \texttt{fail}] = \mathrm{neg}(n)$

For now, we hint that the proof of the claim exploits the two-slot structure in an essential way to transform a non-negligible failure probability in extraction into non-negligible success probability at inverting π. Assuming that the claim holds, it follows readily that

$$\mathsf{sta}^{\mathcal{S}'}(\mathsf{Com}(v)) \cong_c \mathsf{ext\text{-}sta}^{\mathcal{S}'}(\mathsf{Com}(v))$$

STEP 3: *Replace the input to $\mathcal{S}'$ with $\mathsf{Com}(0^n)$.*

Now, we observe that $\mathcal{S}'$ combined with the knowledge extractor on the right runs in expected time $o(T_d)$. This is less than the time it takes to break Com, and thus its output will be indistinguishable whether the input to $\mathcal{S}'$ is $\mathsf{Com}(v)$ or $\mathsf{Com}(0^n)$.

In particular, we have

$$\mathsf{ext\text{-}sta}^{\mathcal{S}'}(\mathsf{Com}(v)) \cong_c \mathsf{ext\text{-}sta}^{\mathcal{S}'}(\mathsf{Com}(0^n))$$

Combining steps 1 and 2, we have

$$\mathsf{mim}^{\mathcal{A}}_{\mathsf{nmCom}}(v) \cong_c \mathsf{ext\text{-}sta}^{\mathcal{S}'}(\mathsf{Com}(v))$$
$$\mathsf{mim}^{\mathcal{A}}_{\mathsf{nmCom}}(0^n) \cong_c \mathsf{ext\text{-}sta}^{\mathcal{S}'}(\mathsf{Com}(0^n))$$

Combining with Step 3 yields $(*)$. □

3.3 Proof of simulation-soundness

We complete the proof of Lemma 2 by establishing the main technical claim. Suppose towards a contradiction that the claim is false, i.e., there is some non-negligible function ϵ such that for all sufficiently large n, there exists some v satisfying

$$\Pr[\mathsf{ext\text{-}sta}^{\mathcal{S}'}(\mathsf{Com}(v)) = \mathtt{fail}] > \epsilon(n)$$

Fix one such n, along with an associated v and identity id. In addition, we may also fix the coin tosses of $\mathcal{S}'$ and some specific $c = \mathsf{Com}(v)$, along some $\tilde{\mathsf{id}}$ on the right, while losing a factor d in the probability ext-sta outputs `fail`. That is, with probability at least $\frac{\epsilon(n)}{d}$ (over the coin tosses of $\mathcal{R}$), the tag on the right is $\tilde{\mathsf{id}}$ and the knowledge extractor outputs witnesses $\pi^{-1}(\tilde{\sigma}_1)$ and $\pi^{-1}(\tilde{\sigma}_2)$. We then construct an adversary $\tilde{\mathcal{A}}$ that for some $j \in \{\tilde{\mathsf{id}}, d-1-\tilde{\mathsf{id}}\}$, inverts π on $\{0,1\}^{\ell_j}$ with probability $\Omega(\frac{\epsilon(n)}{d})$ in time $o(T_j)$, which contradicts the one-wayness of π. Roughly speaking, $\tilde{\mathcal{A}}$ works as follows: on input a challenge $\sigma \in \{0,1\}^{\ell_j}$, simulate the experiment $\mathsf{ext\text{-}sta}^{\mathcal{S}'}(\mathsf{Com}(v))$, and

- if $j = \tilde{\mathsf{id}}$, set $\tilde{\sigma}_1 = \sigma$ and compute $\pi^{-1}(\sigma)$ by extracting the witness from Slot 1; and
- if $j = d - 1 - \tilde{\mathsf{id}}$, set $\tilde{\sigma}_2 = \sigma$ and compute $\pi^{-1}(\sigma)$ by extracting the witness from Slot 2.

Recall that $\mathcal{S}'$ is simply $\mathcal{A}$ with a left execution of $\mathcal{C}'(\mathsf{Com}(v))$ and thus a naive simulation of $\mathcal{S}'$ takes time roughly $T_d \gg T_j$. The bottleneck to an efficient simulation lies in computing each of the messages γ_1, γ_2 in stages 1 and 2 in the computation of $\mathcal{C}'$. We adopt one of three strategies to accomplish this in time $o(T_j)$: compute the message by computing a witness, hardwire the message into the reduction, or argue that we do not need to compute the message for extraction on the right. We consider three representative schedulings of the messages γ_1, γ_2 in relation to the two slots in the right execution. In our analysis we crucially rely on the fact that ext-sta aborts all rewindings that use a different schedule than ext-sta saw in the first simulation τ. Given this property it is sufficient to consider a *static* scheduling. In particular, as the number of possible scheduling is constant, we can WLOG consider a particular fixed scheduling (again at the cost of only a constant loss).

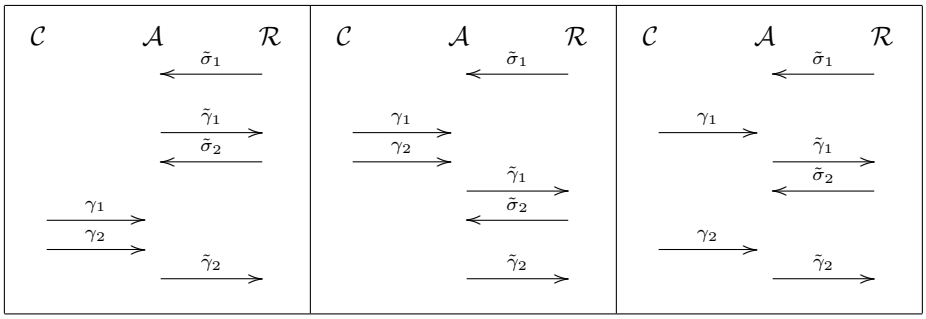

Fig. 3. Three scheduling strategies

Both γ_1, γ_2 are sent after Slot 1 ends. In this case, we construct an adversary that breaks π on $\{0,1\}^{\ell_{\tilde{\mathsf{id}}}}$ by extracting from Slot 1 on the right. Here, we observe that we do not need to compute γ_1 or γ_2. Specifically, we just need to simulate the interaction on the left up to the end of Slot 1. Since both γ_1 and γ_2 are sent after Slot 1, we can complete the simulation on the left in polynomial time.

Both γ_1, γ_2 are sent before Slot 2 begins. In this case, we construct an adversary that breaks π on $\{0,1\}^{\ell_{d-1-\tilde{\mathsf{id}}}}$ by extracting from Slot 2 on the right. Here, we argue that we can fix γ_1 and γ_2. Specifically, we non-uniformly fix a partial transcript of the execution up to the point just before Slot 2 on the right begins, i.e. just before $\tilde{\sigma_2}$ is sent. This fixes all of the left interaction, so there is nothing left to simulate on the left.

γ_1 is contained in Slot 1 and γ_2 in Slot 2. More generally, this covers the case γ_1 is sent before Slot 2 begins and γ_2 is sent after Slot 1 ends. We need to consider two sub-cases:

- $\mathsf{id} < \tilde{\mathsf{id}}$. We construct an adversary that breaks π on $\{0,1\}^{\ell_{\tilde{\mathsf{id}}}}$ by extracting from Slot 1 on the right. Here, we compute γ_1 (by inverting σ_1) and observe that we do not need to compute γ_2. Specifically, we just need to simulate the interaction on the left up to the end of Slot 1 on the right, i.e. up to the point $\tilde{\gamma_1}$ is sent. It suffices to compute $\pi^{-1}(\sigma_1)$ and thus γ_1, which can be done in time $o(T_{\mathsf{id}+1}) \ll T_{\tilde{\mathsf{id}}}$.
- $\mathsf{id} > \tilde{\mathsf{id}}$. We construct an adversary that breaks π on $\{0,1\}^{\ell_{d-\tilde{\mathsf{id}}-1}}$ by extracting from Slot 2 on the right. Here, we fix γ_1 and then compute γ_2 (by inverting σ_2). Specifically, we non-uniformly fix a partial transcript of the execution up to the point just before Slot 2 on the right begins, i.e. just before $\tilde{\sigma_2}$ is sent. This fixes all of the first slot on the left (including γ_1) which we may then hardwire into the reduction. To complete the simulation on the left for Slot 2, it suffices to compute $\pi^{-1}(\sigma_2)$ and thus γ_2, which can be done in time $o(T_{d-\mathsf{id}}) \ll T_{d-\tilde{\mathsf{id}}-1}$.

With probability $\Omega(\frac{\epsilon(n)}{d})$, ext-sta outputs `fail` for one of these schedulings. We may then use $\tilde{\mathcal{A}}$ for that scheduling to derive a contradiction to the one-wayness of π.

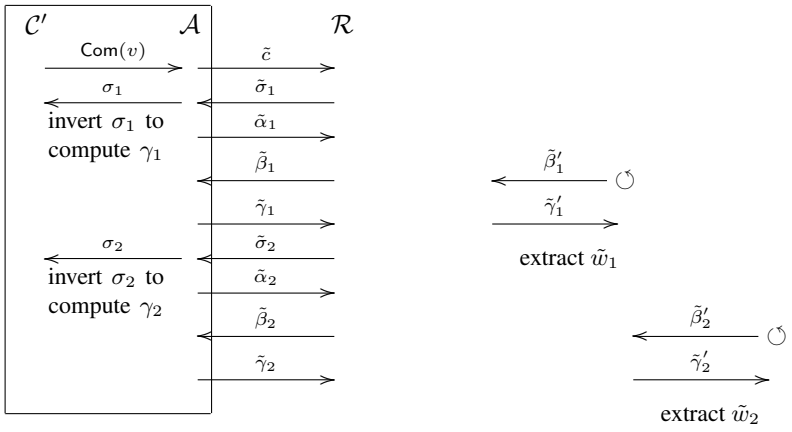

Fig. 4. A pictorial representation of $\mathsf{ext\text{-}sta}^{\mathcal{S}'}(\mathsf{Com}(v))$

Remark 1. We highlight two subtleties in the analysis:

- It is important that in $\tilde{\mathcal{A}}$'s simulation of $\mathsf{ext\text{-}sta}^{\mathcal{S}'}(\mathsf{Com}(v))$, it uses the same witnesses for the relation Λ_{Com} as $\mathcal{C}'$; otherwise, we could have easily solved the problem of efficient simulation by having $\tilde{A}$ use the witness (v, r) for the commitment $\mathsf{Com}(v)$. We cannot appeal to witness-indistinguishability here because we rewind the $\langle \mathcal{P}_{\text{WIPOK}}, \mathcal{V}_{\text{WIPOK}} \rangle$ protocols.
- It is also important for simulation-extractability that σ_2 is sent *after* the completion of Stage 1 in nmCom. This way, we can fix a partial transcript up to the end of Slot 1 while allowing the verifier's challenge σ_2 in Slot 2 to remain undetermined.

4 Non-malleability amplification

In order to apply the non-malleable amplification theorem from [15] to our construction, we first need to modify our construction to satisfy an additional technical requirement, that of non-malleability w.r.t 4-round protocols (to be formalized shortly), which they coin *natural*. [15] also requires that the commitment scheme be *initial-binding*, that is, the first message sent by the sender already determines the value committed to; our commitment scheme clearly satisfies this.

Lemma 3 (Non-malleability amplification [15]). *Let $\langle \mathcal{C}, \mathcal{R} \rangle$ be a $k(n)$-round natural non-malleable commitment scheme for identities of length $t(n)$ with computational complexity $p(n)$. Then, there exists a $15k(n)$-round natural non-malleable commitment scheme for identities of length $2^{t(n)-1}$ with computational complexity $2^{t(n)}p(n) + k(n)poly(n) + poly(n)$.*

Non-malleability w.r.t. k-round protocols. The concept of non-malleability is traditionally only considered in a setting where a man-in-the middle adversary is participating in two (or more) executions of the *same* protocol. We here consider a notion of non-malleability with respect to arbitrary k-round protocols.

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

Definition 2. *Let $(\mathcal{C}, \mathcal{R})$ be a commitment scheme, and B an interactive Turing machine. We say the commitment scheme $(\mathcal{C}, \mathcal{R})$ is* one-many non-malleable w.r.t. *B, if for every probabilistic polynomial-time man-in-the-middle adversary A, and every two sequences $\{y_n^1\}_{n\in N}$ and $\{y_n^2\}_{n\in N}$, such that*

$$\left\{\mathsf{view}_A[\langle B(y_n^1), A(z)\rangle]\right\}_{n\in N, z\in\{0,1\}^*} \approx \left\{\mathsf{view}_A[\langle B(y_n^2), A(z)\rangle]\right\}_{n\in N, z\in\{0,1\}^*}$$

it holds that:

$$\left\{\mathsf{mim}^{B,A}(y_n^1, z)\right\}_{n\in N, z\in\{0,1\}^*} \approx \left\{\mathsf{mim}^{B,A}(y_n^2, z)\right\}_{n\in N, z\in\{0,1\}^*}$$

We say that $(\mathcal{C}, \mathcal{R})$ is one-many non-malleable w.r.t k-round protocols if $(\mathcal{C}, \mathcal{R})$ is one-many non-malleable w.r.t any machine B that interacts with the man-in-the-middle adversary in k rounds.

Modifying our construction. We describe a variant of our construction in Section 3 that is one-many non-malleable w.r.t $(2c-1)$-round protocols for any constant $c > 1$. In addition, the protocol now handles identities of length $c \log\log\log n + O(1)$, although the increase is not necessary for non-malleability amplification. Specifically, we follow the multiple slot approach in [22] to boost the number of slots from 2 to $2c$. On input a tag $\mathsf{id} \in 0, 1, \ldots, d^c - 1$, let $(\mathsf{id}_1, \ldots, \mathsf{id}_c)$ denote the base d representation of id. For $j = 1, 2, \ldots, c$, we will pick a challenge of length ℓ_{id_j} for the $2j - 1$'th slot, and a challenge of length $\ell_{d-1-\mathsf{id}_j}$ for the $2j$'th slot.

The analysis. First, we need to verify that the modified construction remains one-many non-malleable (w.r.t. itself). Indeed, the proof of Lemma 2 and the analysis in Section A

extend in a straight-forward manner to $c > 1$, except in the proof of simulation-soundness, where it is slightly more involved. We will consider two broad classes of scheduling strategies:

- For all $j = 1, 2, \ldots, c$: γ_{2j-1} is contained in Slot $2j - 1$ and γ_{2j} is contained in Slot $2j$.
- There exists some j where one of Slot $2j-1$ or Slot $2j$ contains none of $\gamma_1, \ldots, \gamma_{2c}$.

The previous analysis will still go through, except we now lose a factor $\frac{1}{(dc)^{\Omega(c)}}$ (as opposed to $1/d$ from before) in the probability of inverting the one-way permutation.

Next, we argue that the modified construction is one-many non-malleable w.r.t $(2c-1)$-round protocols. This follows from the fact that we now have $2c$ rewinding slots on the right (c.f. [15]) so that there will always be a slot on the right that does not contain any message from the $(2c-1)$-round protocol executing on the left.

5 Construction from sub-exponential one-way functions

We need to make two modifications to the protocol in Section 3 in order to handle a general one-way function f instead of a one-way permutation π with sub-exponential hardness.

Modifying receiver's challenge. Following [3], we will replace the challenge that the receiver sends at the start of each of the two slots with a 3-round challenge response protocol. This is essentially a cut-and-choose protocol that guarantees that the receiver sends challenges in the range of the one-way function f. Again, we fix some input length ℓ for f corresponding to the desired level of security for the slot.

$\mathcal{R} \rightarrow \mathcal{C}$: Pick s_i^b at random from $\{0,1\}^\ell$ and send $y_i^b = f(s_i^b)$ for $b = 0, 1,\ i = 1, 2, \ldots, n$.

$\mathcal{C} \rightarrow \mathcal{R}$: Send $\mu = (\mu_1, \ldots, \mu_n)$ at random from $\{0,1\}^n$.

$\mathcal{R} \rightarrow \mathcal{C}$: Send $(s_1^{\mu_1}, \ldots, s_n^{\mu_n})$.

$\mathcal{C}$: Verify that for all $i = 1, 2, \ldots, n$: $f(s_i^{\mu_i}) = y_i$.

The sender will then run $\langle \mathcal{P}_{\text{WIPOK}}, \mathcal{V}_{\text{WIPOK}} \rangle$ on the instance $(c, y_1^0, y_1^1, \ldots, y_n^0, y_n^1, \mu)$ w.r.t. the following relation:

$$\Lambda_{\mathsf{Com}} = \{((c, y_1^0, y_1^1, \ldots, y_n^0, y_n^1, \mu), (v, r, i, s)), \mid c = \mathsf{Com}(v; r) \text{ OR } f(s) = y_i^{1-\mu_i}\}$$

The challenge-response protocol has the following properties (cf. [3]):

– With probability $1 - 2^{-n}$ over μ, if the sender accepts at the end of the challenge-response protocol, then there exists a trapdoor witness for the relation Λ_{Com}. Indeed, a trapdoor witness exists unless at most one value in each pair (y_i^0, y_i^1) lies in $f(\{0,1\}^\ell)$, in which case there exists at most one μ for which the sender will not abort.
– It is computationally infeasible for a $2^{O(\ell)}$-time adversary to find a trapdoor witness for the relation Λ_{Com} if f is an exponential one-way function.

Modifying the commitment schemes. We will use Naor's commitment scheme [19] in Com and in $\langle \mathcal{P}_{\text{WIPOK}}, \mathcal{V}_{\text{WIPOK}} \rangle$. Specifically, we will commit v by committing to each bit of v in parallel. We may set the values of $T_0, \ldots, T_d$ as before. The complexity of breaking a T_d-hiding commitment via brute-force is now $\mathrm{poly}(n) \cdot 2^{O((\log T_d)^\kappa))}$ (for some constant $\kappa > 1$ that depends on the seed length of pseudorandom generators from one-way functions in [12]). We can then set $\ell_{d+1} = n^{\Theta(\kappa/\delta)}$ to ensure that $T_{d+1}^{1/2} > \mathrm{poly}(n) \cdot 2^{O((\log T_d)^\kappa))}$.

Acknowledgments. We thank Vinod Vaikuntanathan for an inspiring discussion on non-malleability.

References

[1] B. Barak. How to go beyond the black-box simulation barrier. In *FOCS*, pages 106–115, 2001.
[2] B. Barak. Constant-round coin-tossing with a man in the middle or realizing the shared random string model. In *FOCS*, pages 345–355, 2002.
[3] M. Bellare, M. Jakobsson, and M. Yung. Round-optimal zero-knowledge arguments based on any one-way function. In *EUROCRYPT*, pages 280–305, 1997.
[4] R. Canetti, O. Goldreich, S. Goldwasser, and S. Micali. Resettable zero-knowledge. In *STOC*, pages 235–244, 2000.
[5] G. Di Crescenzo, Y. Ishai, and R. Ostrovsky. Non-interactive and non-malleable commitment. In *Proc. 30th STOC*, pages 141–150, 1998.
[6] D. Dolev, C. Dwork, and M. Naor. Nonmalleable cryptography. *SIAM J. Comput.*, 30(2): 391–437, 2000.
[7] U. Feige and A. Shamir. Zero knowledge proofs of knowledge in two rounds. In *CRYPTO*, pages 526–544, 1989.
[8] O. Goldreich and A. Kahan. How to construct constant-round zero-knowledge proof systems for NP. *J. Cryptology*, 9(3):167–190, 1996.
[9] O. Goldreich and H. Krawczyk. On the composition of zero-knowledge proof systems. *SIAM J. Comput.*, 25(1):169–192, 1996.
[10] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game or a completeness theorem for protocols with honest majority. In *STOC*, pages 218–229, 1987.
[11] O. Goldreich, S. Micali, and A. Wigderson. Proofs that yield nothing but their validity for all languages in NP have zero-knowledge proof systems. *J. ACM*, 38(3):691–729, 1991. Prelim. version in FOCS '86.

[12] J. Håstad, R. Impagliazzo, L. A. Levin, and M. Luby. A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28(4):1364–1396, 1999.
[13] R. Impagliazzo and S. Rudich. Limits on the provable consequences of one-way permutations. In *STOC*, pages 44–61, 1989.
[14] J. Katz, R. Ostrovsky, and A. Smith. Round efficiency of multi-party computation with a dishonest majority. In *EUROCRYPT*, pages 578–595, 2003.
[15] H. Lin and R. Pass. Non-malleability amplification. In *STOC*, pages 189–198, 2009.
[16] H. Lin, R. Pass, and M. Venkitasubramaniam. Concurrent non-malleable commitments from any one-way function. In *TCC*, pages 571–588, 2008.
[17] H. Lin, R. Pass, and M. Venkitasubramaniam. A unified framework for concurrent security: universal composability from stand-alone non-malleability. In *STOC*, pages 179–188, 2009.
[18] M. Liskov, A. Lysyanskaya, S. Micali, L. Reyzin, and A. Smith. Mutually independent commitments. In *ASIACRYPT*, pages 385–401, 2001.
[19] M. Naor. Bit commitment using pseudorandomness. *J. Cryptology*, 4(2):151–158, 1991.
[20] R. Ostrovsky, G. Persiano, and I. Visconti. Simulation-based concurrent non-malleable commitments and decommitments. In *TCC*, pages 91–108, 2009.
[21] O. Pandey, R. Pass, and V. Vaikuntanathan. Adaptive one-way functions and applications. In *CRYPTO*, pages 57–74, 2008.
[22] R. Pass. Bounded-concurrent secure multi-party computation with a dishonest majority. In *STOC*, pages 232–241, 2004.
[23] R. Pass and A. Rosen. Bounded-concurrent secure two-party computation in a constant number of rounds. In *FOCS*, pages 404–413, 2003.
[24] R. Pass and A. Rosen. New and improved constructions of nonmalleable cryptographic protocols. *SIAM J. Comput.*, 38(2):702–752, 2008. Preliminary version in STOC '05.
[25] R. Pass and A. Rosen. Concurrent nonmalleable commitments. *SIAM J. Comput.*, 37(6): 1891–1925, 2008. Preliminary version in FOCS '05.
[26] A. Sahai. Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In *FOCS*, pages 543–553, 1999.

A nmCom is one-many non-malleable

Here, we establish a stronger claim, namely that the protocol nmCom is in fact one-many non-malleable for identities of length $\log\log\log n + O(1)$. We do not need this stronger property, although it is of independent interest. To see why the claim holds, suppose there are m right interactions, where the tags are respectively $(\tilde{\mathsf{id}}_1, \ldots, \tilde{\mathsf{id}}_m)$ and the committed values are respectively $(\tilde{v}_1, \ldots, \tilde{v}_m)$. We modify Step 2 to extract each $\tilde{v}_i$ on the right where $\tilde{\mathsf{id}}_i \neq \mathsf{id}$. As before, we will sample one transcript τ, and then attempt to extract witnesses for each of the m right executions. We need an expected $2m|\Gamma|$ rewindings, $2|\Gamma|$ for each of the m right executions. Next, we will need to show that the probability that the extractor outputs `fail` for any of the m right interactions in negligible. If this probability is at least ϵ, then there is some right interaction for which the extractor outputs `fail` for that interaction with probability at least $\frac{\epsilon}{m}$. Simply repeat the analysis for simulation-soundness in Section 3.3 for this execution (and simulate $\mathcal{R}$ for the other $m-1$ interactions internally).