

Non-Interactive Non-Malleability from Quantum Supremacy

Yael Tauman Kalai
Microsoft Research, MIT
yael@microsoft.com

Dakshita Khurana
Microsoft Research
dakshita.khurana@microsoft.com

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 with respect to 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 with respect to commitment for $\epsilon \log \log n$ tags (for any constant $\epsilon > 0$) into a non-interactive non-malleable commitment with respect to 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.

Contents

1	Introduction	1
1.1	Our Results	2
1.1.1	Non-interactive non-malleable commitments for $O(\log \log n)$ tags.	3
1.1.2	Non-interactive Tag Amplification from NIWIs	4
1.2	Applications and Directions for Future Work	5
1.3	Prior work	6
2	Overview of Techniques	7
2.1	Non-interactive Tag Amplification	7
2.1.1	Tag Amplification using NIWIs: First Stab	7
2.1.2	Overview of Our Compiler	8
2.1.3	More Detailed Protocol Description	10
2.2	Non-Malleable Commitments with respect to Commitment for $\epsilon \log \log n$ Tags	11
3	Definitions	12
3.1	Non-Malleable Commitments w.r.t. Replacement	13
3.2	Non-Malleable Commitments w.r.t. Commitment	16
4	Non-Malleable Commitments for Small Tags	17
4.1	Non-Malleable Commitment Scheme for $\eta \cdot \log \log n$ Tags	17
4.2	Instantiating the Assumption	23
5	Non-Malleability Amplification	24
6	Putting Things Together: Non-Malleable Commitments for All Tags	47
	References	52
A	Two Message Arguments from Quantum Polynomial Hardness	52
A.1	Modified Blum Protocol	52
A.2	Construction of Two-message Arguments	52

1 Introduction

Non-malleability, first introduced by Dolev, Dwork and Naor [DDN91] 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.

Non-interactive Commitments. A non-interactive commitment scheme consists of a commitment algorithm, that on input a message m and randomness r , outputs a string 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 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 constructing non-malleable commitments in the plain model in as few rounds as possible (e.g [DDN91, Bar02, PR05a, PR08, LPV, PPV08, LP09, Wee10, PW10, LP, Goy11, GLOV12, GRRV14, GPR16, COSV16, COSV17, Khu17, LPS17, KS17]). So far, the only known constructions of non-interactive non-malleable commitments (without setup) are the ones by Pandey, Pass and Vaikuntanathan [PPV08], based on a strong non-falsifiable assumption, and Bitansky and Lin [BL18], 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 with respect to commitment.** Intuitively, non-malleability with respect to 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

¹We will sometimes omit explicitly writing the randomness r .

(\tilde{m}, \tilde{r}) such that $\tilde{c} = \text{com}(\tilde{m}; \tilde{r})$. It was shown by [BFMR18] that in the case of non-interactive commitments, non-malleability w.r.t. commitment is equivalent to CCA-security.

- **Non-malleability with respect to replacement.** A weaker, yet natural, notion of malleability is non-malleability with respect to replacement [Goy11]. 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 with respect to 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 [BFMR18] 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 with respect to 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 $\text{Com}(m_b)$.

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.

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 the sub-exponential decisional linear (DLIN) assumption).

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

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

³Non-malleability w.r.t. replacement implies non-malleability w.r.t. opening, as defined by Goyal et al. [GKS16].

novel use of the leakage lemma due to Gentry and Wichs [GW11]. 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.

1.1.1 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 [PPV08, LP12]⁴.

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

To achieve this, we rely on the leveraging technique of Pass and Wee [PW10] 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.

In Section 2.2, we give an overview of the proof of Informal Theorem 1. However, our main conceptual novelty in this part, which we describe next, is the idea of constructing an adaptive or parallel-CCA secure commitment scheme for 2 tags using *quantum supremacy*.

Using Quantum Supremacy. Loosely speaking, in order to construct an adaptive 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 [GHKW17]); and instantiating com_0 as a post-quantum *insecure* commitment (such as one based on factoring or discrete log via Goldreich-Levin [Lev87]), 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 or discrete log, and sub-exponential quantum hardness of LWE or LPN, there exist sub-exponentially secure many-to-many CCA bit commitments for 2 tags.*

⁴We interpret these as CCA-secure commitments for the sake of simplicity of our proof, but in the non-interactive setting [BFMR18], these are equivalent to many-many non-malleable commitments w.r.t. commitment.

Combining this with Informal Theorem 1, we have:

Informal Theorem 2. *Assuming sub-exponential hardness of factoring or discrete log, and sub-exponential quantum hardness of LWE or LPN , there exist many-to-many same-tag non-interactive non-malleable commitments with respect to commitment 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 exist [BL18]. The work of [LPS17] 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 [BL18]. 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 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. Whenever this technique is used, the resulting scheme needs to rely on super-polynomial (and often sub-exponential) hardness assumptions. 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 assumptions. For example, using our ideas, one can appropriately instantiate the protocols in [JKKR17] 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. We give details of this construction in Appendix A.

We also remark that one can substitute the assumption on sub-exponential quantum hardness of LWE with sub-exponentially secure time-lock puzzles [LPS17], or sub-exponentially secure one-way functions [BL18] admitting hardness amplification, to obtain (many-to-many) non-malleable commitments w.r.t. replacement for $\epsilon \log \log n$ tags.

1.1.2 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.

Informal Theorem 3 (Tag Amplification from NIWIs). *Assuming many-to-many same-tag non-malleable commitments with respect to replacement for $\epsilon \log \log n$ tags (for an arbitrarily small constant $\epsilon > 0$) and sub-exponentially secure NIWIs, there exist non-interactive non-malleable commitments with respect to 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 [GOS12], or from derandomization assumptions [BOV07], or assuming indistinguishability obfuscation [BP15]. Interestingly, to prove this theorem, we crucially rely on the Gentry-Wichs leakage lemma [GW11]. We provide a high-level overview of this amplification technique, as well as its proof, in Section 2.1.

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 with respect to commitment for $\epsilon \log \log n$ tags.
- Non-interactive non-malleable commitments with respect to replacement for 2^n tags, additionally assuming sub-exponentially secure NIWIs.

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 motivate and give some additional perspectives on this notion.

- Non-malleable commitments w.r.t. replacement imply the indistinguishability-based notion of non-malleable commitments w.r.t. opening, as defined in [CGM⁺16, GKS16]. Therefore, we obtain 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).*

- It was shown by [BFMR18] that the definitions of non-malleability w.r.t. commitment and (one-to-one) CCA-security are equivalent in the setting of non-interactive commitments. We observe that their proofs also imply equivalence between non-malleability w.r.t. replacement and 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 is allowed to respond with any arbitrary value of its choice.
- Non-interactive non-malleable commitments w.r.t. opening are known to be equivalent to block-wise non-malleable codes [CGM⁺16] 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 5. *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.*

- 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 [GRRV14, COSV16]) 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 [BP04, LPS17], and admitting a non-uniform simulator. Combining these with our non-malleable commitments w.r.t. replacement, we have the following theorem.

Informal Theorem 6. *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.*

Similarly, 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.

- Non-malleable commitments w.r.t. replacement are known to be sufficient for MPC [Goy11]. 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.
- The recent works of [KY18, FF18] 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 [LPS17] constructed non-interactive non-malleable commitments w.r.t. commitment against a restricted class of *uniform* man-in-the-middle adversaries, assuming sub-exponentially secure time-lock puzzles, sub-exponential NIWI and sub-exponential collision-resistant hash functions against uniform adversaries. [BDK⁺18] in very recent independent work construct a significantly weaker object that they call non-interactive quasi-non-malleable commitments, based on well-studied assumptions. These commitments guarantee security against adversaries running in a-priori bounded polynomial time $O(n^c)$, but allow honest parties to run in longer (polynomial) time. These are used as a stepping stone in [BDK⁺18] to build non-malleable codes against bounded polynomial time tampering.

In this paper, our focus is on the non-interactive setting in the plain model against all polynomial-sized adversaries. 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. [PPV08] 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 [BL18] 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. [DDN91, LP09, Wee10, LPS17, BL18]). Of these, only the recent work of [BL18] considers tag amplification in the non-interactive setting against general adversaries. [BL18] 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.

Our tag amplification technique only assumes the existence of a sub-exponentially secure NIWI proof. However, while our tag amplification technique yields commitments that are non-malleable w.r.t. replacement, the one in [BL18] yields commitments that are (concurrent) non-malleable w.r.t. commitment.

2 Overview of Techniques

We now proceed to an informal overview of our techniques. We start by describing our key technical contribution: our tag amplification compiler. We then proceed to describe our new constructions of non-malleable commitments for small tags, which satisfy many-to-many same-tag non-malleability with respect to commitment.

2.1 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 with respect to 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 ⁵. 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 $\tilde{T} = \{\tilde{t}_1, \tilde{t}_2, \dots, \tilde{t}_{\alpha/2}\}$ such that $\tilde{T} \neq T$, there exists at least one index i such that $\tilde{t}_i \notin \{t_1, t_2, \dots, t_{\alpha/2}\}$. 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 $\text{com}_{\tilde{t}_i}(m')$ for a related message m' . This seems to suggest that an adversary cannot generate a *valid* commitment $\text{com}_{\tilde{T}}(\tilde{m})$ 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 [KS17] describes how this problem can be solved using a NIZK, 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-malleable⁶. Instead, in this work, we rely on non-interactive proofs satisfying a weaker hiding property, i.e., witness indistinguishability⁷. This introduces several problems and bottlenecks that do not come up when using NIZKs. In particular, techniques in [KS17] 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.

2.1.1 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

⁵Our actual encoding of T to $\{t_1, t_2, \dots, t_{\alpha/2}\}$ is slightly more sophisticated, but achieves the same effect.

⁶To be precise, they need to rely on the fact that the NIZK is "more secure" than the underlying commitment scheme.

⁷As with NIZKs used in [KS17], we also require our NIWI to be more secure than the underlying commitment, which results in a sub-exponential assumption on the NIWI.

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 $\{r_{i,b}\}_{i \in [\alpha/2], b \in \{0,1\}} \xleftarrow{\$} \{0,1\}^*$, 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. replacement⁸.

Unfortunately, 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) .⁹ 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. replacement¹⁰.

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

2.1.2 Overview of Our Compiler

Our idea is to have each commitment consist of at least $(\ell + 1)$ repetitions (as opposed to only 2), where ℓ is the number of commitments that the adversary can receive (on the left). Additionally, the committer is required to provide a NIWI proof that in all but one of these parallel executions, the messages committed across all $\alpha/2$ tags are identical. That is, in a valid commitment, all $(\ell + 1)$ repetitions commit to the same message.

Now, when we perform the same hybrid argument as above where we change one left commitment in each hybrid, we have the following invariant: in all but one repetition, by the soundness of the NIWI, the messages committed by the MIM are identical for all tags. This helps us argue

⁸On the other hand, if we used a NIZK, the resulting scheme would be many-to-1 non-malleable w.r.t. commitment. Obtaining non-malleability w.r.t. replacement appears to be inherent unless one relies on NIZKs (and a CRS).

⁹This is the standard approach used in all previous work on this topic.

¹⁰This 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.

that the MIM is forced to change the message in *at most one repetition in every hybrid*. We give a more detailed explanation in [Section 2.1.3](#).

Since there are $(\ell + 1)$ repetitions and only ℓ hybrids, we then deduce that messages in at least one repetition must remain unchanged at the end of *all ℓ hybrids*¹¹. 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 2](#) for details.

Technical Bottlenecks. The intuition above seems to imply that the adversary cannot convert a commitment to m to a commitment to a related message m' . However, to formally prove that this construction satisfies the definition of non-malleability w.r.t. replacement¹², we need to argue that there exists an (inefficient) extractor $\mathcal{V}_{\text{Real}}(\tau_{\text{Real}})$ corresponding to the transcript of a “real” experiment with honest messages (m_1, \dots, m_ℓ) , and an (inefficient) extractor $\mathcal{V}_{\text{Ideal}}(\tau_{\text{Ideal}})$ corresponding to the 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 required to output \tilde{m} . Therefore, we need to define distributions $\mathcal{V}_{\text{Real}}(\tau_{\text{Real}})$ and $\mathcal{V}_{\text{Ideal}}(\tau_{\text{Ideal}})$, and ensure that they are 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.¹³ 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 [[GW11](#)] plays a crucial role. This decision of whether to output \tilde{m} or output \perp will be dictated by a leakage function. We will need to “carry over” this leakage function across hybrids by relying on the leakage lemma. 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 in [Section 2.1.3](#), and refer the reader to [Section 5](#) for details of the proof.

¹¹To simplify our proof, we rely on 10ℓ repetitions (instead of $\ell + 1$) repetitions, to ensure that the messages in *most* repetitions remain unchanged.

¹²We refer the reader to [Definition 3](#) for a one-to-one definition, and [Definition 2](#) for a many-to-many definition.

¹³We note that these distributions are indeed indistinguishable if the adversary always generates valid commitments.

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 2.2](#)). We will 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.

2.1.3 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 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 protocol is not many-to-many, because as explained above, 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 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 adversary 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 a hybrid argument. 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 joint distribution of the message \tilde{m} committed by the MIM using small tag $\tilde{t}_1 \notin \{t_{1,1}, t_{1,2}, \dots, t_{1,\alpha/2}\}$ in all k repetitions, cannot change. This follows from the many-to- z non-malleability with respect to 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 his repetitions. This implies that the MIM continues to commit

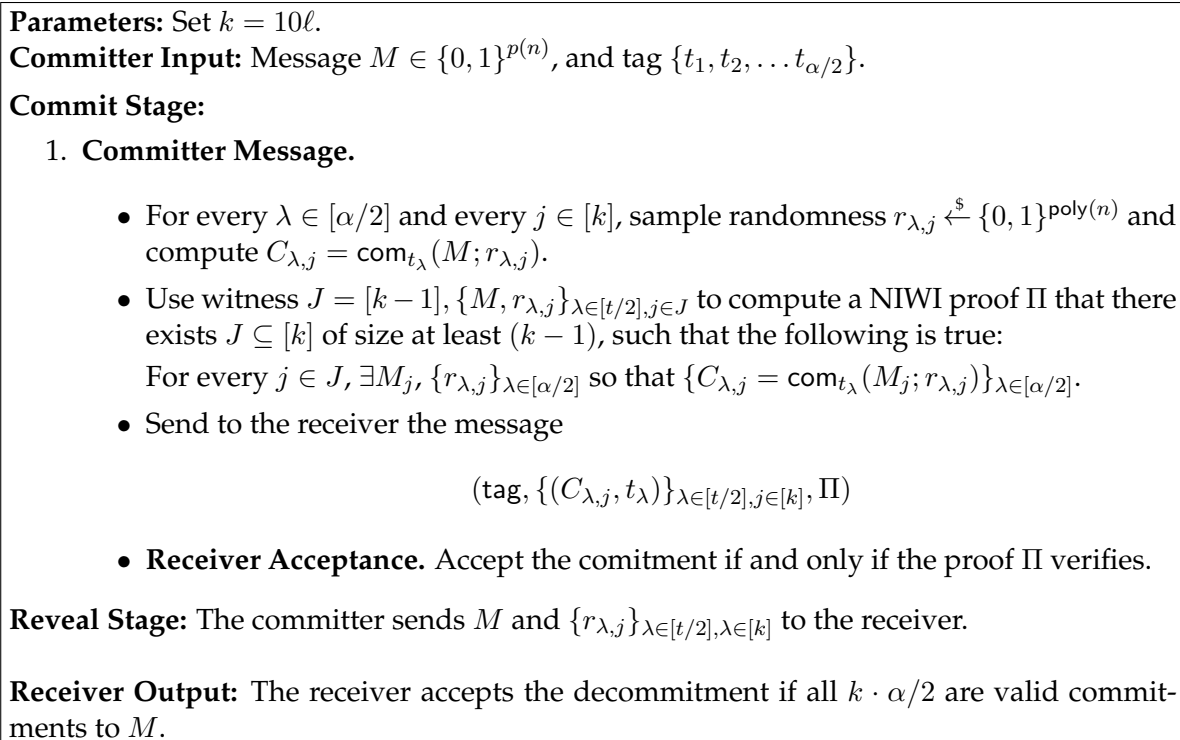


Figure 1: Round-Preserving Tag Amplification

to \tilde{m} in least $(k - 1)$ of the repetitions, *for every tag*. Using a similar argument to the one above, one can argue that on input ℓ commitments to messages $0, 0, m_3, \dots, m_\ell$ (i.e., with the first two messages replaced by 0), the MIM continues to commit to \tilde{m} in at least $(k - 2)$ repetitions, for every tag.

Continuing this way, we observe that the MIM continues to commit to \tilde{m} in at least $(k - \ell)$ repetitions, with respect to 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 \perp , 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 to the leakage lemma [GW11]. We refer the reader to Section 5 for details.

2.2 Non-Malleable Commitments with respect to Commitment for $\epsilon \log \log n$ Tags

To conclude this overview, we give a brief description of how we build a non-malleable commitment for $\epsilon \log \log n$ tags, where $\epsilon > 0$ is a small constant, from sub-exponential adaptive commitments for two tags¹⁴. This uses ideas similar to those in [LPS17, KS17].

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 .

¹⁴Recall that we build adaptive commitments for two tags using quantum supremacy, as discussed in Section 1.1. We refer the reader to Section 4.2 for further details.

We rely on a technique of Pass and Wee [PW10] to show that from any such adaptive commitments, one can 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} .

In order to commit to a bit b with tag $\in [\zeta]$, the committer first XOR secret shares 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 bit b with respect to tag tag , commits to a related bit b' with respect to $\widetilde{\text{tag}} \neq \text{tag}$. 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 with respect to 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. We give a formal construction and proof in [Section 4.1](#).

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} ,

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

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* ([Goy11]). 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 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 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} ¹⁵.

¹⁵Note that this definition explicitly considers auxiliary information z , but is equivalent to one that does not consider z . We explicitly consider z for convenience.

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 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 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 $(\tilde{m}_1, \dots, \tilde{m}_n) = \text{abort}$.

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 $(\tilde{m}_{\text{Ideal},1}, \dots, \tilde{m}_{\text{Ideal},n}) = \text{abort}$.

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

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

(c) *We require:*

$$((c_1, \dots, c_\ell), (\tilde{c}_1, \dots, \tilde{c}_y), z, \mathcal{V}_{\text{Real}}(c_1, \dots, c_\ell, \tilde{c}_1, \dots, \tilde{c}_y)) \approx_c$$

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

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

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

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

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 $\tilde{c}_{0,i} = \text{com}_{\tilde{\text{tag}}}(\tilde{M}_i; \tilde{r}_i)$, set $\tilde{m}_i = \tilde{M}_i$, and otherwise no restrictions are placed on \tilde{m}_i . We require that

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

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

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 $\tilde{c}_{\ell,i} = \text{com}_{\tilde{\text{tag}}}(\tilde{M}_i; \tilde{r}_i)$, set $\tilde{m}_i = \tilde{M}_i$, and otherwise no restrictions are placed on \tilde{m}_i . We require that

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

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:*

$$\begin{aligned} & ((c_{0,1}, \dots, c_{0,\ell}), \mathbf{a}_0, (\tilde{c}_{0,1}, \dots, \tilde{c}_{0,y}), z_0, \mathcal{V}_{\text{Real}}(c_{0,1}, \dots, c_{0,\ell}, \mathbf{a}_0, \tilde{c}_{0,1}, \dots, \tilde{c}_{0,y})) \approx_c \\ & ((c_{\ell,1}, \dots, c_{\ell,\ell}), \mathbf{a}_\ell, (\tilde{c}_{\ell,1}, \dots, \tilde{c}_{\ell,y}), z_\ell, \mathcal{V}_{\text{Ideal}}(c_{\ell,1}, \dots, c_{\ell,\ell}, \mathbf{a}_\ell, \tilde{c}_{\ell,1}, \dots, \tilde{c}_{\ell,y})) \end{aligned}$$

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 1](#) and [Theorem 2](#) show that our constructions also satisfy this stronger definition.*

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

In this section, we consider the stronger definition of non-malleability with respect to commitment [PR05b]. This definition 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 what follows we present a hybrid between these two variants, 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 hybrid definition is used as a stepping stone to achieve our main result.

Definition 4 (ℓ -to-Many Same-tag Non-malleable Commitments with respect to Commitment).

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-many same-tag non-malleable with respect to commitment if the following 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-many Non-malleability.** There exists a (possibly inefficient) function \mathcal{V} such that the following holds:
 - For any $m_1, \dots, m_\ell \in \{0, 1\}^p$, any $r_1, \dots, r_\ell \in \{0, 1\}^{\text{poly}(n)}$, and any $\text{tag}_1, \dots, \text{tag}_\ell \in [N]$, set $c_i = \text{com}_{\text{tag}_i}(m_i; r_i)$ for every $i \in [\ell]$, and let $(\tilde{c}_1, \dots, \tilde{c}_k, z) = \mathcal{A}(c_1, \dots, c_\ell)$. If there exists $\tilde{\text{tag}} \in [N] \setminus \{\text{tag}_i\}_{i \in [\ell]}$ such that $\tilde{c}_1, \dots, \tilde{c}_k$ all use $\tilde{\text{tag}}$, then continue. Otherwise set $(\tilde{m}_1, \dots, \tilde{m}_k) = \text{abort}$. For each $i \in [k]$, 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 set $\tilde{m}_i = \perp$. We require that $\mathcal{V}(c_1, \dots, c_\ell, \tilde{c}_1, \dots, \tilde{c}_k) = (\tilde{m}_1, \dots, \tilde{m}_k)$.
 - For any $r_{\text{Ideal},1}, \dots, r_{\text{Ideal},\ell} \in \{0, 1\}^{\text{poly}(n)}$, any $\text{tag}_1, \dots, \text{tag}_\ell \in [N]$, set $c_{\text{Ideal},i} = \text{com}_{\text{tag}_i}(0^p; r_{\text{Ideal},i})$ for every $i \in [\ell]$, and let $(\tilde{c}_{\text{Ideal},1}, \dots, \tilde{c}_{\text{Ideal},k}, z_{\text{Ideal}}) = \mathcal{A}(c_{\text{Ideal},1}, \dots, c_{\text{Ideal},\ell})$. If there exists $\tilde{\text{tag}} \in [N] \setminus \{\text{tag}_i\}_{i \in [\ell]}$ such that $\tilde{c}_{\text{Ideal},1}, \dots, \tilde{c}_{\text{Ideal},k}$ all use $\tilde{\text{tag}}$, then continue. Otherwise set $(\tilde{m}_1, \dots, \tilde{m}_k) = \text{abort}$. For each $i \in [k]$, 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 $\tilde{c}_i = \text{com}_{\tilde{\text{tag}}}(\tilde{M}_{\text{Ideal},i}, \tilde{r}_{\text{Ideal},i})$, set $\tilde{m}_{\text{Ideal},i} = \tilde{M}_{\text{Ideal},i}$, and otherwise set $\tilde{m}_{\text{Ideal},i} = \perp$. We require that

$$\mathcal{V}(c_{\text{Ideal},1}, \dots, c_{\text{Ideal},\ell}, \tilde{c}_{\text{Ideal},1}, \dots, \tilde{c}_{\text{Ideal},k}) = (\tilde{m}_{\text{Ideal},1}, \dots, \tilde{m}_{\text{Ideal},k}).$$

- For any $m_1, \dots, m_\ell \in \{0, 1\}^p$, any $\text{tag}_1, \dots, \text{tag}_\ell \in [N]$, and for $r_1, \dots, r_\ell \xleftarrow{\$} \{0, 1\}^{\text{poly}(n)}$ and $r_{\text{Ideal},1}, \dots, r_{\text{Ideal},\ell} \xleftarrow{\$} \{0, 1\}^{\text{poly}(n)}$, set $c_i = \text{com}_{\text{tag}_i}(m_i; r_i)$ and $c_{\text{Ideal},i} = \text{com}_{\text{tag}_i}(0^p; r_{\text{Ideal},i})$ for every $i \in [\ell]$. Then,

$$(c_1, \dots, c_\ell, \tilde{c}_1, \dots, \tilde{c}_k, z), \mathcal{V}(c_1, \dots, c_\ell, \tilde{c}_1, \dots, \tilde{c}_k) \approx_c$$

$$(c_{\text{Ideal},1}, \dots, c_{\text{Ideal},\ell}, \tilde{c}_{\text{Ideal},1}, \dots, \tilde{c}_{\text{Ideal},k}, z_{\text{Ideal}}), \mathcal{V}(c_{\text{Ideal},1}, \dots, c_{\text{Ideal},\ell}, \tilde{c}_{\text{Ideal},1}, \dots, \tilde{c}_{\text{Ideal},k}).$$

A hybrid argument from Lin *et. al* [LPV] shows that [Definition 4](#) above is equivalent to the seemingly weaker definition of 1-to-many same-tag non-malleability, and the seemingly stronger definition of many-to-many same-tag non-malleability, both defined below.

Definition 5 (1-to-many Same-tag Non-malleable Commitments with respect to Commitment).

We say that a commitment scheme is 1-to-many same-tag non-malleable with respect to commitment if the scheme satisfies [Definition 4](#) for $\ell(n) = 1$.

Definition 6 (Many-to-many Same-tag Non-malleable Commitments with respect to Commitment).

We say that a commitment scheme is many-to-many same-tag non-malleable with respect to commitment if for every polynomial $\ell(\cdot)$, it is ℓ -to-many same-tag non-malleable w.r.t. commitment (as in [Definition 4](#)).

Remark 2. We note that [Definition 6](#) is strictly stronger than [Definition 3](#). Any ℓ -to-many same-tag non-malleable commitment w.r.t. commitment com is also an ℓ -to-many same-tag auxiliary-input non-malleable commitment w.r.t. replacement for any polynomial ℓ .

This follows by setting $t_V(n) = p + \text{poly}(n)$, i.e., the time required to brute-force break the commitment, such that \mathcal{V} is computable in time at most $\text{poly}(2^{t_V(n)})$. We can then carrying out a hybrid argument (due to Lin *et al.* [LPV]). In the i^{th} hybrid for $i \in [y]$, (c_1, \dots, c_i) are generated as commitments to 0^p , and (c_{i+1}, \dots, c_y) are generated as commitments to (m_{i+1}, \dots, m_y) ; as defined, aux_i is generated as a function of all but the randomness used to generate the i^{th} commitment. Indistinguishability between Hybrid_i and Hybrid_{i+1} follows by non-malleability w.r.t. commitment of com applied to c_i , and by the $T_V(n)$ -indistinguishability of $(\text{aux}_i, \text{aux}_{i+1})$.

4 Non-Malleable Commitments for Small Tags

In this section, we construct a many-to-many same-tag non-malleable commitment scheme with respect to commitment ([Definition 6](#)) for $\eta \cdot \log \log(n)$ tags, for a small constant $\eta > 0$. This scheme can be used to instantiate our compiler from [Section 5](#) (see [Section 6](#) for details).

4.1 Non-Malleable Commitment Scheme for $\eta \cdot \log \log n$ Tags

In this section, we construct a many-to-many same-tag non-malleable commitment scheme w.r.t. commitment ([Definition 6](#)) for $\zeta = \eta \cdot \log \log n$ tags, for a small enough constant $\eta > 0$, based on the following assumption about non-interactive bit commitments.

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 poly(2^{n^δ})-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 poly(2^{n^δ})-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 Section 4.2, we describe how to instantiate this assumption based on specific cryptographic assumptions. Pass and Wee [PW10] showed that Assumption 1 can be used to derive a sequence of commitments, described below.

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 $r' \in \{0, 1\}^{\ell_{1-b,j}}$,

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

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 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} .

Overview of the construction of the commitment sequence. This is only for completeness and is taken from Pass and Wee [PW10]. This sequence of commitments is constructed as follows. For every $b \in \{0, 1\}$ and every $i \in [0, \zeta - 1]$, set $\ell_{b,i+1} = \ell_{b,i}^{1/\delta}$ (where $\delta \in (0, 1)$ is the constant from Assumption 1). Note that if $\eta < \frac{1}{\log 1/\delta}$ then setting $\delta' \triangleq \eta \cdot \log(1/\delta) < 1$, the following holds

$$\ell_{b,0} = (\ell_{b,\zeta})^{\delta^\zeta} = (\ell_{b,\zeta})^{(1/\log n)^{\delta^\zeta}} = n^{(1/\log n)^{\delta^\zeta}} = 2^{\log n^{(1-\delta')}}.$$

For every $i \in [\zeta]$ and $b \in \{0, 1\}$, the commitment $\text{com}_{b,i} : \{0, 1\} \times \{0, 1\}^{\ell_{b,i}(n)} \rightarrow \{0, 1\}^{L(\ell_{b,i}(n))}$ is defined by

$$\text{com}_{b,i}(m; r) = \text{com}_b(m; r) \text{ where } \text{com}_b \text{ is invoked with security parameter } \ell_{b,i}(n).$$

Remark 3. In what follows we think of $\text{com}_b : \{0, 1\} \times \{0, 1\}^* \rightarrow \{0, 1\}^*$, where inputs of length n are mapped to inputs of length $L(n)$.

For every $b \in \{0, 1\}$ and every $i \in [\zeta - 1]$, setting $T_{b,i} = 2^{\ell_{b,i}}$,

$$T_{b,i+1} = 2^{\ell_{b,i+1}} = 2^{\ell_{b,i}^{1/\delta}} = T_{b,i}^{\omega(1)},$$

as desired, where the latter equality follows from the fact that δ is a constant smaller than 1.

For every $b \in \{0, 1\}$ and every $i \in [\zeta]$, consider the algorithm $\mathcal{A}_{b,i}$ that inverts $\text{com}_{b,i}$ by enumerating over all possible $r \in \{0, 1\}^{\ell_{b,i}}$, and outputs \perp if no such r is found. Clearly, $\mathcal{A}_{b,i}$ runs in time at most $T_{b,i} \cdot \text{poly}(n)$. Moreover, for any $j \in [\zeta]$ and any $r' \in \{0, 1\}^{\ell_{1-b,j}}$, the algorithm $\mathcal{A}_{b,i}$, on input $\text{com}_{1-b,j}(r')$ and given oracle access to \mathcal{O}_{1-b} , runs in poly-time and outputs $\mathcal{A}_{b,i}^{\mathcal{O}_{1-b}}(f_{1-b,j}(r'))$. Thus, for every $b \in \{0, 1\}$ and every $i \in [\zeta - 1]$, $\mathcal{A}_{b,i}$ is a $T_{b,i} \cdot \text{poly}(n)$ -size algorithm such that for every $j \in [\zeta]$, every $m_1, m_2 \in \{0, 1\}$, every $r \in \{0, 1\}^{\ell_{b,i}}$ and every $r' \in \{0, 1\}^{\ell_{1-b,j}}$,

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

where the probability is over the randomness of \mathcal{O}_{1-b} , as desired.

By sub-exponential hardness of com_b in the presence of \mathcal{O}_{1-b} (which is implied by Assumption 1), it follows that for $k \geq (i + 1)$, inverting $\text{com}_{b,k}$ given oracle to \mathcal{O}_{1-b} requires circuits of size greater than $\text{poly}(2^{\ell_{b,i+1}}) = \text{poly}(2^{\ell_{b,i}}) = \text{poly}(T_{b,i})$. Therefore, for any $\text{poly}(T_{b,i})$ -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 $r \xleftarrow{\$} \{0, 1\}^{\ell_{b,k}(n)}$ and over the randomness of \mathcal{O}_{1-b} , as desired. The fact that $\ell_{b,0} \geq 2^{\log n^{(1-\delta')}}$ (for some constant $\delta' < 1$), implies that for $i \in [\zeta]$,

$$T_{b,i} \triangleq 2^{\ell_{b,i}} \geq 2^{\log^c n},$$

for every constant $c \in \mathbb{N}$.

Our construction of non-malleable commitments for $\zeta(n)$ tags To commit to a message $m = (m_1, \dots, m_p) \in \{0, 1\}^p$ 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 1. *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 (Definition 6) 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 (Definition 6), 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 (Definition 5), against all $2^{\text{poly}(\log n)}$ -size adversaries. This follows by a hybrid argument of [LPV], which proves that any commitment scheme that satisfies Definition 5 also satisfies Definition 6 (and this hybrid argument also holds for $2^{\text{poly}(\log n)}$ -size adversaries).

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 $m = (m_1, \dots, m_p) \in \{0, 1\}^p$,¹⁶ 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 satisfies Definition 5, 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

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

¹⁶We overload notation, here m_i denotes the i^{th} bit of m , and below each \widetilde{m}_i consists of p bits.

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

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

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

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

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

We use \mathcal{D} to construct a $\text{poly}(T_{1, \zeta - \widetilde{\text{tag}}})$ -size adversary $\mathcal{B}^{\mathcal{O}_0}$ that breaks the hiding property of $\text{com}_{1, \zeta - \text{tag}}$, contradicting the security 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 $w_{\alpha+1} = C$.
4. Let $c = (\text{tag}, \{y_j\}_{j \in [p]}, \{w_j\}_{j \in [p]})$.
5. Obtain $(\widetilde{\text{tag}}, \widetilde{c}_1, \dots, \widetilde{c}_k) = \mathcal{A}(c)$.
6. If $\widetilde{\text{tag}} < \text{tag}$, then output a randomly chosen $b \xleftarrow{\$} \{0, 1\}$.
7. For each $\kappa \in [k]$, do the following:
 - Parse $\widetilde{c}_\kappa = (\widetilde{\text{tag}}, \{\widetilde{y}_j^\kappa\}_{j \in [p]}, \{\widetilde{w}_j^\kappa\}_{j \in [p]})$.
 - For each $j \in [p]$, compute $(\widetilde{a}_j^\kappa, \widetilde{r}_j^\kappa) = \mathcal{A}_{1, \zeta - \widetilde{\text{tag}}}^{\mathcal{O}_0}(\widetilde{y}_j^\kappa)$ and $(\widetilde{a}'_j^\kappa, \widetilde{s}_j^\kappa) = \mathcal{A}_{1, \zeta - \widetilde{\text{tag}}}^{\mathcal{O}_0}(\widetilde{w}_j^\kappa)$.
 - If there exists $j \in [p]$ such that $\widetilde{a}_j^\kappa = \perp$ or $\widetilde{a}'_j^\kappa = \perp$, then set $m_\kappa = \perp$.
 - Else, set $m_\kappa = (m_1^\kappa, m_2^\kappa, \dots, m_p^\kappa)$, where $m_j^\kappa = \widetilde{a}_j^\kappa \oplus \widetilde{a}'_j^\kappa$.

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, $(\widetilde{m}_1, \dots, \widetilde{m}_k)$ are all extracted correctly with overwhelming probability.

8. Compute $e = \mathcal{D}'(c, \widetilde{c}_1, \dots, \widetilde{c}_k, \widetilde{m}_1, \dots, \widetilde{m}_k)$.
9. If $e = 0$, output $b' = a_{\alpha+1}$.
10. If $e = 1$, sample and output a uniformly random bit b' .

By Equation (3), together with the fact that $\widetilde{m}_1, \dots, \widetilde{m}_k$ 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 $a_{\alpha+1} \xleftarrow{\$} \{0, 1\}$ (independently of b), this implies that

$$\begin{aligned} & \Pr[(a_{\alpha+1} \oplus b = 0) \wedge (e = 0) | \widetilde{\text{tag}} > \text{tag}] - \\ & \Pr[(a_{\alpha+1} \oplus b = 1) \wedge (e = 0) | \widetilde{\text{tag}} > \text{tag}] \geq \frac{\Delta}{2(p+1)}. \end{aligned}$$

which in turn implies that

$$\begin{aligned} & \Pr[(b' = b) \wedge (e = 0) | \widetilde{\text{tag}} > \text{tag}] - \\ & \Pr[(b' \neq b) \wedge (e = 0) | \widetilde{\text{tag}} > \text{tag}] \geq \frac{\Delta}{2(p+1)}. \end{aligned}$$

which implies

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

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. \square

4.2 Instantiating the Assumption

Claim 1. Assume the existence of constants $\delta > \delta' > 0$, and two families of non-interactive bit commitments $\mathcal{C}_0 : \{0, 1\} \times \{0, 1\}^n \rightarrow \{0, 1\}^*$ and $\mathcal{C}_1 : \{0, 1\} \times \{0, 1\}^{n^\delta} \rightarrow \{0, 1\}^*$, so that:

- For every $2^{n^{2\delta}}$ -size adversary \mathcal{A} ,

$$|\Pr[\mathcal{A}(\mathcal{C}_0(0; r)) = 0] - \Pr[\mathcal{A}(\mathcal{C}_0(1; r)) = 0]| = \text{negl}(n)$$

over the randomness of sampling $r \xleftarrow{\$} \{0, 1\}^n$.

- There exists a polynomial-size algorithm \mathcal{A}_0 such that for every string c , $\mathcal{A}_0^{\text{BQP}}(c)$ outputs $(b, r) \in \{0, 1\}^{n+1}$, such that

$$\Pr[(c = \mathcal{C}_0(b, r)) \vee (\exists (b', r') \text{ such that } c = \mathcal{C}_0(b'; r'))] = 1 - \text{negl}(n)$$

- For every $2^{n^{\delta'}}$ -size adversary \mathcal{A} ,

$$|\Pr[\mathcal{A}^{\text{BQP}}(\mathcal{C}_1(0; r)) = 0] - \Pr[\mathcal{A}^{\text{BQP}}(\mathcal{C}_1(1; r)) = 0]| = \text{negl}(n)$$

over the randomness of sampling $r \xleftarrow{\$} \{0, 1\}^{n^\delta}$.

Then [Assumption 1](#) follows by defining

$$\text{com}_0(b; r_1, r_2, \dots, r_n) \triangleq \mathcal{C}_0(b; r_1, r_2, \dots, r_n),$$

and

$$\text{com}_1(b; r_1, r_2, \dots, r_n) \triangleq \mathcal{C}_1(b; r_1, r_2, \dots, r_n),$$

and setting $\mathcal{O}_0 = \text{BQP}$ and $\mathcal{O}_1 = \text{DTIME}(2^{n^\delta})$.

Concretely, it suffices to instantiate \mathcal{C}_0 to be any non-interactive commitment scheme whose hiding is based on the sub-exponential hardness of any problem that is invertible given a BQP oracle; for example, subexponential hardness of Factoring or Discrete Log.

Such a commitment can be obtained from any injective one-way function $F_0 : \{0, 1\}^n \rightarrow \{0, 1\}^*$ for which the one-wayness of F_0 is based on the sub-exponential hardness of Factoring (or Discrete Log), and that F_0 is invertible everywhere given access to an oracle that breaks Factoring (or Discrete Log) respectively. Given such an injective one-way function, the bit commitment \mathcal{C}_0 can be defined as follows: On input bit b and randomness $r, s \in \{0, 1\}^n$, output $F_0(r), s, (\langle r, s \rangle \oplus b)$. By the Goldreich-Levin hardcore theorem, the bit b is computationally hidden as long as the function F_0 is one-way.

\mathcal{C}_1 can be instantiated as any commitment where the hiding property holds against sub-exponential quantum adversaries. Such commitments were recently constructed by [GHKW17](#) assuming sub-exponential quantum hardness of lattice-based problems, such as the learning with errors (LWE) problem or the learning parity with noise (LPN) problem. These can also be instantiated based on any injective one-way functions with sub-exponential quantum hardness, by relying on the hardcore bit technique as described above.

Below, we describe how a simplification of the assumption above gives tag-based (many-to-many) adaptive commitments for two tags, based only on sub-exponential hardness of factoring/discrete log, and quantum polynomial hardness of LWE.

Remark 4. Assume the existence of a constant $0 < \delta < 1$ and two families of bit commitments $\mathcal{C}_0 : \{0, 1\} \times \{0, 1\}^{n^{1/\delta}} \rightarrow \{0, 1\}^*$ and $\mathcal{C}_1 : \{0, 1\} \times \{0, 1\}^n \rightarrow \{0, 1\}^*$, such that the following is true:

- For every 2^n -size adversary \mathcal{A} ,

$$|\Pr[\mathcal{A}(\mathcal{C}_0(0; r)) = 0] - \Pr[\mathcal{A}(\mathcal{C}_0(1; r)) = 0]| = \text{negl}(n)$$

over the randomness of sampling $r \xleftarrow{\$} \{0, 1\}^{n^{1/\delta}}$.

- There exists a polynomial-size algorithm \mathcal{A}_0 such that for every string c , $\mathcal{A}_0^{\text{BQP}}(c) \rightarrow (b, r)$ such that

$$\Pr[(c = \mathcal{C}_0(b, r)) \vee (\nexists b', r' \text{ such that } c = \mathcal{C}_0(b'; r'))] = 1 - \text{negl}(n)$$

- For every poly-size adversary \mathcal{A} ,

$$|\Pr[\mathcal{A}^{\text{BQP}}(\mathcal{C}_1(0; r)) = 0] - \Pr[\mathcal{A}^{\text{BQP}}(\mathcal{C}_1(1; r)) = 0]| = \text{negl}(n)$$

over the randomness of sampling $r \xleftarrow{\$} \{0, 1\}^n$.

Then, a tag-based adaptive commitment scheme for two tags, against polynomial-sized adversaries, can be defined as follows: $\text{com}_0 \triangleq \mathcal{C}_0$ and $\text{com}_1 \triangleq \mathcal{C}_1$; $\mathcal{O}_0 = \text{BQP}$ and $\mathcal{O}_1 = \text{DTIME}(2^n)$.

By [PPV08], this implies non-interactive many-to-many non-malleable commitment for two tags. In this case, \mathcal{C}_0 can be instantiated as any non-interactive bit commitment whose hardness is based on the sub-exponential hardness of factoring or discrete log, and \mathcal{C}_1 as a non-interactive bit commitment based on quantum polynomial hardness of LWE or LPN [GHKW17].

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 $T_V \geq \mathcal{T}$.

Theorem 2. 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.

Parameters: Set $k = 10\ell$ and let \mathbb{T} denote the unordered set of all possible subsets of $[t]$ of size $t/2$.

Language L : We define the language $L = \{\{C_{\lambda,j}, s_\lambda\}_{\lambda \in [t/2], j \in [k]} : \exists J \subset [k], |J| = k - 1, \exists \{M_j, r_{\lambda,j}\}_{j \in J, \lambda \in [t/2]} \text{ s.t. } C_{\lambda,j} = \text{Com}_{s_\lambda}(M_j; r_{\lambda,j}) \forall j \in J, \lambda \in [t/2]\}$.

Committer Input: Message $M \in \{0, 1\}^{p(n)}$, and tag $\text{tag} \in [N]$, where $N = \binom{t}{t/2}$.

Commit Stage: To commit to a message M w.r.t. tag tag , do the following:

1. Pick the i^{th} element in \mathbb{T} , for $i = \text{tag}$. Denote this element by $\{s_1, \dots, s_{t/2}\}$.
2. **Committer Message.** For every $\lambda \in [t/2]$ and every $j \in [k]$, sample randomness $r_{\lambda,j} \xleftarrow{\$} \{0, 1\}^{\text{poly}(n)}$ and compute $C_{\lambda,j} = \text{Com}_{s_\lambda}(M; r_{\lambda,j})$. Use witness $J = [k] \setminus \{1\}, \{M, r_{\lambda,j}\}_{\lambda \in [t/2], j \in J}$ to compute a NIWI proof Π for the statement:

$$\{(C_{\lambda,j}, s_\lambda)\}_{\lambda \in [t/2], j \in [k]} \in L.$$

Send to the receiver the message

$$(\text{tag}, \{(C_{\lambda,j}, s_\lambda)\}_{\lambda \in [t/2], j \in [k]}, \Pi)$$

3. **Receiver Acceptance.** The receiver accepts the commitment $(\text{tag}, \{(C_{\lambda,j}, s_\lambda)\}, \Pi)$ if and only if the proof Π is accepted by the verifier of the NIWI system and $\{s_\lambda\}_{\lambda \in [t/2]}$ is the i^{th} element in \mathbb{T} for $i = \text{tag}$.

Reveal Stage: The committer sends the message M and the randomness $\{r_{\lambda,j}\}_{\lambda \in [t/2], j \in [k]}$ to the receiver.

Receiver Output: The receiver verifies that all the commitments were correctly decommitted, and accepts the decommitment $(M, \{r_{\lambda,j}\}_{\lambda \in [t/2], j \in [k]})$ if and only if $\forall \lambda, j \in [t/2] \times [k] : C_{\lambda,j} = \text{Com}_{s_\lambda}(M; r_{\lambda,j})$.

Figure 2: Round-Preserving Tag Amplification

We will later invoke the compiler presented in Figure 2 with the scheme for $\eta \cdot \log \log n$ tags, that we constructed in Section 4, to obtain a scheme for $(\log n)^\epsilon$ tags for some small constant $\epsilon > 0$. Then we will invoke this compiler again on the resulting scheme to obtain a scheme for $2 \log^2 n$, and then a third (and final) time on the resulting scheme to obtain a scheme for $n^{\log n}$ tags. We refer to Section 6 for details.

Proof. Statistical binding of the protocol in Figure 2 follows directly from the statistical binding property of the underlying commitment scheme Com . Computational hiding follows from the hiding of the underlying commitments, and the witness indistinguishability property of Π . Since the proof of non-malleability subsumes hiding, we directly prove non-malleability below. Specifically, we prove auxiliary-input non-malleability w.r.t. replacement for $T'_V = T_V + (n \cdot 2^y)^{c \log n}$, where T_V is the running time of $\mathcal{V}_{\text{Real}}$ and $\mathcal{V}_{\text{Ideal}}$ corresponding to the underlying scheme Com , and where $c \in \mathbb{N}$ is a constant that will be specified towards the end of the proof. As mentioned above, throughout the proof, we assume for simplicity that $T_V > \text{poly}(n \cdot 2^y)$ for all polynomials $\text{poly}(\cdot)$.

To this end, fix any poly-size adversary \mathcal{A} , any messages $m_1, \dots, m_\ell \in \{0, 1\}^{p(n)}$, and any

tags $\text{tag}_1, \dots, \text{tag}_\ell \in [N]$, where $N = \binom{t}{t/2}$. Let $(\text{aux}_1, \dots, \text{aux}_\ell)$ denote any auxiliary information functions according [Definition 3](#). We use Com_s to denote the non-malleable commitment for small tags and com_{tag} to denote the scheme from [Figure 2](#). For every $\beta \in [0, \ell]$ we define τ_β as follows:

- Set $c_{\beta, \alpha} = \text{com}_{\text{tag}_\alpha}(0^p; r_\alpha)$ where $r_\alpha \xleftarrow{s} \{0, 1\}^*$, for every $\alpha \in [1, \beta]$.
- Set $c_{\beta, \alpha} = \text{com}_{\text{tag}_\alpha}(m_\alpha; r_\alpha)$ where $r_\alpha \xleftarrow{s} \{0, 1\}^*$, for every $\alpha \in [\beta + 1, \ell]$.
- Set $\tau_\beta \triangleq (c_{\beta, 1}, \dots, c_{\beta, \ell}, \mathbf{a}_\beta, \tilde{c}_1, \dots, \tilde{c}_y, z_\beta)$ where

$$\mathbf{a}_\beta = \text{aux}_\beta(c_{\beta, 1}, \dots, c_{\beta, \ell}, (0^p)_{\times[\beta-1]}, m_{\beta+1}, \dots, m_\ell, r_1, \dots, r_{\beta-1}, r_{\beta+1}, \dots, r_\ell)$$

$$\text{and } (\tilde{c}_1, \dots, \tilde{c}_y, z_\beta) = \mathcal{A}(c_{\beta, 1}, \dots, c_{\beta, \ell}, \mathbf{a}_\beta).$$

We need to prove that there exist (inefficient) functions $\mathcal{V}_{\text{real}}$ and $\mathcal{V}_{\text{ideal}}$ satisfying the validity condition in [Definition 3](#) such that

$$\left(\tau_0, \mathcal{V}_{\text{Real}}(\tau_0) \right) \approx \left(\tau_\ell, \mathcal{V}_{\text{Ideal}}(\tau_\ell) \right).$$

We defer the description $\mathcal{V}_{\text{Real}}$ and $\mathcal{V}_{\text{Ideal}}$ to the end of this proof. In what follows, we consider an experiment where the PPT man-in-the-middle adversary \mathcal{A} receives input commitments c_i for every $i \in [\ell]$, and define a function that extracts the values committed by \mathcal{A} .

To this end, we parse $\widetilde{\text{tag}}$ output by \mathcal{A} as comprising of an unordered set of small tags, denoted by $\{\tilde{s}_1, \tilde{s}_2, \dots, \tilde{s}_{t/2}\}$, and $\text{tag}_i = \{s_{1,i}, s_{2,i}, \dots, s_{t/2,i}\}$ for $i \in [\ell]$. We note that if $\widetilde{\text{tag}} \notin \{\text{tag}_1, \dots, \text{tag}_\ell\}$, then for every $i \in [\ell]$ there exists at least one index $\tilde{\lambda}_i \in [t/2]$ such that $\tilde{s}_{\tilde{\lambda}_i} \notin \{s_{1,i}, s_{2,i}, \dots, s_{t/2,i}\}$.

As an intermediate step, for every $\beta \in [\ell]$, we define (inefficient) functions $\mathcal{V}_{\beta, \text{real}}(\tau)$ and $\mathcal{V}_{\beta, \text{ideal}}(\tau)$ as follows:

$$\mathcal{V}_{\beta, \text{real}}(\tau) = (\chi_1, \dots, \chi_y)$$

where each $\chi_i = (\tilde{M}_i, \psi_i)$, and \tilde{M}_i is a value extracted from the commitment \tilde{c}_i . We require that for each $i \in [y]$ the extracted value \tilde{M}_i is correct whenever \tilde{c}_i is a valid commitment. More specifically, $\mathcal{V}_{\beta, \text{real}}(\tau)$ is defined as follows:

- Parse $\tau = (c_1, \dots, c_\ell, \mathbf{a}, \tilde{c}_1, \dots, \tilde{c}_y, z)$.
- Parse $c_\beta = (\text{tag}_\beta, \{(C_{\lambda,j}, s_\lambda)\}_{\lambda \in [t/2], j \in [k]}, \Pi)$.
- For each $i \in [y]$, parse $\tilde{c}_i = (\widetilde{\text{tag}}, \{(\tilde{C}_{\lambda,j}^i, \tilde{s}_\lambda)\}_{\lambda \in [t/2], j \in [k]}, \Pi_i)$.
- For every $i \in [y]$ such that Π_i is rejected, set $\chi_i = \text{REJECT}$ (note that this step is efficient).
- For $\widetilde{\text{tag}} = \{\tilde{s}_1, \dots, \tilde{s}_{t/2}\}$, let $\tilde{\lambda}_\beta$ denote the smallest index such that $\tilde{s}_{\tilde{\lambda}_\beta} \notin \{s_{1,\beta}, s_{2,\beta}, \dots, s_{t/2,\beta}\}$.
- Extract (χ_1, \dots, χ_y) from $(\tilde{c}_1, \dots, \tilde{c}_y)$ as follows:
 1. We define auxiliary-input functions $\{\text{au}_{\lambda,j}\}_{\lambda \in [t/2], j \in [k]}$ for the underlying commitment scheme Com , as follows:

(a) For each $j' \in [k]$ define an auxiliary input function $\text{au}_{j'}$, that on input

$$\left(\{C'_{\lambda,j}, s'_\lambda\}_{\lambda \in [t/2], j \in [k]}, \{M'_{\lambda,j}\}_{\lambda \in [t/2], j \in [k] \setminus \{j'\}}, \{r'_{\lambda,j}\}_{\lambda \in [t/2], j \in [k] \setminus \{j'\}} \right)$$

sets

$$X = \{(C'_{\lambda,j}, s'_\lambda)\}_{\lambda \in [t/2], j \in [k]} \text{ and } W = ([k] \setminus \{j'\}, \{M_{\lambda,j}, r_{\lambda,j}\}_{\lambda \in [t/2], j \in [k] \setminus \{j'\}}).$$

If $R_L(X, W) = 1$ then it outputs a NIWI for the statement $X \in L^{17}$, else it outputs \perp .

(b) For all $\lambda \in [t/2]$ and $j \in [k]$, set $\text{au}_{\lambda,j} = \text{au}_j^{18}$.

2. We define an adversary \mathcal{A}' for the underlying commitment Com , that has $\beta, m_1, \dots, m_\ell, \text{tag}_1, \dots, \text{tag}_\ell$ hardwired, and on input $\left(\{(C'_{\lambda,j}, s'_\lambda)\}_{\lambda \in [t/2], j \in [k]}, \text{au}' \right)$ does the following:

- Sample $r_1, \dots, r_\ell \xleftarrow{\$} \{0, 1\}^*$.
- For all $i \in [\ell] \setminus \{\beta\}$, generate c_i using randomness r_i , according to the distribution of $c_{\beta-1,i}$ defined above.
- Set $c_\beta = (\text{tag}_\beta, \{(C'_{\lambda,j}, s'_\lambda)\}_{\lambda \in [t/2], j \in [k]}, \text{au}')$.
- Compute $\mathbf{a}_\beta = \text{aux}_\beta(c_1, \dots, c_\ell, (0^p)_{\times(\beta-1)}, m_{\beta+1}, \dots, m_\ell, r_1, \dots, r_{\beta-1}, r_{\beta+1}, \dots, r_\ell)$.
- Execute $\mathcal{A}(c_1, \dots, c_\ell, \mathbf{a}_\beta)$ to obtain $(\tilde{c}_1, \dots, \tilde{c}_y, z_\beta)$.
- Let $Z = (z_\beta, c_1, \dots, c_{\beta-1}, c_{\beta+1}, \dots, c_\ell, \tilde{c}_1, \dots, \tilde{c}_y)$.
- For each $i \in [y]$, parse $\tilde{c}_i = (\tilde{\text{tag}}, \{(\tilde{C}^i_{\lambda,j}, \tilde{s}_\lambda)\}_{\lambda \in [t/2], j \in [k]}, \Pi_i)$.
- Output $\left(\{(\tilde{C}^i_{\lambda_\beta,j}, \tilde{s}_{\lambda_\beta})\}_{i \in [y], j \in [k]}, Z \right)$.

3. For all $\lambda \in [t/2]$ and $j \in [k]$, set $m_{\lambda,j} = m_\beta$.

4. For all $\lambda \in [t/2]$ and $j \in [k]$, set $s_{\lambda,j} = s'_\lambda$, where $\{s'_\lambda\}_{\lambda \in [t/2]}$ are the input tags.

5. For $\{m_{\lambda,j}\}_{\lambda \in [t/2], j \in [k]}, \{s_{\lambda,j}\}_{\lambda \in [t/2], j \in [k]}, \{\text{au}_{\lambda,j}\}_{\lambda \in [t/2], j \in [k]}$, and \mathcal{A}' defined above, there exist distributions $\mathcal{V}_{\text{Real}}$ and $\mathcal{V}_{\text{Ideal}}$ corresponding to the commitment scheme Com , satisfying [Definition 3](#), which we denote by v_{Real} and v_{Ideal} .¹⁹ Compute

$$\{X^i_{\lambda_\beta,j}\}_{i \in [y], j \in [k]} = v_{\text{Real}} \left(\{C_{\lambda,j}, s_\lambda\}_{\lambda \in [t/2], j \in [k]}, \Pi, \{(\tilde{C}^i_{\lambda_\beta,j}, \tilde{s}_{\lambda_\beta})\}_{i \in [y], j \in [k]} \right).$$

6. For each $i \in [y]$, if there exists a message \tilde{M} and subset $J \subseteq [k]$ where $|J| \geq (\frac{3k}{4} + 1 - \beta)$, such that for all $j \in J$, $X^i_{\lambda_\beta,j} = \tilde{M}$, set $\chi_i = (\tilde{M}, |J|)$, else set $\chi_i = \perp$.

- Output (χ_1, \dots, χ_y) .

We define $\mathcal{V}_{\beta, \text{ideal}}(\tau)$ identically to $\mathcal{V}_{\beta, \text{real}}$, except that in Step 5, it computes

$$\{X^i_{\lambda_\beta,j}\}_{i \in [y], j \in [k]} = v_{\text{Ideal}} \left(\{C_{\lambda,j}, s_\lambda\}_{\lambda \in [t/2], j \in [k]}, \Pi, \{(\tilde{C}^i_{\lambda_\beta,j}, \tilde{s}_{\lambda_\beta})\}_{i \in [y], j \in [k]} \right).$$

We have the following observation about the efficiency of $\mathcal{V}_{\beta, \text{ideal}}$ and $\mathcal{V}_{\beta, \text{real}}$ for $\beta \in [\ell]$.

¹⁷Refer to [Figure 2](#) for a description of the language L .

¹⁸Note that $\text{au}_{\lambda,j}$ as a function gets as input additional messages and randomness, which he ignores.

¹⁹While [Definition 3](#) considers a vector of messages m_1, \dots, m_ℓ , (respectively tags t_1, \dots, t_ℓ and commitments c_1, \dots, c_ℓ), here and in the rest of the proof, we often abuse notation and denote the $t\ell/2$ “left” messages as a $t \times \ell/2$ matrix of messages (and similarly for tags and commitments).

Claim 2. For every $\beta \in [\ell]$, the function $\mathcal{V}_{\beta,\text{ideal}}$ and the function $\mathcal{V}_{\beta,\text{real}}$ are computable in time $T_V(n) + \text{poly}(n)$ where $T_V(n)$ is the time required to compute $\mathcal{V}_{\text{Real}}$ and $\mathcal{V}_{\text{Ideal}}$ for Com.

We will make use of $\mathcal{V}_{1,\text{real}}$ and $\mathcal{V}_{\ell,\text{ideal}}$ to define the functions $\mathcal{V}_{\text{Real}}$ and $\mathcal{V}_{\text{Ideal}}$ respectively. Before describing these functions, we prove some lemmas about the distributions output by the functions $\mathcal{V}_{\beta,\text{real}}, \mathcal{V}_{\beta,\text{ideal}}$ described above.

Lemma 1. For all $\beta \in [\ell]$,

$$\left(\tau_{\beta-1}, \mathcal{V}_{\beta,\text{real}}(\tau_{\beta-1}) \right) \approx_{\text{poly}(n \cdot 2^y)} \left(\tau_{\beta}, \mathcal{V}_{\beta,\text{ideal}}(\tau_{\beta}) \right)$$

where $\tau_{\beta-1}, \tau_{\beta}$ are defined above.

Proof. Fix any $\beta \in [\ell]$. Note that $c_{\beta,\alpha} \equiv c_{\beta-1,\alpha}$ for every $\alpha \neq \beta$, and the only difference between the two distributions is the following.

- $c_{\beta-1,\beta} \stackrel{\$}{\leftarrow} \text{com}_{\text{tag}_{\beta}}(m_{\beta})$ whereas $c_{\beta,\beta} \stackrel{\$}{\leftarrow} \text{com}_{\text{tag}_{\beta}}(0)$.
- $\tau_{\beta-1}$ is defined with auxiliary-input function $\text{aux}_{\beta-1}$, whereas τ_{β} is defined with auxiliary-input function aux_{β} .

We prove this lemma via two hybrid experiments. These hybrids rely on auxiliary-input non-malleability w.r.t. replacement ([Definition 3](#)) of Com for tags in $[t]$, against $\text{poly}(n \cdot 2^y)$ -size machines. We will also rely on the witness indistinguishability of the NIWI against $\text{poly}(T_V)$ -size machines.

In the following hybrids, we start with $(\tau_{\beta-1}, \mathcal{V}_{\beta,\text{real}}(\tau_{\beta-1}))$. We then switch from using $\text{aux}_{\beta-1}$ to using aux_{β} in $\tau_{\beta-1}$. Finally, we switch from generating the β^{th} commitment as a commitment to m_{β} to generating it as a commitment to 0^p . We note the differences between the hybrids by underlining them.

Hybrid₀ : This hybrid is defined as follows:

- Compute $(c_1, \dots, c_{\ell}) := (c_{\beta-1,1}, \dots, c_{\beta-1,\ell})$ as described above, by sampling randomness (r_1, \dots, r_{ℓ}) .
- Compute $\mathbf{a}_{\beta-1} := \text{aux}_{\beta-1}(c_1, \dots, c_{\ell}, (0^p)_{\times(\beta-2)}, m_{\beta}, \dots, m_{\ell}, r_1, \dots, r_{\beta-2}, r_{\beta}, \dots, r_{\ell})$.
- Obtain $(\tilde{c}_1, \dots, \tilde{c}_y, z) = \mathcal{A}(c_1, \dots, c_{\ell}, \mathbf{a}_{\beta-1})$.
- Set $\tau_{\beta-1} = (c_1, \dots, c_{\ell}, \mathbf{a}_{\beta-1}, \tilde{c}_1, \dots, \tilde{c}_y, z)$.
- Output $(\tau_{\beta-1}, \mathcal{V}_{\beta,\text{real}}(\tau_{\beta-1}))$ for $\mathcal{V}_{\beta,\text{real}}$ defined above.

Hybrid₁ : This is the same as the previous hybrid, except that it uses aux_{β} instead of $\text{aux}_{\beta-1}$. Formally, this hybrid is defined as follows:

- Compute $(c_1, \dots, c_{\ell}) := (c_{\beta-1,1}, \dots, c_{\beta-1,\ell})$ as described above, by sampling randomness (r_1, \dots, r_{ℓ}) .
- Compute $\mathbf{a}_{\beta} := \text{aux}_{\beta}(c_1, \dots, c_{\ell}, (0^p)_{\times(\beta-1)}, m_{\beta+1}, \dots, m_{\ell}, r_1, \dots, r_{\beta-1}, r_{\beta+1}, \dots, r_{\ell})$.
- Obtain $(\tilde{c}_1, \dots, \tilde{c}_y, z) = \mathcal{A}(c_1, \dots, c_{\ell}, \mathbf{a}_{\beta})$.
- Set $\tau = (c_1, \dots, c_{\ell}, \mathbf{a}_{\beta}, \tilde{c}_1, \dots, \tilde{c}_y, z)$.
- Output $(\tau, \mathcal{V}_{\beta,\text{real}}(\tau))$ for $\mathcal{V}_{\beta,\text{real}}$ defined above.

Recall that in both Hybrid₀ and Hybrid₁, (c_1, \dots, c_ℓ) are generated according to the distribution $(c_{\beta-1,1}, \dots, c_{\beta-1,\ell})$. By Equation (1) (see Definition 3)

$$(c_{\beta-1,1}, \dots, c_{\beta-1,\ell}, \text{aux}_{\beta-1}(c_{\beta-1,1}, \dots, c_{\beta-1,\ell}, m_1, \dots, m_{\beta-2}, m_\beta, \dots, m_\ell, r_1, \dots, r_{\beta-2}, r_\beta, \dots, r_\ell)) \approx_{\text{poly}(T_V)} \\ (c_{\beta-1,1}, \dots, c_{\beta-1,\ell}, \text{aux}_\beta(c_{\beta-1,1}, \dots, c_{\beta-1,\ell}, m_1, \dots, m_{\beta-1}, m_{\beta+1}, \dots, m_\ell, r_1, \dots, r_{\beta-1}, r_{\beta+1}, \dots, r_\ell)).$$

This, together with the fact that \mathcal{A} runs in polynomial time and $V_{\beta,\text{real}}(\tau)$ can be computed in time $T_V + \text{poly}(n)$, implies that

$$\text{Hybrid}_0 \approx_{\text{poly}(T_V)} \text{Hybrid}_1.$$

Hybrid₂ : This is the same as the previous hybrid, except that it computes c_β as a commitment to 0 (instead of a commitment to m_β). It uses $\mathcal{V}_{\beta,\text{ideal}}$ instead of $\mathcal{V}_{\beta,\text{real}}$ to compute χ_1, \dots, χ_y . Formally, this hybrid is defined as follows:

- Compute $(c_1, \dots, c_\ell) := (c_{\beta,1}, \dots, c_{\beta,\ell})$ as described above, by sampling randomness (r_1, \dots, r_ℓ) .
- Compute $\mathbf{a}_\beta := \text{aux}_\beta(c_1, \dots, c_\ell, (0^p)_{\times(\beta-1)}, m_{\beta+1}, \dots, m_\ell, r_1, \dots, r_{\beta-1}, r_{\beta+1}, \dots, r_\ell)$.
- Obtain $(\tilde{c}_1, \dots, \tilde{c}_y, z) = \mathcal{A}(c_1, \dots, c_\ell, \mathbf{a}_\beta)$.
- Set $\tau_\beta = (c_1, \dots, c_\ell, \mathbf{a}_\beta, \tilde{c}_1, \dots, \tilde{c}_y, z)$.
- Output $(\tau_\beta, \mathcal{V}_{\beta,\text{ideal}}(\tau_\beta))$ for $\mathcal{V}_{\beta,\text{ideal}}$ defined above.

Note that $c_1, \dots, c_{\beta-1}, c_{\beta+1}, \dots, c_\ell$ are generated identically in Hybrid₁ and Hybrid₂. We prove that Hybrid₁ $\approx_{\text{poly}(n \cdot 2^y)}$ Hybrid₂ by $\text{poly}(n \cdot 2^y)$ non-malleability of the underlying commitment scheme Com.

To this end, suppose there exists a $\text{poly}(n \cdot 2^y)$ -time distinguisher that distinguishes Hybrid₁ and Hybrid₂ with advantage $1/\text{poly}(n \cdot 2^y)$ (for infinitely many $n \in \mathbb{N}$). We will use this distinguisher to contradict $kt/2$ -to- ky same-tag auxiliary-input non-malleability w.r.t. replacement of Com against $\text{poly}(n \cdot 2^y)$ -size adversaries, according to Definition 3, as follows.

We define auxiliary-input functions $\{\text{au}_{\lambda,j}\}_{\lambda \in [t/2], j \in [k]}$, an adversary \mathcal{A}' , and messages and tags corresponding to the underlying commitment scheme Com. These are defined identically to the ones defined in the description of $\mathcal{V}_{\beta,\text{real}}$ and $\mathcal{V}_{\beta,\text{ideal}}$. We recall the definitions below.

1. We define $\{\text{au}_{\lambda,j}\}_{\lambda \in [t/2], j \in [k]}$ as follows:

- (a) For each $j' \in [k]$, define an auxiliary input function $\text{au}_{j'}$ that on input

$$(\{C'_{\lambda,j}, s'_\lambda\}_{\lambda \in [t/2], j \in [k]}, \{M'_{\lambda,j}\}_{\lambda \in [t/2], j \in [k] \setminus \{j'\}}, \{r'_{\lambda,j}\}_{\lambda \in [t/2], j \in [k] \setminus \{j'\}}),$$

sets

$$X = \{(C'_{\lambda,j}, s'_\lambda)\}_{\lambda \in [t/2], j \in [k]} \text{ and } W = ([k] \setminus \{j'\}, \{M_{\lambda,j}, r_{\lambda,j}\}_{\lambda \in [t/2], j \in [k] \setminus \{j'\}}).$$

If $R_L(X, W) = 1$, then it outputs a NIWI for the statement $X \in L^{20}$, else it outputs \perp .

- (b) For all $\lambda \in [t/2]$ and $j \in [k]$, set $\text{au}_{\lambda,j} = \text{au}_j^{21}$.

²⁰Refer to Figure 2 for a description of the language L .

²¹Note that $\text{au}_{\lambda,j}$ takes as input additional messages and randomness which he ignores.

2. We define adversary \mathcal{A}' that has $\beta, m_1, \dots, m_\ell, \text{tag}_1, \dots, \text{tag}_\ell$ hardwired, and on input

$$\left(\{(C'_{\lambda,j}, s'_\lambda)\}_{\lambda \in [t/2], j \in [k]}, \mathbf{au}' \right),$$

does the following:

- (a) Sample $r_1, \dots, r_\ell \xleftarrow{\$} \{0, 1\}^*$.
 - (b) For all $i \in [\ell] \setminus \{\beta\}$, generate c_i using randomness r_i , according to the distribution of $c_{\beta-1,i}$ defined above.
 - (c) Set $c_\beta = (\text{tag}_\beta, \{(C'_{\lambda,j}, s'_\lambda)\}_{\lambda \in [t/2], j \in [k]}, \mathbf{au}'$).
 - (d) Compute $\mathbf{a}_\beta = \mathbf{au}_{\times\beta}(c_1, \dots, c_\ell, (0^p)_{\times(\beta-1)}, m_{\beta+1}, \dots, m_\ell, r_1, \dots, r_{\beta-1}, r_{\beta+1}, \dots, r_\ell)$.
 - (e) Execute $\mathcal{A}(c_1, \dots, c_\ell, \mathbf{a}_\beta)$ to obtain $(\tilde{c}_1, \dots, \tilde{c}_y, z)$.
 - (f) Set $Z = (z, c_1, \dots, c_\beta, c_{\beta+1}, \dots, c_\ell, \tilde{c}_1, \dots, \tilde{c}_y)$.
 - (g) For each $i \in [y]$, parse $\tilde{c}_i = (\tilde{\text{tag}}, \{(\tilde{C}^i_{\lambda,j}, \tilde{s}_\lambda)\}_{\lambda \in [t/2], j \in [k]}, \Pi_i)$.
 - (h) Output $\left(\{(\tilde{C}^i_{\lambda_\beta,j}, \tilde{s}_{\lambda_\beta})\}_{i \in [y], j \in [k]}, Z \right)$.
3. For all $\lambda \in [t/2]$ and $j \in [k]$, set $m_{\lambda,j} = m_\beta$.
4. For all $\lambda \in [t/2]$ and $j \in [k]$, set $s_{\lambda,j} = s'_\lambda$, where $\{s'_\lambda\}_{\lambda \in [t/2]}$ are the input tags.

Note that $\{m_{\lambda,j}\}_{\lambda \in [t/2], j \in [k]}, \{s_{\lambda,j}\}_{\lambda \in [t/2], j \in [k]}, \{\mathbf{au}_{\lambda,j}\}_{\lambda \in [t/2], j \in [k]}$, as defined above, satisfy Equation (1) (Definition 3), corresponding to the underlying commitment Com. Therefore, for

$$\{m_{\lambda,j}\}_{\lambda \in [t/2], j \in [k]}, \{s_{\lambda,j}\}_{\lambda \in [t/2], j \in [k]}, \{\mathbf{au}_{\lambda,j}\}_{\lambda \in [t/2], j \in [k]}, \mathcal{A}'$$

there exist distributions $\mathcal{V}_{\text{Real}}$ and $\mathcal{V}_{\text{Ideal}}$, which we denote by v_{Real} and v_{Ideal} , such that:

$$(\tau_0, v_{\text{Real}}(\tau_0)) \approx_{\text{poly}(n \cdot 2^y)} (\tau_{kt/2}, v_{\text{Ideal}}(\tau_{kt/2})) \quad (4)$$

where $\tau_0 = (\{C^0_{\lambda,j}\}_{\lambda \in [t/2], j \in [k]}, \mathbf{au}^0, \mathcal{A}'(\{C^0_{\lambda,j}\}_{\lambda \in [t/2], j \in [k]}, \mathbf{au}^0))$ is such that for every $\lambda \in [t/2], j \in [k]$,

- $r_{\lambda,j} \xleftarrow{\$} \{0, 1\}^*$,
- $C^0_{\lambda,j} = \text{Com}_{\text{tag}_{\lambda,j}}(m_\beta, r_{\lambda,j})$, and
- $\mathbf{au}^0 = \mathbf{au}_{kt/2}(\{C^0_{\lambda,j}\}_{\lambda \in [t/2], j \in [k]}, (m_\beta)_{\times(kt/2-1)}, \{r_{\lambda,j}\}_{(\lambda,j) \in [t/2] \times [k] \setminus \{(t/2,k)\}})$.

Similarly, $\tau_{kt/2} = (\{C^{kt/2}_{\lambda,j}\}_{\lambda \in [t/2], j \in [k]}, \mathbf{au}^{kt/2}, \mathcal{A}'(\{C^{kt/2}_{\lambda,j}\}_{\lambda \in [t/2], j \in [k]}, \mathbf{au}^{kt/2}))$ is such that for every $\lambda \in [t/2], j \in [k]$,

- $r_{\lambda,j} \xleftarrow{\$} \{0, 1\}^*$,
- $C^{kt/2}_{\lambda,j} = \text{Com}_{\text{tag}_{\lambda,j}}(0^p, r_{\lambda,j})$, and
- $\mathbf{au}^{kt/2} = \mathbf{au}_{t/2,k}(\{C^0_{\lambda,j}\}_{\lambda \in [t/2], j \in [k]}, (0)_{\times(kt/2-1)}, \{r_{\lambda,j}\}_{(\lambda,j) \in [t/2] \times [k] \setminus \{(t/2,k)\}})$.

Note that Hybrid_1 outputs

$$(c_{\beta-1,1}, \dots, c_{\beta-1,\ell}, \mathbf{a}_\beta, \tilde{c}_1, \dots, \tilde{c}_y, z, v_{\text{Real}}(\tau_0))$$

which is efficiently computable from

$$(\tau_0, v_{\text{Real}}(\tau_0)).$$

Similarly, Hybrid_2 outputs

$$(c_{\beta,1}, \dots, c_{\beta,\ell}, \mathbf{a}_\beta, \tilde{c}_1, \dots, \tilde{c}_y, z, v_{\text{Ideal}}(\tau_{tk/2}))$$

which is efficiently computable from

$$(\tau_{kt/2}, v_{\text{Ideal}}(\tau_{kt/2})).$$

Therefore, if there exists a $\text{poly}(n \cdot 2^y)$ -time distinguisher \mathcal{D} such that for infinitely many $n \in \mathbb{N}$

$$\left| \Pr[\mathcal{D}(\text{Hybrid}_1) = 1] - \Pr[\mathcal{D}(\text{Hybrid}_2) = 1] \right| \geq \frac{1}{\text{poly}(n \cdot 2^y)},$$

then there exists a $\text{poly}(n \cdot 2^y)$ -time distinguisher \mathcal{D}' such that for infinitely many $n \in \mathbb{N}$

$$\left| \Pr[\mathcal{D}'(\tau_0, v_{\text{Real}}(\tau_0)) = 1] - \Pr[\mathcal{D}'(\tau_{t/2,k}, v_{\text{Ideal}}(\tau_{t/2,k})) = 1] \right| \geq \frac{1}{\text{poly}(n \cdot 2^y)}$$

contradicting Equation (4), as desired.

These two hybrids prove that

$$\left(\tau_{\beta-1}, \mathcal{V}_{\beta,\text{real}}(\tau_{\beta-1}) \right) \approx_{\text{poly}(n \cdot 2^y)} \left(\tau_\beta, \mathcal{V}_{\beta,\text{ideal}}(\tau_\beta) \right),$$

completing the proof of the lemma. □

It is tempting to try to prove that for every $\beta \in [\ell - 1]$,

$$\left(\tau, \mathcal{V}_{\beta,\text{ideal}}(\tau) \right) \approx \left(\tau, \mathcal{V}_{\beta+1,\text{real}}(\tau) \right), \quad (5)$$

and then use Lemma 1, together with a standard hybrid argument, to argue that

$$\left(\tau_0, \mathcal{V}_{1,\text{real}}(\tau_0) \right) \approx \left(\tau_\ell, \mathcal{V}_{\beta,\text{ideal}}(\tau_\ell) \right),$$

and thus define $\mathcal{V}_{\text{Real}}$ according to $(\tau_0, \mathcal{V}_{1,\text{real}}(\tau_0))$, while converting each (\tilde{M}_i, ψ_i) to \tilde{M}_i ; and similarly, define $\mathcal{V}_{\text{Ideal}}$ according to $(\tau_\ell, \mathcal{V}_{\beta,\text{ideal}}(\tau_\ell))$, while converting each (\tilde{M}_i, ψ_i) to \tilde{M}_i .

Unfortunately, Equation (5) is not necessarily true. Equation (5) is true if we assume the adversary always generates valid commitments. However, for a general adversary, it may be the case that $\mathcal{V}_{\beta,\text{ideal}}(\tau)$ outputs \perp , whereas $\mathcal{V}_{\beta+1,\text{real}}(\tau)$ does not. The reason is two-fold: First, $\mathcal{V}_{\beta+1,\text{real}}$ extracts the commitments corresponding to $\tilde{s}_{\lambda_{\beta+1}}$ whereas $\mathcal{V}_{\beta,\text{ideal}}$ extracts the commitments corresponding to $\tilde{s}_{\lambda_\beta}$. Second, $\mathcal{V}_{\beta+1,\text{real}}$ is “more forgiving” than $\mathcal{V}_{\beta,\text{ideal}}$.

Jumping ahead, to make the distributions $\mathcal{V}_{\beta,\text{ideal}}(\tau)$ and $\mathcal{V}_{\beta+1,\text{real}}(\tau)$ indistinguishable, we need to sometimes convert the output of $\mathcal{V}_{\beta+1,\text{real}}(\tau)$ to \perp . We do this by using the leakage lemma due to Gentry and Wichs [GW11], as follows.

In the next part of the proof, we fix any $\beta \in [\ell - 1]$ and consider the distributions $\mathcal{V}_{\beta, \text{ideal}}(\tau_\beta)$ and $\mathcal{V}_{\beta+1, \text{real}}(\tau_\beta)$.

In what follows, we consider an arbitrary transcript τ , and for any $\beta \in [\ell - 1]$ and any $i \in [y]$, we denote the output of $\mathcal{V}_{\beta, \text{ideal}}(\tau)$ by (χ_1, \dots, χ_y) , and denote by $\mathcal{V}_{\beta, \text{ideal}}^i(\tau)$ the value χ_i . Similarly, we denote the output of $\mathcal{V}_{\beta+1, \text{real}}(\tau)$ by (χ_1, \dots, χ_y) , and denote by $\mathcal{V}_{\beta+1, \text{real}}^i(\tau)$ the value χ_i . Note that the output χ_i of $\mathcal{V}_{\beta, \text{ideal}}^i$ only depends on values $\{\widetilde{M}_{\lambda_\beta, j}\}_{j \in [k]}$ corresponding to the i^{th} right commitments with (small) tag $\widetilde{s}_{\lambda_\beta} \notin \{s_{1, \beta}, \dots, s_{\ell, \beta}\}$, whereas the output χ_i of $\mathcal{V}_{\beta+1, \text{real}}^i$ only depends on the values $\{\widetilde{M}_{\lambda_{\beta+1}, j}\}_{j \in [k]}$ corresponding to the i^{th} commitments with (small) tag $\widetilde{s}_{\lambda_{\beta+1}} \notin \{s_{1, \beta+1}, \dots, s_{\ell, \beta+1}\}$.

Note that since REJECT is efficiently computable, if $\mathcal{V}_{\beta, \text{ideal}}^i(\tau) = \text{REJECT}$ then $\mathcal{V}_{\beta+1, \text{real}}^i(\tau) = \text{REJECT}$. Moreover, if Π_i is an accepting proof, then by the soundness of the NIWI proof system,²² $\{\widetilde{M}_{\lambda_\beta, j}\}_{j \in [k]}$ and $\{\widetilde{M}_{\lambda_{\beta+1}, j}\}_{j \in [k]}$ in the i^{th} commitment share at least $(k - 1)$ elements. Therefore, by the soundness of the NIWI, the following holds for all τ , where $\tau = (c_1, \dots, c_\ell, \widetilde{c}_1, \dots, \widetilde{c}_y, z)$ denotes any transcript of an execution.

1. $\mathcal{V}_{\beta, \text{ideal}}^i(\tau) = (\widetilde{M}, \psi)$ for $\psi \in \left[\left(\frac{3k}{4} + 1 - \beta \right), k \right] \implies$
 $\mathcal{V}_{\beta+1, \text{real}}^i(\tau) = (\widetilde{M}, \psi')$ where $\psi' \in \{\psi - 1, \psi, \psi + 1\}$ and $\psi' \leq k$.
2. $\mathcal{V}_{\beta, \text{ideal}}^i(\tau) = \perp \implies \mathcal{V}_{\beta+1, \text{real}}^i(\tau) \in \{\perp\} \cup \{(\widetilde{M}, \psi')\}_{\widetilde{M} \in \{0, 1\}^{p(n)}}$ for $\psi' \in \left\{ \left(\frac{3k}{4} - \beta \right), \left(\frac{3k}{4} - \beta + 1 \right) \right\}$.

Next we prove that if $\mathcal{V}_{\beta, \text{ideal}}^i(\tau) \neq \perp$, then one can efficiently compute $\mathcal{V}_{\beta+1, \text{real}}^i(\tau)$ from $\mathcal{V}_{\beta, \text{ideal}}^i(\tau)$ given only 2 additional bits of (inefficient) leakage. To this end, we define an (inefficient) leakage function $\pi_\beta^i(\cdot)$ with range $\{0, -1, 1, *\}$, and an efficient function f^i such that for every transcript τ ,

$$\mathcal{V}_{\beta+1, \text{real}}^i(\tau) = f^i(\mathcal{V}_{\beta, \text{ideal}}^i(\tau), \pi_\beta^i(\tau)) \quad (6)$$

whenever $\pi_\beta^i(\tau) \neq *$.

Intuitively, the leakage value 0 indicates that $\mathcal{V}_{\beta+1, \text{real}}^i(\tau)$ is identical to $\mathcal{V}_{\beta, \text{ideal}}^i(\tau)$. The leakage values $\{-1, 1\}$ indicate that $\mathcal{V}_{\beta, \text{ideal}}^i(\tau)$ contains some message and value ψ , and $\mathcal{V}_{\beta+1, \text{real}}^i(\tau)$ contains the same message and $\psi \pm 1$. The leakage value $*$ indicates failure. We formally define π_β^i and f^i in [Figure 3](#).

Note that the leakage π_β consists of $2y$ bits. Recall that by [Lemma 1](#), the distributions $(\tau_{\beta-1}, \mathcal{V}_{\beta, \text{real}}(\tau_{\beta-1}))$ and $(\tau_\beta, \mathcal{V}_{\beta, \text{ideal}}(\tau_\beta))$ are $\text{poly}(n \cdot 2^y)$ -indistinguishable. Therefore, we can rely on the Gentry-Wichs leakage lemma [\[GW11\]](#) (stated below, and adapted to our regime) to argue that there exists an augmentation function $\widehat{\pi}_{\beta-1}$ such that

$$\left(\tau_{\beta-1}, \mathcal{V}_{\beta, \text{real}}(\tau_{\beta-1}), \pi_{\beta-1}(\tau_{\beta-1}) \right) \approx \left(\tau_\beta, \mathcal{V}_{\beta, \text{ideal}}(\tau_\beta), \widehat{\pi}_{\beta-1}(\tau_\beta) \right).$$

Lemma 2. [\[GW11, JP14, CLP15\]](#) Let X and Y be any two distributions that cannot be distinguished with advantage more than ϵ by machines of size T . Then, for any leakage function π that outputs γ bits, there exists a function $\widehat{\pi}$ such that

$$\left(x, \pi(x) \right) \text{ and } \left(y, \widehat{\pi}(y) \right) \text{ for } x \leftarrow X \text{ and } y \leftarrow Y$$

²²For simplicity, we rely on deterministically verifiable NIWIs. Known NIWIs [\[BOV07, GOS12\]](#) satisfy this condition.




Figure 3: The definition of π_β and f .

cannot be distinguished with advantage more than 2ϵ , by machines of size $\left(\frac{\epsilon^4 T}{2^{\gamma \cdot \log T}}\right)^{\frac{1}{3}}$.²³

Furthermore, if $X \approx_{\text{poly}(T)} Y$ and $H_\infty(X), H_\infty(Y) \geq \omega(\log(T) + \log^2(n))$, then in time $2^{O(\gamma)} \cdot (n \cdot T)^{\log n}$ one can compute $\widehat{\pi}$ for which

$$(x, \pi(x)) \approx_{\text{poly}(T/2^\gamma)} (y, \widehat{\pi}(y)).$$

We next describe how we use [Lemma 1](#) and [Lemma 2](#) to complete the proof of [Theorem 2](#). Recall that in order to show that the scheme in [Figure 2](#) satisfies [Definition 3](#), we need to define $\mathcal{V}_{\text{real}}$ and $\mathcal{V}_{\text{ideal}}$ and prove that $(\tau_0, \mathcal{V}_{\text{real}}(\tau_0)) \approx (\tau_\ell, \mathcal{V}_{\text{ideal}}(\tau_\ell))$.

First attempt at defining $\mathcal{V}_{\text{Real}}$ and $\mathcal{V}_{\text{Ideal}}$.

For simplicity, we assume that $\pi_\beta^i(\tau_\beta) \neq *$ for every $i \in [y]$ and every $\beta \in [1, \ell]$.

Fix any $\beta \in [\ell]$. By [Lemma 1](#),

$$(\tau_\beta, \mathcal{V}_{\beta,\text{ideal}}(\tau_\beta)) \approx_{\text{poly}(n \cdot 2^\beta)} (\tau_{\beta-1}, \mathcal{V}_{\beta,\text{real}}(\tau_{\beta-1})). \quad (7)$$

By Equation (6), with our simplifying assumption,

$$\mathcal{V}_{\beta,\text{real}}(\tau_{\beta-1}) = f(\mathcal{V}_{\beta-1,\text{ideal}}(\tau_{\beta-1}), \pi_{\beta-1}(\tau_{\beta-1})). \quad (8)$$

²³Note that this statement is not asymptotic. In the next sentence we state an asymptotic version that suffices for our applications.

Combining Equation (7) and Equation (8),

$$\left(\tau_\beta, \mathcal{V}_{\beta, \text{ideal}}(\tau_\beta)\right) \approx_{\text{poly}(n \cdot 2y)} \left(\tau_{\beta-1}, f\left(\mathcal{V}_{\beta-1, \text{ideal}}(\tau_{\beta-1}), \pi_{\beta-1}(\tau_{\beta-1})\right)\right). \quad (9)$$

Setting $\beta = \ell$ in Equation (9),

$$\left(\tau_\ell, \mathcal{V}_{\ell, \text{ideal}}(\tau_\ell)\right) \approx_{\text{poly}(n \cdot 2y)} \left(\tau_{\ell-1}, f\left(\mathcal{V}_{\ell-1, \text{ideal}}(\tau_{\ell-1}), \pi_{\ell-1}(\tau_{\ell-1})\right)\right). \quad (10)$$

We will denote the left hand side of Equation (10) by $J_\ell(\tau_\ell)$, and the right hand side by $J_{\ell-1}(\tau_{\ell-1})$. Thus,

$$J_\ell(\tau_\ell) \approx_{\text{poly}(n \cdot 2y)} J_{\ell-1}(\tau_{\ell-1}).$$

Next, we will derive a sequence $J_{\ell-2}(\tau_{\ell-2}), J_{\ell-3}(\tau_{\ell-3}), \dots, J_1(\tau_1), J_0(\tau_0)$, and prove that for every $\alpha \in [\ell - 1, \dots, 0]$,

$$J_{\alpha+1}(\tau_{\alpha+1}) \approx J_\alpha(\tau_\alpha).$$

To this end, setting $\beta = \ell - 1$ in Equation (9),

$$\left(\tau_{\ell-1}, \mathcal{V}_{\ell-1, \text{ideal}}(\tau_{\ell-1})\right) \approx_{\text{poly}(n \cdot 2y)} \left(\tau_{\ell-2}, f\left(\mathcal{V}_{\ell-2, \text{ideal}}(\tau_{\ell-2}), \pi_{\ell-2}(\tau_{\ell-2})\right)\right). \quad (11)$$

Since the output of $\pi_{\ell-1}$ consists of $2y$ bits, set $\hat{\pi}_{\ell-1}$ according to the leakage lemma (Lemma 2) for Equation (11) so that:

$$\left(\tau_{\ell-1}, \mathcal{V}_{\ell-1, \text{ideal}}(\tau_{\ell-1}), \pi_{\ell-1}(\tau_{\ell-1})\right) \approx \left(\tau_{\ell-2}, f\left(\mathcal{V}_{\ell-2, \text{ideal}}(\tau_{\ell-2}), \pi_{\ell-2}(\tau_{\ell-2})\right), \hat{\pi}_{\ell-1}(\tau_{\ell-2})\right). \quad (12)$$

Remark 5. We note that according to Lemma 2, the leakage $\hat{\pi}_{\ell-1}$ is a function of $(\tau_{\ell-2}, \mathcal{V}_{\ell-2, \text{ideal}}(\tau_{\ell-2}))$ and not only of $\tau_{\ell-2}$. However, to avoid cluttering of notation, we denote $\hat{\pi}_{\ell-1}(\tau_{\ell-2}, \mathcal{V}_{\ell-2, \text{ideal}}(\tau_{\ell-2}))$ by $\hat{\pi}_{\ell-1}(\tau_{\ell-2})$. We use similar notation throughout the proof.

Because f is an efficient function, we can apply it to Equation (12) so that:

$$\left(\tau_{\ell-1}, f\left(\mathcal{V}_{\ell-1, \text{ideal}}(\tau_{\ell-1}), \pi_{\ell-1}(\tau_{\ell-1})\right)\right) \approx \left(\tau_{\ell-2}, f\left(f\left(\mathcal{V}_{\ell-2, \text{ideal}}(\tau_{\ell-2}), \pi_{\ell-2}(\tau_{\ell-2})\right), \hat{\pi}_{\ell-1}(\tau_{\ell-2})\right)\right). \quad (13)$$

Define $\eta_{\ell-2}(\cdot)$ and $h(\cdot)$ such that for all τ and all $i \in [y]$ ²⁴,

$$\text{If } \pi_{\ell-2}^i(\tau) = * \text{ or } \hat{\pi}_{\ell-1}^i(\tau) = *, \text{ set } \eta_{\ell-2}^i(\tau) = *.$$

$$\text{Else, set } \eta_{\ell-2}^i(\tau) = \pi_{\ell-2}^i(\tau) + \hat{\pi}_{\ell-1}^i(\tau).$$

$$h^i((\widetilde{M}, \psi), \eta) = (\widetilde{M}, \psi + \eta) \text{ if } \eta \neq *, \text{ and } h^i(\perp, \cdot) = h^i(\cdot, *) = \perp.$$

Note that

$$f\left(f\left(\mathcal{V}_{\ell-2, \text{ideal}}(\tau_{\ell-2}), \pi_{\ell-2}(\tau_{\ell-2})\right), \hat{\pi}_{\ell-1}(\tau_{\ell-2})\right) = h\left(\mathcal{V}_{\ell-2, \text{ideal}}(\tau_{\ell-2}), \eta_{\ell-2}(\tau_{\ell-2})\right)$$

Substituting into Equation (13), we have:

$$\left(\tau_{\ell-1}, f\left(\mathcal{V}_{\ell-1, \text{ideal}}(\tau_{\ell-1}), \pi_{\ell-1}(\tau_{\ell-1})\right)\right) \approx \left(\tau_{\ell-2}, h\left(\mathcal{V}_{\ell-2, \text{ideal}}(\tau_{\ell-2}), \eta_{\ell-2}(\tau_{\ell-2})\right)\right). \quad (14)$$

We denote the right hand side of Equation (14) by $J_{\ell-2}(\tau_{\ell-2})$.

²⁴We note that h is identical to f , except that it operates over a larger leakage domain.

Note that the left hand side of the Equation (14) is $J_{\ell-1}(\tau_{\ell-1})$, and thus

$$J_{\ell-1}(\tau_{\ell-1}) \approx J_{\ell-2}(\tau_{\ell-2}).$$

Next, setting $\beta = (\ell - 2)$ in Equation (9) we have:

$$\left(\tau_{\ell-2}, \mathcal{V}_{\ell-2, \text{ideal}}(\tau_{\ell-2})\right) \approx_{\text{poly}(n \cdot 2^y)} \left(\tau_{\ell-3}, f(\mathcal{V}_{\ell-3, \text{ideal}}(\tau_{\ell-3}), \pi_{\ell-3}(\tau_{\ell-3}))\right). \quad (15)$$

Since the output of $\eta_{\ell-2}$ consists of $O(y)$ bits, set $\hat{\eta}_{\ell-2}$ according to the leakage lemma (Lemma 2) for Equation (15) so that:

$$\left(\tau_{\ell-2}, \mathcal{V}_{\ell-2, \text{ideal}}(\tau_{\ell-2}), \eta_{\ell-2}(\tau_{\ell-2})\right) \approx \left(\tau_{\ell-3}, f(\mathcal{V}_{\ell-3, \text{ideal}}(\tau_{\ell-3}), \pi_{\ell-3}(\tau_{\ell-3})), \hat{\eta}_{\ell-2}(\tau_{\ell-3})\right). \quad (16)$$

Because h is an efficient function, applying it to both sides of the equation above:

$$\left(\tau_{\ell-2}, h(\mathcal{V}_{\ell-2, \text{ideal}}(\tau_{\ell-2}), \eta_{\ell-2}(\tau_{\ell-2}))\right) \approx \left(\tau_{\ell-3}, h(f(\mathcal{V}_{\ell-3, \text{ideal}}(\tau_{\ell-3}), \pi_{\ell-3}(\tau_{\ell-3})), \hat{\eta}_{\ell-2}(\tau_{\ell-3}))\right). \quad (17)$$

Define $\eta_{\ell-3}(\cdot)$ such that for all τ and all $i \in [y]$,

$$\text{If } \pi_{\ell-3}^i(\tau) = * \text{ or } \hat{\eta}_{\ell-2}^i(\tau) = *, \text{ set } \eta_{\ell-3}^i(\tau) = *,$$

$$\text{Else set } \eta_{\ell-3}^i(\tau) = \pi_{\ell-3}^i(\tau) + \hat{\eta}_{\ell-2}^i(\tau).$$

$$\text{Note that } h^i((\widetilde{M}, \psi), \eta) = (\widetilde{M}, \psi + \eta) \text{ if } \eta \neq *, \text{ and } h^i(\perp, \cdot) = h^i(\cdot, *) = \perp.$$

Therefore,

$$h(f(\mathcal{V}_{\ell-3, \text{ideal}}(\tau_{\ell-3}), \pi_{\ell-3}(\tau_{\ell-3})), \hat{\eta}_{\ell-2}(\tau_{\ell-3})) = h(\mathcal{V}_{\ell-3, \text{ideal}}(\tau_{\ell-3}), \eta_{\ell-3}(\tau_{\ell-3})).$$

Substituting into Equation (17), we have:

$$\left(\tau_{\ell-2}, h(\mathcal{V}_{\ell-2, \text{ideal}}(\tau_{\ell-2}), \eta_{\ell-2}(\tau_{\ell-2}))\right) \approx \left(\tau_{\ell-3}, h(\mathcal{V}_{\ell-3, \text{ideal}}(\tau_{\ell-3}), \eta_{\ell-3}(\tau_{\ell-3}))\right) \quad (18)$$

We denote the right hand side of Equation (18) by $J_{\ell-3}(\tau_{\ell-3})$.

Note that the left hand side of the Equation (18) is $J_{\ell-2}(\tau_{\ell-2})$, and thus

$$J_{\ell-2}(\tau_{\ell-2}) \approx J_{\ell-3}(\tau_{\ell-3}).$$

Similarly, for all $\alpha \in [1, \ell - 4]$, we define $\eta_\alpha(\cdot)$ that outputs $O(y \cdot \log \ell)$ bits, and define

$$J_\alpha(\tau_\alpha) \triangleq h(\mathcal{V}_{\alpha, \text{ideal}}(\tau_\alpha), \eta_\alpha(\tau_\alpha)).$$

The same argument as above implies that for every $\alpha \in [\ell - 1]$,

$$(\tau_{\alpha+1}, J_{\alpha+1}(\tau_{\alpha+1})) \approx (\tau_\alpha, J_\alpha(\tau_\alpha)).$$

This follows by setting $\beta = (\alpha + 1)$ in Equation (9), and then applying the sequence of Equations (16, 17, 18) with $(\alpha + 1)$ instead of $(\ell - 2)$.

Therefore, we conclude that

$$(\tau_\ell, J_\ell(\tau_\ell)) \approx (\tau_1, J_1(\tau_1)). \quad (19)$$

Recall that we need to prove that

$$(\tau_\ell, J_\ell(\tau_\ell)) \approx (\tau_0, J_0(\tau_0)),$$

for some J_0 , which we define next. To this end, by [Lemma 1](#) (for $\beta = 1$):

$$\left(\tau_1, \mathcal{V}_{1,\text{ideal}}(\tau_1)\right) \approx_{\text{poly}(n \cdot 2^y)} \left(\tau_0, \mathcal{V}_{1,\text{real}}(\tau_0)\right). \quad (20)$$

Set $\hat{\eta}_1$ according to the leakage lemma ([Lemma 2](#)) for Equation (20),

$$\left(\tau_1, \mathcal{V}_{1,\text{ideal}}(\tau_1), \eta_1(\tau_1)\right) \approx \left(\tau_0, \mathcal{V}_{1,\text{real}}(\tau_0), \hat{\eta}_1(\tau_0)\right)$$

Because h is an efficient function,

$$\left(\tau_1, h(\mathcal{V}_{1,\text{ideal}}(\tau_1), \eta_1(\tau_1))\right) \approx \left(\tau_0, h(\mathcal{V}_{1,\text{real}}(\tau_0), \hat{\eta}_1(\tau_0))\right). \quad (21)$$

Thus,

$$\left(\tau_1, J_1(\tau_1)\right) \approx \left(\tau_0, h(\mathcal{V}_{1,\text{real}}(\tau_0), \hat{\eta}_1(\tau_0))\right) \triangleq \left(\tau_0, J_0(\tau_0)\right), \quad (22)$$

which together with Equation (19), implies that

$$\left(\tau_\ell, J_\ell(\tau_\ell)\right) \approx \left(\tau_0, J_0(\tau_0)\right), \quad (23)$$

as desired.

Thus, under the assumption that $(\pi_\beta^i \neq *)$ for all $i \in [y]$ and all $\beta \in [1, \ell - 1]$, we can define F as a function that on input a tuple (χ_1, \dots, χ_y) outputs $(\tilde{m}_1, \dots, \tilde{m}_y)$ such that for all $i \in [y]$, if $\chi_i = (\tilde{M}, \psi)$ for some $\tilde{M} \in \{0, 1\}^p$ and some $\psi \geq 0$, then $\tilde{m}_i = \tilde{M}$, else $\tilde{m}_i = \perp$. We can then define,

$$\mathcal{V}_{\text{Real}}(\tau_0) = F(J_0(\tau_0))$$

and

$$\mathcal{V}_{\text{Ideal}}(\tau_\ell) = F(J_\ell(\tau_\ell))$$

and by Equation (23) conclude that

$$\left(\tau_0, \mathcal{V}_{\text{Real}}(\tau_0)\right) \approx \left(\tau_\ell, \mathcal{V}_{\text{Ideal}}(\tau_\ell)\right),$$

as desired.

However, the assumption that $\pi_\beta^i \neq *$ for all $i \in [y]$ and all $\beta \in [1, \ell - 1]$, is not necessarily true. In what follows, we remove this assumption. To that end, we need to define the distributions $\mathcal{V}_{\text{Real}}$ and $\mathcal{V}_{\text{Ideal}}$ differently, to ensure indistinguishability even when for some $\beta \in [1, \ell - 1]$ and some $i \in [y]$, $\pi_\beta^i = *$.

Defining $\mathcal{V}_{\text{Real}}$ and $\mathcal{V}_{\text{Ideal}}$ in the general case.

Note that if $\pi_\beta^i(\tau_\beta) \neq *$ for all $i \in [y]$, then Equation (8) holds. However, if $\pi_\beta^i(\tau_\beta) = *$ for some $i \in [y]$, then Equation (8) does not hold. In this case,

$$\mathcal{V}_{\beta+1,\text{real}}^i(\tau_\beta) = \left(\tilde{M}, \frac{3k}{4} - \beta + b\right) \text{ for some } \tilde{M} \text{ and some } b \in \{0, 1\}, \text{ whereas}$$

$$f^i(\mathcal{V}_{\beta,\text{ideal}}^i(\tau_\beta), \pi_\beta^i(\tau_\beta)) = \perp.$$

We next define an efficient function g and inefficient leakage functions ϕ_β for every $\beta \in [1, \ell - 1]$, such that g applied to $\mathcal{V}_{\beta+1}$ sometimes changes a value $(\widetilde{M}, \frac{3k}{4} - \beta + b)$ to \perp , ensuring that

$$g(\mathcal{V}_{\beta+1, \text{real}}(\tau_\beta), \phi_\beta(\tau_\beta)) = f(\mathcal{V}_{\beta, \text{ideal}}(\tau_\beta), \pi_\beta(\tau_\beta)). \quad (24)$$

To this end, we define the leakage function $\phi_\beta = (\phi_\beta^1, \dots, \phi_\beta^y)$ such that for every $i \in [y]$ and every transcript τ ,

- $(\phi_\beta^i(\tau) = *) \iff (\pi_\beta^i(\tau) = *)$.
- $\phi_\beta^i(\tau) = 0$ otherwise.

Define the function g that takes input x and $\Phi \in \{0, *\}$ such that for all x ,

- $g(x, 0) = x$,
- $g(x, *) = \perp$.

For any $\mathcal{X} = (\chi_1, \dots, \chi_y)$ and any leakage value $\Phi \in \{0, *\}^y$, define $\chi'_i \triangleq g(\chi_i, \Phi_i)$ and set $g(\mathcal{X}, \Phi) = (\chi'_1, \dots, \chi'_y)$ ²⁵. Note that for every $\beta \in [1, \ell - 1]$, the functions g and ϕ_β defined above, satisfy Equation (24), as desired.

Recall that by Lemma 1, for every $\beta \in [0, \ell - 1]$,

$$\left(\tau_{\beta+1}, \mathcal{V}_{\beta+1, \text{ideal}}(\tau_{\beta+1}), \right) \approx_{\text{poly}(n \cdot 2^y)} \left(\tau_\beta, \mathcal{V}_{\beta+1, \text{real}}(\tau_\beta) \right).$$

We set $\widehat{\phi}_\beta$ according to the leakage lemma (Lemma 2)²⁶, so that

$$\left(\tau_{\beta+1}, \mathcal{V}_{\beta+1, \text{ideal}}(\tau_{\beta+1}), \widehat{\phi}_\beta(\tau_{\beta+1}) \right) \approx_{\text{poly}(n \cdot 2^y)} \left(\tau_\beta, \mathcal{V}_{\beta+1, \text{real}}(\tau_\beta), \phi_\beta(\tau_\beta) \right), \quad (25)$$

and rely on the fact that g is efficiently computable, to conclude that for all $\beta \in [1, \ell - 1]$:

$$\begin{aligned} & \left(\tau_{\beta+1}, g(\mathcal{V}_{\beta+1, \text{ideal}}(\tau_{\beta+1}), \widehat{\phi}_\beta(\tau_{\beta+1})) \right) \approx_{\text{poly}(n \cdot 2^y)} \\ & \left(\tau_\beta, g(\mathcal{V}_{\beta+1, \text{real}}(\tau_\beta), \phi_\beta(\tau_\beta)) \right) = \\ & \left(\tau_\beta, f(\mathcal{V}_{\beta, \text{ideal}}(\tau_\beta), \pi_\beta(\tau_\beta)) \right). \end{aligned} \quad (26)$$

Setting $\beta = \ell - 1$ in Equation (26), we have:

$$\left(\tau_\ell, g(\mathcal{V}_{\ell, \text{ideal}}(\tau_\ell), \widehat{\phi}_{\ell-1}(\tau_\ell)) \right) \approx_{\text{poly}(n \cdot 2^y)} \left(\tau_{\ell-1}, f(\mathcal{V}_{\ell-1, \text{ideal}}(\tau_{\ell-1}), \pi_{\ell-1}(\tau_{\ell-1})) \right). \quad (27)$$

Setting $\beta = (\ell - 2)$ in Equation (26),

$$\left(\tau_{\ell-1}, g(\mathcal{V}_{\ell-1, \text{ideal}}(\tau_{\ell-1}), \widehat{\phi}_{\ell-2}(\tau_{\ell-1})) \right) \approx_{\text{poly}(n \cdot 2^y)} \left(\tau_{\ell-2}, f(\mathcal{V}_{\ell-2, \text{ideal}}(\tau_{\ell-2}), \pi_{\ell-2}(\tau_{\ell-2})) \right) \quad (28)$$

²⁵Here, we overload notation and allow g to take as input an individual pair χ or a tuple of pairs (χ_1, \dots, χ_y) .

²⁶Note that when applying Lemma 2, we have a $\text{poly}(n \cdot 2^y)$ loss in security. Therefore, this usage cannot be made repeatedly, more than a constant number of times. This is the first application, and we will use this lemma three times.

Since $\pi_{\ell-1}$ outputs $2y$ bits, we will set $\widehat{\pi}_{\ell-1}$ according to the leakage lemma (Lemma 2)²⁷ for Equation (28) such that

$$\left(\tau_{\ell-1}, g\left(\mathcal{V}_{\ell-1,\text{ideal}}(\tau_{\ell-1}), \widehat{\phi}_{\ell-2}(\tau_{\ell-1})\right), \pi_{\ell-1}(\tau_{\ell-1}) \right) \approx_{\text{poly}(n \cdot 2^y)} \left(\tau_{\ell-2}, f\left(\mathcal{V}_{\ell-2,\text{ideal}}(\tau_{\ell-2}), \pi_{\ell-2}(\tau_{\ell-2})\right), \widehat{\pi}_{\ell-1}(\tau_{\ell-2}) \right)$$

Because the function f is efficient, we conclude that

$$\begin{aligned} & \left(\tau_{\ell-1}, f\left(g\left(\mathcal{V}_{\ell-1,\text{ideal}}(\tau_{\ell-1}), \widehat{\phi}_{\ell-2}(\tau_{\ell-1})\right), \pi_{\ell-1}(\tau_{\ell-1})\right) \right) \approx_{\text{poly}(n \cdot 2^y)} \\ & \left(\tau_{\ell-2}, f\left(f\left(\mathcal{V}_{\ell-2,\text{ideal}}(\tau_{\ell-2}), \pi_{\ell-2}(\tau_{\ell-2})\right), \widehat{\pi}_{\ell-1}(\tau_{\ell-2})\right) \right) \end{aligned}$$

We define $\eta_{\ell-2}$ exactly as before, such that for all τ and all $i \in [y]$,

$$\eta_{\ell-2}^i(\tau) = * \iff \pi_{\ell-2}^i(\tau) = * \text{ or } \widehat{\pi}_{\ell-1}^i(\tau) = *$$

$$\text{Else, } \eta_{\ell-2}^i(\tau) = \pi_{\ell-2}^i(\tau) + \widehat{\pi}_{\ell-1}^i(\tau).$$

Recall that $h^i((\widetilde{M}, \psi), \eta) = (\widetilde{M}, \psi + \eta)$ if $\eta \neq *$, and $h^i(\perp, \cdot) = h^i(\cdot, *) = \perp$.

By definition,

$$f\left(f\left(\mathcal{V}_{\ell-2,\text{ideal}}(\tau_{\ell-2}), \pi_{\ell-2}(\tau_{\ell-2})\right), \widehat{\pi}_{\ell-1}(\tau_{\ell-2})\right) = h\left(\mathcal{V}_{\ell-2,\text{ideal}}(\tau_{\ell-2}), \eta_{\ell-2}(\tau_{\ell-2})\right).$$

Thus,

$$\left(\tau_{\ell-1}, f\left(g\left(\mathcal{V}_{\ell-1,\text{ideal}}(\tau_{\ell-1}), \widehat{\phi}_{\ell-2}(\tau_{\ell-1})\right), \pi_{\ell-1}(\tau_{\ell-1})\right) \right) \approx_{\text{poly}(n \cdot 2^y)} \left(\tau_{\ell-2}, h\left(\mathcal{V}_{\ell-2,\text{ideal}}(\tau_{\ell-2}), \eta_{\ell-2}(\tau_{\ell-2})\right) \right) \quad (29)$$

Similarly, setting $\beta = (\ell - 3)$ in Equation (26),

$$\left(\tau_{\ell-2}, g\left(\mathcal{V}_{\ell-2,\text{ideal}}(\tau_{\ell-2}), \widehat{\phi}_{\ell-3}(\tau_{\ell-2})\right) \right) \approx_{\text{poly}(n \cdot 2^y)} \left(\tau_{\ell-3}, f\left(\mathcal{V}_{\ell-3,\text{ideal}}(\tau_{\ell-3}), \pi_{\ell-3}(\tau_{\ell-3})\right) \right) \quad (30)$$

Since $\eta_{\ell-2}$ outputs $O(y)$ bits, we will set $\widehat{\eta}_{\ell-2}$ according to the leakage lemma (Lemma 2)²⁸ for Equation (30) such that

$$\left(\tau_{\ell-2}, g\left(\mathcal{V}_{\ell-2,\text{ideal}}(\tau_{\ell-2}), \widehat{\phi}_{\ell-3}(\tau_{\ell-2})\right), \eta_{\ell-2}(\tau_{\ell-2}) \right) \approx_{\text{poly}(n \cdot 2^y)} \left(\tau_{\ell-3}, f\left(\mathcal{V}_{\ell-3,\text{ideal}}(\tau_{\ell-3}), \pi_{\ell-3}(\tau_{\ell-3})\right), \widehat{\eta}_{\ell-2}(\tau_{\ell-3}) \right) \quad (31)$$

Because the function h is efficient, we conclude that

$$\begin{aligned} & \left(\tau_{\ell-2}, h\left(g\left(\mathcal{V}_{\ell-2,\text{ideal}}(\tau_{\ell-2}), \widehat{\phi}_{\ell-3}(\tau_{\ell-2})\right), \eta_{\ell-2}(\tau_{\ell-2})\right) \right) \approx_{\text{poly}(n \cdot 2^y)} \\ & \left(\tau_{\ell-3}, h\left(f\left(\mathcal{V}_{\ell-3,\text{ideal}}(\tau_{\ell-3}), \pi_{\ell-3}(\tau_{\ell-3})\right), \widehat{\eta}_{\ell-2}(\tau_{\ell-3})\right) \right) \end{aligned} \quad (32)$$

²⁷This is the second sequential application of the leakage lemma.

²⁸This is the second sequential application of the leakage lemma.

We define $\eta_{\ell-3}$ exactly as before such that for all τ and all $i \in [y]$,

$$\eta_{\ell-3} = * \iff \pi_{\ell-3}^i = * \text{ or } \widehat{\eta}_{\ell-2}^i(\tau) = *, \text{ and}$$

$$\eta_{\ell-3}^i(\tau) = \pi_{\ell-3}^i(\tau) + \widehat{\eta}_{\ell-2}^i(\tau).$$

Recall that $h^i((\widetilde{M}, \psi), \eta) = (\widetilde{M}, \psi + \eta)$ if $\eta \neq *$, and $h^i(\perp, \cdot) = h^i(\cdot, *) = \perp$.

By definition,

$$\left(\tau_{\ell-3}, h\left(f\left(\mathcal{V}_{\ell-3, \text{ideal}}(\tau_{\ell-3}), \pi_{\ell-3}(\tau_{\ell-3}), \widehat{\eta}_{\ell-2}(\tau_{\ell-3})\right)\right) \right) = \left(\tau_{\ell-3}, h\left(\mathcal{V}_{\ell-3, \text{ideal}}(\tau_{\ell-3}), \eta_{\ell-3}(\tau_{\ell-3})\right) \right) \quad (33)$$

Therefore,

$$\left(\tau_{\ell-2}, h\left(g\left(\mathcal{V}_{\ell-2, \text{ideal}}(\tau_{\ell-2}), \widehat{\phi}_{\ell-3}(\tau_{\ell-2}), \eta_{\ell-2}(\tau_{\ell-2})\right)\right) \right) \approx_{\text{poly}(n \cdot 2^y)} \left(\tau_{\ell-3}, h\left(\mathcal{V}_{\ell-3, \text{ideal}}(\tau_{\ell-3}), \eta_{\ell-3}(\tau_{\ell-3})\right) \right) \quad (34)$$

Similarly, for all $\alpha \in [\ell - 4]$, we can define η_α that outputs $O(y \log \ell)$ bits and prove that:

$$\left(\tau_{\alpha+1}, h\left(g\left(\mathcal{V}_{\alpha+1, \text{ideal}}(\tau_{\alpha+1}), \widehat{\phi}_\alpha(\tau_{\alpha+1}), \eta_{\alpha+1}(\tau_{\alpha+1})\right)\right) \right) \approx_{\text{poly}(n \cdot 2^y)} \left(\tau_\alpha, h\left(\mathcal{V}_{\alpha, \text{ideal}}(\tau_\alpha), \eta_\alpha(\tau_\alpha)\right) \right) \quad (35)$$

This follows by setting $\beta = \alpha$ in Equation (26) and then applying the sequence of Equations (31, 32, 33) with $(\alpha + 1)$ instead of $(\ell - 2)$.

Our next step will be to replace the left hand side of Equation (35) with

$$\left(\tau_{\alpha+1}, g\left(h\left(\mathcal{V}_{\alpha+1, \text{ideal}}(\tau_{\alpha+1}), \eta_{\alpha+1}(\tau_{\alpha+1}), \widehat{\phi}_\alpha(\tau_{\alpha+1})\right)\right) \right),$$

i.e., with h and g interchanged. To this end, we prove the following claim about the commutativity of the functions g and h .

Claim 3. For any leakage values $\Phi \in \{0, *\}^y$ and $\Pi \in ([-\ell, \ell] \cup *)^y$, and any \mathcal{X} ,

$$g(h(\mathcal{X}, \Pi), \Phi) = h(g(\mathcal{X}, \Phi), \Pi). \quad (36)$$

Proof. We prove this via the following exhaustive case analysis, for each $i \in [y]$:

- If $\Phi^i = 0$ then g is the identity function, and hence

$$g(h^i(\chi_i, \Pi^i), \Phi^i) = h^i(g(\chi_i, \Phi^i), \Pi^i).$$

- If $\Phi^i = *$ then

$$g(h^i(\chi_i, \Pi^i), \Phi^i) = h^i(g(\chi_i, \Phi^i), \Pi^i) = \perp,$$

where the fact that $g(h^i(\chi_i, \Pi^i), \Phi^i) = \perp$ follows from the fact that $g(\cdot, *) = \perp$, and the fact that $h^i(g(\chi_i, \Phi^i), \Pi^i) = \perp$ follows from the fact that $g(\cdot, *) = \perp$ and the fact that $h^i(\perp, \cdot) = \perp$ (where the latter follows from the definition of h^i together with the fact that $f(\perp, \cdot) = \perp$).

This completes the proof of the claim. \square

Therefore, defining $\eta_\ell \equiv 0^y$, $\eta_{\ell-1} \triangleq \pi_{\ell-1}$, and $\eta_{\ell-2}, \dots, \eta_1$ as defined above, it holds that for every $\alpha \in [\ell - 1]$:

$$\left(\tau_{\alpha+1}, g\left(h(\mathcal{V}_{\alpha+1, \text{ideal}}(\tau_{\alpha+1}), \eta_{\alpha+1}(\tau_{\alpha+1})), \widehat{\phi}_\alpha(\tau_{\alpha+1})\right) \right) \approx_{\text{poly}(n \cdot 2^y)} \left(\tau_\alpha, h(\mathcal{V}_{\alpha, \text{ideal}}(\tau_\alpha), \eta_\alpha(\tau_\alpha)) \right) \quad (37)$$

We define $\widehat{\Phi}_1 \triangleq \widehat{\phi}_1$. We would like the right hand side of the equation above for α to be the same as the left hand side applied with $\alpha - 1$, since then could denote the left hand side by $\left(\tau_{\alpha+1}, J_{\alpha+1}(\tau_{\alpha+1}) \right)$ and the right hand side by $\left(\tau_\alpha, J_\alpha(\tau_\alpha) \right)$, we use a standard hybrid argument to argue that

$$\left(\tau_\ell, J_\ell(\tau_\ell) \right) \approx \left(\tau_1, J_1(\tau_1) \right).$$

In what follows, we modify Equation (37) to have this desired structure.

To this end, substitute $\alpha = 1$ in Equation (37), to obtain

$$\left(\tau_2, g\left(h(\mathcal{V}_{2, \text{ideal}}(\tau_2), \eta_2(\tau_2)), \widehat{\phi}_1(\tau_2)\right) \right) \approx_{\text{poly}(n \cdot 2^y)} \left(\tau_1, h(\mathcal{V}_{1, \text{ideal}}(\tau_1), \eta_1(\tau_1)) \right)$$

Similarly, substitute $\alpha = 2$ in Equation (37), to obtain

$$\left(\tau_3, g\left(h(\mathcal{V}_{3, \text{ideal}}(\tau_3), \eta_3(\tau_3)), \widehat{\phi}_2(\tau_3)\right) \right) \approx_{\text{poly}(n \cdot 2^y)} \left(\tau_2, h(\mathcal{V}_{2, \text{ideal}}(\tau_2), \eta_2(\tau_2)) \right)$$

Set $\widehat{\Phi}_1$ according to the leakage lemma (Lemma 2)²⁹ such that:

$$\left(\tau_3, g\left(h(\mathcal{V}_{3, \text{ideal}}(\tau_3), \eta_3(\tau_3)), \widehat{\phi}_2(\tau_3)\right), \widehat{\Phi}_1(\tau_3) \right) \approx_{\text{poly}(n \cdot 2^y)} \left(\tau_2, h(\mathcal{V}_{2, \text{ideal}}(\tau_2), \eta_2(\tau_2)), \widehat{\Phi}_1(\tau_2) \right) \quad (38)$$

Because g is an efficient function,

$$\left(\tau_3, g\left(g\left(h(\mathcal{V}_{3, \text{ideal}}(\tau_3), \eta_3(\tau_3)), \widehat{\phi}_2(\tau_3)\right), \widehat{\Phi}_1(\tau_3)\right) \right) \approx_{\text{poly}(n \cdot 2^y)} \left(\tau_2, g\left(h(\mathcal{V}_{2, \text{ideal}}(\tau_2), \eta_2(\tau_2)), \widehat{\Phi}_1(\tau_2)\right) \right) \quad (39)$$

At this point, we will define $\widehat{\Phi}_2(\cdot)$ so that for every τ and every $i \in [y]$,

$$\left(\widehat{\Phi}_2^i(\tau) = * \right) \text{ if } \left(\widehat{\phi}_2^i(\tau) = * \text{ or } \widehat{\Phi}_1^i(\tau) = * \right), \text{ and } 0 \text{ otherwise.}$$

Note that

$$g\left(g\left(h(\mathcal{V}_{3, \text{ideal}}(\tau_3), \eta_3(\tau_3)), \widehat{\phi}_2(\tau_3)\right), \widehat{\Phi}_1(\tau_3)\right) = g\left(h(\mathcal{V}_{3, \text{ideal}}(\tau_3), \eta_3(\tau_3)), \widehat{\Phi}_2(\tau_3)\right)$$

Substituting this in Equation (39),

$$\left(\tau_3, g\left(h(\mathcal{V}_{3, \text{ideal}}(\tau_3), \eta_3(\tau_3)), \widehat{\Phi}_2(\tau_3)\right) \right) \approx_{\text{poly}(n \cdot 2^y)} \left(\tau_2, g\left(h(\mathcal{V}_{2, \text{ideal}}(\tau_2), \eta_2(\tau_2)), \widehat{\Phi}_1(\tau_2)\right) \right)$$

²⁹This is the third sequential application of the leakage lemma.

Similarly, starting with Equation (37) and applying Lemma 2, we can define $\widehat{\Phi}_3$ such that:

$$\left(\tau_4, g\left(h(\mathcal{V}_{4,\text{ideal}}(\tau_4), \eta_4(\tau_4)), \widehat{\Phi}_3(\tau_4)\right) \right) \approx_{\text{poly}(n \cdot 2^y)} \left(\tau_3, g\left(h(\mathcal{V}_{3,\text{ideal}}(\tau_3), \eta_3(\tau_3)), \widehat{\Phi}_2(\tau_3)\right) \right) \quad (40)$$

This follows by substituting $\alpha = 3$ in Equation (37), and then applying Equations (38) and (39), with 3 instead of 2, and relying on $\widehat{\Phi}_2$ instead of $\widehat{\Phi}_1$.

Similarly, we define $\widehat{\Phi}_4, \dots, \widehat{\Phi}_{\ell-1}$. Setting

$$J_1(\tau_1) = h(\mathcal{V}_{1,\text{ideal}}(\tau_1), \eta_1(\tau_1)),$$

and

$$J_\alpha(\tau_\alpha) = g\left(h(\mathcal{V}_{\alpha,\text{ideal}}(\tau_\alpha), \eta_\alpha(\tau_\alpha)), \widehat{\Phi}_{\alpha-1}(\tau_\alpha)\right) \text{ for } \alpha \in [2, \ell],$$

we conclude that for all $\alpha \in [2, \ell]$:

$$(\tau_\alpha, J_\alpha(\tau_\alpha)) \approx (\tau_{\alpha-1}, J_{\alpha-1}(\tau_{\alpha-1})).$$

This follows by starting with Equation (37), and then applying Equations (38) and (39), with α instead of 3.

This implies

$$(\tau_1, J_1(\tau_1)) \approx (\tau_\ell, J_\ell(\tau_\ell))$$

Recall that by Lemma 1 we have,

$$(\tau_1, \mathcal{V}_{1,\text{ideal}}(\tau_1)) \approx_{\text{poly}(n \cdot 2^y)} (\tau_0, \mathcal{V}_{1,\text{real}}(\tau_0)) \quad (41)$$

We set $\widehat{\eta}_1$ according to the leakage lemma (Lemma 2) such that:

$$(\tau_1, \mathcal{V}_{1,\text{ideal}}(\tau_1), \eta_1(\tau_1)) \approx_{\text{poly}(n \cdot 2^y)} (\tau_0, \mathcal{V}_{1,\text{real}}(\tau_0), \widehat{\eta}_1(\tau_0))$$

This, together with the fact that h is efficiently computable, implies that

$$\left(\tau_1, h\left(\mathcal{V}_{1,\text{ideal}}(\tau_1), \eta_1(\tau_1)\right) \right) \approx_{\text{poly}(n \cdot 2^y)} \left(\tau_0, h\left(\mathcal{V}_{1,\text{real}}(\tau_0), \widehat{\eta}_1(\tau_0)\right) \right)$$

Define $J_0(\tau_0) = h(\mathcal{V}_{1,\text{real}}(\tau_0), \widehat{\eta}_1(\tau_0))$, then we have that for all $\alpha \in [1, \ell]$,

$$(\tau_\alpha, J_\alpha(\tau_\alpha)) \approx (\tau_{\alpha-1}, J_{\alpha-1}(\tau_{\alpha-1})) \quad (42)$$

We define F as a function that on input a tuple (χ_1, \dots, χ_y) outputs $(\widetilde{m}_1, \dots, \widetilde{m}_y)$ such that for all $i \in [y]$, if $\chi_i = (\widetilde{M}, \psi)$ for some $\widetilde{M} \in \{0, 1\}^p$ and some $\psi \geq 0$, then $\widetilde{m}_i = \widetilde{M}$, else $\widetilde{m}_i = \perp$. We then define $\mathcal{V}_{\text{Real}}$ as

$$\mathcal{V}_{\text{Real}}(\tau_0) = F\left(J_0(\tau_0)\right) = F\left(h\left(\mathcal{V}_{1,\text{real}}(\tau_0), \widehat{\eta}_1(\tau_0)\right)\right)$$

and $\mathcal{V}_{\text{Ideal}}$ as

$$\mathcal{V}_{\text{Ideal}}(\tau_\ell) = F\left(J_\ell(\tau_\ell)\right) = F\left(g\left(\mathcal{V}_{\ell,\text{ideal}}, \widehat{\Phi}_{\ell-1}(\tau_\ell)\right)\right)$$

which by Equation (42), implies that

$$(\tau_0, \mathcal{V}_{\text{Real}}(\tau_0)) \approx (\tau_\ell, \mathcal{V}_{\text{Ideal}}(\tau_\ell))$$

Efficiency of computing $\mathcal{V}_{\text{Real}}$ and $\mathcal{V}_{\text{Ideal}}$. Note that $\mathcal{V}_{\text{Real}}(\tau_0) = F\left(h\left(\mathcal{V}_{1,\text{real}}(\tau_0), \widehat{\eta}_1(\tau_0)\right)\right)$, where F and h are efficiently computable, and $\mathcal{V}_{1,\text{real}}$ is computable in time T_V . Moreover, $\widehat{\eta}_1$ outputs a string of length at most $O(y \log \ell)$, and is defined as the result of applying the leakage lemma (Lemma 2) to Equation (41). Therefore, setting $\gamma = O(y \log \ell)$ and $T = n \cdot 2^y$ in Lemma 2, the function $\widehat{\eta}_1$ is computable in time $(n \cdot 2^y)^{O(\log n)}$.

Similarly, $\mathcal{V}_{\text{Ideal}}(\tau_\ell) = F\left(g\left(\mathcal{V}_{\ell,\text{ideal}}, \widehat{\Phi}_{\ell-1}(\tau_\ell)\right)\right)$, where F and g are efficiently computable, and $\mathcal{V}_{\ell,\text{ideal}}$ is computable in time T_V . We next argue that $\widehat{\Phi}_{\ell-1}$ is computable in time $(n \cdot 2^y)^{O(\log n)}$. To this end, recall that $\widehat{\Phi}_{\ell-1} = \widehat{\phi}_{\ell-1} + \widehat{\widehat{\Phi}}_{\ell-2}$, and thus it suffices to bound the running time of $\widehat{\phi}_{\ell-1}$ and $\widehat{\widehat{\Phi}}_{\ell-2}$. Recall that $\widehat{\phi}_{\ell-1}$ outputs a string of length $O(\log y)$, and is defined as the result of applying the leakage lemma (Lemma 2) to Equation (37). Therefore, setting $\gamma = O(y)$ and $T = n \cdot 2^y$ the function $\widehat{\phi}_{\ell-1}$ is computable in time $(n \cdot 2^y)^{O(\log n)}$. Similarly, $\widehat{\widehat{\Phi}}_{\ell-2}$ outputs a string of length $O(\log y)$, and is again defined as the result of applying the leakage lemma (Lemma 2). Therefore, setting $\gamma = O(y)$ and $T = n \cdot 2^y$ in Lemma 2, the function $\widehat{\widehat{\Phi}}_{\ell-2}$ is computable in time $(n \cdot 2^y)^{O(\log n)}$. Therefore, $\widehat{\Phi}_{\ell-1}$ is computable in time $(n \cdot 2^y)^{O(\log n)}$.

We thus conclude that $\mathcal{V}_{\text{Real}}$ and $\mathcal{V}_{\text{Ideal}}$ are computable in time $T'_V = T_V + (n \cdot 2^y)^{O(\log n)}$, as desired.

Validity of $\mathcal{V}_{\text{Real}}$ and $\mathcal{V}_{\text{Ideal}}$. We now prove that $\mathcal{V}_{\text{Real}}$ and $\mathcal{V}_{\text{Ideal}}$ satisfy the validity condition in Definition 3.

Proving that $\mathcal{V}_{\text{Real}}$ satisfies the validity condition.

Recall that $\mathcal{V}_{\text{Real}}(\tau_0) = F\left(h\left(\mathcal{V}_{1,\text{Real}}(\tau_0), \widehat{\eta}_1(\tau_0)\right)\right)$. First, we note that for every $\beta \in [\ell]$, $\mathcal{V}_{\beta,\text{Real}}$ and $\mathcal{V}_{\beta,\text{Ideal}}$ invoke v_{Real} or v_{Ideal} respectively on the underlying commitments corresponding to adversary's small tags $\widetilde{\lambda}_\beta$ in all the right sessions. Therefore by construction, for every $\beta \in [\ell]$ and every $i \in [y]$, and for every transcript τ in which the i^{th} commitment \widetilde{c}_i is valid (i.e., is of the form $\widetilde{c}_i = \text{com}(\widetilde{M}; r)$ for some \widetilde{M} and r), it holds that $\mathcal{V}_{\beta,\text{Real}}^i(\tau)$ outputs (\widetilde{M}, k) . Hence, to prove that $\mathcal{V}_{\text{Real}}$ satisfies the validity condition, it suffices to prove that for every $i \in [y]$,

$$\Pr\left[\left(\mathcal{V}_{1,\text{Real}}^i(\tau_0) = (\widetilde{M}, k)\right) \wedge \left(\widehat{\eta}_1^i(\tau_0) = *\right)\right] = \text{negl}(n)$$

We prove the following stronger equation, for every $i \in [y]$:

$$\Pr\left[\left(\mathcal{V}_{1,\text{Real}}^i(\tau_0) \neq \perp\right) \wedge \left(\widehat{\eta}_1^i(\tau_0) = *\right)\right] = \text{negl}(n)$$

Recall that for any $\beta \in [\ell - 1]$, by definition of π_β , for all $i \in [y]$,

$$\pi_\beta^i(\tau_\beta) = * \implies \mathcal{V}_{\beta,\text{Ideal}}^i(\tau_\beta) = \perp. \quad (43)$$

Setting $\beta = [\ell - 1]$ in Equation (43), for all $i \in [y]$, $\pi_{\ell-1}^i(\tau_{\ell-1}) = * \implies \mathcal{V}_{\ell-1,\text{Ideal}}^i(\tau_{\ell-1}) = \perp$.

Recall that $\widetilde{\eta}_{\ell-1} \triangleq \pi_{\ell-1}$. Since $\widehat{\eta}_{\ell-1}$ is computed based on $\eta_{\ell-1}$ by applying the leakage lemma (refer to Equation (12)), we have that for all $i \in [y]$,

$$\Pr\left[\left(f^i\left(\mathcal{V}_{\ell-2,\text{Ideal}}^i(\tau_{\ell-2}), \pi_{\ell-2}^i(\tau_{\ell-2})\right) \neq \perp\right) \wedge \left(\widehat{\eta}_{\ell-1}^i(\tau_{\ell-2}) = *\right)\right] = \text{negl}(n).$$

Next, we prove that for every β , there exists a negligible function μ_β , such that for every $\beta \in [\ell - 3]$, if

$$\Pr \left[(f^i(\mathcal{V}_{\beta+1, \text{Ideal}}^i(\tau_{\beta+1}), \pi_\beta^i(\tau_{\beta+1})) \neq \perp) \wedge (\widehat{\eta}_{\beta+2}^i(\tau_{\beta+1}) = *) \right] = \epsilon$$

then

$$\Pr \left[(f^i(\mathcal{V}_{\beta, \text{Ideal}}^i(\tau_\beta), \pi_\beta^i(\tau_\beta)) \neq \perp) \wedge (\widehat{\eta}_{\beta+1}^i(\tau_\beta) = *) \right] \leq \epsilon + \mu_\beta(n).$$

Fix any $\beta \in [\ell - 3]$. By definition, $\eta_{\beta+1}^i(\tau) = * \iff (\widehat{\eta}_{\beta+2}^i(\tau) = *) \text{ or } (\pi_{\beta+1}^i(\tau) = *)$. Therefore,

$$\begin{aligned} & \Pr \left[(\mathcal{V}_{\beta+1, \text{Ideal}}^i(\tau_{\beta+1}) \neq \perp) \wedge (\eta_{\beta+1}^i(\tau_{\beta+1}) = *) \right] \\ &= \Pr \left[(\mathcal{V}_{\beta+1, \text{Ideal}}^i(\tau_{\beta+1}) \neq \perp) \wedge (\widehat{\eta}_{\beta+2}^i(\tau_{\beta+1}) = *) \wedge (\pi_{\beta+1}^i(\tau_{\beta+1}) \neq *) \right] \\ &+ \Pr \left[(\mathcal{V}_{\beta+1, \text{Ideal}}^i(\tau_{\beta+1}) \neq \perp) \wedge (\pi_{\beta+1}^i(\tau_{\beta+1}) = *) \right] \\ &= \Pr \left[(\mathcal{V}_{\beta+1, \text{Ideal}}^i(\tau_{\beta+1}) \neq \perp) \wedge (\widehat{\eta}_{\beta+2}^i(\tau_{\beta+1}) = *) \wedge (\pi_{\beta+1}^i(\tau_{\beta+1}) \neq *) \right] + 0 \\ &= \Pr \left[f^i(\mathcal{V}_{\beta+1, \text{Ideal}}^i(\tau_{\beta+1}), \pi_{\beta+1}^i(\tau_{\beta+1})) \neq \perp \wedge (\widehat{\eta}_{\beta+2}^i(\tau_{\beta+1}) = *) \wedge (\pi_{\beta+1}^i(\tau_{\beta+1}) \neq *) \right] = \epsilon \end{aligned}$$

Since $\widehat{\eta}_{\beta+1}$ is computed based on $\eta_{\beta+1}$ by applying the leakage lemma to Equation (9) (w.r.t. $\beta + 1$), we have that for all $i \in [y]$,

$$\Pr \left[(f^i(\mathcal{V}_{\beta, \text{Ideal}}^i(\tau_\beta), \pi_\beta^i(\tau_\beta)) \neq \perp) \wedge (\widehat{\eta}_{\beta+1}^i(\tau_\beta) = *) \right] \leq \epsilon + \mu_\beta(n),$$

for some negligible function μ_β , as required. This implies that for every $\beta \in [1, \ell - 3]$,

$$\Pr \left[(f^i(\mathcal{V}_{\beta, \text{Ideal}}^i(\tau_\beta), \pi_\beta^i(\tau_\beta)) \neq \perp) \wedge (\widehat{\eta}_{\beta+1}^i(\tau_\beta) = *) \right] = \text{negl}(n)$$

Specifically, setting $\beta = 1$, we have that for every $i \in [y]$,

$$\Pr \left[f^i(\mathcal{V}_{1, \text{Ideal}}^i(\tau_1), \pi_1^i(\tau_1)) \neq \perp \wedge (\widehat{\eta}_2^i(\tau_1) = *) \right] = \text{negl}(n)$$

Since by definition, $\eta_1^i(\tau) = * \iff (\widehat{\eta}_2^i(\tau) = *) \text{ or } (\pi_1^i(\tau) = *)$, this implies:

$$\begin{aligned} & \Pr \left[(\mathcal{V}_{1, \text{Ideal}}^i(\tau_1) \neq \perp) \wedge (\eta_1^i(\tau_1) = *) \right] \\ &= \Pr \left[(\mathcal{V}_{1, \text{Ideal}}^i(\tau_1) \neq \perp) \wedge (\widehat{\eta}_2^i(\tau_1) = *) \wedge (\pi_1^i(\tau_1) \neq *) \right] + \Pr \left[(\mathcal{V}_{1, \text{Ideal}}^i(\tau_1) \neq \perp) \wedge (\pi_1^i(\tau_1) = *) \right] \\ &= \Pr \left[(\mathcal{V}_{1, \text{Ideal}}^i(\tau_1) \neq \perp) \wedge (\widehat{\eta}_2^i(\tau_1) = *) \wedge (\pi_1^i(\tau_1) \neq *) \right] + 0 \\ &= \Pr \left[(f(\mathcal{V}_{1, \text{Ideal}}^i(\tau_1), \pi_1^i(\tau_1)) \neq \perp) \wedge (\widehat{\eta}_2^i(\tau_1) = *) \wedge (\pi_1^i(\tau_1) \neq *) \right] + 0 \\ &= \text{negl}(n) \end{aligned}$$

Since $\widehat{\eta}_1$ is derived from η_1 by applying the leakage lemma to Equation (41), we have that for every $i \in [y]$,

$$\Pr \left[(\mathcal{V}_{1, \text{Real}}^i(\tau_0) \neq \perp) \wedge (\widehat{\eta}_1^i(\tau_0) = *) \right] = \text{negl}(n),$$

as desired.

Proving that $\mathcal{V}_{\text{ideal}}$ satisfies the validity condition.

Recall that $\mathcal{V}_{\text{ideal}}(\tau_\ell) = F\left(g\left(\mathcal{V}_{\ell,\text{ideal}}(\tau_\ell), \widehat{\Phi}_{\ell-1}(\tau_\ell)\right)\right)$. First, we note that for every $\beta \in [\ell]$, $\mathcal{V}_{\beta,\text{Real}}$ and $\mathcal{V}_{\beta,\text{Ideal}}$ invoke v_{Real} or v_{Ideal} respectively on the underlying commitments corresponding to adversary's small tags $\tilde{\lambda}_\beta$ in all the right sessions. Therefore by construction, for every $\beta \in [\ell]$ and every $i \in [y]$, and for every transcript τ in which the i^{th} commitment \tilde{c}_i is valid (i.e., is of the form $\tilde{c}_i = \text{com}(\widetilde{M}; r)$ for some \widetilde{M} and r), it holds that $\mathcal{V}_{\beta,\text{Ideal}}^i(\tau)$ outputs (\widetilde{M}, k) . Hence, to prove that $\mathcal{V}_{\text{ideal}}$ satisfies the validity condition, it suffices to prove that for every $i \in [y]$,

$$\Pr\left[\left(\mathcal{V}_{\ell,\text{Ideal}}^i(\tau_\ell) \in \{(\widetilde{M}, k)\}_{\widetilde{M} \in \{0,1\}^p}\right) \wedge \left(\widehat{\Phi}_{\ell-1}^i(\tau_\ell) = *\right)\right] = \text{negl}(n).$$

First, we recall that for every $\beta \in [\ell - 1]$ and every $i \in [y]$, by definition

$$(\phi_\beta^i(\tau_\beta) = *) \iff (\pi_\beta^i(\tau_\beta) = *)$$

Moreover, if $(\pi_\beta^i(\tau_\beta) = *)$ then

$$\mathcal{V}_{\beta+1,\text{real}}^i(\tau_\beta) = \left(\widetilde{M}, \frac{3k}{4} - \beta + b\right) \text{ for some } \widetilde{M} \in \{0,1\}^p, b \in \{0,1\}.$$

Therefore, if $(\phi_\beta^i(\tau_\beta) = *)$ then

$$\mathcal{V}_{\beta+1,\text{real}}^i(\tau_\beta) = \left(\widetilde{M}, \frac{3k}{4} - \beta + b\right) \text{ for some } \widetilde{M} \in \{0,1\}^p, b \in \{0,1\}.$$

Recall that for $\beta \in [\ell - 1]$, by Equation (25),

$$(\tau_{\beta+1}, \mathcal{V}_{\beta+1,\text{ideal}}(\tau_{\beta+1}), \widehat{\Phi}_\beta(\tau_{\beta+1})) \approx (\tau_\beta, \mathcal{V}_{\beta+1,\text{real}}(\tau_\beta), \phi_\beta(\tau_\beta)).$$

Therefore for every $\beta \in [\ell - 1]$ and every $i \in [y]$,

$$\Pr\left[\left(\mathcal{V}_{\beta+1,\text{ideal}}^i(\tau_{\beta+1}) \notin \{(\widetilde{M}, \psi)\}_{\widetilde{M} \in \{0,1\}^p, \psi \in \{\frac{3k}{4} - \beta, \frac{3k}{4} - \beta + 1\}}\right) \wedge \left(\widehat{\Phi}_\beta^i(\tau_{\beta+1}) = *\right)\right] = \text{negl}(n). \quad (44)$$

Recall that for every $i \in [y]$, h^i on input $((\widetilde{M}, \psi), \alpha)$, for every α in the range of $\eta_{\beta+1}^i(\cdot)$ outputs either \perp (if $\alpha = *$) or (\widetilde{M}, ψ') , where $\psi' \in [\psi - \ell + \beta + 1, \psi + \ell - \beta - 1]$ ³⁰. Therefore, if

$$\mathcal{V}_{\beta+1,\text{ideal}}^i(\tau_{\beta+1}) \in \left\{ \left(\widetilde{M}, \frac{3k}{4} - \beta\right), \left(\widetilde{M}, \frac{3k}{4} - \beta + 1\right) \right\}$$

then

$$h^i(\mathcal{V}_{\beta+1,\text{ideal}}^i(\tau_{\beta+1}), \eta_{\beta+1}^i(\tau_{\beta+1})) = * \text{ or,}$$

$$h^i(\mathcal{V}_{\beta+1,\text{ideal}}^i(\tau_{\beta+1}), \eta_{\beta+1}^i(\tau_{\beta+1})) = (\widetilde{M}, \psi') \text{ where } \psi' \in \left[\frac{3k}{4} - \ell + 1, \frac{3k}{4} + \ell \right]$$

which implies that for every $i \in [y]$, $\beta \in [\ell - 1]$, there exists a negligible function μ such that

$$\Pr\left[\left(h^i(\mathcal{V}_{\beta+1,\text{ideal}}^i(\tau_{\beta+1}), \eta_{\beta+1}^i(\tau_{\beta+1})) \notin \{*\} \cup \{(\widetilde{M}, \psi)\}_{\widetilde{M} \in \{0,1\}^p, \psi \in \{\frac{3k}{4} - \ell + 1, \frac{3k}{4} + \ell\}}\right) \wedge \left(\widehat{\Phi}_\beta^i(\tau_{\beta+1}) = *\right)\right] = \mu(n). \quad (45)$$

³⁰This is because $\eta_{\beta+1}^i$ outputs values in the range $[-\ell + (\beta + 1), \ell - (\beta + 1)]$.

Recall that by Equation (37), for all $\beta \in [\ell - 1]$:

$$\left(\tau_{\beta+1}, g\left(h(\mathcal{V}_{\beta+1,\text{ideal}}(\tau_{\beta+1}), \eta_{\beta+1}(\tau_{\beta+1})), \widehat{\phi}_{\beta}(\tau_{\beta+1})\right) \right) \approx_{\text{poly}(n \cdot 2^y)} \left(\tau_{\beta}, h(\mathcal{V}_{\beta,\text{ideal}}(\tau_{\beta}), \eta_{\beta}(\tau_{\beta})) \right)$$

Therefore taking $\beta = 2$ and applying the leakage lemma, we get that $\widehat{\Phi}_1 \triangleq \widehat{\phi}_1$ is such that:

$$\left(\left(\tau_3, g\left(h(\mathcal{V}_{3,\text{ideal}}(\tau_3), \eta_3(\tau_3)), \widehat{\phi}_2(\tau_3)\right) \right), \widehat{\phi}_1(\tau_3) \right) \approx_{\text{poly}(n \cdot 2^y)} \left(\left(\tau_2, h(\mathcal{V}_{2,\text{ideal}}(\tau_2), \eta_2(\tau_2)) \right), \widehat{\phi}_1(\tau_2) \right)$$

Combining this with Equation (45), we have that for every $i \in [y]$,

$$\Pr \left[\left(g\left(h^i(\mathcal{V}_{3,\text{ideal}}^i(\tau_3), \eta_3^i(\tau_3)), \widehat{\phi}_2^i(\tau_3)\right) \notin \{*\} \cup \{(\widetilde{M}, \psi)\}_{\widetilde{M} \in \{0,1\}^p, \psi \in [\frac{3k}{4} - \ell + 1, \frac{3k}{4} + \ell]} \right) \wedge \left(\widehat{\Phi}_1^i(\tau_3) = * \right) \right] = \text{negl}(n)$$

Next, we prove that for every $\beta \in [3, \ell - 1]$ there exists a negligible function μ_{β} such that if:

$$\Pr \left[\left(g\left(h^i(\mathcal{V}_{\beta,\text{ideal}}^i(\tau_{\beta}), \eta_{\beta}^i(\tau_{\beta})), \widehat{\phi}_{\beta-1}^i(\tau_{\beta})\right) \notin \{*\} \cup \{(\widetilde{M}, \psi)\}_{\widetilde{M} \in \{0,1\}^p, \psi \in [\frac{3k}{4} - \ell + 1, \frac{3k}{4} + \ell]} \right) \wedge \left(\widehat{\Phi}_{\beta-2}^i(\tau_{\beta}) = * \right) \right] = \varepsilon$$

then

$$\Pr \left[\left(g\left(h^i(\mathcal{V}_{\beta+1,\text{ideal}}^i(\tau_{\beta+1}), \eta_{\beta+1}^i(\tau_{\beta+1})), \widehat{\phi}_{\beta}^i(\tau_{\beta+1})\right) \notin \{*\} \cup \{(\widetilde{M}, \psi)\}_{\widetilde{M} \in \{0,1\}^p, \psi \in [\frac{3k}{4} - \ell + 1, \frac{3k}{4} + \ell]} \right) \wedge \left(\widehat{\Phi}_{\beta-1}^i(\tau_{\beta+1}) = * \right) \right] \leq \varepsilon + \mu_{\beta}(n)$$

To this end, fix any $\beta \in [3, \ell - 1]$.

Since $\widehat{\Phi}_{\beta-1}^i(\tau) = * \iff (\widehat{\phi}_{\beta-1}^i = * \text{ OR } \widehat{\Phi}_{\beta-2}^i = *)$, we have for every $i \in [y]$,

$$\begin{aligned} & \Pr \left[\left(h^i(\mathcal{V}_{\beta,\text{ideal}}^i(\tau_{\beta}), \eta_{\beta}^i(\tau_{\beta})) \notin \{*\} \cup \{(\widetilde{M}, \psi)\}_{\widetilde{M} \in \{0,1\}^p, \psi \in [\frac{3k}{4} - \ell + 1, \frac{3k}{4} + \ell]} \right) \wedge \left(\widehat{\Phi}_{\beta-1}^i(\tau_{\beta}) = * \right) \right] \\ &= \Pr \left[h^i(\mathcal{V}_{\beta,\text{ideal}}^i(\tau_{\beta}), \eta_{\beta}^i(\tau_{\beta})) \notin \{*\} \cup \{(\widetilde{M}, \psi)\}_{\widetilde{M} \in \{0,1\}^p, \psi \in [\frac{3k}{4} - \ell + 1, \frac{3k}{4} + \ell]} \wedge \left(\widehat{\phi}_{\beta-1}^i(\tau_{\beta}) = * \right) \right] \\ &+ \Pr \left[h^i(\mathcal{V}_{\beta,\text{ideal}}^i(\tau_{\beta}), \eta_{\beta}^i(\tau_{\beta})) \notin \{*\} \cup \{(\widetilde{M}, \psi)\}_{\widetilde{M} \in \{0,1\}^p, \psi \in [\frac{3k}{4} - \ell + 1, \frac{3k}{4} + \ell]} \wedge \left(\widehat{\Phi}_{\beta-2}^i(\tau_{\beta}) = * \right) \wedge \left(\widehat{\phi}_{\beta-1}^i(\tau_{\beta}) \neq * \right) \right] \\ &= \Pr \left[h^i(\mathcal{V}_{\beta,\text{ideal}}^i(\tau_{\beta}), \eta_{\beta}^i(\tau_{\beta})) \notin \{*\} \cup \{(\widetilde{M}, \psi)\}_{\widetilde{M} \in \{0,1\}^p, \psi \in [\frac{3k}{4} - \ell + 1, \frac{3k}{4} + \ell]} \wedge \left(\widehat{\phi}_{\beta-1}^i(\tau_{\beta}) = * \right) \right] \\ &+ \Pr \left[g\left(h^i(\mathcal{V}_{\beta,\text{ideal}}^i(\tau_{\beta}), \eta_{\beta}^i(\tau_{\beta})), \widehat{\phi}_{\beta-1}^i(\tau_{\beta})\right) \notin \{*\} \cup \{(\widetilde{M}, \psi)\}_{\widetilde{M} \in \{0,1\}^p, \psi \in [\frac{3k}{4} - \ell + 1, \frac{3k}{4} + \ell]} \wedge \left(\widehat{\Phi}_{\beta-2}^i(\tau_{\beta}) = * \right) \right. \\ &\quad \left. \wedge \left(\widehat{\phi}_{\beta-1}^i(\tau_{\beta}) \neq * \right) \right] \\ &\leq \mu(n) + \varepsilon. \end{aligned}$$

Since $\widehat{\Phi}_{\beta-1}$ is obtained by applying the leakage lemma (Lemma 2) to Equation (37), we have that for every $i \in [y]$, there exists a negligible function μ'_β such that

$$\Pr \left[\left(g \left(h^i \left(\mathcal{V}_{\beta+1, \text{ideal}}^i(\tau_{\beta+1}), \eta_{\beta+1}^i(\tau_{\beta+1}) \right), \widehat{\phi}_{\beta}^i(\tau_{\beta+1}) \right) \notin \{*\} \cup \{(\widetilde{M}, \psi)\}_{\widetilde{M} \in \{0,1\}^p, \psi \in \left[\frac{3k}{4} - \ell + 1, \frac{3k}{4} + \ell \right]} \right) \wedge \left(\widehat{\Phi}_{\beta-1}^i(\tau_{\beta+1}) = * \right) \right] \leq \varepsilon + \mu(n) + \mu'_\beta(n)$$

Setting $\mu_\beta = \mu + \mu'_\beta$, we have

$$\Pr \left[\left(g \left(h^i \left(\mathcal{V}_{\beta+1, \text{ideal}}^i(\tau_{\beta+1}), \eta_{\beta+1}^i(\tau_{\beta+1}) \right), \widehat{\phi}_{\beta}^i(\tau_{\beta+1}) \right) \notin \{*\} \cup \{(\widetilde{M}, \psi)\}_{\widetilde{M} \in \{0,1\}^p, \psi \in \left[\frac{3k}{4} - \ell + 1, \frac{3k}{4} + \ell \right]} \right) \wedge \left(\widehat{\Phi}_{\beta-1}^i(\tau_{\beta+1}) = * \right) \right] \leq \varepsilon + \mu_\beta(n)$$

as desired.

This implies that (setting $\beta = \ell - 1$):

$$\Pr \left[g \left(h^i \left(\mathcal{V}_{\ell, \text{ideal}}^i(\tau_\ell), \eta_\ell^i(\tau_\ell) \right), \widehat{\phi}_{\ell-1}^i(\tau_\ell) \right) \notin \{*\} \cup \{(\widetilde{M}, \psi)\}_{\widetilde{M} \in \{0,1\}^p, \psi \in \left[\frac{3k}{4} - \ell + 1, \frac{3k}{4} + \ell \right]} \right) \wedge \left(\widehat{\Phi}_{\ell-2}^i(\tau_\ell) = * \right) \right] = \text{negl}(n).$$

Since $\widehat{\Phi}_{\ell-1}^i(\tau) = * \iff (\widehat{\phi}_{\ell-1}^i = * \text{ OR } \widehat{\Phi}_{\ell-2}^i = *)$, we have for every $i \in [y]$,

$$\begin{aligned} & \Pr \left[\left(h^i \left(\mathcal{V}_{\ell, \text{ideal}}^i(\tau_\ell), \eta_\ell^i(\tau_\ell) \right) \notin \{*\} \cup \{(\widetilde{M}, \psi)\}_{\widetilde{M} \in \{0,1\}^p, \psi \in \left[\frac{3k}{4} - \ell + 1, \frac{3k}{4} + \ell \right]} \right) \wedge \left(\widehat{\Phi}_{\ell-1}^i(\tau_\ell) = * \right) \right] \\ &= \Pr \left[\left(h^i \left(\mathcal{V}_{\ell, \text{ideal}}^i(\tau_\ell), \eta_\ell^i(\tau_\ell) \right) \notin \{*\} \cup \{(\widetilde{M}, \psi)\}_{\widetilde{M} \in \{0,1\}^p, \psi \in \left[\frac{3k}{4} - \ell + 1, \frac{3k}{4} + \ell \right]} \right) \wedge \left(\widehat{\phi}_{\ell-1}^i(\tau_\ell) = * \right) \right] \\ &+ \Pr \left[\left(h^i \left(\mathcal{V}_{\ell, \text{ideal}}^i(\tau_\ell), \eta_\ell^i(\tau_\ell) \right) \notin \{*\} \cup \{(\widetilde{M}, \psi)\}_{\widetilde{M} \in \{0,1\}^p, \psi \in \left[\frac{3k}{4} - \ell + 1, \frac{3k}{4} + \ell \right]} \right) \wedge \left(\widehat{\Phi}_{\ell-2}^i(\tau_\beta) = * \right) \wedge \left(\widehat{\phi}_{\ell-1}^i(\tau_\ell) \neq * \right) \right] \\ &= \Pr \left[\left(h^i \left(\mathcal{V}_{\ell, \text{ideal}}^i(\tau_\ell), \eta_\ell^i(\tau_\ell) \right) \notin \{*\} \cup \{(\widetilde{M}, \psi)\}_{\widetilde{M} \in \{0,1\}^p, \psi \in \left[\frac{3k}{4} - \ell + 1, \frac{3k}{4} + \ell \right]} \right) \wedge \left(\widehat{\phi}_{\ell-1}^i(\tau_\ell) = * \right) \right] \\ &+ \Pr \left[g \left(h^i \left(\mathcal{V}_{\ell, \text{ideal}}^i(\tau_\ell), \eta_\ell^i(\tau_\ell) \right), \widehat{\phi}_{\ell-1}^i(\tau_\ell) \right) \notin \{*\} \cup \{(\widetilde{M}, \psi)\}_{\widetilde{M} \in \{0,1\}^p, \psi \in \left[\frac{3k}{4} - \ell + 1, \frac{3k}{4} + \ell \right]} \right) \wedge \left(\widehat{\Phi}_{\ell-2}^i(\tau_\beta) = * \right) \right] \\ &\wedge \left(\widehat{\phi}_{\ell-1}^i(\tau_\ell) \neq * \right) \\ &= \mu(n) + \text{negl}(n) = \text{negl}(n). \end{aligned}$$

Since $\eta_\ell := 0^y$, this implies:

$$\Pr \left[\left(\mathcal{V}_{\ell, \text{ideal}}^i(\tau_\ell) \notin \{*\} \cup \{(\widetilde{M}, \psi)\}_{\widetilde{M} \in \{0,1\}^p, \psi \in \left[\frac{3k}{4} - \ell, \frac{3k}{4} + \ell \right]} \right) \wedge \left(\widehat{\Phi}_{\ell-1}^i(\tau_\ell) = * \right) \right] = \text{negl}(n)$$

which implies

$$\Pr \left[\left(\mathcal{V}_{\ell, \text{ideal}}^i(\tau_\ell) \notin \{(\widetilde{M}, k)\}_{\widetilde{M} \in \{0,1\}^p} \right) \wedge \left(\widehat{\Phi}_{\ell-1}^i(\tau_\ell) = * \right) \right] = \text{negl}(n),$$

as desired. This completes the proof of validity, and also concludes the proof of the theorem. \square

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 3. *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 satisfying [Definition 4](#), 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 1](#), assuming sub-exponential hardness of factoring or discrete log, and sub-exponential quantum hardness of LWE. By [Remark 2](#), this scheme also satisfies [Definition 3](#) for any polynomials ℓ and y .

- 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 2](#), 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 2](#) 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 2](#), 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 2](#) 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 2](#), we note that $z = 10\ell y$ and that $\mathcal{C}_{[\log^2 n]}$ 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 2](#) 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 (randomized) 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 that Com_{vk} satisfies [Definition 1](#).

□

Acknowledgements

The second author is very grateful to Amit Sahai for many discussions, going back to 2015, about the question of building non-interactive non-malleable commitments from “incomparable one-way functions”, which was very similar to the notion of tag-based adaptive commitments for two tags.

References

- [Bar02] Boaz Barak. Constant-Round Coin-Tossing with a Man in the Middle or Realizing the Shared Random String Model. In *FOCS 2002*, pages 345–355, 2002. 1
- [BD18] Zvika Brakerski and Nico Döttling. Two-message statistically sender-private OT from LWE. In Amos Beimel and Stefan Dziembowski, editors, *Theory of Cryptography - 16th International Conference, TCC 2018, Panaji, India, November 11-14, 2018, Proceedings, Part*

- II*, volume 11240 of *Lecture Notes in Computer Science*, pages 370–390. Springer, 2018. [52](#), [53](#)
- [BDK⁺18] Marshall Ball, Dana Dachman-Soled, Mukul Kulkarni, Huijia Lin, and Tal Malkin. Non-malleable codes against bounded polynomial time tampering. *IACR Cryptology ePrint Archive*, 2018:1015, 2018. [6](#)
- [BFMR18] Brandon Broadnax, Valerie Fetzer, Jörn Müller-Quade, and Andy Rupp. Non-malleability vs. cca-security: The case of commitments. In Michel Abdalla and Ricardo Dahab, editors, *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*, volume 10770 of *Lecture Notes in Computer Science*, pages 312–337. Springer, 2018. [2](#), [3](#), [5](#)
- [BGI⁺17] Saikrishna Badrinarayanan, Sanjam Garg, Yuval Ishai, Amit Sahai, and Akshay Wadia. Two-message witness indistinguishability and secure computation in the plain model from new assumptions. *IACR Cryptology ePrint Archive*, 2017:433, 2017. [52](#), [54](#)
- [BL18] Nir Bitansky and Huijia Lin. One-message zero knowledge and non-malleable commitments. *IACR Cryptology ePrint Archive*, 2018:613, 2018. [1](#), [4](#), [6](#)
- [BOV07] Boaz Barak, Shien Jin Ong, and Salil P. Vadhan. Derandomization in cryptography. *SIAM J. Comput.*, 37(2):380–400, 2007. [4](#), [32](#)
- [BP04] Mihir Bellare and Adriana Palacio. The knowledge-of-exponent assumptions and 3-round zero-knowledge protocols. In *CRPTO 2004*, pages 273–289, 2004. [5](#)
- [BP15] Nir Bitansky and Omer Paneth. Zaps and non-interactive witness indistinguishability from indistinguishability obfuscation. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *Theory of Cryptography - 12th Theory of Cryptography Conference, TCC 2015, Warsaw, Poland, March 23-25, 2015, Proceedings, Part II*, volume 9015 of *Lecture Notes in Computer Science*, pages 401–427. Springer, 2015. [4](#)
- [CGM⁺16] Nishanth Chandran, Vipul Goyal, Pratyay Mukherjee, Omkant Pandey, and Jalaj Upadhyay. Block-wise non-malleable codes. In Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi, editors, *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy*, volume 55 of *LIPIcs*, pages 31:1–31:14. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016. [5](#)
- [CLP15] Kai-Min Chung, Edward Lui, and Rafael Pass. From weak to strong zero-knowledge and applications. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *Theory of Cryptography - 12th Theory of Cryptography Conference, TCC 2015, Warsaw, Poland, March 23-25, 2015, Proceedings, Part I*, volume 9014 of *Lecture Notes in Computer Science*, pages 66–92. Springer, 2015. [32](#)
- [COSV16] Michele Ciampi, Rafail Ostrovsky, Luisa Siniscalchi, and Ivan Visconti. Concurrent non-malleable commitments (and more) in 3 rounds. In Matthew Robshaw and Jonathan Katz, editors, *Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part III*, volume 9816 of *Lecture Notes in Computer Science*, pages 270–299. Springer, 2016. [1](#), [5](#)

- [COSV17] Michele Ciampi, Rafail Ostrovsky, Luisa Siniscalchi, and Ivan Visconti. Four-round concurrent non-malleable commitments from one-way functions. In *Annual International Cryptology Conference*, pages 127–157. Springer, 2017. 1
- [DDN91] Danny Dolev, Cynthia Dwork, and Moni Naor. Non-Malleable Cryptography (Extended Abstract). In *STOC 1991*, 1991. 1, 6
- [FF18] Peter Fenteany and Benjamin Fuller. Non-malleable digital lockers. Cryptology ePrint Archive, Report 2018/957, 2018. <https://eprint.iacr.org/2018/957>. 6
- [GHKW17] Rishab Goyal, Susan Hohenberger, Venkata Koppula, and Brent Waters. A generic approach to constructing and proving verifiable random functions. In Kalai and Reyzin [KR17], pages 537–566. 3, 23, 24
- [GKS16] Vipul Goyal, Dakshita Khurana, and Amit Sahai. Breaking the three round barrier for non-malleable commitments. In *FOCS*, 2016. 2, 5
- [GLOV12] Vipul Goyal, Chen-Kuei Lee, Rafail Ostrovsky, and Ivan Visconti. Constructing non-malleable commitments: A black-box approach. In *FOCS*, 2012. 1
- [GOS12] Jens Groth, Rafail Ostrovsky, and Amit Sahai. New techniques for noninteractive zero-knowledge. *J. ACM*, 59(3):11:1–11:35, 2012. 4, 32
- [Goy11] Vipul Goyal. Constant Round Non-malleable Protocols Using One-way Functions. In *STOC 2011*, pages 695–704. ACM, 2011. 1, 2, 6, 13
- [GPR16] Vipul Goyal, Omkant Pandey, and Silas Richelson. Textbook non-malleable commitments. In *STOC*, pages 1128–1141, New York, NY, USA, 2016. ACM. 1
- [GRRV14] Vipul Goyal, Silas Richelson, Alon Rosen, and Margarita Vald. An algebraic approach to non-malleability. In *FOCS 2014*, pages 41–50, 2014. 1, 5
- [GW11] Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In Lance Fortnow and Salil P. Vadhan, editors, *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pages 99–108. ACM, 2011. 3, 4, 9, 11, 31, 32
- [JKKR17] Abhishek Jain, Yael Tauman Kalai, Dakshita Khurana, and Ron Rothblum. Distinguisher-dependent simulation in two rounds and its applications. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part II*, volume 10402 of *Lecture Notes in Computer Science*, pages 158–189. Springer, 2017. 4, 52, 54
- [JP14] Dimitar Jetchev and Krzysztof Pietrzak. How to fake auxiliary input. In Yehuda Lindell, editor, *Theory of Cryptography - 11th Theory of Cryptography Conference, TCC 2014, San Diego, CA, USA, February 24-26, 2014. Proceedings*, volume 8349 of *Lecture Notes in Computer Science*, pages 566–590. Springer, 2014. 32
- [Khu17] Dakshita Khurana. Round optimal concurrent non-malleability from polynomial hardness. In Kalai and Reyzin [KR17], pages 139–171. 1

- [KR17] Yael Kalai and Leonid Reyzin, editors. *Theory of Cryptography - 15th International Conference, TCC 2017, Baltimore, MD, USA, November 12-15, 2017, Proceedings, Part II*, volume 10678 of *Lecture Notes in Computer Science*. Springer, 2017. 50
- [KS17] Dakshita Khurana and Amit Sahai. Birthday simulation from exponential hardness: 2 round non-malleable commitments and 3 round gap zk. 2017. 1, 7, 11
- [KY18] Ilan Komargodski and Eylon Yogev. Another step towards realizing random oracles: Non-malleable point obfuscation. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part I*, volume 10820 of *Lecture Notes in Computer Science*, pages 259–279. Springer, 2018. 6
- [Lev87] Leonid A. Levin. One-way functions and pseudorandom generators. *Combinatorica*, 7(4):357–363, 1987. 3, 52
- [LP] Huijia Lin and Rafael Pass. Constant-round Non-malleable Commitments from Any One-way Function. In *STOC 2011*, pages 705–714. 1
- [LP09] Huijia Lin and Rafael Pass. Non-malleability Amplification. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC '09*, pages 189–198, 2009. 1, 6
- [LP12] Huijia Lin and Rafael Pass. Black-box constructions of composable protocols without set-up. In Reihaneh Safavi-Naini and Ran Canetti, editors, *Advances in Cryptology - CRYPTO 2012 - 32nd Annual Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2012. Proceedings*, volume 7417 of *Lecture Notes in Computer Science*, pages 461–478. Springer, 2012. 3
- [LPS17] Huijia Lin, Rafael Pass, and Pratik Soni. 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>. 1, 4, 5, 6, 11
- [LPV] Huijia Lin, Rafael Pass, and Muthuramakrishnan Venkitasubramaniam. Concurrent Non-malleable Commitments from Any One-Way Function. In *TCC 2008*, pages 571–588. 1, 17, 20
- [PPV08] Omkant Pandey, Rafael Pass, and Vinod Vaikuntanathan. Adaptive One-Way Functions and Applications. In *Advances in Cryptology — CRYPTO '08*, pages 57–74, 2008. 1, 3, 6, 24
- [PR05a] Rafael Pass and Alon Rosen. Concurrent Non-Malleable Commitments. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science, FOCS '05*, pages 563–572, 2005. 1
- [PR05b] Rafael Pass and Alon Rosen. New and improved constructions of non-malleable cryptographic protocols. In *STOC 2005*, pages 533–542, 2005. 16
- [PR08] Rafael Pass and Alon Rosen. New and Improved Constructions of Nonmalleable Cryptographic Protocols. *SIAM J. Comput.*, 38(2):702–752, 2008. 1

- [PW10] Rafael Pass and Hoeteck Wee. Constant-round non-malleable commitments from sub-exponential one-way functions. In *EUROCRYPT 2010*, pages 638–655, 2010. [1](#), [3](#), [12](#), [18](#), [19](#)
- [Wee10] Hoeteck Wee. Black-box, round-efficient secure computation via non-malleability amplification. In *FOCS 2010*, pages 531–540, 2010. [1](#), [6](#)

A Two Message Arguments from Quantum Polynomial Hardness

In this section, we construct a two-message argument system assuming quantum (polynomial) hardness of LWE and polynomial hardness of factoring or discrete log. Recall that an argument system is a proof system where soundness holds only against computationally bounded (polynomial-size) provers.

The argument system we construct satisfies zero-knowledge with super-polynomial simulation and witness indistinguishability, as well as weak zero-knowledge, strong witness indistinguishability and witness hiding against delayed-input verifiers.

This argument system is the same as the one in [JKKR17], except that while the [JKKR17] protocol relies on any commitment and any two-message OT with security against unbounded receivers and *super-polynomial senders*, we rely on commitments based on factoring or discrete log [Lev87] and OT based on LWE [BD18] with security against unbounded receivers and quantum polynomial senders. This allows us to use quantum supremacy instead of complexity leveraging to argue soundness of the protocol.

A.1 Modified Blum Protocol

We use a specific instantiation of the Blum Σ -protocol for Graph Hamiltonicity. Recall that in the Blum protocol, the prover commits to a random permutation and the permuted graph. If the verifier’s challenge bit is 0, the prover decommits to the permutation, and if the bit is 1, decommits to a Hamiltonian cycle in this permuted graph. In our setting, instead of using *any* commitment, we will use a statistically binding, quantum-breakable non-interactive commitment³¹. Quantum breakable commitments are statistically binding (string) commitments that are computationally hiding against polynomial-size receivers, but where the value committed to by any malicious PPT committer can be extracted in quantum polynomial time. As described in Section 4.2, such a scheme exists based on any one-way function that is invertible in BQP, such as one based on factoring or discrete log.

This protocol has soundness at least $\frac{1}{2} - \text{negl}(n)$ against quantum polynomial-size provers, and satisfies honest-verifier zero-knowledge against PPT verifiers.

A.2 Construction of Two-message Arguments

Our two-message argument is essentially identical to the one in [JKKR17, BGI⁺17]³². Recall that these protocols use two-message oblivious transfer (OT) with super-polynomial security to convert the Blum Σ -protocol into a two-message argument system. We will use the modified Blum protocol above together with polynomially secure OT which is based on LWE .

³¹These can be non-interactive or 2-message commitment; we rely on non-interactive commitments for simplicity.

³²The only difference is that we use specific LWE -based OT and specific commitments, whereas they rely on super-polynomially secure OT and any commitment.

Let $\text{OT}_Q = (\text{OT}_{Q,1}, \text{OT}_{Q,2})$ denote a two-message bit oblivious transfer protocol. Let $\text{OT}_{Q,1}(b)$ denote the first (receiver) message of this OT protocol with receiver input b , and let $\text{OT}_{Q,2}(m_0, m_1)$ denote the second (sender) message of the OT protocol with sender input bits m_0, m_1 ³³. We require this protocol to satisfy the following properties:

- The distributions $\text{OT}_{Q,1}(0)$ and $\text{OT}_{Q,1}(1)$ are computationally indistinguishable.
- For every unbounded malicious receiver that outputs first message o_1 , there exists a bit b such that for every pair of messages (m_0, m_1) , the distributions $\text{OT}_{Q,2}(m_0, m_1)$ and $\text{OT}_{Q,2}(m_b, m_b)$ are statistically indistinguishable.

This can be instantiated assuming quantum polynomial hardness of LWE [BD18].

Let $\Sigma = (a, e, z)$ denote the three messages of the modified Blum protocol, where a denotes the message of our underlying (quantum-breakable) commitment. We will perform a parallel repetition of this protocol, thus for each $i \in [n]$, (a_i, e_i, z_i) are messages corresponding to an underlying modified Blum protocol with a single-bit challenge (i.e., where $e_i \in \{0, 1\}$).

Two-Message Argument

- **Verifier Message:**
 - For every $i \in [n]$ do the following:
 - * Sample challenge e_i uniformly at random.
 - * Compute $o_{1,i} = \text{OT}_{Q,1,i}(e_i)$.
 - Send $\{o_{1,i}\}_{i \in [n]}$ in parallel.
- **Prover Message:**
 - Obtain input $x \in L$, witness w such that $R_L(x, w) = 1$.
 - For every $i \in [n]$, do:
 - * Emulate the honest prover of the Blum Protocol by sampling commitment a_i and compute answers z_i^0, z_i^1 corresponding to verifier challenges 0 and 1 respectively.
 - * Compute $o_{2,i} = \text{OT}_{Q,2,i}(z_i^0, z_i^1)$.
 - Send $\{a_i, o_{2,i}\}_{i \in [n]}$.
- **Verifier Output:** For every $i \in [n]$ verifier V recovers z_i as the output of $(\text{OT}_{Q,1,i}, \text{OT}_{Q,2,i})$ for $i \in [n]$, and outputs accept if for all $i \in [n]$, $(a_i, e_i, z_i)_{i \in [n]}$ is an accepting transcript of the underlying modified Blum protocol.

Figure 4: Two Message Argument System for NP

Theorem 4. *Assuming quantum polynomial hardness of LWE, and polynomial hardness of factoring or discrete log, there exist two-message arguments with adaptive soundness against polynomial-sized quantum provers, and satisfying:*

³³We note that the second (sender) message also depends on the first (receiver) message; we omit this dependence here for succinctness.

- *Witness Indistinguishability and ZK with superpolynomial simulation against all malicious verifiers*
- *Distributional weak zero-knowledge, strong WI and witness hiding against delayed-input verifiers.*

Proof. (Sketch) In this proof sketch, we focus on proving soundness.

We note that the proof of secrecy follows directly from previous work [BGI⁺17, JKKR17]. This is because the protocol in [BGI⁺17, JKKR17] is proven secure assuming *any* two-round malicious OT with security against unbounded receivers, and *any* commitment; and we only change it to rely on a *specific* instantiation of two-round malicious OT and commitments.

To prove soundness, [BGI⁺17, JKKR17] relied on the fact that OT hides the receiver choice bit from *super-polynomial senders*. We show that super-polynomial assumptions are not required if we assume that the receiver's choice bit is hidden from quantum polynomial-sized machines.

We show that if the underlying commitment is quantum-breakable and if OT_Q hides the receiver choice bit from polynomial-sized quantum malicious senders, the protocol in Figure 4 satisfies adaptive soundness against polynomial-sized quantum provers.

We will prove this by contradiction. Fix any cheating quantum prover P^* and any $p(n) = \frac{1}{\text{poly}(n)}$, such that that with probability at least $p(n)$ over the randomness of sampling verifier message, for infinitely many $n \in \mathbb{N}$, P^* outputs $x \notin L$ together with an accepting transcript for x according to the protocol in Figure 4. We will construct a quantum polynomial-size reduction R^{P^*} that on input $\{o_{1,i} \stackrel{\$}{\leftarrow} \text{OT}_{1,i}(e_i)\}_{i \in [n]}$ for $e = e_1 e_2 \dots e_n \stackrel{\$}{\leftarrow} \{0, 1\}^n$, with oracle access to P^* , outputs \tilde{e} such that with probability at least $\frac{1}{p(n)}$, $\tilde{e} = e$ (contradicting the receiver security of the OT).

The reduction R on input $\{o_{1,i}\}_{i \in [n]}$ does the following:

- Send $\{o_{1,i}\}_{i \in [n]}$ to P^* and obtain $\{a_i\}_{i \in [n]}$.
- Break the quantum breakable commitments $\{a_i\}_{i \in [n]}$ to obtain $\{v_i\}_{i \in [n]}$.
- Compute the string \tilde{e} as follows:
 1. For all $i \in [n]$, if v_i consists of (π, G') such that $G' = \pi(G)$, set $\tilde{e}_i = 0$.
 2. Else set $\tilde{e}_i = 1$.
- Output $\tilde{e} = \tilde{e}_1, \dots, \tilde{e}_n$.

Recall that P^* breaks soundness with probability $p = \frac{1}{\text{poly}(n)}$ for infinitely many $n \in \mathbb{N}$. This means that over the randomness of sampling e , P^* outputs $x \notin L$ and a_i^* , $\text{OT}_2(z_i^0, z_i^1)$ for $i \in [n]$ that cause the verifier to accept with probability p .

Moreover, whenever $x \notin L$ (that is, the corresponding Graph G does not contain any Hamiltonian Cycle) and the transcript is accepting, then for every $i \in [n]$, either

- The committed graph is a correctly permuted variant of G . Since the transcript is accepting, this implies that $e_i = 0$, and note that in this case we set $\tilde{e}_i = 0$, or
- The committed graph consists of a Hamiltonian cycle. Since the transcript is accepting, this implies that $e_i = 1$, and note that in this case we set $\tilde{e}_i = 1$.

Therefore, in any accepting transcript where $x \notin L$, it must hold that $\tilde{e} = e$.

Since P^* outputs $x \notin L$ and an accepting proof for x with probability at least $p(n)$ (for infinitely many $n \in \mathbb{N}$), it follows that $\tilde{e} = e$ with probability at least $p(n)$ for infinitely many $n \in \mathbb{N}$, which contradicts receiver input-hiding security of the OT against polynomial-sized quantum machines. This completes the proof of soundness. \square