

Non-Interactive Non-Malleability from Quantum Supremacy

Yael Tauman Kalai ^{*} and Dakshita Khurana ^{**}

Abstract. We construct non-interactive non-malleable commitments without setup in the plain model, under well-studied assumptions.

First, we construct non-interactive non-malleable commitments w.r.t. commitment for $\epsilon \log \log n$ tags for a small constant $\epsilon > 0$, under the following assumptions:

1. Sub-exponential hardness of factoring or discrete log.
2. Quantum sub-exponential hardness of learning with errors (LWE).

Second, as our key technical contribution, we introduce a new tag amplification technique. We show how to convert any non-interactive non-malleable commitment w.r.t. commitment for $\epsilon \log \log n$ tags (for any constant $\epsilon > 0$) into a non-interactive non-malleable commitment w.r.t. replacement for 2^n tags. This part only assumes the existence of sub-exponentially secure non-interactive witness indistinguishable (NIWI) proofs, which can be based on sub-exponential security of the decisional linear assumption.

Interestingly, for the tag amplification technique, we crucially rely on the leakage lemma due to Gentry and Wichs (STOC 2011). For the construction of non-malleable commitments for $\epsilon \log \log n$ tags, we rely on quantum supremacy. This use of quantum supremacy in classical cryptography is novel, and we believe it will have future applications. We provide one such application to two-message witness indistinguishable (WI) arguments from (quantum) polynomial hardness assumptions.

1 Introduction

Non-malleability, first introduced by Dolev, Dwork and Naor [11] aims to counter the ubiquitous problem of man-in-the-middle (MIM) attacks on cryptographic protocols. A MIM adversary participates in two or more instantiations of a protocol, trying to use information obtained in one execution to breach security in the other protocol execution. A non-malleable protocol should ensure that such an adversary gains no advantage. A long-standing problem in this area has been to build non-malleable protocols, without any additional setup or rounds of interaction. In this paper, we develop techniques to address this question based on well-studied assumptions. We focus on a core non-malleable primitive – a commitment scheme.

^{*} Microsoft Research and MIT, Cambridge, USA. Email: yael@microsoft.com

^{**} Microsoft Research, Cambridge, USA and UIUC, Urbana-Champaign, USA. Email: dakshita@illinois.edu

Non-interactive Commitments. A non-interactive commitment scheme consists of a commitment algorithm, that on input a message m and randomness r , outputs a commitment to m , which is denoted by $\text{com}(m; r)$ ¹. A commitment scheme is required to be both binding and hiding. The (statistical) binding requirement asserts that a commitment cannot be opened to two different messages $m \neq m'$, namely, there do not exist $m \neq m'$ and randomness r, r' such that $\text{com}(m; r) = \text{com}(m'; r')$. The (computational) hiding property asserts that for any two messages, m and m' (of the same length), the distributions $\text{com}(m)$ and $\text{com}(m')$ are computationally indistinguishable. We note that one could also consider computational binding and statistical hiding, however such commitment schemes are known to require at least two rounds of interaction when dealing with non-uniform adversaries. The focus of this work is on the non-interactive setting.

Non-interactive non-malleable commitments. Loosely speaking, a commitment scheme is said to be non-malleable if no MIM adversary, given a commitment $\text{com}(m)$, can efficiently generate a commitment $\text{com}(m')$, such that the message m' is related to the original message m .

Non-malleable commitments are among the core building blocks of various cryptographic protocols such as coin-flipping, secure auctions, electronic voting, general multi-party computation (MPC) protocols, and non-malleable proof systems. Therefore, they have a direct impact on the round complexity of such protocols. For example, many constructions of concurrent MPC against Byzantine adversaries are bottlenecked by the round complexity of non-malleable commitments.

As such, there has been a long line of work on obtaining constructions of non-malleable commitments in the plain model in as few rounds as possible (e.g [11,2,32,34,30,31,27,36,35,26,15,17,19,18,9,10,23,29,24]). So far, the only known constructions of non-interactive non-malleable commitments (without setup) are the ones by Pandey, Pass and Vaikuntanathan [31], based on a strong non-falsifiable assumption, and Bitansky and Lin [5], based on a relatively new assumption about sub-exponential incompressible functions. We elaborate on these related works in Section 1.3.

Indeed, constructing non-interactive non-malleable commitment schemes (without setup) from standard assumptions, has been a long standing open problem and is the focus of this work. Three primary flavours of non-malleability have been considered in the literature:

- **Non-malleability w.r.t. commitment.** Intuitively, non-malleability w.r.t. commitment, which is the strongest of the three definitions, requires that for any two messages $m_0, m_1 \in \{0, 1\}^p$, the distributions $(\text{Com}(m_0), \tilde{m}_0)$ and $(\text{Com}(m_1), \tilde{m}_1)$ are computationally indistinguishable. Here \tilde{m}_b is the message committed to by the MIM given $\text{Com}(m_b)$, and is set to \perp if the adversary given $\text{Com}(m_b)$ outputs \tilde{c} for which there do not exist any (\tilde{m}, \tilde{r}) such that $\tilde{c} = \text{com}(\tilde{m}; \tilde{r})$. Another definition that is considered in the literature is

¹ We will sometimes omit explicitly writing the randomness r .

that of CCA-security for commitment schemes. It is known [7] that in the case of non-interactive commitments, non-malleability w.r.t. commitment is equivalent to (one-to-one) CCA-security.

- **Non-malleability w.r.t. replacement.** A weaker, yet natural, notion of malleability is non-malleability w.r.t. replacement [15]. This requires that for any two messages $m_0, m_1 \in \{0, 1\}^p$, the distributions $(\text{Com}(m_0), \tilde{m}_0)$ and $(\text{Com}(m_1), \tilde{m}_1)$ are indistinguishable *whenever* $\tilde{m}_0, \tilde{m}_1 \neq \perp$.² This is exactly like non-malleability w.r.t. commitment, except that the adversary is allowed to perform “selective abort” attacks, where the event that the adversary committed to an invalid message, is allowed to be correlated with the honest message. This guarantees that a man-in-the-middle adversary cannot commit to *valid* messages that are related to the message committed in an honest protocol. We observe that the proofs in [7] demonstrate that non-interactive non-malleability w.r.t. replacement is equivalent to a weaker form of CCA-security. We further elaborate upon this in Section 1.2.
- **Non-malleability w.r.t. opening.** This is an even weaker³, yet natural notion, which requires that for any two messages m_0, m_1 , the joint distribution of $(\text{Com}(m_0), \tilde{m}_0)$ and $(\text{Com}(m_1), \tilde{m}_1)$ are indistinguishable *whenever* $\tilde{m}_0, \tilde{m}_1 \neq \perp$, where \tilde{m}_b is the message *opened* by the MIM given $\text{Com}(m_b)$. The crucial difference from both the previous definitions is that \tilde{m}_0, \tilde{m}_1 represent the messages opened by the adversary, as opposed to the messages committed. Informally, this allows an adversary to commit to a message that is related to an honest message, as long the adversary is unable to convincingly open these commitments.

This work focuses on the first two definitions. We also note that all non-malleable commitment schemes assume that parties have “tags” (or id’s), and require non-malleability to hold whenever the adversary is trying to commit w.r.t. **tag** that is different from an honest **tag**. We differentiate between the following two settings:

- One-to-one setting, where the man-in-the-middle (MIM) gets a single committed message and generates a single commitment.
- Many-to-many (concurrent) setting, where the MIM receives many commitments and is allowed to generate many commitments. Here, the guarantee is that for any two sets of committed messages sent to the MIM, the joint distribution of these committed messages and the messages that the MIM commits to, are indistinguishable.

In this work, we focus on the one-to-one definition. But as a stepping stone, we define and construct many-to-many *same-tag* non-malleable commitments. This is similar to the many-to-many notion, except that it restricts the MIM to use the same tag in all commitments that he outputs.

² As earlier, \tilde{m}_b denotes the message committed to by the MIM given $\text{Com}(m_b)$.

³ Non-malleability w.r.t. replacement implies non-malleability w.r.t. opening, as defined by Goyal et al. [16].

1.1 Our Results

In this paper, we first construct non-malleable commitments w.r.t. commitment for $\epsilon \log \log n$ tags (for some small constant $\epsilon > 0$) in the many-to-many same-tag setting, based on well-studied hardness assumptions, which we elaborate on below. Then we present a general “tag amplification” compiler that converts any non-malleable commitment w.r.t. replacement with $\epsilon \log \log n$ tags in the many-to-many same tag setting, into a non-malleable commitment w.r.t. replacement with 2^n tags in the one-to-one setting, assuming sub-exponential NIWI (which can in turn be based on sub-exponential decisional linear (DLIN)).

For the first result, our contribution is primarily conceptual, and relies on using *quantum supremacy*. Our second result contains the bulk of the technical difficulty. In this part, we make a novel use of the leakage lemma due to Gentry and Wichs [13]. The use of the leakage lemma in this context is surprising, since a-priori the problem of non-malleability seems quite unrelated to leakage. In what follows, we state our results in more detail.

Non-interactive non-malleable commitments for $O(\log \log n)$ tags. We construct non-interactive non-malleable commitments w.r.t. commitment for $\epsilon \log \log n$ tags (for a small constant $\epsilon > 0$) assuming:

- Sub-exponential hardness of factoring or discrete log.
- Sub-exponential hardness of learning with errors (LWE) or learning parity with noise (LPN) against quantum circuits.

More generally, we construct non-malleable commitments w.r.t. commitment for $\epsilon \log \log n$ tags from any sub-exponentially secure bit commitment for 2 tags (denoted by com_0 and com_1), for which the hiding property of com_0 holds even given an oracle that breaks com_1 , and similarly the hiding property of com_1 holds even given an oracle that breaks com_0 . Such commitments are known as *adaptive* or *CCA-secure* commitments [31,28], and imply many-to-many non-interactive non-malleable commitments w.r.t. commitment.

Informal Theorem 1 *Assuming the existence of sub-exponentially CCA-secure many-to-many non-interactive bit commitments for 2 tags, there exist many-to-many same-tag non-interactive non-malleable string commitments w.r.t. commitment for $\epsilon \log \log n$ tags (for a small constant $\epsilon > 0$).*

To achieve this, we start with the leveraging technique of Pass and Wee [35] that allows us to construct, from any sub-exponentially secure non-interactive commitment, a series of $\epsilon \log \log n$ commitments, each harder than the previous one. But this only provides hardness in one direction, and in particular does not even yield commitments for 2 tags that are non-malleable w.r.t. each other.

Our main conceptual novelty in this part, which we describe next, is the idea of constructing a CCA secure commitment scheme for 2 tags using *quantum supremacy*. Later, we describe how we can carefully combine this insight with the technique of [35] to obtain non-malleable commitments for $\epsilon \log \log n$ tags.

Using Quantum Supremacy. Loosely speaking, in order to construct a CCA-secure commitment for 2 tags, we need two axes of hardness: One axis in which com_0 is harder than com_1 , and the other in which com_1 is harder than com_0 .

We build such an axis by relying on quantum supremacy, which is the ability of quantum computers to solve problems (such as factoring) that are believed to be hard for classical computers. Namely, we construct two commitment algorithms com_0 and com_1 such that for quantum algorithms, breaking com_1 is harder than breaking com_0 , and yet for classical algorithms, breaking com_0 is harder than breaking com_1 .

This is achieved by instantiating com_1 as a post-quantum secure commitment (such as one based on LWE or LPN [14]); and instantiating com_0 as a post-quantum *insecure* commitment (such as one based on factoring or discrete log), albeit with a much larger security parameter. Now, given a BQP oracle, com_1 is secure but com_0 is not; at the same time, classical machines can break com_1 faster than they can break com_0 . We prove the following claim:

Informal Claim 1 *Assuming sub-exponential hardness of factoring/discrete log and sub-exponential quantum hardness of LWE/LPN , there exist sub-exponentially CCA secure many-to-many non-interactive commitments for 2 tags.*

Combining this with Informal Theorem 1, we have:

Informal Theorem 2 *Assuming sub-exponential hardness of factoring/discrete log, and sub-exponential quantum hardness of LWE/LPN , many-to-many same-tag non-interactive non-malleable commitments w.r.t. commitment exist for $\epsilon \log \log n$ tags, for a small constant $\epsilon > 0$.*

Prior to this work, obtaining non-interactive non-malleable commitments w.r.t. commitment, even for just two tags, required the non-standard assumption that there exist sub-exponential incompressible one-way functions, and either sub-exponentially secure time-lock puzzles or sub-exponentially secure one-way functions admitting hardness amplification [5]. The work of [29] constructed non-interactive non-malleable commitments *w.r.t. extraction* (which is similar to w.r.t. replacement) for $O(\log \log n)$ tags assuming sub-exponentially secure time-lock puzzles or sub-exponentially secure one-way functions that admit hardness amplification [5]. We show that non-interactive non-malleable commitments w.r.t. commitment for $\epsilon \log \log n$ tags (in fact, even parallel CCA commitments for 2 tags) can be constructed based on much more well-studied assumptions than previously known.

We also remark that one can substitute the assumption on sub-exponential quantum hardness of LWE with sub-exponentially secure time-lock puzzles [29], or sub-exponentially secure one-way functions [5] admitting hardness amplification, to obtain (many-to-many) non-malleable commitments w.r.t. replacement for $\epsilon \log \log n$ tags.

We believe that this idea of using quantum supremacy may have other applications in classical cryptography. In particular, the technique of complexity

leveraging, which breaks hardness of one primitive while retaining hardness of another, is extensively used in cryptography. Typically, when this technique is used, the resulting scheme relies on super-polynomial (and often sub-exponential) hardness. We believe that in several such applications, the complexity leveraging technique can be replaced with quantum supremacy, thus converting such super-polynomial hardness assumptions to quantum polynomial hardness. For example, using our ideas, one can appropriately instantiate the protocols in [21] to obtain two-message witness indistinguishable protocols based on quantum-polynomial hardness of LWE , and polynomially hard one-way functions (such as those based on factoring or discrete log) that are invertible in BQP .

Non-interactive Tag Amplification from NIWIs Our more involved technical contribution is a non-interactive tag amplification technique that relies only on sub-exponentially secure non-interactive witness indistinguishable (NIWI) proofs for NP .

Informal Theorem 3 (Tag Amplification from NIWIs) *Assuming many-to-many same-tag non-malleable commitments w.r.t. replacement for $\epsilon \log \log n$ tags (for an arbitrarily small constant $\epsilon > 0$) and sub-exponentially secure NIWIs for NP , there exist non-interactive non-malleable commitments w.r.t. replacement for 2^n tags.*

We note that sub-exponentially secure NIWIs can be constructed assuming the sub-exponential hardness of the decisional linear problem [20], or from derandomization assumptions [3], or assuming indistinguishability obfuscation [6]. Interestingly, to prove this theorem, we crucially rely on the Gentry-Wichs leakage lemma [13]. We provide a high-level overview of this amplification technique, as well as its proof, in Section 2.2. To summarize, assuming sub-exponential hardness of factoring or discrete log, as well as sub-exponential quantum hardness of LWE or LPN , there exist:

- Non-interactive non-malleable commitments w.r.t. commitment for $\epsilon \log \log n$ tags.
- Non-interactive non-malleable commitments w.r.t. replacement for 2^n tags, additionally assuming sub-exponentially secure NIWIs for NP .

1.2 Applications and Directions for Future Work

As mentioned above, our final result (for 2^n tags) satisfies non-malleability w.r.t. replacement. In what follows, we give applications of this notion. Prior to our work, these were only known under strong non-standard assumptions [5].

Applications to Other Notions of Commitment

- **Non-malleability w.r.t. Opening.** As previously mentioned, non-malleable commitments w.r.t. replacement imply non-malleable commitments w.r.t. opening, as defined in [8,16]. Therefore, we obtain the first non-interactive non-malleable commitments w.r.t. opening from well-studied assumptions.

Informal Theorem 4 *Assuming sub-exponential hardness of discrete log or factoring, sub-exponential quantum hardness of LWE or LPN, and sub-exponentially secure NIWIs, there exist non-interactive non-malleable commitments w.r.t. opening (for 2^n tags).*

- **CCA Secure Commitments.** It was observed by [7] that the definitions of (one-to-one) non-malleability w.r.t. commitment and (one-to-one) CCA-security are equivalent in the non-interactive setting. We observe that in a similar way, non-malleability w.r.t. replacement implies a weaker notion of one-to-one CCA-security, where if the adversary queries the CCA oracle with a commitment to an invalid value, the oracle self-destructs.
- **Restricted Adversaries.** When restricted to adversaries that only output valid commitments, the notions of non-malleability w.r.t. replacement and non-malleability w.r.t. commitment are equivalent. Therefore, non-malleable commitments w.r.t. replacement can be combined with an appropriate ZK proof of validity of the commitment (as is implicit in [19,9]) to obtain non-malleable commitments w.r.t. commitment. For instance, (sub-exponential) NIWI and (sub-exponential) keyless collision resistant hash functions against *uniform adversaries* are known to imply one-message zero-knowledge with soundness against uniform (sub-exponential time) adversaries [4,29], and admitting a non-uniform simulator. Combining these with our non-malleable commitments w.r.t. replacement, we have the following theorem.

Informal Theorem 5 *Assuming sub-exponential hardness of discrete log or factoring against non-uniform adversaries, sub-exponential quantum hardness of LWE or LPN against non-uniform adversaries, sub-exponentially secure keyless collision-resistant hash functions against uniform adversaries, and sub-exponentially secure NIWIs against uniform adversaries, there exist non-interactive non-malleable commitments w.r.t. commitment against uniform adversaries.*

In a similar way, our commitment with can be appended with one-message ZK arguments of validity of the commitments, against any restricted class of adversaries, to yield non-malleable commitments w.r.t. commitment, against the same restricted classes of adversaries.

Other Applications

- **Upgrading NIZKs.** Non-interactive non-malleable commitments can also be used to upgrade NIZKs to satisfy a form of simulation soundness, without modifying the CRS. Informally, to give a simulation sound NIZK for statement x with witness w , the prover can generate a non-malleable commitment to w with tag x , and provide a (standard) NIZK proof that the commitment is a valid commitment to a witness for x . Note that non-malleability with respect to replacement suffices for this application because the NIZK can be used to provide a proof of validity of the commitment.

- **Block-wise Non-Malleable Codes.** Non-interactive non-malleable commitments w.r.t. opening are known to be equivalent to block-wise non-malleable codes [8] with two blocks. Block-wise non-malleable codes are a strengthening of the notion of split-state non-malleable codes. Using our result, we obtain the first block-wise non-malleable codes that only require two blocks (or states), based on well-studied assumptions.

Informal Theorem 6 *Assuming sub-exponential hardness of discrete log or factoring, sub-exponential quantum hardness of LWE or LPN, and sub-exponentially secure NIWIs, there exist 2-block blockwise non-malleable codes.*

Directions for Future Work

- **MPC.** Non-malleable commitments w.r.t. replacement are known to be sufficient for MPC [15]. We believe that our constructions of non-malleable commitments w.r.t. replacement will help obtain constructions of two-message concurrent secure computation against malicious adversaries (with super-polynomial simulation) from well-studied assumptions. A detailed exploration is beyond the scope of this work.
- **Non-Malleable Cryptographic Primitives.** The recent works of [25,12] give constructions of non-malleable point obfuscation and non-malleable digital lockers from strong variants of the DDH assumption. We believe that our commitments will find applications to achieving non-malleability in context of witness encryption, obfuscation and many other inherently non-interactive primitives, based on well-studied assumptions.

1.3 Prior work

The work of [29] constructed non-interactive non-malleable commitments w.r.t. commitment against a restricted class of *uniform* adversaries, assuming sub-exponentially secure time-lock puzzles, sub-exponential NIWI and sub-exponential collision-resistant hash functions against uniform adversaries. A very recent independent work [1] constructs an object called non-interactive quasi-non-malleable commitment (w.r.t. commitment), based on well-studied assumptions. This guarantees security against adversaries running in a-priori bounded polynomial time $O(n^c)$, but allows honest parties to run in longer (polynomial) time.

In this paper, our focus is on the non-interactive setting in the plain model against non-uniform adversaries with arbitrary polynomial running time. In this setting, constructions of non-malleable commitments have remained elusive, except based on non-standard assumptions. In particular, prior to our work, there were only two known constructions, described below.

Pandey et al. [31] constructed non-interactive concurrent non-malleable commitments w.r.t. commitment, starting from a non-falsifiable assumption, that already incorporates a strong form of non-malleability called *adaptive* injective one-way functions. Very recently, Bitansky and Lin [5] constructed concurrent non-interactive non-malleable commitments w.r.t. commitment, based on the

(relatively new, non-standard) assumption that there exist sub-exponential incompressible functions, sub-exponentially secure NIWI proofs, and either sub-exponential injective one-way functions that admit hardness amplification or sub-exponential time-lock puzzles.

Non-Interactive Tag Amplification. Tag amplification has been extensively studied in the non-malleability literature (e.g. [11,27,36,29,5]). Of these, only the recent work of [5] considers tag amplification in the non-interactive setting against general adversaries. They make a relatively non-standard assumption about the existence of sub-exponential incompressible one-way functions, in addition to assuming the existence of a sub-exponentially secure NIWI proofs. Using this incompressibility assumption, they construct a variant of one-message ZK proofs with weak soundness guarantees, and they use this variant of ZK to emulate techniques used in prior work for tag amplification.

On the other hand, our tag amplification technique only assumes the existence of a sub-exponentially secure NIWI proof, and is therefore substantially different from prior techniques for tag amplification (all of which crucially required ZK). However, while our tag amplification technique yields commitments that are non-malleable w.r.t. replacement, the one in [5] yields commitments that are (concurrent) non-malleable w.r.t. commitment.

2 Overview of Our Techniques

We now provide an informal overview of our techniques.

2.1 Non-Malleable Commitments w.r.t. Commitment for $\epsilon \log \log n$ Tags

As discussed earlier, we realize sub-exponential adaptive commitments for two tags based on sub-exponential quantum hardness of LWE/LPN and sub-exponential hardness of factoring/discrete log. We now describe how we use these to obtain non-malleable commitments for a small number of tags ($\epsilon \log \log n$ tags where $\epsilon > 0$ is a small constant), which satisfy many-to-many same-tag non-malleability w.r.t. commitment. We give a formal construction of non-malleable commitments for $\epsilon \log \log n$ tags, and its proof in Section 4.

Assume the existence of adaptive commitments $\text{com}_0, \text{com}_1$, and oracles $\mathcal{O}_0, \mathcal{O}_1$ such that com_0 is *sub-exponentially* hard to invert given oracle \mathcal{O}_1 , but com_1 is invertible in the presence of \mathcal{O}_1 . Similarly, com_1 is *sub-exponentially* hard to invert given oracle \mathcal{O}_0 , but com_0 is invertible in the presence of \mathcal{O}_0 .

We show that from any such adaptive commitments, one can use complexity leveraging to derive a sequence of (bit) commitments $\{\text{com}_{d,i}\}_{d \in \{0,1\}, i \in [\zeta]}$, where $\zeta = \epsilon \log \log n$ for a small constant $0 < \epsilon < 1$, and where

$$\text{com}_{d,i} : \{0, 1\} \times \{0, 1\}^{\ell_{d,i}(n)} \rightarrow \{0, 1\}^*$$

such that for each $d \in \{0, 1\}$,

$$\ell_{d,1} = \omega(\log n) < \ell_{d,2} < \dots < \ell_{d,\zeta-1} < \ell_{d,\zeta} \triangleq n$$

and for every $i, j, k \in [\zeta]$ for which $k > i$, inverting $\text{com}_{d,k}$ relative to the oracle \mathcal{O}_{1-d} requires more time than jointly inverting $\text{com}_{d,i}$ and $\text{com}_{1-d,j}$, relative to the oracle \mathcal{O}_{1-d} . A variant of this technique was used by Pass and Wee [35].

Construction. In order to commit to a bit b with $\text{tag} \in [\zeta]$, the committer first XOR secret shares the bit b to obtain two shares b_1 and b_2 . The commitment to b simply consists of $(\text{com}_{0,\text{tag}}(b_1), \text{com}_{1,\zeta-\text{tag}}(b_2))$.

Analysis. Suppose there exists a MIM (adversary) that on input a commitment to a bit b w.r.t. tag tag , commits to a related bit b' w.r.t. $\widetilde{\text{tag}} \neq \text{tag}$. We have the following possibilities:

- If $\text{tag} > \widetilde{\text{tag}}$, then breaking $\text{com}_{0,\text{tag}}$ relative to oracle \mathcal{O}_1 is harder than jointly breaking $\text{com}_{0,\widetilde{\text{tag}}}$ and $\text{com}_{1,\zeta-\widetilde{\text{tag}}}$ relative to \mathcal{O}_1 .
- If $\text{tag} < \widetilde{\text{tag}}$, then breaking $\text{com}_{1,\zeta-\text{tag}}$ relative to \mathcal{O}_0 is harder than jointly breaking $\text{com}_{0,\widetilde{\text{tag}}}$ and $\text{com}_{1,\zeta-\widetilde{\text{tag}}}$ relative to \mathcal{O}_0 .

In the first case, we extract the bit b' committed by the MIM by jointly breaking $\text{com}_{0,\widetilde{\text{tag}}}$ and $\text{com}_{1,\zeta-\widetilde{\text{tag}}}$ relative to \mathcal{O}_1 , and if b' is related to b , we get a contradiction to the hardness of breaking $\text{com}_{0,\text{tag}}$ relative to \mathcal{O}_1 . We can use a similar argument in the second case.

We also observe that we can allow the MIM to generate an arbitrary number of commitments on the right with the same $\widetilde{\text{tag}}$, and rely on the same assumptions to argue that the joint distribution of bits committed by the MIM (in many right commitments) remains independent of the honest bit. This gives us many-to-many same tag non-malleable commitments w.r.t. commitment for $\epsilon \log \log n$ tags. For simplicity, we only focused on bit commitments in this overview. However it is easy to extend this construction to obtain string commitments for $\epsilon \log \log n$ tags, based on sub-exponential adaptive bit commitments for two tags.

2.2 Non-interactive Tag Amplification

Our starting point is the following basic idea. Start with a non-malleable commitment scheme com for tags in $[\alpha]$ where $\alpha \leq \text{poly}(n)$, and obtain a scheme Com for tags in $[2^{\alpha/2}]$, as follows: To commit to a message m w.r.t. a tag T , first compute $\{t_1, t_2, \dots, t_{\alpha/2}\}$, such that each $t_i = (i||T_i)$ where T_i denotes the i^{th} bit of T^4 . Let

$$\text{Com}_T(m) \triangleq \{\text{com}_{t_i}(m)\}_{i \in [\alpha/2]}.$$

Note that for any two tags $T = \{t_1, t_2, \dots, t_{\alpha/2}\}$ and $\widetilde{T} = \{\widetilde{t}_1, \widetilde{t}_2, \dots, \widetilde{t}_{\alpha/2}\}$ such that $\widetilde{T} \neq T$, there exists at least one index i such that $\widetilde{t}_i \notin \{t_1, t_2, \dots, t_{\alpha/2}\}$. Therefore, if the underlying com is $\alpha/2$ -to-1 non-malleable, then given $\text{Com}_T(m) = \{\text{com}_{t_i}(m)\}_{i \in [\alpha/2]}$, it should be hard to generate $\text{com}_{\widetilde{t}_i}(m')$ for a related message m' . Therefore, an adversary cannot generate a *valid* commitment $\text{com}_{\widetilde{T}}(\widetilde{m})$

⁴ Our actual encoding of T to $\{t_1, t_2, \dots, t_{\alpha/2}\}$ is slightly more sophisticated, but achieves the same effect.

to a related message \tilde{m} , i.e., that the resulting scheme is non-malleable w.r.t. replacement.

However, the security of this scheme completely breaks down even if the adversary receives *two commitments*. Specifically, an adversary that receives two commitments $\text{Com}_T(m)$ and $\text{Com}_{T'}(m)$ with different tags $T = \{t_1, t_2, \dots, t_{\alpha/2}\}$ and $T' = \{t'_1, t'_2, \dots, t'_{\alpha/2}\}$, can easily output $\text{Com}_{\tilde{T}}(m)$, where $\tilde{T} = \{t_1, \dots, t_{\alpha/4}, t'_{\alpha/4+1}, \dots, t'_{\alpha/2}\}$. In other words, the resulting scheme *does not satisfy* many-to-1 non-malleability (or even 2-to-1 non-malleability), and is only non-malleable in the 1-to-1 setting.

Thus, using this idea we can go from $\eta \log \log n$ tags to $2^{\frac{\eta}{2} \log \log n} = \log^{\frac{\eta}{2}} n$ tags, but cannot continue further, since this compiler uses an underlying commitment which is many-to-one non-malleable (or more specifically, $\alpha/2$ -to-1 non-malleable).

The blueprint in Khurana and Sahai [24] describes how this problem can be solved using a NIZK argument, which requires the existence of a common random string (which we want to avoid). Namely, they show that if we append to the commitment $C = \{\text{com}_{t_i}(m)\}_{i \in [\alpha/2]}$ a NIZK proof that all these $\alpha/2$ commitments com_{t_i} are to the same message m , then one can indeed prove that this resulting scheme is many-to-one non-malleable⁵. Instead, in this work, we rely on non-interactive proofs satisfying a weaker hiding property, i.e., witness indistinguishability⁶. This introduces several problems that do not come up when using NIZKs. In particular, techniques in [24] rely on the reduction’s ability to generate “simulated” proofs, a notion that is not applicable when using NIWIs. We discuss these barriers in further detail below.

Tag Amplification using NIWIs: First Stab. While NIWI proofs have been extremely useful in a wide variety of cryptographic settings, they often become meaningless when trying to prove NP statements that have a single witness, such as the one described above. Typically, NIWI proofs are only useful for statements that have at least two independent witnesses.

One can create a statement with two independent witnesses by repeating the blueprint twice in parallel. Namely, commit to a message m by computing $C_1 = \{\text{com}_{t_i}(m; r_{i,1})\}_{i \in [\alpha/2]}$, $C_2 = \{\text{com}_{t_i}(m; r_{i,2})\}_{i \in [\alpha/2]}$ where $\{r_{i,b}\}_{i \in [\alpha/2], b \in \{0,1\}} \stackrel{\$}{\leftarrow} \{0,1\}^*$, and add a NIWI proving that all the commitments, in either C_1 or C_2 , are to the same message.

Indeed, one can easily prove that if the underlying scheme for α tags is $(\alpha/2)$ -to-1 non-malleable, then the resulting scheme is one-to-one non-malleable w.r.t. replacement (which was the case even before we started using NIWIs)⁷. Unfortu-

⁵ To be precise, they need to rely on the fact that the NIZK is “more secure” than the underlying commitment scheme.

⁶ As with NIZKs used in [24], we also require our NIWI to be more secure than the underlying commitment, which results in a sub-exponential assumption on the NIWI.

⁷ On the other hand, if we used a NIZK, the resulting scheme would be many-to-1 non-malleable w.r.t. commitment.

nately, it is not clear if the resulting scheme satisfies even 2-to-1 non-malleability (w.r.t. replacement). Roughly speaking, the problem is as follows. For simplicity, consider a MIM that obtains commitments which are both commitments to m_1 or both to m_2 , and tries to copy m_1 (or m_2). A natural approach to rule out such a MIM would be to rely on an intermediate hybrid, in which the MIM obtains a commitment to (m_1, m_2) .⁸ Unfortunately, we have no way to use a hybrid argument to rule out a MIM that does the following:

- In the first hybrid, on input commitments to (m_1, m_1) , outputs a (valid) commitment to m_1 .
- In the intermediate hybrid, on input commitments to (m_1, m_2) , outputs an invalid commitment where the first repetition in the MIM’s commitment consists of all commitments to m_1 , and the second repetition consists of all commitments to m_2 , and these commitments are accompanied with an accepting NIWI proof.
- In the final hybrid, on input commitments to (m_2, m_2) , outputs a (valid) commitment to m_2 .

The problem is that neither of the two pairs of adjacent hybrids can be used to get a contradiction to the one-to-one non-malleability, because neither are violating the non-malleability criterium w.r.t. replacement⁹.

However, as we already noted above, many-to-one non-malleability is essential if we want to use the compiler again. In fact, it may seem like the NIWIs were not useful at all, since we could get one-to-one non-malleability even for the basic scheme described at the beginning of this overview, which did not require any NIWI (or NIZK). While at first, this approach seems to be inherently problematic, we will now describe how we can nevertheless rely on NIWIs to obtain our desired compiler, as follows

Overview of Our Compiler. Our idea is to have each commitment consist of $(\ell + 1)$ repetitions (as opposed to only 2), where ℓ is the number of commitments that the adversary can receive (on the left).

Namely, our new (outer) commitment scheme will consist of a matrix of (inner) commitments corresponding to the underlying small tag commitment scheme. This matrix contains $(\ell + 1)$ rows, corresponding to each of the repetitions, and $(\alpha/2)$ columns, corresponding to the small tags of the underlying scheme. The honest committer generates *all* $(\alpha/2) \cdot (\ell + 1)$ inner commitments to the same message (with independent randomness). Additionally, the committer is required to provide a NIWI proof that ℓ out of the $(\ell + 1)$ rows satisfy the following property: The message committed using the inner commitment scheme across all $\alpha/2$ tags is identical for this row (but this message is not required to be identical across different rows).

⁸ This is the standard approach used in all previous work on this topic.

⁹ This problem can be avoided by relying on NIZKs which would prevent the MIM from behaving as in the intermediate hybrid. However, we cannot rely on NIZKs because they require a CRS.

Now, let us perform the same hybrid argument as above, where in the j^{th} hybrid, we change the j^{th} left outer commitment from a commitment to m_1 to a commitment to m_2 . Then, for the outer commitment output by the MIM, which is an $(\ell + 1) \times (\alpha/2)$ matrix of inner commitments, the following must be true.

1. Recall that at least one small tag of the MIM differs from *every* small tag used in the j^{th} left outer commitment. Therefore, by the non-malleability of the underlying commitment, the value committed by the MIM *in all inner commitments accross at least one column* (corresponding to this differing small tag) does not change.
2. Moreover, by the soundness of the NIWI provided by the MIM, at least ℓ of the rows satisfy the following property: the values committed across all $\alpha/2$ tags is identical for this row.

Combining (1) and (2) implies that the values committed by the MIM across at least ℓ rows do not change. In other words, the MIM may change the values committed in *at most one row in every hybrid*.

But since there are $(\ell + 1)$ rows and only ℓ hybrids, we deduce that there exists at least one row for which the messages remained unchanged at the end of *all ℓ hybrids*¹⁰. Therefore, no adversary can commit to a *valid* message that is related to the messages committed to in the left executions.

We show that this compiler works even if the underlying scheme is non-malleable w.r.t. replacement (as opposed to being non-malleable w.r.t. commitment). However, there is a loss in parameters when applying this compiler, i.e., the compiler converts any ℓ -to- z non-malleable commitment w.r.t. replacement into an ℓ' -to- z' non-malleable commitment w.r.t. replacement, where ℓ' and z' are smaller than ℓ and z . We do not discuss exact parameter constraints here, but refer the reader to Theorem 3 for details. We will give a more detailed explanation in Section 2.2.

Technical Bottlenecks. The intuition above seems to imply that the adversary cannot convert a commitment to m into a commitment to a related message m' . Proving this formally requires overcoming many technical difficulties. Specifically, the definition of non-malleability w.r.t. replacement¹¹, requires that there exist an (inefficient) extractor $\mathcal{V}_{\text{Real}}$ that extracts the message committed by the adversary from a transcript of a “real” experiment with honest messages (m_1, \dots, m_ℓ) , and an (inefficient) extractor $\mathcal{V}_{\text{Ideal}}$ that extracts the message committed by the adversary from a transcript of an “ideal” experiment with honest messages $(0, 0, \dots, 0)$, such that the joint distribution of the view of the MIM in the real experiment and the values output by $\mathcal{V}_{\text{Real}}$, is indistinguishable from the joint distribution of the view of the MIM in the ideal experiment and the values output by $\mathcal{V}_{\text{Ideal}}$. Furthermore, whenever the MIM generates a “valid” commitment \tilde{c} to a message \tilde{m} in either the real or ideal experiment, $\mathcal{V}_{\text{Real}}$ and $\mathcal{V}_{\text{Ideal}}$ are

¹⁰ To simplify our proof, we rely on 10ℓ repetitions (instead of $\ell + 1$) repetitions, to ensure that the messages in *most* repetitions remain unchanged.

¹¹ We refer the reader to Definition 3 for a one-to-one definition, and Definition 2 for a many-to-many definition.

required to output \tilde{m} . Whenever the message committed by the MIM is invalid, we impose no restrictions on the output of $\mathcal{V}_{\text{Real}}$ and $\mathcal{V}_{\text{Ideal}}$. To formally prove security, we will need to define these extractors $\mathcal{V}_{\text{Real}}$ and $\mathcal{V}_{\text{Ideal}}$, and ensure that their output distributions remain indistinguishable.

It is tempting to define $\mathcal{V}_{\text{Real}}$ and $\mathcal{V}_{\text{Ideal}}$ to output \tilde{M} corresponding to the MIM’s commitment string \tilde{c} , if there exists \tilde{r} such that $\tilde{c} = \text{com}(\tilde{M}, \tilde{r})$, and otherwise output \perp . However, as observed by the intuition above, these distributions will not necessarily be indistinguishable.¹² Namely, the adversary may generate valid commitments when given commitments to m and commit to \perp when given commitments to 0.

Intuitively, to make these distributions indistinguishable, we will introduce some “slack”, and sometimes output a valid message even though the adversary did not commit to a “perfectly valid” message. The question is the following: Suppose that the adversary outputs a commitment that is “close to” being a valid commitment to a message \tilde{m} . Should the extractors $\mathcal{V}_{\text{Real}}$ or $\mathcal{V}_{\text{Ideal}}$ output \tilde{m} or output \perp ? This is precisely where the leakage lemma of Gentry and Wichs [13] plays a crucial role. More specifically, we define a function π that outputs the decision bit of whether to output \perp , or to output one of the extracted messages (and also specifies which of the extracted messages should be output). This function is inefficient.

Now informally, the leakage lemma states that for every two indistinguishable distributions (X, Y) and every unbounded leakage function π , there exists a relatively efficient simulator that outputs a leakage π' such that $(X, \pi(X))$ is indistinguishable from $(Y, \pi'(Y))$.

In our context, the decision of whether to output \tilde{m} or output \perp in any particular hybrid will be dictated by the leakage lemma. Specifically, we will rely on the lemma where X and Y correspond to the view of the MIM in two consecutive hybrids, and where π is the leakage function described above. The leakage lemma will help us “carry over” this leakage across indistinguishable hybrids in a relatively efficient manner. The proof of non-malleability of this amplification step is the primary technical contribution of our paper.

There are many additional technical subtleties that were not discussed. For instance, in order to argue that the compiler can be applied several times, we work with a strong variant of non-malleability w.r.t. replacement (which only strengthens our final result). We give a more detailed protocol description towards the end of this section, and also refer the reader to Section 5 for details of the construction.

Putting Things Together. We now describe how we use this compiler to obtain our final result, i.e. non-malleable commitments for 2^n tags. Our starting point is our scheme for $\eta \log \log n$ tags which is many-to-many same-tag non-malleable w.r.t. commitment, and in particular is many-to-many same-tag non-malleable w.r.t. replacement (we give an overview of this scheme in Section 2.1). We will

¹² We note that these distributions are indeed indistinguishable if the adversary always generates valid commitments.

use the compiler above *three times*: First we convert the scheme for $\eta \log \log n$ tags into a scheme for $\log^{\eta/2}(n)$ tags, then we convert the resulting scheme for $\log^{\eta/2}(n)$ tags into a scheme for $2^{\log^\epsilon n}$ tags (for a small constant $\epsilon > 0$). We apply the compiler one final time to the scheme for $2^{\log^\epsilon n}$ tags to get a scheme for $\omega(n^{\log n})$ tags.

We note that it is not clear that we can run the compiler on itself many times, since every time we run the compiler, there is a loss in parameters. However, we set parameters carefully so that this nevertheless goes through.

To go from $n^{\log n}$ tags to 2^n tags, we use the (standard) idea of relying on sub-exponentially secure signatures. Specifically, to commit to a message m with tag $T \in [2^n]$, we generate a random pair sk, vk of signing and verification keys for the underlying scheme, where the verification key is of length $\log^2 n$ bits. We use vk as our “small” tag for the non-malleable commitment, and sign the larger tag $T \in [2^n]$ with sk . The security of this construction follows by the (sub-exponential) unforgeability of the underlying signature scheme. We refer the reader to Section 6 for more details.

More Detailed Protocol Description. Finally, to help the reader navigate our tag amplification protocol, we now give a slightly more detailed description of our protocol, and the intuition for non-malleability. As mentioned above, to commit to a message m with tag $T = \{t_1, t_2, \dots, t_{\alpha/2}\}$, the committer commits to the message $k = 10\ell$ times in parallel with tags $\{t_1, t_2, \dots, t_{\alpha/2}\}$, using fresh randomness each time.

Our protocol is described informally in Figure 1. Note that the resulting commitment is not many-to-many, because as explained above, even for ℓ -to-1 non-malleability, the size of the resulting commitment grows linearly with ℓ .

Roughly, we prove that if our underlying commitment scheme com is many-to- z non-malleable w.r.t. replacement, and is secure against 2^y -sized adversaries, then the resulting scheme is ℓ -to- y same-tag non-malleable w.r.t. replacement, for any y and ℓ such that $\ell \cdot y < \frac{z}{10}$. We require the NIWI to be WI against $\text{poly}(T)$ -time adversaries, where T is the time required to brute-force break com .

Intuition for Non-Malleability. For simplicity, let us consider a MIM that on input ℓ commitments, with corresponding tags T_1, T_2, \dots, T_ℓ , outputs a single commitment \tilde{c} with tag \tilde{T} (in our actual proof, the MIM is allowed to output multiple commitments, albeit using the same tag).

We need to argue that the MIM on input ℓ commitments to messages m_1, \dots, m_ℓ cannot output a *valid* commitment to a related message \tilde{m} . As eluded to earlier, this is done via hybrids. Let us suppose for contradiction that on input commitments to m_1, \dots, m_ℓ , the adversary outputs a valid commitment to \tilde{m} .

We consider a hybrid where the first honest commitment (on the left) is generated as a commitment to 0 (but the rest are commitments to m_2, \dots, m_ℓ). Letting $T_1 := \{t_{1,1}, t_{1,2}, \dots, t_{1,\alpha/2}\}$, one can argue that the distribution of the message \tilde{m} committed by the MIM in the column corresponding small tag $\tilde{t}_1 \notin \{t_{1,1}, t_{1,2}, \dots, t_{1,\alpha/2}\}$ cannot change in all k rows. This follows from the many-to- z

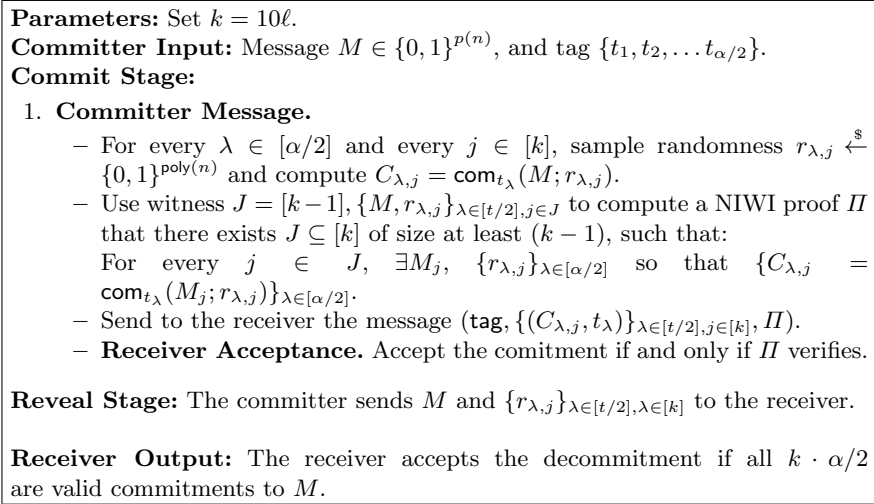


Fig. 1. Round-Preserving Tag Amplification

non-malleability w.r.t. commitment of com for $z \geq k$ (and relying on the fact that NIWI is hard against $\text{poly}(T)$ -time adversaries).

Furthermore, by the soundness of the MIM's NIWI, this implies that the MIM continues to commit to \tilde{m} in at least $(k-1)$ of the rows. This implies that the MIM continues to commit to \tilde{m} in least $(k-1)$ of the rows, *for every tag*. the MIM continues to commit to \tilde{m} in at least $(k-2)$ of the rows, for every tag.

Continuing this way, we observe that the MIM continues to commit to \tilde{m} in at least $(k-\ell)$ of the rows, w.r.t. every tag, on input ℓ commitments to messages $(0, 0, \dots, 0)$. Therefore on input $(0, 0, \dots, 0)$, the MIM either continues to commit to \tilde{m} or commits to an invalid value, and therefore, \tilde{m} must be unrelated to m . This is the key intuition for the security of our scheme.

As explained above, the actual analysis of indistinguishability of the *joint distribution* of the protocol transcript and messages committed by the MIM is quite involved, and requires multiple careful applications of the leakage lemma [13]. We refer the reader to Section 5 for details.

3 Definitions

Let n denote the security parameter. In all our definitions, the input message to the commitment scheme will be sampled from $\{0, 1\}^p$ for a polynomially bounded function $p = p$.

For any $T = T(n)$, we use $\mathcal{X} \approx_{\text{poly}(T(n))} \mathcal{Y}$ to denote two distributions such that for every $(T(n))^{O(1)}$ -size distinguisher \mathcal{D} ,

$$\Pr[\mathcal{D}(x) = 1 | x \xleftarrow{\$} \mathcal{X}] - \Pr[\mathcal{D}(x) = 1 | x \xleftarrow{\$} \mathcal{Y}] = \text{negl}(n).$$

We denote by $\mathcal{X} \approx \mathcal{Y}$, the event that $\mathcal{X} \approx_{\text{poly}(n)} \mathcal{Y}$.

3.1 Non-Malleable Commitments w.r.t. Replacement

In this section, we present the main definition of non-malleability that we achieve, which is known as *non-malleability w.r.t. replacement* ([15]). This definition is weaker than the original definition of non-malleability, which is known as *non-malleability with respect to commitment* (and is formally defined in Section 3.2).

Non-malleability considers a man-in-the-middle that receives a commitment to a message $m \in \{0, 1\}^p$ and generates a new commitment \tilde{c} . We say that the man-in-the-middle commits to \perp if there does not exist any (\tilde{m}, \tilde{r}) such that $\tilde{c} = \text{com}(\tilde{m}; \tilde{r})$. Intuitively, the definition of non-malleability with respect to commitment requires that for any two messages $m_0, m_1 \in \{0, 1\}^p$, the joint distributions of $(\text{Com}(m_0), \tilde{m}_0)$ and $(\text{Com}(m_1), \tilde{m}_1)$ are indistinguishable, where \tilde{m}_b is the message committed to by the MIM given $\text{Com}(m_b)$. The definition of non-malleability w.r.t. replacement (that we achieve) intuitively requires this to hold only conditioned on $\tilde{m}_0, \tilde{m}_1 \neq \perp$.

We emphasize that we consider the case where the MIM gets a single committed message and generates a single commitment. This is known as the “one-to-one” definition. A stronger definition is the “many-to-many” definition (also known as concurrent non-malleability), where the MIM receives many commitments and is allowed to generate many commitments, and the guarantee is that for any two sets of messages committed to and sent to the MIM, the joint distribution of these commitments and the messages committed to by the MIM, are indistinguishable.

Definition 1 (Non-Malleable Commitments w.r.t. Replacement). *A non-interactive non-malleable (one-to-one) string commitment scheme with N tags consists of a probabilistic poly-time algorithm \mathcal{C} , that takes as input a message $m \in \{0, 1\}^p$, randomness $r \in \{0, 1\}^{\text{poly}(n)}$, and a $\text{tag} \in [N]$, and outputs a commitment $\text{com}_{\text{tag}}(m; r)$. It is said to be non-malleable w.r.t. replacement if the following two properties hold:*

1. **Statistical binding.** *There do not exist $m_0, m_1 \in \{0, 1\}^p$, $r_0, r_1 \in \{0, 1\}^{\text{poly}(n)}$ and $\text{tag}_0, \text{tag}_1 \in [N]$ such that $m_0 \neq m_1$ and $\text{com}_{\text{tag}_0}(m_0; r_0) = \text{com}_{\text{tag}_1}(m_1; r_1)$.*
2. **One-to-One Non-malleability.** *For any poly-size adversary \mathcal{A} , any $m \in \{0, 1\}^p$ and any $\text{tag} \in [N]$, there exist (possibly inefficient) functions $\mathcal{V}_{\text{Real}}$ and $\mathcal{V}_{\text{Ideal}}$ such that the following holds:*
 - (a) *Sample $r \xleftarrow{\$} \{0, 1\}^{\text{poly}(n)}$ and set $c = \text{com}_{\text{tag}}(m; r)$. Let $(\tilde{c}, z) = \mathcal{A}(c)$. If there exists $\tilde{\text{tag}} \in [N] \setminus \{\text{tag}\}$, $\tilde{M} \in \{0, 1\}^{p(n)}$ and $\tilde{r} \in \{0, 1\}^{\text{poly}(n)}$ such that $\tilde{c} = \text{com}_{\tilde{\text{tag}}}(\tilde{M}; \tilde{r})$ then $\tilde{m} = \tilde{M}$, otherwise no restrictions are placed on \tilde{m} . We require that*

$$\Pr[\mathcal{V}_{\text{Real}}(c, \tilde{c}) = \tilde{m}] = 1 - \text{negl}(n).$$
 - (b) *Sample $r_{\text{Ideal}} \xleftarrow{\$} \{0, 1\}^{\text{poly}(n)}$ and set $c_{\text{Ideal}} = \text{com}_{\text{tag}}(0^p; r_{\text{Ideal}})$. Let $(\tilde{c}_{\text{Ideal}}, z_{\text{Ideal}}) = \mathcal{A}(c_{\text{Ideal}})$. If there exists $\tilde{\text{tag}} \in [N] \setminus \{\text{tag}\}$, $\tilde{M}_{\text{Ideal}} \in \{0, 1\}^{p(n)}$ and $\tilde{r}_{\text{Ideal}} \in$*

$\{0, 1\}^{\text{poly}(n)}$ such that $\tilde{c}_{\text{Ideal}} = \text{com}_{\tilde{\text{tag}}}(\tilde{M}_{\text{Ideal}}; \tilde{r}_{\text{Ideal}})$ then $\tilde{m}_{\text{Ideal}} = \tilde{M}_{\text{Ideal}}$, otherwise no restrictions are placed on \tilde{m}_{Ideal} . We require that

$$\Pr[\mathcal{V}_{\text{Ideal}}(c_{\text{Ideal}}, \tilde{c}_{\text{Ideal}}) = \tilde{m}_{\text{Ideal}}] = 1 - \text{negl}(n).$$

(c) We require:

$$(c, \tilde{c}, z, \mathcal{V}_{\text{Real}}(c, \tilde{c})) \approx_c (c_{\text{Ideal}}, \tilde{c}_{\text{Ideal}}, z_{\text{Ideal}}, \mathcal{V}_{\text{Ideal}}(c_{\text{Ideal}}, \tilde{c}_{\text{Ideal}})).$$

over the randomness of sampling r, r_{Ideal} ¹³.

We next present an (intermediate) security definition that we use as a stepping stone to achieve our main result. This is a many-to-many version of Definition 1, that restricts the adversary to use the same tag in all commitments that he outputs.

Definition 2 (ℓ -to- y Same-tag Non-malleable Commitments w.r.t. Replacement). A non-interactive non-malleable commitment scheme with N tags consists of a probabilistic poly-time algorithm \mathcal{C} , that takes as input a message $m \in \{0, 1\}^p$, randomness $r \in \{0, 1\}^{\text{poly}(n)}$, and a $\text{tag} \in [N]$, and outputs a commitment $\text{com}_{\text{tag}}(m; r)$. It is said to be ℓ -to- y same-tag non-malleable w.r.t. replacement for polynomials $\ell(\cdot)$ and $y(\cdot)$, if the following two properties hold:

1. **Statistical binding.** There do not exist $m_0, m_1 \in \{0, 1\}^p$, $r_0, r_1 \in \{0, 1\}^{\text{poly}(n)}$ and $\text{tag}_0, \text{tag}_1 \in [N]$ such that $m_0 \neq m_1$ and $\text{com}_{\text{tag}_0}(m_0; r_0) = \text{com}_{\text{tag}_1}(m_1; r_1)$.
2. **ℓ -to- y Non-malleability.** For any poly-size adversary \mathcal{A} , any $m_1, \dots, m_\ell \in \{0, 1\}^p$, and any $\text{tag}_1, \dots, \text{tag}_\ell \in [N]$, there exist (possibly inefficient) functions $\mathcal{V}_{\text{Real}}$ and $\mathcal{V}_{\text{Ideal}}$ such that the following holds:

(a) Sample $r_1, \dots, r_\ell \xleftarrow{\$} \{0, 1\}^{\text{poly}(n)}$, set $c_i = \text{com}_{\text{tag}_i}(m_i; r_i)$ for every $i \in [\ell]$, and let $(\tilde{c}_1, \dots, \tilde{c}_y, z) = \mathcal{A}(c_1, \dots, c_\ell)$.

If there exists $\tilde{\text{tag}} \in [N] \setminus \{\text{tag}_i\}_{i \in [\ell]}$ such that $\tilde{c}_1, \dots, \tilde{c}_y$ all use $\tilde{\text{tag}}$, then continue. Otherwise set $(\tilde{m}_1, \dots, \tilde{m}_y) = \text{abort}$.

For each $i \in [y]$, if there exists $\tilde{M}_i \in \{0, 1\}^p$ and $\tilde{r}_i \in \{0, 1\}^{\text{poly}(n)}$ for which $\tilde{c}_i = \text{com}_{\tilde{\text{tag}}}(\tilde{M}_i; \tilde{r}_i)$, set $\tilde{m}_i = \tilde{M}_i$, and otherwise no restrictions are placed on \tilde{m}_i . We require that

$$\Pr[\mathcal{V}_{\text{Real}}(c_1, \dots, c_\ell, \tilde{c}_1, \dots, \tilde{c}_y) = (\tilde{m}_1, \dots, \tilde{m}_y)] = 1 - \text{negl}(n)$$

(b) Sample $r_{\text{Ideal}, 1}, \dots, r_{\text{Ideal}, \ell} \xleftarrow{\$} \{0, 1\}^{\text{poly}(n)}$, set $c_{\text{Ideal}, i} = \text{com}_{\text{tag}_i}(0^p; r_{\text{Ideal}, i})$ for every $i \in [\ell]$, and let $(\tilde{c}_{\text{Ideal}, 1}, \dots, \tilde{c}_{\text{Ideal}, y}, z_{\text{Ideal}}) = \mathcal{A}(c_{\text{Ideal}, 1}, \dots, c_{\text{Ideal}, \ell})$. If there exists $\tilde{\text{tag}} \in [N] \setminus \{\text{tag}_i\}_{i \in [\ell]}$ such that $\tilde{c}_{\text{Ideal}, 1}, \dots, \tilde{c}_{\text{Ideal}, y}$ all use $\tilde{\text{tag}}$, then continue. Otherwise set $(\tilde{m}_{\text{Ideal}, 1}, \dots, \tilde{m}_{\text{Ideal}, y}) = \text{abort}$.

For each $i \in [y]$, if there exists $\tilde{M}_{\text{Ideal}, i} \in \{0, 1\}^p$ and $\tilde{r}_{\text{Ideal}, i} \in \{0, 1\}^{\text{poly}(n)}$ for which $\tilde{c}_{\text{Ideal}, i} = \text{com}_{\tilde{\text{tag}}}(\tilde{M}_{\text{Ideal}, i}; \tilde{r}_{\text{Ideal}, i})$, set $\tilde{m}_{\text{Ideal}, i} = \tilde{M}_{\text{Ideal}, i}$, and otherwise no restrictions are placed on $\tilde{m}_{\text{Ideal}, i}$. We require that

$$\begin{aligned} \Pr[\mathcal{V}_{\text{Ideal}}(c_{\text{Ideal}, 1}, \dots, c_{\text{Ideal}, \ell}, \tilde{c}_{\text{Ideal}, 1}, \dots, \tilde{c}_{\text{Ideal}, y}) = (\tilde{m}_{\text{Ideal}, 1}, \dots, \tilde{m}_{\text{Ideal}, y})] \\ = 1 - \text{negl}(n) \end{aligned}$$

¹³ Note that this definition explicitly considers auxiliary information z , but is equivalent to one that does not consider z . We explicitly consider z for convenience.

(c) We require:

$$\begin{aligned} & ((c_1, \dots, c_\ell), (\tilde{c}_1, \dots, \tilde{c}_y), z, \mathcal{V}_{\text{Real}}(c_1, \dots, c_\ell, \tilde{c}_1, \dots, \tilde{c}_y)) \approx_c \\ & ((c_{\text{Ideal},1}, \dots, c_{\text{Ideal},\ell}), (\tilde{c}_{\text{Ideal},1}, \dots, \tilde{c}_{\text{Ideal},y}), z_{\text{Ideal}}, \\ & \mathcal{V}_{\text{Ideal}}(c_{\text{Ideal},1}, \dots, c_{\text{Ideal},\ell}, \tilde{c}_{\text{Ideal},1}, \dots, \tilde{c}_{\text{Ideal},y})) \end{aligned}$$

over the randomness of sampling r_1, \dots, r_ℓ and $r_{\text{Ideal},1}, \dots, r_{\text{Ideal},\ell}$.

In what follows, we define a slight strengthening of ℓ -to- y same-tag non-malleability w.r.t. replacement. Namely, in the definition below we allow the MIM to *obtain as input* some restricted auxiliary information on the honest messages and randomness.

Definition 3 (ℓ -to- y Same-tag Auxiliary-Input Non-malleable Commitments w.r.t. Replacement). A non-interactive non-malleable commitment scheme with N tags consists of a probabilistic poly-time algorithm \mathcal{C} , that takes as input a message $m \in \{0, 1\}^p$, randomness $r \in \{0, 1\}^{\text{poly}(n)}$, and a tag $\text{tag} \in [N]$, and outputs a commitment $\text{com}_{\text{tag}}(m; r)$. It is said to be ℓ -to- y same-tag auxiliary-input non-malleable w.r.t. replacement for polynomials $\ell(\cdot)$ and $y(\cdot)$, if the following two properties hold:

1. **Statistical binding.** There do not exist $m_0, m_1 \in \{0, 1\}^p$, $r_0, r_1 \in \{0, 1\}^{\text{poly}(n)}$ and $\text{tag}_0, \text{tag}_1 \in [N]$ such that $m_0 \neq m_1$ and $\text{com}_{\text{tag}_0}(m_0; r_0) = \text{com}_{\text{tag}_1}(m_1; r_1)$.
2. **ℓ -to- y Non-malleability.** There exists a function $t_V : \mathbb{N} \rightarrow \mathbb{N}$ such that the following holds.

Fix any messages $m_1, \dots, m_\ell \in \{0, 1\}^p$, any $\text{tag}_1, \dots, \text{tag}_\ell$, and any efficient auxiliary input functions $\text{aux}_1, \text{aux}_2, \dots, \text{aux}_\ell$, where for every $i \in [\ell]$, aux_i takes as input the commitments (c_1, \dots, c_ℓ) together with the messages and randomness used to compute $(c_1, \dots, c_{i-1}, c_{i+1}, \dots, c_\ell)$. Set $T_V(n) = 2^{t_V(n)}$. For every $\beta \in [\ell]$ define ℓ commitments $c_{\beta,1}, \dots, c_{\beta,\ell}$, where $c_{\beta,i} = \text{com}_{\text{tag}_i}(0^p; r_i)$ for every $i \in [\beta]$, and $c_{\beta,i} = \text{com}_{\text{tag}_i}(m_i; r_i)$ for every $i \in [\beta + 1, \ell]$, where $r_1, \dots, r_\ell \stackrel{\$}{\leftarrow} \{0, 1\}^{\text{poly}(n)}$.

Suppose that for every $\beta \in [0, \ell - 1]$,

$$\begin{aligned} & (c_{\beta,1}, \dots, c_{\beta,\ell}, \text{aux}_\beta(c_{\beta,1}, \dots, c_{\beta,\ell}, (0^p)_{\times(\beta-1)}, m_{\beta+1}, \dots, m_\ell, \\ & r_1, \dots, r_{\beta-1}, r_{\beta+1}, \dots, r_\ell)) \approx_{T_V(n)} \\ & (c_{\beta,1}, \dots, c_{\beta,\ell}, \text{aux}_{\beta+1}(c_{\beta,1}, \dots, c_{\beta,\ell}, (0^p)_{\times(\beta)}, m_{\beta+2}, \dots, m_\ell, \\ & r_1, \dots, r_\beta, r_{\beta+2}, \dots, r_\ell)) \end{aligned} \quad (1)$$

where $\text{aux}_0 \triangleq \text{aux}_\ell$.

Fix any polynomial-size adversary \mathcal{A} , and for every $\beta \in [0, \ell]$ let

$$\begin{aligned} & (\tilde{c}_{\beta,1}, \dots, \tilde{c}_{\beta,y}, z_\beta) = \\ & \mathcal{A}(c_{\beta,1}, \dots, c_{\beta,\ell}, \text{aux}_\beta(c_{\beta,1}, \dots, c_{\beta,\ell}, (0^p)_{\times(\beta-1)}, m_{\beta+1}, \dots, m_\ell, \\ & r_1, \dots, r_{\beta-1}, r_{\beta+1}, \dots, r_\ell)). \end{aligned}$$

We require that there exist (possibly inefficient) functions $\mathcal{V}_{\text{Real}}$ and $\mathcal{V}_{\text{Ideal}}$, each computable in time $T_V(n)$, such that:

(a) If there exists $\widetilde{\text{tag}} \in [N] \setminus \{\text{tag}_i\}_{i \in [\ell]}$ such that $\widetilde{c}_{0,1}, \dots, \widetilde{c}_{0,y}$ all use tag $\widetilde{\text{tag}}$, then continue. Otherwise set $(\widetilde{m}_1, \dots, \widetilde{m}_n) = \text{abort}$.

For each $i \in [y]$, if there exists $\widetilde{M}_i \in \{0,1\}^p$ and $\widetilde{r}_i \in \{0,1\}^{\text{poly}(n)}$ for which $\widetilde{c}_{0,i} = \text{com}_{\widetilde{\text{tag}}}(\widetilde{M}_i; \widetilde{r}_i)$, set $\widetilde{m}_i = \widetilde{M}_i$, and otherwise no restrictions are placed on \widetilde{m}_i . We require that

$$\Pr[\mathcal{V}_{\text{Real}}(c_{0,1}, \dots, c_{0,\ell}, \mathbf{a}_0, \widetilde{c}_{0,1}, \dots, \widetilde{c}_{0,y}) = (\widetilde{m}_1, \dots, \widetilde{m}_y)] = 1 - \text{negl}(n)$$

where $\mathbf{a}_0 = \text{aux}_0(c_{0,1}, \dots, c_{0,\ell}, m_1, \dots, m_{\ell-1}, r_1, \dots, r_{\ell-1})$.

(b) If there exists $\widetilde{\text{tag}} \in [N] \setminus \{\text{tag}_i\}_{i \in [\ell]}$ such that $\widetilde{c}_{\ell,1}, \dots, \widetilde{c}_{\ell,y}$ all use tag $\widetilde{\text{tag}}$, then continue. Otherwise set $(\widetilde{m}_1, \dots, \widetilde{m}_n) = \text{abort}$.

For each $i \in [y]$, if there exists $\widetilde{M}_i \in \{0,1\}^p$ and $\widetilde{r}_i \in \{0,1\}^{\text{poly}(n)}$ for which $\widetilde{c}_{\ell,i} = \text{com}_{\widetilde{\text{tag}}}(\widetilde{M}_i; \widetilde{r}_i)$, set $\widetilde{m}_i = \widetilde{M}_i$, and otherwise no restrictions are placed on \widetilde{m}_i . We require that

$$\Pr[\mathcal{V}_{\text{Ideal}}(c_{\ell,1}, \dots, c_{\ell,\ell}, \mathbf{a}_\ell, \widetilde{c}_{\ell,1}, \dots, \widetilde{c}_{\ell,y}) = (\widetilde{m}_1, \dots, \widetilde{m}_y)] = 1 - \text{negl}(n)$$

where $\mathbf{a}_\ell = \text{aux}_\ell(c_{\ell,1}, \dots, c_{\ell,\ell}, (0^p)_{\times(\ell-1)}, r_1, \dots, r_{\ell-1})$.

(c) We require:

$$((c_{0,1}, \dots, c_{0,\ell}), \mathbf{a}_0, (\widetilde{c}_{0,1}, \dots, \widetilde{c}_{0,y}), z_0, \mathcal{V}_{\text{Real}}(c_{0,1}, \dots, c_{0,\ell}, \mathbf{a}_0, \widetilde{c}_{0,1}, \dots, \widetilde{c}_{0,y})) \approx_c$$

$$((c_{\ell,1}, \dots, c_{\ell,\ell}), \mathbf{a}_\ell, (\widetilde{c}_{\ell,1}, \dots, \widetilde{c}_{\ell,y}), z_\ell, \mathcal{V}_{\text{Ideal}}(c_{\ell,1}, \dots, c_{\ell,\ell}, \mathbf{a}_\ell, \widetilde{c}_{\ell,1}, \dots, \widetilde{c}_{\ell,y}))$$

over the randomness of sampling $c_{0,1}, \dots, c_{0,\ell}$ and $c_{\ell,1}, \dots, c_{\ell,\ell}$.

Remark 1. One can strengthen these definitions, to require non-malleability to hold for any two sets of messages (m_1^1, \dots, m_ℓ^1) and (m_1^2, \dots, m_ℓ^2) , such that $\mathcal{V}_{\text{Real}}$ (as before) considers an experiment where the honest committer generates commitments to (m_1^1, \dots, m_ℓ^1) , whereas $\mathcal{V}_{\text{Ideal}}$ considers an experiment where the honest committer generates commitments to (m_1^2, \dots, m_ℓ^2) (instead of generating commitments to 0s). The proofs of Theorem 2 and Theorem 3 show that our constructions also satisfy this stronger definition.

3.2 Non-Malleable Commitments w.r.t. Commitment

We also consider the stronger definition of non-malleability with respect to commitment [33]. This definition is standard in the literature; it is sometimes considered in the many-to-many setting (known as concurrent non-malleability), where the adversary (man-in-the-middle) receives many commitments “on the left” and generates many commitments “on the right”. It is also sometimes considered in the one-to-one setting, where the man-in-the-middle receives a single commitment “on the left” and generates a single commitment “on the right”. In this paper, we use a variant where we require the MIM to use the same tag in all “right” commitments, and we refer to this as the many-to- k same-tag variant. This definition is used as a stepping stone to achieve our main result, and is omitted from this version due to space constraints.

4 Non-Malleable Commitments for Small Tags

In this section, we construct a many-to-many same-tag non-malleable commitment scheme w.r.t. commitment for $\zeta = \eta \cdot \log \log n$ tags, for a small enough constant $\eta > 0$, based on the following assumption.

Assumption 1 *There exist non-interactive bit commitments $\text{com}_0 : \{0, 1\} \times \{0, 1\}^n \rightarrow \{0, 1\}^{L(n)}$ and $\text{com}_1 : \{0, 1\} \times \{0, 1\}^n \rightarrow \{0, 1\}^{L(n)}$ with the following properties.*

1. **There exists an oracle relative to which com_0 is *sub-exponentially hiding*, but com_1 is extractable.** *There exists an (inefficient, possibly randomized) oracle \mathcal{O}_1 and a poly-size algorithm \mathcal{A}_1 such that for every $n \in \mathbb{N}$ and every $(m, r) \in \{0, 1\} \times \{0, 1\}^n$,*

$$\Pr[\mathcal{A}_1^{\mathcal{O}_1}(\text{com}_1(m; r)) = (m, r)] = 1 - \text{negl}(n).$$

where the probability is over the randomness of \mathcal{O}_1 . Moreover, on input any string c for which $\exists(m, r)$ such that $c = \text{com}_1(m; r)$, we require that $\mathcal{A}_1^{\mathcal{O}_1}$ output \perp .

Yet, there exists a constant $\delta > 0$ such that for every $n \in \mathbb{N}$, every $\text{poly}(2^{n^\delta})$ -size adversary \mathcal{A} , and every pair of messages m_1 and m_2 in $\{0, 1\}$,

$$\left| \Pr[\mathcal{A}^{\mathcal{O}_1}(\text{com}_0(m_1; r)) = 1] - \Pr[\mathcal{A}^{\mathcal{O}_1}(\text{com}_0(m_2; r)) = 1] \right| = \text{negl}(n),$$

where the probability is over $r \xleftarrow{\$} \{0, 1\}^n$ and over the randomness of \mathcal{O}_1 .

2. **There exists an oracle relative to which com_1 is *sub-exponentially hard to invert* but com_0 is invertible.** *There exists an (inefficient, possibly randomized) oracle \mathcal{O}_0 and a poly-size algorithm \mathcal{A}_0 such that for every $n \in \mathbb{N}$ and every $(m, r) \in \{0, 1\} \times \{0, 1\}^n$,*

$$\Pr[\mathcal{A}_0^{\mathcal{O}_0}(\text{com}_0(m; r)) = (m, r)] = 1 - \text{negl}(n)$$

where the probability is over the randomness of \mathcal{O}_0 . Moreover, on input any string c for which $\exists(m, r)$ such that $c = \text{com}_0(m; r)$, we require that $\mathcal{A}_0^{\mathcal{O}_0}$ output \perp .

Yet, there exists a constant $\delta > 0$ such that for every $n \in \mathbb{N}$, every $\text{poly}(2^{n^\delta})$ -size adversary \mathcal{A} , and every pair of messages m_1 and m_2 in $\{0, 1\}$,

$$\left| \Pr[\mathcal{A}^{\mathcal{O}_0}(\text{com}_1(m_1; r)) = 1] - \Pr[\mathcal{A}^{\mathcal{O}_0}(\text{com}_1(m_2; r)) = 1] \right| = \text{negl}(n),$$

where the probability is over $r \xleftarrow{\$} \{0, 1\}^n$ and over the randomness of \mathcal{O}_0 .

In the full version, we formally show that it suffices to instantiate com_0 as any commitment whose hiding is based on the sub-exponential hardness of factoring/discrete log or any other problem that is invertible given a BQP oracle,

and it suffices to instantiate com_1 as any commitment whose hiding holds against sub-exponential quantum adversaries.

We note that Assumption 1 can be used to derive a sequence of commitments, described below [35].

There exist inefficient (possibly randomized) oracles $\mathcal{O}_0, \mathcal{O}_1$, a small constant $\eta > 0$, and a sequence $\{\text{com}_{b,i}\}_{b \in \{0,1\}, i \in [\zeta]}$ of commitment functions, where $\zeta = \eta \cdot \log \log(n)$ and

$$\text{com}_{b,i} : \{0, 1\} \times \{0, 1\}^{\ell_{b,i}(n)} \rightarrow \{0, 1\}^{L(\ell_{b,i}(n))}$$

such that for each $b \in \{0, 1\}$,

$$\ell_{b,1} = \omega(\log n^{\log \log n}) < \ell_{b,2} < \dots < \ell_{b,\zeta-1} < \ell_{b,\zeta} \triangleq n$$

and for every $i, j, k \in [\zeta]$ such that $k > i$, inverting $\text{com}_{b,k}$ relative to the oracle \mathcal{O}_{1-b} requires more time than jointly inverting $\text{com}_{b,i}$ and $\text{com}_{1-b,j}$ relative to the oracle \mathcal{O}_{1-b} .

Formally, for every $b \in \{0, 1\}$ and every $i \in [\zeta - 1]$ there exists a $T_{b,i} \cdot \text{poly}(n)$ -size algorithm $\mathcal{A}_{b,i}$ such that for every $j \in [\zeta]$, every messages $m_1, m_2 \in \{0, 1\}$, every $r \in \{0, 1\}^{\ell_{b,i}}$ and $r' \in \{0, 1\}^{\ell_{1-b,j}}$,

$$\begin{aligned} \Pr \left[(\mathcal{A}_{b,i}^{\mathcal{O}_{1-b}}(\text{com}_{b,i}(m_1; r)) = (m_1, r)) \wedge (\mathcal{A}_{b,i}^{\mathcal{O}_{1-b}}(\text{com}_{1-b,j}(m_2; r')) = (m_2, r')) \right] \\ = 1 - \text{negl}(n), \end{aligned}$$

where the probability is over the randomness of \mathcal{O}_{1-b} . Moreover, on input any element outside the range of $\text{com}_{b,i}$ or $\text{com}_{1-b,j}$, $\mathcal{A}_{b,i}^{\mathcal{O}_{1-b}}$ outputs \perp .

Yet, for every $\text{poly}(T_{b,i})$ -size adversary \mathcal{A} and every $k > i$,

$$\left| \Pr[\mathcal{A}^{\mathcal{O}_{1-b}}(\text{com}_{b,k}(m_1; r)) = 1] - \Pr[\mathcal{A}^{\mathcal{O}_{1-b}}(\text{com}_{b,k}(m_2; r)) = 1] \right| = \text{negl}(n),$$

where the probability is over $r \leftarrow \{0, 1\}^{\ell_{b,k}(n)}$ and over the randomness of \mathcal{O}_{1-b} . An overview of the construction of this sequence of commitments, following the technique of [35], can be found in the full version of the paper.

Our construction of non-malleable commitments for $\zeta(n)$ tags To commit to a message $m = (m_1, \dots, m_p) \in \{0, 1\}$ with respect to tag , using randomness $(r_i, s_i, a_i)_{i \in [p]}$, where for every $i \in [p]$, $r_i, s_i \xleftarrow{\$} \{0, 1\}^{\ell_{0,\text{tag}}} \times \{0, 1\}^{\ell_{1,\zeta-\text{tag}}}$ and $a_i \xleftarrow{\$} \{0, 1\}$, our commitment algorithm is defined by:

$$\begin{aligned} \text{Com}_{\text{tag}} \left(m; (r_i, s_i, a_i)_{i \in [p]} \right) \\ = \left(\text{tag}, (\text{com}_{0,\text{tag}}(a_i; r_i))_{i \in [p]}, (\text{com}_{1,\zeta-\text{tag}}(m_i \oplus a_i; s_i))_{i \in [p]} \right). \end{aligned}$$

Theorem 2. *If Assumption 1 holds, then there exists a constant $\eta > 0$ such that Com_{tag} is a non-interactive many-to-many same-tag non-malleable commitment scheme w.r.t. commitment for $\zeta = \eta \cdot \log \log(n)$ tags, against all $2^{\text{poly}(\log n)}$ -size adversaries.*

Proof. The fact that Com is statistically binding follows from the fact that $\text{com}_{b,i}$ are all statistically binding, which in turn follows from the fact that com_0 and com_1 are statistically binding. We next argue that Com is many-to-many same-tag non-malleable w.r.t. commitment against all $2^{\text{poly}(\log n)}$ -size adversaries. To this end, it suffices to prove that it is 1-to-many same-tag non-malleable w.r.t. commitment against all $2^{\text{poly}(\log n)}$ -size adversaries. This follows by a hybrid argument of [30], which proves that any commitment scheme that satisfies the one-to-many definition also satisfies the many-to-many definition.

To prove non-malleability, fix a $2^{\text{poly}(\log n)}$ -size adversary \mathcal{A} , and fix any $k \leq \text{poly}(n)$. Given a message $m = (m_1, \dots, m_p) \in \{0, 1\}^p$,¹⁴ we consider the following distribution:

Choose at random $b \xleftarrow{\$} \{0, 1\}$ and $R \xleftarrow{\$} \{0, 1\}^{\text{poly}(n)}$. If $b = 0$ then let $c = \text{Com}_{\text{tag}}(0^p; R)$. If $b = 1$ then let $c = \text{Com}_{\text{tag}}(m; R)$. Let

$$(\widetilde{\text{tag}}, \widetilde{c}_1, \dots, \widetilde{c}_k) = \mathcal{A}(c).$$

Consider the joint distribution

$$(c, \widetilde{c}_1, \dots, \widetilde{c}_k, \widetilde{m}_1, \dots, \widetilde{m}_k),$$

where for every $i \in [k]$, if there exists $\widetilde{M}_i \in \{0, 1\}^p$ and randomness $R_i \in \{0, 1\}^{\text{poly}(n)}$ such that $\widetilde{c}_i = \text{com}_{\widetilde{\text{tag}}}(\widetilde{M}_i, R_i)$, then $\widetilde{m}_i = \widetilde{M}_i$; else $\widetilde{m}_i = \perp$.

To prove that this construction is secure, it suffices to prove that for every $2^{\text{poly}(\log n)}$ -size adversary \mathcal{D} and every message m ,

$$\Pr[\mathcal{D}(c, \widetilde{c}_1, \dots, \widetilde{c}_k, \widetilde{m}_1, \dots, \widetilde{m}_k) = b] = \frac{1}{2} + \text{negl}(n).$$

To prove this, it suffices to show that for every $2^{\text{poly}(\log n)}$ -size adversary \mathcal{D} and every message m , if $\Pr[\widetilde{\text{tag}} > \text{tag}] \geq \frac{1}{\text{poly}(n)}$ for some polynomial $\text{poly}(\cdot)$, then

$$\Pr[\mathcal{D}(c, \widetilde{c}_1, \dots, \widetilde{c}_k, \widetilde{m}_1, \dots, \widetilde{m}_k) = b | \widetilde{\text{tag}} > \text{tag}] = \frac{1}{2} + \text{negl}(n),$$

and if $\Pr[\widetilde{\text{tag}} < \text{tag}] \geq \frac{1}{\text{poly}(n)}$ for some polynomial $\text{poly}(\cdot)$, then

$$\Pr[\mathcal{D}(c, \widetilde{c}_1, \dots, \widetilde{c}_k, \widetilde{m}_1, \dots, \widetilde{m}_k) = b | \widetilde{\text{tag}} < \text{tag}] = \frac{1}{2} + \text{negl}(n).$$

Suppose that $\Pr[\widetilde{\text{tag}} > \text{tag}] = \widehat{p} = \frac{1}{\text{poly}(n)}$. Note that $\widetilde{\text{tag}} > \text{tag}$ implies, $\zeta - \widetilde{\text{tag}} < \zeta - \text{tag}$. Suppose for the sake of contradiction that there exists a $2^{\text{poly}(\log n)}$ -size distinguisher \mathcal{D} and a non-negligible function Δ such that

$$\Pr[\mathcal{D}(c, \widetilde{c}_1, \dots, \widetilde{c}_k, \widetilde{m}_1, \dots, \widetilde{m}_k) = b | \widetilde{\text{tag}} > \text{tag}] \geq \frac{1}{2} + \Delta. \quad (2)$$

¹⁴ We overload notation, here m_i denotes the i^{th} bit of m , and below each \widetilde{m}_i consists of p bits.

Consider the following hybrid distributions H_0, \dots, H_p , where H_α is defined by choosing $m' = (m_1, \dots, m_\alpha, 0, \dots, 0) \in \{0, 1\}^p$ and setting $c = \text{Com}_{\text{tag}}(m'; r)$ for a randomly chosen $r \xleftarrow{\$} \{0, 1\}^{\text{poly}(n)}$.

By a standard hybrid argument, we conclude that there exists $\alpha \in \{0, 1, \dots, p\}$ and a $2^{\text{poly}(\log n)}$ -size distinguisher \mathcal{D}' such that

$$\Pr[\mathcal{D}'(c, \tilde{c}_1, \dots, \tilde{c}_k, \tilde{m}_1, \dots, \tilde{m}_k | \widetilde{\text{tag}} > \text{tag}, H_\alpha) = 0] - \Pr[\mathcal{D}'(c, \tilde{c}_1, \dots, \tilde{c}_k, \tilde{m}_1, \dots, \tilde{m}_k | \widetilde{\text{tag}} > \text{tag}, H_{\alpha+1}) = 0] \geq \frac{\Delta}{p+1}. \quad (3)$$

Note that this implies that $\tilde{m}_{\alpha+1} = 1$, since otherwise H_α and $H_{\alpha+1}$ are identical.

We use \mathcal{D} to construct a $\text{poly}(T_{1, \xi - \widetilde{\text{tag}}})$ -size adversary $\mathcal{B}^{\mathcal{O}_0}$ that breaks the hiding property of $\text{com}_{1, \zeta - \text{tag}}$. Recall that

$$\begin{aligned} & \text{Com}_{\text{tag}}\left(m; (r_i, s_i, a_i)_{i \in [p]}\right) \\ &= \left(\text{tag}, (\text{com}_{0, \text{tag}}(a_i; r_i))_{i \in [p]}, (\text{com}_{1, \zeta - \text{tag}}(m_i \oplus a_i; s_i))_{i \in [p]}\right). \end{aligned}$$

Fix any $\text{tag} \in [\zeta]$. The algorithm $\mathcal{B}^{\mathcal{O}_0}$, given input a string C in the range of $\text{com}_{1, \zeta - \text{tag}}$, and oracle access to \mathcal{D} does the following:

1. For each $j \in [p]$ sample $r_j \xleftarrow{\$} \{0, 1\}^{\ell_{0, \text{tag}}}$ and compute $y_j = \text{com}_{0, \text{tag}}(a_j; r_j)$.
2. For each $j \in [\alpha] \cup [\alpha + 2, p]$, sample $s_j \xleftarrow{\$} \{0, 1\}^{\ell_{1, \zeta - \text{tag}}}$ and compute $w_j = \text{com}_{1, \zeta - \text{tag}}(m_j \oplus a_j; s_j)$.
3. Let $w_{\alpha+1} = C$, $c = (\text{tag}, \{y_j\}_{j \in [p]}, \{w_j\}_{j \in [p]})$. Set $(\widetilde{\text{tag}}, \tilde{c}_1, \dots, \tilde{c}_k) = \mathcal{A}(c)$.
4. If $\widetilde{\text{tag}} < \text{tag}$, then output a randomly chosen $b \xleftarrow{\$} \{0, 1\}$.
5. For each $\kappa \in [k]$, do the following:
 - Parse $\tilde{c}_\kappa = (\widetilde{\text{tag}}, \{\tilde{y}_j^\kappa\}_{j \in [p]}, \{\tilde{w}_j^\kappa\}_{j \in [p]})$.
 - For each $j \in [p]$, compute $(\tilde{a}_j^\kappa, \tilde{r}_j^\kappa) = \mathcal{A}_{1, \zeta - \widetilde{\text{tag}}}^{\mathcal{O}_0}(\tilde{y}_j^\kappa)$ and $(\tilde{a}'_j^\kappa, \tilde{s}_j^\kappa) = \mathcal{A}_{1, \zeta - \widetilde{\text{tag}}}^{\mathcal{O}_0}(\tilde{w}_j^\kappa)$.
 - If there exists $j \in [p]$ such that $\tilde{a}_j^\kappa = \perp$ or $\tilde{a}'_j^\kappa = \perp$, then set $m_\kappa = \perp$.
 - Else, set $m_\kappa = (m_1^\kappa, m_2^\kappa, \dots, m_p^\kappa)$, where $m_j^\kappa = \tilde{a}_j^\kappa \oplus \tilde{a}'_j^\kappa$.

Recall that for every $\widetilde{\text{tag}} \in [\zeta]$, $\mathcal{A}_{1, \zeta - \widetilde{\text{tag}}}^{\mathcal{O}_0}$ is a $T_{1, \zeta - \widetilde{\text{tag}}} \cdot \text{poly}(n)$ -size oracle-aided algorithm that:

- Inverts $\text{com}_{0, \widetilde{\text{tag}}}$ on *any* element in the image of $\text{com}_{0, \widetilde{\text{tag}}}$ with overwhelming probability (over the randomness of \mathcal{O}_0), and outputs \perp on input any element outside the image of $\text{com}_{0, \widetilde{\text{tag}}}$.
- Inverts $\text{com}_{1, \zeta - \widetilde{\text{tag}}}$ on *any* element in the image of $\text{com}_{1, \zeta - \widetilde{\text{tag}}}$ with overwhelming probability (over the randomness of \mathcal{O}_0), and outputs \perp on input any element outside the image of $\text{com}_{1, \zeta - \widetilde{\text{tag}}}$.

Therefore, $(\tilde{m}_1, \dots, \tilde{m}_k)$ are extracted correctly w.h.p.

6. Compute $e = \mathcal{D}'(c, \tilde{c}_1, \dots, \tilde{c}_k, \tilde{m}_1, \dots, \tilde{m}_k)$.
7. If $e = 0$, output $b' = a_{\alpha+1}$. If $e = 1$, output uniformly random b' .

By Equation (3), together with the fact that $\tilde{m}_1, \dots, \tilde{m}_k$ were computed correctly with overwhelming probability,

$$\begin{aligned} & \Pr[e = 0 | (\widetilde{\text{tag}} > \text{tag}) \wedge (a_{\alpha+1} \oplus b = 0)] - \\ & \Pr[e = 0 | (\widetilde{\text{tag}} > \text{tag}) \wedge (a_{\alpha+1} \oplus b = 1)] \geq \frac{\Delta}{p+1}. \end{aligned}$$

Since $a_{\alpha+1} \stackrel{s}{\leftarrow} \{0, 1\}$ (independently of b), this implies that

$$\Pr[(b' = b) \wedge (e = 0) | \widetilde{\text{tag}} > \text{tag}] = \frac{1}{2} \Pr[(e = 0) | \widetilde{\text{tag}} > \text{tag}] + \frac{\Delta}{4(p+1)}$$

Also note that we sample b' uniformly at random if $e = 1$. Therefore,

$$\begin{aligned} & \Pr[(b' = b) \wedge (e = 1) | \widetilde{\text{tag}} > \text{tag}] - \\ & \Pr[(b' \neq b) \wedge (e = 1) | \widetilde{\text{tag}} > \text{tag}] = 0 \end{aligned}$$

which implies

$$\Pr[(b' = b) \wedge (e = 1) | \widetilde{\text{tag}} > \text{tag}] = \frac{1}{2} \Pr[(e = 1) | \widetilde{\text{tag}} > \text{tag}]$$

This implies that

$$\Pr[\mathcal{B}^{\mathcal{O}_0}(\text{com}_{1, \zeta - \widetilde{\text{tag}}}(b)) = b | \widetilde{\text{tag}} > \text{tag}] \geq \frac{1}{2} + \frac{\Delta}{4(p+1)} - \text{negl}(n),$$

contradicting Assumption 1. The case where $\Pr[\widetilde{\text{tag}} < \text{tag}] = \frac{1}{\text{poly}(n)}$, is identical to the previous case, with the roles of com_0 and com_1 reversed, thus we omit the proof. This completes the proof of non-malleability.

5 Non-Malleability Amplification

In this section, we present a non-interactive amplification technique to bootstrap non-malleable commitments for small tags into non-malleable commitments for large tags. We present a compiler that converts any $5\ell t$ -to- z same-tag *auxiliary-input* non-malleable commitment scheme *w.r.t. replacement* (Definition 3) for tags in $[t]$ into an ℓ -to- y same-tag *auxiliary-input* non-malleable commitment scheme *w.r.t. replacement* (Definition 3) for tags in $\left[\binom{t}{t/2} \right]$, for any y and any ℓ such that $\ell y \leq \frac{z}{10}$. We describe our compiler in Figure 2. We emphasize that the size of the resulting commitment scheme grows linearly with ℓ .

We denote the commitment scheme for tags in $[t]$ by Com . We require the scheme Com to be secure against \mathcal{T} -size adversaries, for $\mathcal{T} = \text{poly}(n \cdot 2^y)$.

Let $T_V : \mathbb{N} \rightarrow \mathbb{N}$ denote the time bound associated with Com (i.e., the time required to compute $\mathcal{V}_{\text{Real}}$ and $\mathcal{V}_{\text{Ideal}}$). Our compiler assumes the existence of a NIWI (non-interactive witness indistinguishable) proof system, where witness indistinguishability holds against $\text{poly}(T_V, \mathcal{T})$ -size adversaries. From now, we assume for simplicity (and without loss of generality) that $T_V \geq \mathcal{T}$.

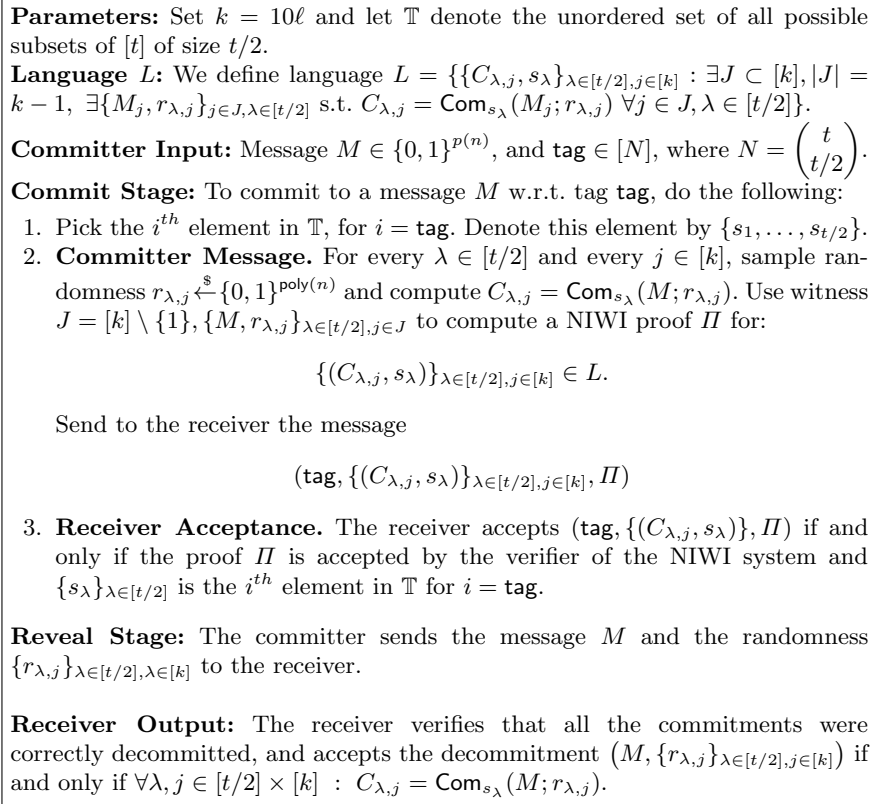


Fig. 2. Round-Preserving Tag Amplification

Theorem 3. *For any polynomials y, z, ℓ and t , where $\ell y \leq \frac{z}{10}$, assuming Com is $5\ell t$ -to- z same-tag auxiliary-input non-malleable w.r.t. replacement (Definition 3) for tags in $[t]$ against $\text{poly}(n \cdot 2^y)$ -size adversaries, and assuming sub-exponentially secure NIWI, the scheme in Figure 2 is ℓ -to- y same-tag auxiliary-input non-malleable w.r.t. replacement (Definition 3) for tags in $\left[\binom{t}{t/2} \right]$ against polynomial size adversaries.*

An overview of the intuition for this construction was provided in Section 2.2. In the formal proof (please refer to the full version of the paper [22]), for every $\beta \in [\ell]$, we define τ_β as the transcript generated by the MIM when the first β left commitments are to 0, and the remaining are to $m_{\beta+1}, \dots, m_\ell$. We then build a sequence of extractors $\mathcal{V}_{\beta, \text{real}}$ and $\mathcal{V}_{\beta, \text{ideal}}$ for $\beta \in [\ell]$, where V_{real} roughly corresponds to $\mathcal{V}_{1, \text{real}}$ and V_{ideal} to $\mathcal{V}_{\ell, \text{ideal}}$. These are such that the joint distribution $(\tau_\beta, \mathcal{V}_{\beta, \text{ideal}}(\tau_\beta)) \approx_c (\tau_{\beta-1}, \mathcal{V}_{\beta, \text{real}}(\tau_{\beta-1}))$. Roughly, we also define a $2y$ -bit inefficient leakage function π_β and an efficient function f such that for every τ ,

$\mathcal{V}_{\beta,\text{real}}(\tau) = f(\mathcal{V}_{\beta-1,\text{ideal}}(\tau), \pi_{\beta} - 1(\tau))$. Combining these equations implies that for every $\beta \in [2, \ell]$:

$$(\tau_{\beta}, \mathcal{V}_{\beta,\text{ideal}}(\tau_{\beta})) \approx_c \tau_{\beta-1}, f(\mathcal{V}_{\beta-1,\text{ideal}}(\tau_{\beta-1}), \pi_{\beta} - 1(\tau_{\beta-1}))$$

We then use the leakage lemma to simulate leakage $\widehat{\pi}_{\beta-1}$ such that

$$\begin{aligned} (\tau_{\beta}, \mathcal{V}_{\beta,\text{ideal}}(\tau_{\beta})) &\approx_c \tau_{\beta-1}, f(\mathcal{V}_{\beta-1,\text{ideal}}(\tau_{\beta-1}), \pi_{\beta-1}(\tau_{\beta-1})) \\ &\approx_c \tau_{\beta-2}, f(\mathcal{V}_{\beta-2,\text{ideal}}(\tau_{\beta-2}), \pi_{\beta-2}(\tau_{\beta-2})), \widehat{\pi}_{\beta-1}(\tau_{\beta-2}) \approx_c \tau_{\beta-3} \dots \end{aligned}$$

Continuing this way, we obtain efficiently simulatable leakage η and an efficiently computable function F such that $\tau_{\ell}, \mathcal{V}_{\ell,\text{ideal}}(\tau_{\ell}) \approx_c \tau_0, F(\mathcal{V}_{1,\text{real}}(\tau_0), \eta(\tau_0))$. This allows us to set $\mathcal{V}_{\text{real}}$ as $F(\mathcal{V}_{1,\text{real}}(\tau_0), \eta(\tau_0))$ while preserving indistinguishability. We refer the reader to the full version for a detailed proof.

6 Putting Things Together: Non-Malleable Commitments for All Tags

In this section, we describe how one can combine results from Section 4 and Section 5 to obtain our main result.

Theorem 4. *There exists a non-interactive non-malleable commitment w.r.t. replacement satisfying Definition 1, assuming the following:*

- Sub-exponential hardness of factoring or discrete log.
- Sub-exponential quantum hardness of LWE.
- Sub-exponential non-interactive witness indistinguishable (NIWI) proofs.

Proof. To obtain this theorem, we apply the following sequence of steps:

- Let $\mathcal{C}_{[\eta \log \log n]}$ denote a many-to-many same-tag non-malleable commitment w.r.t. commitment for $\eta \log \log n$ tags where $0 < \eta < 1$, secure against $2^{\text{poly} \log n}$ -size adversaries. Such a scheme is constructed in Theorem 2, assuming sub-exponential hardness of factoring or discrete log, and sub-exponential quantum hardness of LWE.
- Apply the compiler in Section 5 to $\mathcal{C}_{[\eta \log \log n]}$. Specifically, setting $y = \log^3 n, \ell = \log^3 n, z = \log^7 n, t = \eta \log \log n$ in Theorem 3, we note that $z \geq 10\ell y$ and $\mathcal{C}_{[\eta \log \log n]}$ is $5\ell t$ -to- z same-tag auxiliary-input non-malleable w.r.t. replacement against $\text{poly}(n \cdot 2^y)$ -size adversaries. Therefore, Theorem 3 gives a $(\log^3 n)$ -to- $(\log^3 n)$ same-tag auxiliary-input non-malleable commitment w.r.t. replacement satisfying Definition 3, for $\log^{\epsilon} n$ tags, (for a small constant $\epsilon > 0$), against polynomial-size adversaries. Denote this resulting scheme by $\mathcal{C}_{[\log^{\epsilon} n]}$.
- Apply the compiler in Section 5 once again, this time to $\mathcal{C}_{[\log^{\epsilon} n]}$. Specifically, setting $y = 10, \ell = 10 \log^2 n, z = 1000 \log^2 n, t = \log^{\epsilon} n$ in Theorem 3, we note that $z = 10\ell y$ and that $\mathcal{C}_{[\log^{\epsilon} n]}$ is $5\ell t$ -to- z same-tag auxiliary-input non-malleable w.r.t. replacement against $\text{poly}(n \cdot 2^y)$ -size adversaries.

Therefore, Theorem 3 gives a $10 \log^2 n$ -to-10 same-tag auxiliary-input non-malleable commitment w.r.t. replacement satisfying Definition 3, for $2 \log^2 n$ tags, against polynomial-size adversaries. Denote this resulting scheme by $\mathcal{C}_{[2 \log^2 n]}$.

- Apply the compiler in Section 5 one final time, this time to $\mathcal{C}_{[2 \log^2 n]}$. Specifically, setting $\ell = y = 1, z = 10, t = 2 \log^2 n$ in Theorem 3, we note that $z = 10\ell y$ and that $\mathcal{C}_{[2 \log^2 n]}$ is $5\ell t$ -to- z same-tag auxiliary-input non-malleable w.r.t. replacement against $\text{poly}(n \cdot 2^y)$ -size adversaries.

Therefore, Theorem 3 gives a 1-to-1 auxiliary-input non-malleable commitment w.r.t. replacement satisfying Definition 3, for $n^{\log n}$ tags, against polynomial-size adversaries. Denote this resulting scheme by $\mathcal{C}_{[n^{\log n}]}$.

- Next, assume the existence of a sub-exponentially secure digital signature scheme. More specifically, assume the existence of a signature scheme such that poly-size adversary cannot forge signatures w.r.t. verification keys of size $\log^2 n$ (except with negligible probability). Such a scheme is implied by sub-exponential one-way functions. Denote the keys for such a scheme by (vk, sk) , the setup algorithm by $\text{Setup}(1^\lambda)$ and the signing algorithm by $\text{Sign}(sk, \cdot)$.

Then starting with a non-malleable commitment scheme (w.r.t. replacement) according to Definition 1 for tags in $[n^{\log n}]$ (denoted by $\mathcal{C}_{[n^{\log n}]}$), we build non-malleable commitments for tags in $[2^n]$, satisfying Definition 1 as follows:

To commit to message m with tag $T \in [2^n]$, sample $(vk, sk) \xleftarrow{\$} \text{Setup}(1^{\log^2 n})$, compute a commitment $c \leftarrow \text{Com}_{vk}(m)$, and a signature $\sigma \leftarrow \text{Sign}(sk, T)$. Output (vk, c, σ) . Here $\text{Com}_{vk}(\cdot)$ denotes the commitment algorithm of $\mathcal{C}_{[n^{\log n}]}$ corresponding to tag vk , and we note that $|vk| = \log^2 n$ bits.

For every PPT man-in-the-middle \mathcal{A} that outputs $(\widetilde{vk}, \widetilde{c}, \widetilde{\sigma})$, one of the following holds.

- Either $\widetilde{vk} = vk$, in which case by unforgeability of the signature scheme, if $\widetilde{T} \neq T$ then $\widetilde{\sigma}$ does not verify.
- Or $\widetilde{vk} \neq vk$, in which case the message committed to in \widetilde{c} is “unrelated” to the message committed to in c , i.e., it satisfies the non-malleability condition of Definition 1, since we assume Com_{vk} satisfies Definition 1.

References

1. Ball, M., Dachman-Soled, D., Kulkarni, M., Lin, H., Malkin, T.: Non-malleable codes against bounded polynomial time tampering. IACR Cryptology ePrint Archive 2018, 1015 (2018), <https://eprint.iacr.org/2018/1015>
2. Barak, B.: Constant-Round Coin-Tossing with a Man in the Middle or Realizing the Shared Random String Model. In: FOCS 2002. pp. 345–355 (2002)
3. Barak, B., Ong, S.J., Vadhan, S.P.: Derandomization in cryptography. SIAM J. Comput. 37(2), 380–400 (2007), <https://doi.org/10.1137/050641958>
4. Bellare, M., Palacio, A.: The knowledge-of-exponent assumptions and 3-round zero-knowledge protocols. In: CRPTO 2004. pp. 273–289 (2004)

5. Bitansky, N., Lin, H.: One-message zero knowledge and non-malleable commitments. *IACR Cryptology ePrint Archive* 2018, 613 (2018), <https://eprint.iacr.org/2018/613>
6. Bitansky, N., Paneth, O.: Zaps and non-interactive witness indistinguishability from indistinguishability obfuscation. In: *Theory of Cryptography - 12th Theory of Cryptography Conference, TCC 2015*. pp. 401–427 (2015), https://doi.org/10.1007/978-3-662-46497-7_16
7. Broadnax, B., Fetzer, V., Müller-Quade, J., Rupp, A.: Non-malleability vs. cca-security: The case of commitments. In: *Public-Key Cryptography - PKC 2018 - 21st IACR International Conference on Practice and Theory of Public-Key Cryptography, Rio de Janeiro, Brazil, March 25-29, 2018, Proceedings, Part II*. pp. 312–337 (2018), https://doi.org/10.1007/978-3-319-76581-5_11
8. Chandran, N., Goyal, V., Mukherjee, P., Pandey, O., Upadhyay, J.: Block-wise non-malleable codes. In: *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy*. pp. 31:1–31:14 (2016), <https://doi.org/10.4230/LIPIcs.ICALP.2016.31>
9. Ciampi, M., Ostrovsky, R., Siniscalchi, L., Visconti, I.: Concurrent non-malleable commitments (and more) in 3 rounds. In: *Advances in Cryptology - CRYPTO 2016*. pp. 270–299 (2016), https://doi.org/10.1007/978-3-662-53015-3_10
10. Ciampi, M., Ostrovsky, R., Siniscalchi, L., Visconti, I.: Four-round concurrent non-malleable commitments from one-way functions. In: *Annual International Cryptology Conference*. pp. 127–157. Springer (2017)
11. Dolev, D., Dwork, C., Naor, M.: Non-Malleable Cryptography (Extended Abstract). In: *STOC 1991* (1991)
12. Fenteany, P., Fuller, B.: Non-malleable digital lockers. *Cryptology ePrint Archive, Report 2018/957* (2018), <https://eprint.iacr.org/2018/957>
13. Gentry, C., Wichs, D.: Separating succinct non-interactive arguments from all falsifiable assumptions. In: Fortnow, L., Vadhan, S.P. (eds.) *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*. pp. 99–108. ACM (2011), <http://doi.acm.org/10.1145/1993636.1993651>
14. Goyal, R., Hohenberger, S., Koppula, V., Waters, B.: A generic approach to constructing and proving verifiable random functions. In: Kalai, Y., Reyzin, L. (eds.) *Theory of Cryptography - 15th International Conference, TCC 2017. Lecture Notes in Computer Science, vol. 10678*, pp. 537–566. Springer (2017), https://doi.org/10.1007/978-3-319-70503-3_18
15. Goyal, V.: Constant Round Non-malleable Protocols Using One-way Functions. In: *STOC 2011*. pp. 695–704. ACM (2011)
16. Goyal, V., Khurana, D., Sahai, A.: Breaking the three round barrier for non-malleable commitments. In: *FOCS (2016)*
17. Goyal, V., Lee, C.K., Ostrovsky, R., Visconti, I.: Constructing non-malleable commitments: A black-box approach. In: *FOCS (2012)*
18. Goyal, V., Pandey, O., Richelson, S.: Textbook non-malleable commitments. In: *STOC*. pp. 1128–1141. ACM, New York, NY, USA (2016), <http://doi.acm.org/10.1145/2897518.2897657>
19. Goyal, V., Richelson, S., Rosen, A., Vald, M.: An algebraic approach to non-malleability. In: *FOCS 2014*. pp. 41–50 (2014)
20. Groth, J., Ostrovsky, R., Sahai, A.: New techniques for noninteractive zero-knowledge. *J. ACM* 59(3), 11:1–11:35 (2012), <http://doi.acm.org/10.1145/2220357.2220358>

21. Jain, A., Kalai, Y.T., Khurana, D., Rothblum, R.: Distinguisher-dependent simulation in two rounds and its applications. In: *Advances in Cryptology - CRYPTO 2017*. pp. 158–189 (2017), https://doi.org/10.1007/978-3-319-63715-0_6
22. Kalai, Y., Khurana, D.: Non-interactive non-malleability from quantum supremacy. *Electronic Colloquium on Computational Complexity (ECCC)* 25, 203 (2018), <https://eccc.weizmann.ac.il/report/2018/203>
23. Khurana, D.: Round optimal concurrent non-malleability from polynomial hardness. In: *Theory of Cryptography - 15th International Conference, TCC 2017, Baltimore, MD, USA, November 12-15, 2017, Proceedings, Part II*. pp. 139–171 (2017), https://doi.org/10.1007/978-3-319-70503-3_5
24. Khurana, D., Sahai, A.: How to achieve non-malleability in one or two rounds. In: *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*. pp. 564–575 (2017), <https://doi.org/10.1109/FOCS.2017.58>
25. Komargodski, I., Yogev, E.: Another step towards realizing random oracles: Non-malleable point obfuscation. In: Nielsen, J.B., Rijmen, V. (eds.) *Advances in Cryptology - EUROCRYPT 2018. Lecture Notes in Computer Science*, vol. 10820, pp. 259–279. Springer (2018), https://doi.org/10.1007/978-3-319-78381-9_10
26. Lin, H., Pass, R.: Constant-round Non-malleable Commitments from Any One-way Function. In: *STOC 2011*. pp. 705–714
27. Lin, H., Pass, R.: Non-malleability Amplification. In: *Proceedings of the 41st Annual ACM Symposium on Theory of Computing*. pp. 189–198. *STOC '09* (2009)
28. Lin, H., Pass, R.: Black-box constructions of composable protocols without setup. In: Safavi-Naini, R., Canetti, R. (eds.) *Advances in Cryptology - CRYPTO 2012. Lecture Notes in Computer Science*, vol. 7417, pp. 461–478. Springer (2012), https://doi.org/10.1007/978-3-642-32009-5_27
29. Lin, H., Pass, R., Soni, P.: Two-round and non-interactive concurrent non-malleable commitments from time-lock puzzles. *Cryptology ePrint Archive, Report 2017/273* (2017), <http://eprint.iacr.org/2017/273>
30. Lin, H., Pass, R., Venkatasubramanian, M.: Concurrent Non-malleable Commitments from Any One-Way Function. In: *TCC 2008*. pp. 571–588
31. Pandey, O., Pass, R., Vaikuntanathan, V.: Adaptive One-Way Functions and Applications. In: *Advances in Cryptology — CRYPTO '08*. pp. 57–74 (2008)
32. Pass, R., Rosen, A.: Concurrent Non-Malleable Commitments. In: *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*. pp. 563–572. *FOCS '05* (2005)
33. Pass, R., Rosen, A.: New and improved constructions of non-malleable cryptographic protocols. In: *STOC 2005*. pp. 533–542 (2005)
34. Pass, R., Rosen, A.: New and Improved Constructions of Nonmalleable Cryptographic Protocols. *SIAM J. Comput.* 38(2), 702–752 (2008)
35. Pass, R., Wee, H.: Constant-round non-malleable commitments from sub-exponential one-way functions. In: *EUROCRYPT 2010*. pp. 638–655 (2010)
36. Wee, H.: Black-box, round-efficient secure computation via non-malleability amplification. In: *FOCS 2010*. pp. 531–540 (2010)