

Non-malleable Commitments against Quantum Attacks

Nir Bitansky*

Huijia Lin[†]

Omri Shmueli[‡]

Abstract

We construct, under standard hardness assumptions, the first non-malleable commitments secure against quantum attacks. Our commitments are statistically binding and satisfy the standard notion of *non-malleability with respect to commitment*. We obtain the following instantiations:

- A $\log^*(\lambda)$ -round classical protocol based on quantum fully-homomorphic encryption and the quantum hardness of Learning with Errors.
- A polynomial-round classical protocol based on post-quantum oblivious transfer.
- A polynomial-round quantum protocol based on post-quantum one-way functions.

Previously, non-malleable commitments with quantum security were only known against a restricted class of adversaries known as *synchronizing adversaries*. At the heart of our results is a general technique that allows to modularly obtain non-malleable commitments from any extractable commitment protocol, obliviously of the underlying extraction strategy (black-box or non-black-box), round complexity, and whether communication is quantum or classical. The transformation preserves the quantum security of the underlying extractable commitments, and is new even in the classical setting.

*Tel Aviv University, nirbitan@tau.ac.il. Member of the Check Point Institute of Information Security. Supported by ISF grants 18/484 and 19/2137, by Len Blavatnik and the Blavatnik Family Foundation, by the Blavatnik Interdisciplinary Cyber Research Center at Tel Aviv University, and by the European Union Horizon 2020 Research and Innovation Program via ERC Project REACT (Grant 756482).

[†]University of Washington, rachel@cs.washington.edu. Supported by NSF grant CNS-1936825 (CAREER), CNS-2026774, a Hellman Fellowship, a JP Morgan AI Research Award, a Simons Collaboration grant on the Theory of Algorithmic Fairness.

[‡]Tel Aviv University, omrishmueli@mail.tau.ac.il. Supported by ISF grants 18/484 and 19/2137, by Len Blavatnik and the Blavatnik Family Foundation, and by the European Union Horizon 2020 Research and Innovation Program via ERC Project REACT (Grant 756482).

Contents

1	Introduction	1
2	Technical Overview	4
2.1	Understanding the Challenges	4
2.2	Leveraging Extractable Commitments in the Non-Synchronizing Setting	6
2.3	Tag Amplification	10
3	Preliminaries	13
3.1	Interactive Protocols, Witness Indistinguishability, and Zero Knowledge	14
3.1.1	Witness Indistinguishability	14
3.1.2	Quantum Zero-Knowledge Protocols	15
3.2	Commitments	15
3.3	Quantum Rewinding Lemma	16
3.4	Non-Malleable Commitments	17
3.5	Committed Value Oracle	18
3.6	Extractable Commitments	18
3.6.1	r -Robustness.	19
3.6.2	First-message binding.	19
4	Two-sided Extraction	20
4.1	Two-sided Extractor	20
4.2	Proof of Two-sided Extraction	21
4.3	Non-Interleaving Extraction	22
5	Post-quantum Non-malleable Commitment For Few Tags	23
5.1	Ideally-Scheduled Block Commitments	24
5.2	Adversaries with Predetermined Ideal Schedule	26
5.3	Proof of Proposition 5.1	27
5.4	Robustness	30
6	Tag Amplification	31
6.1	Analysis of The Tag Amplification	32

1 Introduction

Commitments are one of the most basic cryptographic primitives. They enable a sender to commit to a string to be *opened* at a later stage. As long as the commitment is not opened, it is *hiding* — efficient receivers learn nothing about the committed value. Furthermore, the commitment is *statistically binding* — with overwhelming probability, the commitment can be opened to a single, information-theoretically determined value in the commitment phase. While these basic security guarantees go a long way in terms of applications, they do not always suffice. In particular, they do not prevent a *man-in-the-middle* adversary from receiving a commitment to a given value v from one party and trying to send to another party a commitment to a related value, say $v - 1$ (without knowing the committed value v at all).

Such attacks are called "mauling attacks" and in some settings could be devastating. For instance, consider the scenario where a city opens a bidding process for the construction of a new city hall. Companies are instructed to commit to their proposed bid using a commitment scheme, and these commitments are opened at the end of the bidding period. If the scheme is "malleable", company A may manage to underbid company B , by covertly mauling B 's commitment to create their own commitment to a lower bid. More generally, ensuring independence of private values is vital in many applications of commitments, such as coin tossing, federated learning, and collaborative computation over private data.

In their seminal work, Dolev, Dwork and Naor introduced the concept of *non-malleable commitments* to protect against mauling attacks [DDN03]. They guarantee that the value \tilde{v} a man-in-the-middle adversary commits to is computationally independent of the value v in the commitment it receives (unless the man-in-the-middle simply "copies", by relaying messages between the honest sender and receiver it interacts with, in which case $\tilde{v} = v$). From its onset, the study of non-malleable cryptography has put stress on achieving solutions without any reliance on trusted parties or any form of trusted setup, and solutions that hold when honest parties may not even be aware of the existence of a man-in-the-middle, and the way it manipulates the messages they send over time. The latter is particularly important in applications where the man-in-the-middle acts "in the dark". For instance, in the aforementioned example, company A may not be aware of the competing company B .

Since their conception, non-malleable commitments have indeed proved to be a useful and versatile building block for ensuring independence of values. They have been used in coin-tossing protocols, secure multiparty computation protocols, non-malleable proof systems (zero-knowledge, witness indistinguishability, multi-prover interactive proofs), and more. Techniques developed for non-malleable commitments are also useful for building non-malleable codes, non-malleable extractors (and two source extractors), and non-malleable time-lock puzzles. The work of [DDN03] constructed the first non-malleable commitments against classical adversaries based on one-way functions. Since then, a plethora of constructions have been proposed achieving different, sometimes optimal, tradeoffs between round-complexity, efficiency, and underlying assumptions (c.f. [Bar02, PR05a, PR05b, LPV08a, LPV09, LP09, PPV08, PW10, Wee10, Goy11, LP11, GLOV12, COSV16, GPR16a, GKS16, COSV17, Khu17, KS17, LPS17, BL18, KK19, GR19, GKLW20]).

Non-Malleability Against Quantum Adversaries. In contrast to the comprehensive understanding of non-malleability in the classical setting, our understanding of non-malleability against quantum adversaries is very much lacking. The threat of quantum attacks has prompted the development of post-quantum cryptography, and yet despite its important role in cryptography, post-quantum non-malleability has yet to catch up. In this work, we construct, under standard assumptions, the first non-malleable commitments with post-quantum security, namely, the hiding and non-malleability properties hold even against efficient quantum adversaries (and binding continues to be information theoretic).

Prior to our work, post-quantum non-malleable commitments were not known under any assumption. Partial progress was made by Agrawal, Bartusek, Goyal, Khurana, and Malavolta [ABG⁺20] who, assuming super-polynomial quantum hardness of Learning With Errors, construct post-quantum non-malleable commitments against a restricted class of adversaries known as *synchronizing adversaries*. A

synchronizing adversary is limited as follows: When acting as a man-in-the-middle between a sender and a receiver, it is bound to synchronize its interactions with the honest parties; namely, when it receives the i -th message from the sender, it immediately sends the i -th message to the receiver and vice versa. Such synchronicity may often not exist for example due to network’s asynchronicity, lack of synchronized clocks, or concurrent executions where parties are unaware of the existence of other executions. Enforcing synchronizing behaviour in general requires a trusted setup (like a broadcast channel) and coordination among parties to enforce message ordering.

The gold standard of non-malleability (since its introduction in [DDN03]) requires handling general, *non-synchronizing* adversaries, who can arbitrarily schedule messages in the two interactions (without awareness of the sender and receiver). In this work, for the first time, we achieve this gold standard non-malleability in the post-quantum setting. As we shall explain later on, the challenge stems from the fact that classical techniques previously used to obtain non-malleability against non-synchronizing adversaries (e.g., as robust extraction [LP09], simulation extractability [PR05a, PR05b] and so on) do not generally apply in the quantum setting. This is due to basic quantum phenomena such as *unclonability* [WZ82] and *state disturbance* [FP96].

In what follows, we first elaborate on our results, and then move on to discuss our techniques.

Our Results in More Detail

We construct statistically binding non-malleable commitments against quantum non-synchronizing adversaries, under standard hardness assumptions. Our main result is a modular construction of post-quantum non-malleable commitments from post-quantum extractable commitments. The latter is a statistically binding commitment protocol that is extractable in the following sense: There exists an efficient quantum extractor-simulator, which given the code of any quantum sender, can simulate its (arbitrary) quantum output, while extracting the committed value.

Theorem 1.1 (Informal). *Assuming k -round post-quantum extractable commitments, there exist $k^{O(1)} \log^* \lambda$ -round post-quantum non-malleable commitments, where λ is the security parameter.*

By default, by post-quantum we mean protocols that can be executed by classical parties, but which are secure against quantum adversaries. In particular, starting from a post-quantum classical extractable commitment we obtain a post-quantum classical non-malleable commitment. However, our result is in fact more general — we can also rely on quantum extractable commitments (i.e. where honest parties can use quantum computation and communication), provided that the committed value is defined by a classical statistically binding commitment. In this case, the resulting non-malleable commitment will also be quantum. Our construction is modular, in particular, it does not depend on the specific extraction strategy used by the underlying extractable commitments, may it be rewinding-based or non-black-box. We therefore obtain several immediate instantiations. First, using the constant-round post-quantum extractable commitments from [BS20], we obtain $O(\log^* \lambda)$ -round post-quantum non-malleable commitments under the same standard assumptions: Quantum Learning with Errors (QLWE) and Quantum Fully-Homomorphic Encryption (QFHE) [Mah18, Bra18].

Corollary 1.1. *Assuming QFHE and QLWE, there exist $O(\log^* \lambda)$ -round post-quantum non-malleable commitments with classical parties.*

More generally, post-quantum k -round extractable commitments are known from any post-quantum, k_1 -round zero-knowledge argument and k_2 -round (semi-honest) oblivious transfer, provided that $k \geq k_1 + k_2 + O(1)$ [HSS15, LN11, BS20]. Since zero knowledge proofs are known from post-quantum one-way functions [Wat09], which in turn follows from post-quantum oblivious transfer, we obtain the following corollary.

Corollary 1.2. *Assuming post-quantum oblivious transfer (with classical parties), there exist $\text{poly}(\lambda)$ -round post-quantum non-malleable commitments with classical parties.*

We note that, in fact, the round complexity of Watrous’ zero-knowledge proof [Wat09] can be reduced to any super constant function $r(\lambda) = \omega(1)$, using standard techniques.¹ Accordingly, we can improve the round complexity to $O(r)$ rounds for any $r = \omega(1)$, provided $O(r)$ -round post-quantum oblivious transfer. However, in this case, although negligible, the soundness error of the corresponding zero-knowledge proofs is relatively large λ^{-r} , and the resulting commitments will thus be at most λ^r -secure (for instance, a poly-logarithmic-round protocol, would be at most quasi-polynomially secure). In contrast, the protocols in Corollaries 1.1, 1.2, and 1.3 below preserve the security of the underlying primitives (up to polynomial loss).

As noted above, our construction in fact works even for extractable commitments with quantum parties, as long as the committed value is determined by a classical statistically binding commitment scheme (given as part of the communication). This property is satisfied by the recent construction of extractable commitments from post-quantum one-way functions by [GLSV21].

Corollary 1.3. *Assuming post-quantum one-way functions, there exist $\text{poly}(\lambda)$ -round quantum-secure non-malleable commitments with quantum parties.*

Our Techniques in a Nutshell. Following the common approach in the literature, we focus on the notion of *tag-based non-malleable commitments* [DDN03], where every commitment is associated with a string called a tag, representing the identity of the party to which the commitment is sent. Indeed, when the number of tags is exponential this notion is known to be equivalent to the aforementioned basic notion of non-malleable commitments. Non-malleability according to this notion asserts that a man-in-the-middle (MIM) (quantum) adversary \mathcal{A} , which interacting with a sender Sen on the *left*, receives a commitment to a value v under one tag tg , cannot interact with a receiver Rec on the *right*, sending it a commitment under a different tag $\tilde{\text{tg}} \neq \text{tg}$ to a related value \tilde{v} .

A bit more formally, we require that for every efficient quantum MIM \mathcal{A} , and two values v_0, v_1 , the output state of \mathcal{A} and its committed value \tilde{v} , are computationally indistinguishable (where \mathcal{A} can choose the tags tg , $\tilde{\text{tg}}$, and \tilde{v} is set to \perp if $\text{tg} = \tilde{\text{tg}}$). Importantly, the MIM adversary \mathcal{A} is non-synchronizing: it can schedule the messages in the left and right interactions in an arbitrary order.

Toward our construction, we take the following common two-step approach:

- **Step 1 – Base commitment for few tags:** we first construct a $k^{O(\tau)}$ -round protocol for a constant number τ of tags from a k -round extractable commitment protocol.
- **Step 2 – Non-malleability amplification:** we present a non-malleability amplification procedure that transforms t -bit tag protocols to $2^{\Omega(t)}$ -bit tag protocols for any $t \in [3, O(\log \lambda)]$. Applying this procedure recursively on the above base protocol $O(\log^* \lambda)$ times allows obtaining 2^λ -tag non-malleability.

As mentioned, this 2-step approach is common in the literature. However, obtaining these two building blocks in the setting of quantum adversaries is not known and requires new ideas.

Base Non-Malleable Commitments. We give a generic transformation from any extractable commitment to a non-malleable commitment with a *constant* number of tags. The main challenge here is dealing with *general* extractable commitments, whereas prior constructions of classically-secure non-malleable

¹For instance, by starting from Blum’s Hamiltonicity protocol [Blu87] that has constant soundness, repeating it in parallel logarithmically many times, and then sequentially $\omega(1)$ times.

commitments heavily leverage special properties of the extraction procedure in specific extractable commitments, which do not hold in the quantum setting.²

In a nutshell, our new technique uses a general extractable commitment scheme to commit to *different sets* of secret shares of the committed value, and then arrange the commitments according to a carefully chosen scheduling, determined by the tag. We then prove, by a combinatorial argument, that whenever the tags are different, one can extract the value the man-in-the-middle commits to, while keeping the commitment it receives hiding, which then implies non-malleability. Along the way, we also develop techniques for working with general extractable commitments in the quantum setting, such as, two-sided extraction, robust extraction, and non-interleaving extraction. In light of the general difficulty of working with extraction in the quantum setting, we believe that our techniques may be helpful for future works.

Non-Malleability Amplification. Classical tag-amplification procedures (where one transforms a non-malleable commitment scheme with respect to a small number of tags to one with a large number of tags) are common in the literature [DDN03, LP11, KS17, LPS17] but do not apply in the quantum setting. The work of [ABG⁺20] showed an amplification procedure secure against synchronizing quantum adversaries. We give the first amplification procedure that is secure against non-synchronizing quantum adversaries. Like in the case of constructing base non-malleable commitments, the procedure can be based on any k -round post-quantum extractable commitment, and incurs an additive $O(k)$ round overhead. We believe this post-quantum amplification procedure is of independent interest.

2 Technical Overview

We now give an overview of the main ideas behind our construction. Following the convention in the non-malleability literature, we refer to the interaction between Sen and \mathcal{A} as the left interaction/commitment, and that between Rec and \mathcal{A} the right interaction/commitment. Similarly, we refer to v , tg (and \tilde{v} , $\tilde{\text{tg}}$) as the left (and right) committed values or tag.

2.1 Understanding the Challenges

Before presenting our base commitments, we explain the main challenges that arise in the quantum setting. First, we recall a basic approach toward proving non-malleability in the classical setting *via extraction*. Here the basic idea is to provide a reduction that given a MIM adversary \mathcal{A} , can efficiently extract the value \tilde{v} that \mathcal{A} commits to on the right. Accordingly, if the MIM \mathcal{A} manages to maul the commitment to v on the left and commit to a related value \tilde{v} on the right, the reduction will gain information about v , and be able to break the *hiding* of the commitment.

The Difficulty in MIM Extraction. Extractable commitments allow for efficient extraction from adversarial senders in the *stand-alone setting*. Such extraction is traditionally done by either means of rewinding, or more generally using the sender’s code. In the MIM setting, where \mathcal{A} acts as a sender on the right, while acting as a receiver on the left, extraction from \mathcal{A} is much more challenging. The problem is that the interaction of \mathcal{A} with the receiver Rec on the right may occur concurrently to its interaction with the sender Sen on the left. This means that a reduction attempting to rewind \mathcal{A} to extract the right committed value, may effectively also need to rewind the sender Sen on the left. (This may happen for example if, when the reduction rewinds \mathcal{A} and sends \mathcal{A} a new message, \mathcal{A} also sends a new message in the left commitment and expects a reply from Sen before proceeding in the right commitment.) In such a case, extraction does not generally work — the “actual” sender of the right commitment is essentially the MIM \mathcal{A} combined with the sender Sen on the left. However, the reduction does not possess the code of

²For example, 3-round special-sound proofs used in [DDN03, LP09, LP11, Goy11], or Barak’s non-black-box extraction strategy in [PR05b, PR05a]. These schemes, which are used as building blocks in previous classical non-malleable commitments, are only known to be extractable against classical adversaries, but not quantum ones.

Sen, specifically, it does not possess its randomness. The challenge is to perform such extraction without access to the secret randomness of the sender on the left, and thus without compromising the hiding of the left commitment.

Indeed, classical non-malleable commitments tend to require more than *plain extractable commitments*. A long array of works (c.f., [DDN03, PR05b, PR05a, LP09, PW10, LP11, Goy11]) design various *safe extraction techniques*, which guarantee extraction on the right without compromising hiding of the left committed value. These safe-extraction techniques rely on properties of specific protocols and extraction strategies, rather than general (stand-alone) extractable commitments. For instance, the protocols of [DDN03, LP09, LP11, Goy11, GPR16b] rely on three-message witness-indistinguishable protocols satisfying an extraction guarantee known as *special soundness*, whereas the protocols in [PR05b, PR05a] rely on the specific structure of Barak’s non-black-box zero knowledge protocol.

The Quantum Barrier. The (safe) extraction techniques used to obtain non malleability in the classical setting fail in the quantum setting. For once, rewinding does not generally work. We cannot record the adversary’s quantum state between rewinding attempts due to the no-cloning theorem [WZ82]. Also, we cannot simply measure between rewindings, as this disturbs that the adversary’s state [FP96]. In this case, even if we do extract, we may not be able to faithfully simulate the adversary’s output state in the protocol³. Similarly, non-black-box techniques do not generally apply. For instance, it is unclear how to apply Barak’s non-black-box simulation technique [Bar02], due to the lack of *universal arguments* [BG08] for quantum computations (this is just to mention one difficulty in using Barak’s strategy in the quantum setting).

The difficulty of applying classical proof techniques in the setting of quantum adversaries is indeed a well known phenomena, and in some settings, quantum proof techniques have been successfully developed to circumvent this difficulty. Perhaps the most famous example of this is in the context of zero-knowledge simulation. Here Watrous [Wat09] shows that in certain settings quantum rewinding is possible and used it to obtain zero-knowledge protocols. Several other rewinding techniques enable extraction, but disturb the adversary’s state in the process [Unr12, CCY20, CMSZ21]. Alternatively, several recent works [AP19, BS20, ABG⁺20] obtain constant round zero-knowledge via non-black-box quantum techniques, using quantum FHE (and assuming LWE). As mentioned in the introduction, these post-quantum zero-knowledge protocols, combined with post-quantum oblivious transfer, do lead to constructions of post-quantum extractable commitments, *but only in the stand-alone setting* (rather than the MIM setting). In particular, known post-quantum extractable commitments do not satisfy the specific properties that the classical safe-extraction techniques require.

Given the above state of affairs, in this work, we aim to construct post-quantum non-malleable commitments modularly based on *any* post-quantum extractable commitment. Such a transformation would immediately yield post-quantum non-malleable commitments based on existing extractable commitments. More so, future advances in quantum extractable commitments, for instance, weakened assumptions or reduced round complexity, would directly translate to improvements in post-quantum non-malleable commitments. Furthermore, the equivalence between extractability and non-malleability is interesting on its own from a theoretical perspective.

The Synchronizing Setting. As observed in [ABG⁺20], if restricted to synchronizing adversaries, such a modular construction exists using ideas from early works [CR87, DDN03]: When committing under a tag $tg \in [\tau]$ for $\tau \leq \lambda$, in every round $i \neq tg$ send an empty message, and in round tg , send an extractable commitment to the value v . Indeed, in the synchronizing setting, a commitment on the left under tag tg would never interleave with the commitment on the right under tag $\tilde{tg} \neq tg$. Thus, safe-extraction opportunities come for free, circumventing the real challenge in achieving non-malleability. It is not hard to see, however, that in the non-synchronizing setting, this approach would

³Recall that non-malleability requires that the joint distribution of the output state of the adversary and the committed value are indistinguishable regardless of the committed value on the left. Hence the reduction needs to extract the committed value without disturbing the state of the quantum adversary.

completely fail as the adversary can always align the extractable commitment on the right with that on the left. The work of [ABG⁺20] further constructed constant-round non-malleable commitments for a super-constant number of tags, based on mildly super-polynomial security of quantum FHE and LWE. The non-malleability of the new protocol, however, still relies on the synchronization of the left and right commitments.

2.2 Leveraging Extractable Commitments in the Non-Synchronizing Setting

We design a base protocol for a constant number of tags that, using *any* (post-quantum) extractable commitment scheme. The protocol guarantees extraction on the right while preserving hiding on the left, even against a quantum non-synchronizing MIM adversary. In this overview, we explain our base commitments in three steps:

- First, we introduce our basic idea in the simplified *one-sided* non-malleability setting where the MIM is restricted to choose a smaller tag on the right than the tag on the left, $\tilde{t}g < tg$.
- Then, we extend the basic idea to the general setting where the MIM may also choose a right tag that is larger $\tilde{t}g > tg$. We illustrate the main ideas here under the simplifying assumption of a certain honest behavior of the adversary.
- Finally, we show how to remove the simplifying assumption on the adversary.

Step 1: One-sided Non-malleability Let us first consider a MIM adversary that given a commitment on the left under tag tg , produces a commitment on the right under a smaller tag $\tilde{t}g < tg$. In our commitment, the sender first secret shares the value v to be committed into shares u_1, \dots, u_n . It then sequentially sends extractable commitments to each of the shares u_1, \dots, u_n – we refer to the entire batch of these sequential extractable commitments as a *block-commitment to v* . The binding and hiding of this protocol follow directly from those of the underlying extractable commitment. We focus on non-malleability.

To achieve non-malleability, the number of shares n is chosen as a function of the tag tg . The goal is to guarantee that in every execution where the tag $\tilde{t}g$ on the right is smaller than the tag tg on the left, there will exist, on the *left*, a commitment to one of the shares u_i that is *free* in the sense that it does not interleave with the interaction on the right; namely, during the commitment to u_i on the left, no message is sent in the right execution (see Figure 1). Before explaining how freeness is achieved, let us explain how we use it to establish non-malleability.

Extracting While Preserving Hiding and First-Message Binding. To argue non-malleability, we show that we can efficiently extract *all* shares $\tilde{u}_1, \dots, \tilde{u}_{\tilde{n}}$ on the right, while preserving the hiding of the free share u_i on the left, and by the security of secret sharing, also the hiding of the committed value v .

Freeness guarantees that almost all commitments on the right do not interleave with the commitment to u_i on the left, more precisely, a single commitment on the right could be “split” by the commitment to u_i on the left (as in Figure 1), which prevents extraction of that right split commitment. To deal with this, we rely on extractable commitments that are *first-message binding*; namely their first sender message fixes the value of the commitment. This gives rise to a simple extraction strategy: for any commitment on the right, where the first sender’s message is sent before the free commitment (on the left), we can extract the corresponding share *non-uniformly*; for the commitments where the first sender’s message occurs afterwards, we use the efficient extractor. Accordingly, we get a non-uniform reduction to the hiding of the free extractable commitment on the left.

We observe that any extractable commitment can be made first-message binding without any additional assumptions, and while increasing round complexity by at most a constant factor. For simplicity

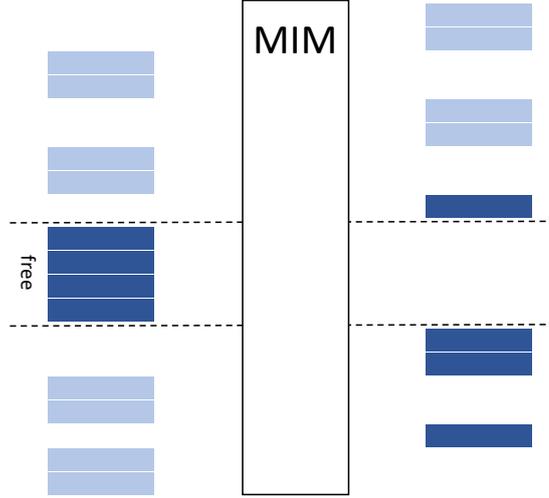


Figure 1: Freeness Example. Each share commitment has 4 messages and there are $n = 3$ shares on the left, and $\tilde{n} = 2$ shares on the right. The second commitment on the left is free. Note that it splits the second commitment on the right.

we describe how to achieve this assuming also non-interactive commitments.⁴ We append to the original extractable commitment a first message where the sender sends a non-interactive commitment to the committed value and add at the end a zero-knowledge argument that this commitment is consistent with the commitment in the original extractable commitment. Extractability follows from the extractability of the original scheme and soundness of the argument, whereas hiding follows from that of the original scheme and the zero knowledge property. We note that (post-quantum) zero-knowledge arguments follow from (post-quantum) extractable commitments with a constant round complexity overhead (see e.g. [BS20]).

Guaranteeing Freeness. To achieve the required freeness property, it suffices to guarantee that whenever $\tilde{\text{tg}} < \text{tg}$, the number of shares $n(\text{tg})$ (and hence the number of extractable commitments) on the left is larger than the total number of messages on the right, which is $k \cdot n(\tilde{\text{tg}})$, where k is the number of messages in each extractable commitment. Accordingly, we choose $n(\text{tg}) = (k + 1)^{\text{tg}}$.

Step 2: Dealing with General Adversaries. The above commitment does not prevent mauling of commitments under tag tg to commitments under tags $\tilde{\text{tg}} > \text{tg}$. To deal with general adversaries, we invoke the above idea again in reverse order. That is, the sender now secret shares the value v twice independently: once to n shares u_1, \dots, u_n , and again to \tilde{n} shares $\bar{u}_1, \dots, \bar{u}_{\tilde{n}}$. It then sequentially sends extractable commitments to the shares $u_1, \dots, u_n, \bar{u}_1, \dots, \bar{u}_{\tilde{n}}$, that is, sending two sequential block-commitments to v . To understand the basic idea, we assume for simplicity, in this step, that the MIM attacker always commits to shares of the same value \tilde{v} in the two block-commitments on the right (in Step 3, we will remove this assumption using zero-knowledge arguments).

Our goal now is to set the number of shares $n(\cdot), \tilde{n}(\cdot)$, based on the tags, to guarantee that there exists a block-commitment on the right with respect to which there exist two extractable commitments to shares u_i and $\bar{u}_{\tilde{i}}$ on the left (one from each left block-commitment) that are *free*. This means we can extract every share from that right block-commitment, while keeping the shares u_i and $\bar{u}_{\tilde{i}}$, and hence the left committed value, hidden. We say that the corresponding block-commitment on the right is *ideally scheduled* (see Figure 2).

Once we establish the existence of an ideally scheduled block, we can prove non-malleability using a non-uniform reduction to the hiding of the extractable commitments to u_i and $\bar{u}_{\tilde{i}}$ similar to the one we

⁴In the body, we observe that Naor commitments [Nao91], which can be obtained from (post-quantum) one-way functions, and thus also from any commitment, are in fact sufficient.

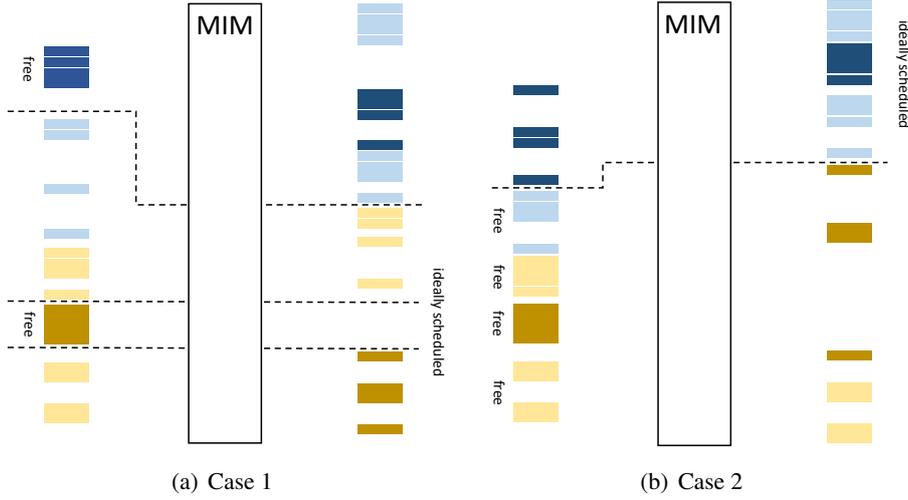


Figure 2: Examples of an ideally scheduled block of shares (on the right). The first block of share commitments is colored in (light/dark) blue and the second in (light/dark) yellow. We mark the commitments on the left that are free with respect to the ideally scheduled block.

used in the first step. Since we are only able to extract from one of the two block-commitments on the right, it is important that both commit to the same value \tilde{v} , and thus our reduction would work, regardless of which one of the two it is able to extract from. Before we explain how to enforce this using ZK in Step 3, we explain how the existence of an ideally scheduled block is established.

Guaranteeing an Ideally Scheduled Block-Commitment. We prove that by setting the parameters n, \bar{n} appropriately, an ideally scheduled block of shares always exists. For this purpose we generalize the combinatorial argument from before. Concretely, we set n, \bar{n} to guarantee that:

1. *Either*, the number of shares $n = n(\text{tg})$ in the first left block-commitment is larger than the total number of messages $k \cdot n(\tilde{\text{tg}})$ in the first right block-commitment,
2. *Or*, the number of shares $\bar{n} = \bar{n}(\text{tg})$ in the second left block-commitment is larger than the total number of messages $k \cdot \bar{n}(\tilde{\text{tg}})$ in the second right block-commitment.

In addition, we require that n, \bar{n} are both at least 2. These conditions can be satisfied for example by setting $n = (k + 1)^{\text{tg}}, \bar{n} = (k + 1)^{\tau - \text{tg}} + 1$, where τ is the total number of tags (namely, $\text{tg} \in [\tau]$).

To see why the above is sufficient, let us assume for instance that Condition 2 of the two above conditions holds (at this point, both are treated symmetrically). We consider two cases:

- **Case 1 (depicted in Figure 2(a))** : the commitment to share u_1 (i.e., the first share of the first block-commitment) on the left ends before the second block-commitment starts on the right. In this case, the commitment to u_1 on the left is free with respect to the second block-commitment on the right. Furthermore, since Condition 2 holds, (by the argument in Step 1,) there also exists a commitment to a share \bar{u}_i (in the second block-commitment) on the left that is also free with respect to the second block-commitment on the right. Accordingly, the second block-commitment on the right is ideally scheduled.
- **Case 2 (depicted in Figure 2(b))** : the commitment to share u_1 on the left ends after the second block of share commitments starts on the right. In this case, the commitments to shares $u_2, \dots, u_n, \bar{u}_1, \dots, \bar{u}_{\bar{n}}$ on the left are all free with respect to the first block-commitment on the right, and thus it is ideally scheduled. (We use the fact that $n \geq 2$, to deduce that a free share u_2 indeed exists.)

Step 3: Use ZK to Ensure Consistency of Right Block-Commitments. Recall that in the last step, we made the simplifying assumption that the MIM adversary always commits to the same value \tilde{v} in the two right block-commitments. The expected approach to removing this assumption, would be to require that the sender gives a (post-quantum) zero-knowledge argument that such consistency indeed holds.

While the soundness of the argument guarantees the required consistency on the right, the addition of a zero knowledge proof brings about new challenges in the reduction of non-malleability to hiding on the left, due to non-synchronizing adversaries. Indeed, in the proof of non-malleability, before using the hiding of the extractable commitments on the left, we must use the zero knowledge property on the left to argue that the proof does not compromise the hidden shares. The problem is that the zero-knowledge argument on the left might interleave with our ideally scheduled block-commitment on the right, and thus with our extraction procedure. For instance, if the extractor wants to rewind the MIM, it might have to rewind the zero knowledge prover on the left, which is not possible. More generally, there could be a circular dependency: The zero-knowledge simulation needs to be applied to the verifier's code which depends on the extractor's code; however, extraction needs to be applied to the sender's code which depends on the simulator's code.

To circumvent this difficulty, we would like to guarantee that an ideally scheduled block-commitment would also be free of the zero knowledge messages on the left, namely, during its execution, no zero knowledge messages should be sent in the left execution (see Figure 3). Indeed, if this is the case, then we can apply the zero knowledge simulator to the verifier that when needed runs the extractor on the right *in its head*. Note that since the the right block-commitment is free from zero knowledge messages on the left, the code of the extractor, and induced verifier, is independent of the simulator's code, breaking the circularity.

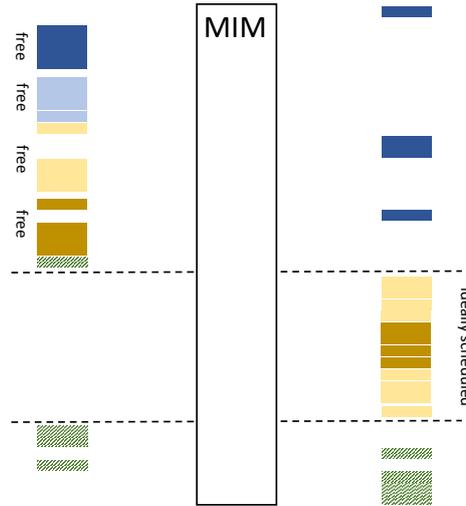


Figure 3: The zero knowledge argument on the left is colored in green. The ideally scheduled block of shares on the right is required to be free of any zero knowledge messages (as well as satisfy the same conditions as before).

Guaranteeing (the Stronger Form of) Ideal Scheduling. To achieve the stronger form of ideal scheduling, we augment the protocol yet again. Specifically, we repeat sequentially for $\ell + 1$ times the second block-commitment to shares $\bar{u}_1, \dots, \bar{u}_{\bar{n}}$, where ℓ is the number of rounds in the zero knowledge protocol. We now require that there is a block-commitment I among the $\ell + 2$ right block-commitments (one of u_1, \dots, u_n , and $\ell + 1$ of $\bar{u}_1, \dots, \bar{u}_{\bar{n}}$) that is ideally scheduled in the following stronger sense:

1. There exist shares u_i and $\bar{u}_{\bar{i}}$ such that *all* commitments to these shares (one to u_i and $\ell + 1$ ones to $\bar{u}_{\bar{i}}$) on the left, are free of the I 'th right block-commitment.
2. The I 'th right block-commitment is free of the zero knowledge argument on the left.

We provide a more involved combinatorial argument (and choice of parameters n, \bar{n}) showing that an ideally scheduled right block-commitment I always exists. Concretely, we set n, \bar{n} to guarantee that:

1. *Either*, the number of shares $n = n(\text{tg})$ in the first left block-commitment as well as the number of shares $\bar{n} = \bar{n}(\text{tg})$ in each of the left block-commitments $2, \dots, \ell + 2$ are *both* larger than the total number of messages $k \cdot n(\tilde{\text{tg}})$ in the first right block-commitment.
2. *Or*, the number of shares $\bar{n} = \bar{n}(\text{tg})$ in each of the left block-commitments $2, \dots, \ell + 2$ is larger than the total number of messages $k \cdot \bar{n}(\tilde{\text{tg}})$ in each of the right block-commitments $2, \dots, \ell + 2$.

Again, we also require that n, \bar{n} are both at least 2. The above conditions can be satisfied for example by setting $n = (k + 1)^{\text{tg}}, \bar{n} = (k + 1)^{2\tau - \text{tg}} + 1$, where τ is the total number of tags. The above two conditions can no longer be treated symmetrically as before. We explain separately, how each one of them implies the existence of an ideally scheduled block on the right (in the stronger sense defined above).

- **Case 1 (applies for either one of the two conditions):** the first block-commitment on the right ends after the knowledge argument on the left had started. In this case, block commitments $2, \dots, \ell + 2$ on the right do not interleave with any of the block commitments on the left. Thus, we only need to establish that one of them does not interleave with the zero knowledge argument on the left. This follows from the fact that there are $\ell + 1$ of them, but only ℓ messages in the zero knowledge argument.
- **Case 2:** Condition 1 holds, but Case 1 above does not hold. First, since Case 1 does not hold, the first right block commitment does not interleave the zero knowledge argument on the left (which only starts after this block commitments ends). Accordingly, it is left to establish that there exist share commitments u_i in left block commitment 1 and \bar{u}_i in each of the left block commitments $2, \dots, \ell + 2$ that are free with respect to the first right block commitment. This is where we use Condition 1 — since the number of messages in this right block is strictly smaller than the number of shares n, \bar{n} in each left block, the required free share commitments are guaranteed to exist.
- **Case 3 (applies for either one of the two conditions):** the commitment to share u_1 on the left ends after the second block of share commitments starts on the right. In this case, the commitments to shares $u_2, \dots, u_n, \bar{u}_1, \dots, \bar{u}_{\bar{n}}$, as well as the zero knowledge argument on the left are all free with respect to the first block-commitment on the right, and thus it is ideally scheduled. (This case is similar to the simplified case depicted in Figure 2(a).)
- **Case 4:** Condition 2 holds, but Case 3 above does not hold. First, since Case 3 does not hold, all the right block commitments $2, \dots, \ell + 2$ do not interleave with the commitment to share u_1 in the first left block commitment. Furthermore, one of these right blocks $\text{blk} \in \{2, \dots, \ell + 2\}$ does not interleave with the zero knowledge argument on the left (which consists of ℓ messages). To deduce that blk is ideally schedule, it is left to show that there is a free share \bar{u}_i in each of the left blocks $2, \dots, \ell + 2$. Here we invoke Condition 2 — the number of messages in blk is strictly smaller than the number of shares \bar{n} in each of the left blocks $2, \dots, \ell + 2$, the required free share commitments are again guaranteed to exist.

2.3 Tag Amplification

We now briefly overview the tag amplification process, which takes a non-malleable commitment $\langle \text{Sen}, \text{Rec} \rangle$ for $t \in [3, O(\log \lambda)]$ bit tags and transforms it into $\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle$ for $T = 2^{t-1}$ bit tags. The amplification procedure is an adaptation of existing procedures from the literature mostly similar to [KS17, ABG⁺20] which in turn is based on that of [DDN03]; however, unlike the first of the two, it

relies on polynomial hardness assumptions, and avoids complexity leveraging, and unlike the second, it works against non-synchronizing adversaries and not only synchronizing ones.

The basic way that previous amplification schemes work is as follows: to commit to a value v , under a tag $\hat{\text{tg}} \in \{0, 1\}^T$ for $T = 2^{t-1}$, we consider $t-1$ tags of the form $\text{tg}_i = (i, \hat{\text{tg}}[i]) \in \{0, 1\}^t$ corresponding to the base scheme (here $\hat{\text{tg}}[i]$ is the i -th bit of $\hat{\text{tg}}$). The committer then sends $t-1$ commitments to the value v in parallel under each one of the tags tg_i , using the base protocol $\langle \text{Sen}, \text{Rec} \rangle$. Finally, a proof that all $t-1$ commitments are consistent is added.

The basic idea behind the transformation is that if all the commitments are consistent, then in order to maul a commitment to value v under tag $\hat{\text{tg}}$ to a commitment to a related value \tilde{v} under tag $\hat{\text{tg}}' \neq \hat{\text{tg}}$, the MIM must create a commitment to \tilde{v} using the base protocol under tag $\text{tg}'_i = (i, \hat{\text{tg}}'[i])$ for every $i \in [t-1]$, by potentially mauling from some of the left commitments to v under tags $\{\text{tg}_i = (i, \hat{\text{tg}}[i])\}_{i \in [t-1]}$. However, the fact that $\hat{\text{tg}} \neq \hat{\text{tg}}'$ means that they differ on at least one bit, that is, $\hat{\text{tg}}[j] \neq \hat{\text{tg}}'[j]$ for some j . Thus, tag $\text{tg}'_j = (j, \hat{\text{tg}}'[j])$ on the right is different from all the tags $\{\text{tg}_i = (i, \hat{\text{tg}}[i])\}_{j \in [t-1]}$ on the left. By the non-malleability of the base protocol, the value committed to under tag $\text{tg}'_j = (j, \hat{\text{tg}}'[j])$ on the right must be independent of the value v committed under tags $\{\text{tg}_i = (i, \hat{\text{tg}}[i])\}$ on the left. Given additionally that the values committed to in all base commitments on the right are the same, the non-malleability of $\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle$ with respect to $\hat{\text{tg}}, \hat{\text{tg}}'$ then follows.

In the setting of synchronizing MIM adversaries the above intuition can be formalized as expected, when the proof of consistency is instantiated with a zero-knowledge argument. In the more general setting of non-synchronizing adversaries, things become more subtle. Specifically, if the zero knowledge argument on the left interleaves with the non-malleable commitments on the right, then it is not clear how to leverage the non-malleability of the base protocol $\langle \text{Sen}, \text{Rec} \rangle$. (More specifically, we need to apply zero-knowledge simulation on the code of the verifier, which however might depend on the honest receiver Rec 's code. Then, we can no longer reduce to the non-malleability of the base protocol.)

To overcome this difficulty, we rely on the Feige-Lapidot-Shamir trapdoor paradigm [FLS99]. The first receiver message in our protocol sets up a trapdoor (a solution to a hard problem), and the final proof of consistency is a witness indistinguishable (WI) proof that either: (1) the $t-1$ commitments are consistent, or (2) the sender “knows” the trapdoor (where formally knowledge is enforced using an extractable commitment). The idea behind the FLS paradigm is that the trapdoor cannot be obtained by a sender running the protocol, and thus the validity of assertion (1) is guaranteed on right. In contrast, we would like to ensure that the reduction of non-malleability to hiding on the left would be able to obtain the trapdoor and use it in order to simulate the WI proof.

We can show that the reduction can indeed do this, but only provided certain scheduling conditions. Specifically, the trapdoor on the left should be set up before the non-malleable commitment on the right occurs. In this case, we can non-uniformly obtain the witness. To deal with the other case, we augment the protocol yet again, adding a plain non-interactive commitment to the committed value v between the trapdoor set up phase and the non-malleable commitment phase. In case the non-malleable commitment on the right starts before the trapdoor set up on the left, then in particular the plain commitment on the right occurs before any commitment was made on the left. In this case, we have a direct reduction from non-malleability to hiding, which non-uniformly obtains the value of the plain commitment on the right (this is akin to our earlier use of “first-message binding”). We refer the reader to Figure 5 for the amplification scheme and Section 6 for the proof.

Robustness. One challenge in the proof above is that even in the case that we can obtain the trapdoor witness on the left, it is not immediate that non-malleability holds when the commitments on the right interleave with the proof. For this, we require that the base non-malleable commitment satisfies an extra property known as r -robustness [LP12]. This property essentially says that the committed value on the right can be extracted without rewinding an arbitrary r -message protocol (the WI proof in our case) executed concurrently. This allows to switch the witness used in the WI on the left, and argue that the

right committed value stays the same after the switch.

We show that our base protocol (described in Section 2.2) is indeed robust for an appropriate choice of parameters. We further show that the tag amplification transformation described here, preserves r -robustness. (See more details in Section 4.)

Two-Sided Extraction via Watrous' Rewinding Lemma. One challenge in our analysis of both the base protocol and the tag amplification procedure is that the adversary's scheduling of messages is *adaptive*. In particular, even though the protocol's design guarantees that executions always contain certain *extraction opportunities*, we do not know ahead of time when they will occur. This is not a problem in the classical setting, where one can typically run first the so called main thread to identify the extraction opportunities and then rewind back to extract. However, such rewinding in the quantum setting might disturb the adversary's state.

The analysis of our base scheme circumvents this difficulty by showing a reduction to adversaries that commit ahead of time to the timing of the so called extraction opportunities. This reduction strongly relies on the fact that the definition of non-malleability is an indistinguishability-based definition. In contrast, r -robustness is a simulation based definition — it requires a simulator that given the code of the MIM adversary can extract on the right, while interacting with an r -message protocol on the left. Let us briefly explain the difficulty in this setting.

To achieve r -robustness, we make sure there are more than r extraction opportunities on the right. Consider a simplified scenario where the MIM gives $r + 1$ extractable commitments, and we want to extract from the "free" extractable commitment that does not interleave with any of the r left messages — we refer to this as *non-interleaving extraction*. The difficulty is that the simulator does not know which extractable commitment would be "free". If the simulator starts an extractable commitment without applying the extractor, it might miss the sole extraction opportunity. On the other hand, if it always applies the extractor, extraction may halt when the adversary expects a message on the left, and the simulator should give up extraction but still faithfully simulate the left and right interactions from here. To resolve this conundrum, we need the extractor of an extractable commitment protocol to be able to interchangeably simulate two types of interactions, ones that will eventually constitute an extraction opportunity and ones that will turn out not to be extractable due to the adversary's scheduling.

Toward this, we prove a *two-sided simulation lemma* for extractable commitments. This lemma shows that we can always enhance the extractor so that in case the sender in the commitment prematurely aborts, not only can we simulate the sender's state at that point, *but also the state of the receiver* (in case of abort, extraction is not required); otherwise, the extractor simulates the sender's state and extracts the committed value as usual (without simulating the state of the receiver). Using this two-sided extractor we can deal with cases where a commitment on the right turns out not to be extractable due to scheduled messages on the left by viewing this event as a premature abort, and then using the simulated state of the receiver to faithfully continue the interaction (without extracting).

The proof of the lemma is inspired by [BS20] and uses the fact that up to the point of abort a real execution and an execution simulated by the extractor are indistinguishable. Our two-sided extractor first tosses a random coin to decide whether to simulate with extraction or to honestly simulate the receiver anticipating an abort; if the guess failed, it tries again (the expected number of trials is negligibly close to two). While this works smoothly in the classical setting, in the quantum setting it should be done with care, as rewinding without state disturbance is typically a problem. In this specific setting, however, we meet the conditions of Watrous' *quantum rewinding lemma* [Wat09] — our extractor is guaranteed to succeed with probability close to $1/2$, obliviously of the quantum internal state of the adversarial sender.

Given the general difficulty of dealing with extraction and rewinding in the quantum setting, our techniques, two-sided extraction, non-interleaving extraction, and robust extraction, may be helpful for future works.

3 Preliminaries

We rely on standard notions of classical Turing machines and Boolean circuits:

- A PPT algorithm is a probabilistic polynomial-time Turing machine.
- For a PPT algorithm M , we denote by $M(x; r)$ the output of M on input x and random coins r . For such an algorithm and any input x , we write $m \in M(x)$ to denote the fact that m is in the support of $M(x; \cdot)$.

We follow standard notions from quantum computation.

- A QPT algorithm is a quantum polynomial-time Turing machine.
- An interactive algorithm M , in a two-party setting, has input divided into two registers and output divided into two registers. For the input, one register I_m is for an input message from the other party, and a second register I_a is an auxiliary input that acts as an inner state of the party. For the output, one register O_m is for a message to be sent to the other party, and another register O_a is again for auxiliary output that acts again as an inner state. For a quantum interactive algorithm M , both input and output registers are quantum.

The Adversarial Model. Throughout, efficient adversaries are modeled as quantum circuits with non-uniform quantum advice (i.e. quantum auxiliary input). Formally, a *polynomial-size adversary* $\mathcal{A} = \{\mathcal{A}_\lambda, \rho_\lambda\}_{\lambda \in \mathbb{N}}$, consists of a polynomial-size non-uniform sequence of quantum circuits $\{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$, and a sequence of polynomial-size mixed quantum states $\{\rho_\lambda\}_{\lambda \in \mathbb{N}}$.

For an interactive quantum adversary in a classical protocol, it can be assumed without loss of generality that its output message register is always measured in the computational basis at the end of computation. This assumption is indeed without the loss of generality, because whenever a quantum state is sent through a classical channel then qubits decohere and are effectively measured in the computational basis.

Indistinguishability in the Quantum Setting.

- Let $f : \mathbb{N} \rightarrow [0, 1]$ be a function.
 - f is negligible if for every constant $c \in \mathbb{N}$ there exists $N \in \mathbb{N}$ such that for all $n > N$, $f(n) < n^{-c}$.
 - f is noticeable if there exists $c \in \mathbb{N}, N \in \mathbb{N}$ such that for every $n \geq N$, $f(n) \geq n^{-c}$.
 - f is overwhelming if it is of the form $1 - \mu(n)$, for a negligible function μ .
- We may consider random variables over bit strings or over quantum states. This will be clear from the context.
- For two random variables X and Y supported on quantum states, quantum distinguisher circuit D with, quantum auxiliary input ρ , and $\mu \in [0, 1]$, we write $X \approx_{D, \rho, \mu} Y$ if

$$|\Pr[D(X; \rho) = 1] - \Pr[D(Y; \rho) = 1]| \leq \mu.$$

- Two ensembles of random variables $\mathcal{X} = \{X_i\}_{\lambda \in \mathbb{N}, i \in I_\lambda}$, $\mathcal{Y} = \{Y_i\}_{\lambda \in \mathbb{N}, i \in I_\lambda}$ over the same set of indices $I = \cup_{\lambda \in \mathbb{N}} I_\lambda$ are said to be *computationally indistinguishable*, denoted by $\mathcal{X} \approx_c \mathcal{Y}$, if for every polynomial-size quantum distinguisher $D = \{D_\lambda, \rho_\lambda\}_{\lambda \in \mathbb{N}}$ there exists a negligible function $\mu(\cdot)$ such that for all $\lambda \in \mathbb{N}, i \in I_\lambda$,

$$X_i \approx_{D_\lambda, \rho_\lambda, \mu(\lambda)} Y_i .$$

- The trace distance between two distributions X, Y supported over quantum states, denoted $\text{TD}(X, Y)$, is a generalization of statistical distance to the quantum setting and represents the maximal distinguishing advantage between two distributions supported over quantum states, by unbounded quantum algorithms. We thus say that ensembles $\mathcal{X} = \{X_i\}_{\lambda \in \mathbb{N}, i \in I_\lambda}, \mathcal{Y} = \{Y_i\}_{\lambda \in \mathbb{N}, i \in I_\lambda}$, supported over quantum states, are statistically indistinguishable (and write $\mathcal{X} \approx_s \mathcal{Y}$), if there exists a negligible function $\mu(\cdot)$ such that for all $\lambda \in \mathbb{N}, i \in I_\lambda$,

$$\text{TD}(X_i, Y_i) \leq \mu(\lambda) .$$

In what follows, we introduce the cryptographic tools used in this work. By default, all algorithms are classical and efficient unless stated otherwise, and security holds against polynomial-size non-uniform quantum adversaries with quantum advice.

3.1 Interactive Protocols, Witness Indistinguishability, and Zero Knowledge

We define proof and argument systems that are secure against quantum adversaries. In what follows, we denote by (P, V) a protocol between two parties P and V . For common input x , we denote by $\text{OUT}_V\langle P, V \rangle(x)$ the output of V in the protocol. For honest verifiers, this output will be a single bit indicating acceptance or rejection of the proof. Malicious quantum verifiers may have arbitrary quantum output (which is formally captured by the verifier outputting its inner quantum state).

Definition 3.1 (Classical Proof and Argument Systems for NP). *Let (P, V) be a protocol with an honest PPT prover P and an honest PPT verifier V for a language $\mathcal{L} \in \text{NP}$, satisfying:*

1. **Perfect Completeness:** *For any $\lambda \in \mathbb{N}, x \in \mathcal{L} \cap \{0, 1\}^\lambda, w \in \mathcal{R}_\mathcal{L}(x)$,*

$$\Pr[\text{OUT}_V\langle P(w), V \rangle(x) = 1] = 1 .$$

2. **Soundness:** *The protocol satisfies one of the following.*

- **Computational Soundness:** *For any quantum polynomial-size prover $P^* = \{P_\lambda^*, \rho_\lambda\}_{\lambda \in \mathbb{N}}$, there exists a negligible function $\mu(\cdot)$ such that for any security parameter $\lambda \in \mathbb{N}$ and any $x \in \{0, 1\}^\lambda \setminus \mathcal{L}$,*

$$\Pr[\text{OUT}_V\langle P_\lambda^*(\rho_\lambda), V \rangle(x) = 1] \leq \mu(\lambda) .$$

A protocol with computational soundness is called an argument.

- **Statistical Soundness:** *There exists a negligible function $\mu(\cdot)$, such that for any (unbounded) prover P^* , any security parameter $\lambda \in \mathbb{N}$, and any $x \in \{0, 1\}^\lambda \setminus \mathcal{L}$,*

$$\Pr[\text{OUT}_V\langle P^*, V \rangle(x) = 1] \leq \mu(\lambda) .$$

A protocol with statistical soundness is called a proof.

3.1.1 Witness Indistinguishability

We rely on classical constant-round (public-coin) proof systems for NP that are witness-indistinguishable; that is, proofs that use different witnesses (for the same statement) are computationally indistinguishable for quantum attackers.

Definition 3.2 (WI Proof System for NP). A classical protocol proof system (P, V) for a language $\mathcal{L} \in NP$ (as in Definition 3.1) is witness-indistinguishable if it satisfies:

Witness Indistinguishability: For every quantum polynomial-size verifier $V^* = \{V_\lambda^*, \rho_\lambda\}_\lambda$,

$$\{\text{OUT}_{V_\lambda^*} \langle P(w_0), V_\lambda^*(\rho_\lambda) \rangle(x)\}_{\lambda, x, w_0, w_1} \approx_c \{\text{OUT}_{V_\lambda^*} \langle P(w_1), V_\lambda^*(\rho_\lambda) \rangle(x)\}_{\lambda, x, w_0, w_1},$$

where $\lambda \in \mathbb{N}$, $x \in \mathcal{L} \cap \{0, 1\}^\lambda$, and $w_0, w_1 \in \mathcal{R}_\mathcal{L}(x)$ are witnesses for x .

Instantiations. 3-message, public-coin classical proof systems with WI follow from classical zero-knowledge proof systems such as the parallel repetition of the 3-coloring protocol [GMW91], which is in turn based on non-interactive perfectly-binding commitments. For the proof system to be WI against quantum attacks, we need the non-interactive commitments to be computationally hiding against quantum adversaries, which can be instantiated for example from QLWE.

3.1.2 Quantum Zero-Knowledge Protocols

We next define post-quantum zero-knowledge classical protocols for NP.

Definition 3.3 (Post-Quantum Zero-Knowledge Classical Protocol). Let (P, V) be a classical protocol (argument or proof) for a language $\mathcal{L} \in NP$ as in Definition 3.1. The protocol is quantum zero-knowledge if it satisfies:

Quantum Zero Knowledge: There exists a quantum polynomial-time simulator Sim , such that for any quantum polynomial-size verifier $V^* = \{V_\lambda^*, \rho_\lambda\}_{\lambda \in \mathbb{N}}$,

$$\{\text{OUT}_{V_\lambda^*} \langle P(w), V_\lambda^*(\rho_\lambda) \rangle(x)\}_{\lambda, x, w} \approx_c \{\text{Sim}(x, V_\lambda^*, \rho_\lambda)\}_{\lambda, x, w},$$

where $\lambda \in \mathbb{N}$, $x \in \mathcal{L} \cap \{0, 1\}^\lambda$, $w \in \mathcal{R}_\mathcal{L}(x)$.

Instantiations. [Wat09] showed how to construct a classical ZK proof system for NP with a polynomial number of rounds. [BS20] show how to construct a classical ZK argument system for NP with a constant number of rounds.

3.2 Commitments

Roughly speaking, A commitment scheme enables a party, called the committer, to commit itself to a value to another party, the receiver. At first the value is hidden from the receiver; this property is called *hiding*. At a later stage when the commitment is opened, it can only reveal a single value as determined in the committing phase; this property is called *binding*. First we define the structure of a commitment scheme.

Definition 3.4 (Commitment Schemes). A commitment scheme is an interactive protocol $\langle \text{Sen}, \text{Rec} \rangle$ with the following properties:

1. Both the committer Sen and the receiver Rec are PPT machines.
2. The commitment scheme has two stages: a commit stage and a reveal stage. In both stages, Sen and Rec receive a security parameter 1^λ as common input. Sen additionally receives a private input $v \in \{0, 1\}^{\text{poly}(\lambda)}$ that is the string to be committed.
3. The commit stage results in a joint output c , called the commitment, a private output for Sen , d , called the decommitment string and a private output for Rec , b .

Without loss of generality, c can be the full transcript of the interaction between Sen and Rec . b is the bit verdict of Rec indicating whether it accepts the commitment c or not; a commitment is accepting if Rec accepts (i.e., $b = 1$) at the end of the commit stage.

4. In the reveal stage, committer Sen sends the pair (v, d) to the receiver Rec, and decides to accept or reject the decommitment (c, v, d) deterministically. We consider the restricted case where the receiver Rec does not keep state from the commit stage to the reveal stage. Then, let $\text{open}_{\langle \text{Sen}, \text{Rec} \rangle}$ denote the Boolean function that corresponds to the verdict of the receiver Rec, that is, $\text{open}(c, v, d) = 1$, if and only if Rec accepts the decommitment (c, v, d) .

A commitment c is valid if and only if there exists a pair (v, d) such that $\text{open}(c, v, d) = 1$.

If Sen and Rec do not deviate from the protocol, then Rec should accept (with probability 1) during the reveal stage.

Next we define the binding and hiding property of a commitment scheme.

Definition 3.5 (Binding). A commitment scheme $\langle \text{Sen}, \text{Rec} \rangle$ is statistically (resp. computationally) binding if for every (resp. quantum polynomial-size) malicious committer Sen^* , there exists a negligible function ν such that Sen^* succeeds in the following game with probability at most $\nu(\lambda)$:

On security parameter 1^λ , Sen^* first interacts with Rec in the commit stage to produce commitment c . Then Sen^* outputs two decommitments (c, v_0, d_0) and (c, v_1, d_1) , and succeeds if $v_1 \neq v_2$ and Rec accepts both decommitments.

The commitment scheme is perfectly binding if no machine Sen^* can ever succeed at the above game.

Definition 3.6 (Hiding). A commitment scheme $\langle \text{Sen}, \text{Rec} \rangle$ is computationally hiding if for every quantum polynomial-size receiver $\text{Rec}^* = \{\text{Rec}_\lambda^*, \rho_\lambda\}_{\lambda \in \mathbb{N}}$ and polynomial $\ell(\cdot)$,

$$\begin{aligned} & \{\text{OUT}_{\text{Rec}_\lambda^*} \langle \text{Sen}(v_0), \text{Rec}_\lambda^*(\rho_\lambda) \rangle (1^\lambda) \}_{\lambda, v_0, v_1} \\ & \approx_c \{\text{OUT}_{\text{Rec}_\lambda^*} \langle \text{Sen}(v_1), \text{Rec}_\lambda^*(\rho_\lambda) \rangle (1^\lambda) \}_{\lambda, v_0, v_1} , \end{aligned}$$

where $\lambda \in \mathbb{N}$, $v_0, v_1 \in \{0, 1\}^{\ell(\lambda)}$.

In the sequel of the paper, a commitment scheme always refers to a statistically-binding commitment.

3.3 Quantum Rewinding Lemma

We use Lemma 9 from [Wat09], which constructs a quantum algorithm for amplifying the success probability of quantum sampler circuits under some conditions. The exact version of the lemma is taken verbatim from [BS20].

Lemma 3.1 (Lemma 9, [Wat09]). *There is a quantum algorithm R that gets as input:*

- A general quantum circuit Q with n input qubits that outputs a classical bit b and an additional m output qubits.
- An n -qubit state $|\psi\rangle$.
- A number $t \in \mathbb{N}$.

R executes in time $t \cdot \text{poly}(|Q|)$ and outputs a distribution over m -qubit states $D_\psi := R(Q, |\psi\rangle, t)$ with the following guarantees.

For an n -qubit state $|\psi\rangle$, denote by Q_ψ the conditional distribution of the output distribution $Q(|\psi\rangle)$, conditioned on $b = 0$, and denote by $p(\psi)$ the probability that $b = 0$. If there exist $p_0, q \in (0, 1)$, $\varepsilon \in (0, \frac{1}{2})$ such that:

- Amplification executes for enough time: $t \geq \frac{\log(1/\varepsilon)}{4 \cdot p_0(1-p_0)}$,

- *There is some minimal probability that $b = 0$: For every n -qubit state $|\psi\rangle$, $p_0 \leq p(\psi)$,*
- *$p(\psi)$ is input-independent, up to ε distance: For every n -qubit state $|\psi\rangle$, $|p(\psi) - q| < \varepsilon$, and*
- *q is closer to $\frac{1}{2}$: $p_0(1 - p_0) \leq q(1 - q)$,*

then for every n -qubit state $|\psi\rangle$,

$$\text{TD}(\mathbf{Q}_\psi, D_\psi) \leq 4\sqrt{\varepsilon} \frac{\log(1/\varepsilon)}{p_0(1 - p_0)} .$$

3.4 Non-Malleable Commitments

Standard commitment schemes are defined in 3.2. Let $\langle \text{Sen}, \text{Rec} \rangle$ be a commitment scheme. In an interaction between a malicious sender Sen^* and honest receiver Rec , we say that Sen^* is *non-aborting* if the Rec accepts (i.e., outputs 1) at the end of the commitment stage. Let $\text{open}_{\langle \text{Sen}, \text{Rec} \rangle}(c, v, d)$ be the function for verifying decommitments of $\langle \text{Sen}, \text{Rec} \rangle$. Define the following value function:

$$\text{val}(c) = \begin{cases} v & \text{if } \exists \text{ unique } v \text{ s.t. } \exists d, \text{open}_{\langle \text{Sen}, \text{Rec} \rangle}(c, v, d) = 1 \\ \perp & \text{otherwise} \end{cases}$$

A commitment c is *valid* if $\text{val}(c) \neq \perp$, and otherwise *invalid*.

Tag-based Commitment Scheme. Following [DDN03, PR05b], we consider *tag-based commitment schemes* where, in addition to the security parameter, the sender and the receiver also receive a “tag”—a.k.a. the identity— tg as common input.

We recall the definition of non-malleability from [LPV08b], adapted to quantum polynomial-size man-in-the-middle adversaries.

Let $\langle \text{Sen}, \text{Rec} \rangle$ be a tag-based commitment scheme, and let $\lambda \in \mathbb{N}$ be a security parameter. Consider a man-in-the-middle (MIM) adversary \mathcal{A} that participates in one left and one right interactions simultaneously. In the left interactions the MIM adversary \mathcal{A} , on auxiliary quantum state ρ , interacts with Sen , receiving commitments to value v , using a tag $\text{tg} \in [T]$ of its choice. In the right interactions \mathcal{A} interacts with Rec attempting to commit to a related value \tilde{v} , again using a tag $\tilde{\text{tg}}$ of length t of its choice. If the right commitment is invalid, or $\tilde{\text{tg}} = \text{tg}$, set $\tilde{v}_i = \perp$ —i.e., choosing the same tags in the left and right interactions is considered invalid. Let $\text{mim}_{\langle \text{Sen}, \text{Rec} \rangle}(\mathcal{A}, \rho, v)$ denote a random variable that describes the value \tilde{v} along with the quantum output of $\mathcal{A}(\rho)$ at the end of the interaction where Sen commits to v on the left.

Definition 3.7. A commitment scheme $\langle \text{Sen}, \text{Rec} \rangle$ is said to be non-malleable if for every quantum polynomial-size man-in-the-middle adversary $A = \{A_\lambda, \rho_\lambda\}_{\lambda \in \mathbb{N}}$ and a polynomial $\ell : \mathbb{N} \rightarrow \mathbb{N}$,

$$\{\text{mim}_{\langle \text{Sen}, \text{Rec} \rangle}(A_\lambda, \rho_\lambda, v)\}_{\lambda, v, v'} \approx_c \{\text{mim}_{\langle \text{Sen}, \text{Rec} \rangle}(A_\lambda, \rho_\lambda, v')\}_{\lambda, v, v'} ,$$

where $\lambda \in \mathbb{N}$ is the security parameter and $v, v' \in \{0, 1\}^{\ell(\lambda)}$ are two committed values by the honest sender.

Generally, the distributions in the MIM experiment include a quantum algorithm with a quantum auxiliary state. A standard strengthening of indistinguishability definitions for distributions of the above-mentioned type is to let the distinguisher prepare an entangled register, which is entangled with the register that contains the auxiliary state of the quantum algorithm in the distribution. In our specific case of MIM distributions the stronger definition (defined below) is equivalent, as we prove next.

Definition 3.8 (Stronger Definition of Non-Malleability). A commitment scheme $\langle \text{Sen}, \text{Rec} \rangle$ is said to be non-malleable (with respect to entanglement) if for every quantum polynomial-size man-in-the-middle adversary $A = \{A_\lambda\}_{\lambda \in \mathbb{N}}$ that can obtain a quantum auxiliary state, a polynomial-size quantum state $\sigma = \{\sigma_\lambda\}_{\lambda \in \mathbb{N}}$ of size at least what A obtains, and a polynomial $\ell : \mathbb{N} \rightarrow \mathbb{N}$,

$$\{\text{mim}_{\langle \text{Sen}, \text{Rec} \rangle}(A_\lambda, \sigma_{1,\lambda}, v), \sigma_{2,\lambda}\}_{\lambda, v, v'} \approx_c \{\text{mim}_{\langle \text{Sen}, \text{Rec} \rangle}(A_\lambda, \sigma_{1,\lambda}, v'), \sigma_{2,\lambda}\}_{\lambda, v, v'} ,$$

where $\lambda \in \mathbb{N}$ is the security parameter, $v, v' \in \{0, 1\}^{\ell(\lambda)}$ are two committed values by the honest sender and σ_1 is the first register of the state σ such that it is in the size of the auxiliary state for A and σ_2 is the rest of the state.

Claim 3.1. Any commitment scheme $\langle \text{Sen}, \text{Rec} \rangle$ satisfying security definition 3.7 also satisfy security definition 3.8.

Proof. Assume $\langle \text{Sen}, \text{Rec} \rangle$ is secure with respect to Definition 3.7 and assume toward contradiction that it is not secure with respect to Definition 3.8. Let $A = \{A_\lambda\}_{\lambda \in \mathbb{N}}$ a MIM adversary and let $D = \{D_\lambda, \sigma_\lambda\}$ a distinguisher that distinguishes between,

$$\{\text{mim}_{\langle \text{Sen}, \text{Rec} \rangle}(A_\lambda, \sigma_{1,\lambda}, v), \sigma_{2,\lambda}\}_{\lambda, v, v'} , \quad \{\text{mim}_{\langle \text{Sen}, \text{Rec} \rangle}(A_\lambda, \sigma_{1,\lambda}, v'), \sigma_{2,\lambda}\}_{\lambda, v, v'} ,$$

for some v, v' . Consider A' a new MIM adversary: A' has quantum auxiliary state σ . The MIM execution of A' is to run A with auxiliary state σ_1 , and keep the rest of σ , which we denote by σ_2 , untouched on the side. D can thus distinguish between the distributions

$$\{\text{mim}_{\langle \text{Sen}, \text{Rec} \rangle}(A'_\lambda, \sigma_\lambda, v)\}_{\lambda, v, v'} , \quad \{\text{mim}_{\langle \text{Sen}, \text{Rec} \rangle}(A'_\lambda, \sigma_\lambda, v')\}_{\lambda, v, v'} ,$$

in contradiction to the security of $\langle \text{Sen}, \text{Rec} \rangle$ with respect to Definition 3.7. □

3.5 Committed Value Oracle

Let $\langle \text{Sen}, \text{Rec} \rangle$ be a (possibly tag-based) commitment scheme. A sequential committed-value oracle $\mathcal{O}^\infty[\langle \text{Sen}, \text{Rec} \rangle]$ of $\langle \text{Sen}, \text{Rec} \rangle$ acts as follows in interaction with a sender Sen^* : it interacts with Sen^* in many *sequential* sessions; in each session,

- it participates with Sen^* in the commit phase of $\langle \text{Sen}, \text{Rec} \rangle$ as the honest receiver Rec (using a tag chosen adaptively by Sen^*), obtaining a commitment c , and
- if Sen^* is *non-aborting* in the commit phase and sends request break, it returns $\text{val}(c)$.

The single-session oracle $\mathcal{O}^1[\langle \text{Sen}, \text{Rec} \rangle]$ is similar to \mathcal{O}^∞ , except that it interacts with the adversary in a single session.

Throughout, when the commitment scheme is clear from the context, we write $\mathcal{O}^\infty, \mathcal{O}^1$ for simplicity.

3.6 Extractable Commitments

We define the standard notion of post-quantum extractable commitments along with several enhancements of this notion.

Definition 3.9. Let $\langle \text{ExCom.Sen}, \text{ExCom.Rec} \rangle$ be a (possibly tag-based) commitment scheme and \mathcal{O}^1 its (single-session) committed value oracle. We say that $\langle \text{ExCom.Sen}, \text{ExCom.Rec} \rangle$ is extractable if there exists a QPT simulator Sim^1 , such that, for every quantum polynomial-size sender $\text{Sen}^* = \{\text{Sen}_\lambda^*, \rho_\lambda\}_{\lambda \in \mathbb{N}}$,

$$\left\{ \text{OUT}_{\text{Sen}_\lambda^*} \left(\text{Sen}_\lambda^{*\mathcal{O}^1}(\rho_\lambda) \right) \right\}_{\lambda \in \mathbb{N}} \approx_c \left\{ \text{Sim}^1(\text{Sen}_\lambda^*, \rho_\lambda) \right\}_{\lambda \in \mathbb{N}} .$$

Sequential Extraction. We analogously define sequential extractability.

Definition 3.10. Let $\langle \text{ExCom.Sen}, \text{ExCom.Rec} \rangle$ be a (possibly tag-based) commitment scheme and \mathcal{O}^∞ its sequential committed value oracle. We say that $\langle \text{ExCom.Sen}, \text{ExCom.Rec} \rangle$ is sequentially extractable if there exists a QPT simulator Sim^∞ , such that, for every quantum polynomial-size sender $\text{Sen}^* = \{\text{Sen}_\lambda^*, \rho_\lambda\}_{\lambda \in \mathbb{N}}$,

$$\left\{ \text{OUT}_{\text{Sen}_\lambda^*} \left(\text{Sen}_\lambda^{*\mathcal{O}^\infty}(\rho_\lambda) \right) \right\}_{\lambda \in \mathbb{N}} \approx_c \left\{ \text{Sim}^\infty(\text{Sen}_\lambda^*, \rho_\lambda) \right\}_{\lambda \in \mathbb{N}} .$$

Constructions of post-quantum extractable commitment from the literature (e.g. [HSS15, LN11, BS20]) address the single-session oracle. However, a standard proof shows that sequential extraction follows.

Lemma 3.2. Any extractable commitment is sequentially extractable.

In fact, later on we prove a stronger Lemma 4.1, which implies the above as a corollary. We now proceed to defining r -robustness and first-message binding which are two strengtheners to the above standard definition of extractable commitments.

3.6.1 r -Robustness.

The work of [LP12] introduced the notion of r -robustness w.r.t. committed value oracle, following similar notions of r -robustness introduced in [CLP16, LP09]. We here recall their definition, adapted to working with quantum polynomial-size adversaries. Let $\langle \text{Sen}, \text{Rec} \rangle$ be a (possibly tag-based) commitment scheme. Consider a man-in-the-middle adversary that participates in an *arbitrary* left interaction with a *limited number r of rounds*, while having access to the committed value oracle $\mathcal{O}^\infty[\langle \text{Sen}, \text{Rec} \rangle]$. Roughly speaking, $\langle \text{Sen}, \text{Rec} \rangle$ is r -robust if the output of \mathcal{A} in any r -round interaction, with access to the oracle $\mathcal{O}^\infty[\langle \text{Sen}, \text{Rec} \rangle]$, can be simulated without the oracle. In other words, having access to the oracle does not help the adversary in breaking the security in any r -round protocol much.

Definition 3.11 (r -robust extraction). Let $\langle \text{Sen}, \text{Rec} \rangle$ be a (possibly tag-based) commitment scheme. We say that $\langle \text{Sen}, \text{Rec} \rangle$ is r -robust w.r.t. the committed-value oracle, if there exists a QPT simulator Sim_r , such that, for every QPT adversary $\mathcal{A} = \{A_\lambda, \rho_\lambda\}_{\lambda \in \mathbb{N}}$, the following holds:

- *Simulation:* For every PPT r -round machine B ,

$$\left\{ \text{OUT}_{A_\lambda} \langle B(z, 1^\lambda), A_\lambda^{\mathcal{O}^\infty[\langle \text{Sen}, \text{Rec} \rangle]}(\rho_\lambda) \rangle \right\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*} \approx_c \left\{ \text{OUT}_{\text{Sim}} \langle B(z, 1^\lambda), \text{Sim}_r(A_\lambda, \rho_\lambda) \rangle \right\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*} .$$

3.6.2 First-message binding.

We define an additional property of extractable commitments which will come in handy later in the construction of post-quantum non-malleable commitments. The property, which we call first-message binding, asserts that the first message of the sender determines the committed value. Additionally, if the first message in the extractable commitment protocol is a receiver message, then the extractor simulates it honestly, in particular, independently of the malicious sender's circuit.

Definition 3.12. Let $\langle \text{ExCom.Sen}, \text{ExCom.Rec} \rangle$ be an extractable commitment scheme. We say that the scheme has first-message binding if:

1. With overwhelming probability over the choice of the honest receiver randomness, the first sender message in the protocol fixes the committed value.

2. If the first message in the protocol is a receiver message, in a simulated session, the extractor ExCom.Ext samples this message by invoking the honest receiver (independently of the malicious sender circuit).

We observe that every extractable commitment can easily be turned into an extractable commitment with first-message binding.

Lemma 3.3. *Let $\langle \text{ExCom.Sen}, \text{ExCom.Rec} \rangle$ be an extractable commitment scheme. Then there exists an extractable commitment scheme $\langle \text{Sen}, \text{Rec} \rangle$ with first-message binding. Furthermore, the sequential extractor Sim^∞ for the scheme also satisfies Property 2 in the above definition.*

Proofsketch. To achieve first message binding, we append to the beginning of protocol $\langle \text{ExCom.Sen}, \text{ExCom.Rec} \rangle$ a statistically-binding commitment with (at most) two messages (such as Naor’s commitment [Nao91], which can be based on post-quantum OWFs, which in turn follow from commitments). We then run the extractable commitment protocol $\langle \text{ExCom.Sen}, \text{ExCom.Rec} \rangle$, and end with a zero-knowledge argument of consistency between the committed value in the initial commitment and the extractable commitment.

The first sender message is now statistically binding by the statistical binding of the initial commitment and the soundness of the zero knowledge argument. The corresponding extractor does not simulate nor extracts from the first initial commitment, but only simulates and extracts from the the original extractable commitment. The same applies for a sequential extractor. \square

4 Two-sided Extraction

In this section, we prove a *two-sided extraction lemma* for any extractable commitment. We then use it to prove a *non-interleaving extraction lemma*, which we later rely on.

4.1 Two-sided Extractor

We define the following variant \mathcal{O}_\perp^1 of the committed value oracle \mathcal{O}^1 . Recall that \mathcal{O}^1 participates in a session of the commit phase of $\langle \text{ExCom.Sen}, \text{ExCom.Rec} \rangle$ with Sen^* , acting as the honest receiver ExCom.Rec . If Sen^* is non-aborting in the commit phase and requests break, \mathcal{O}^1 returns the value $\text{val}(c)$ committed in the produced commitment c .

\mathcal{O}_\perp^1 does the same, except that in the case that Sen^* aborts, it sends back the internal state of the honest receiver ExCom.Rec in that session. That is,

$$\mathcal{O}_\perp^1 \text{ returns } \begin{cases} \text{internal state of ExCom.Rec} & \text{if Sen}^* \text{ aborts} \\ \text{val}(c) & \text{if Sen}^* \text{ is non-aborting in } c \text{ and requests break} \\ \text{nothing} & \text{otherwise} \end{cases}$$

We prove that every extractable commitment satisfies such *two-sided extractability*:

Claim 4.1. *Let $\langle \text{ExCom.Sen}, \text{ExCom.Rec} \rangle$ be an extractable commitment scheme and \mathcal{O}_\perp^1 its enhanced committed value oracle. There exists a QPT simulator Sim_\perp^1 , such that, for every quantum polynomial-size sender $\text{Sen}^* = \{\text{Sen}_\lambda^*, \rho_\lambda\}_{\lambda \in \mathbb{N}}$, the following two ensembles are computationally indistinguishable,*

$$\left\{ \text{OUT}_{\text{Sen}_\lambda^*} \left(\text{Sen}_\lambda^{*\mathcal{O}_\perp^1}(\rho_\lambda) \right) \right\}_{\lambda \in \mathbb{N}} \approx_c \left\{ \text{Sim}_\perp^1(\text{Sen}_\lambda^*, \rho_\lambda) \right\}_{\lambda \in \mathbb{N}} .$$

4.2 Proof of Two-sided Extraction

We now prove Claim 4.1. To construct the simulator Sim_{\perp}^1 , we first construct two simulators, Sim_a and Sim_{na} , for aborting and non-aborting executions respectively. More precisely,

Simulating an aborting execution: Consider a modified adversary $\text{Sen}_a^{*\mathcal{O}_{\perp}^1}(\rho)$ that proceeds identically to $\text{Sen}^{*\mathcal{O}_{\perp}^1}(\rho)$, except that after sending a commitment of $\langle \text{ExCom.Sen}, \text{ExCom.Rec} \rangle$ to \mathcal{O}_{\perp}^1 , if the commitment is non-aborting (i.e. the receiver outputs 1 at the end of the interaction) then it terminates and outputs fail. If the commitment is aborting, it proceeds as Sen^* does – receives the state of ExCom.Rec and outputs its arbitrary quantum output $|\phi\rangle$. We construct Sim_a that on input (Sen^*, ρ) simulates the output state of $\text{Sen}_a^{*\mathcal{O}_{\perp}^1}(\rho)$ as follows:

- $\text{Sim}_a(\text{Sen}^*, \rho)$ runs $\text{Sen}^*(\rho)$ with an honest receiver ExCom.Rec in the commit phase; let ϕ_R be the state of ExCom.Rec at the end of the commit phase.
- If the commit phase is non-aborting, it outputs fail.
- Else it returns ϕ_R to Sen^* , and outputs the output state ϕ of Sen^* .

Observe that whenever the commitment is aborting, Sim_a emulates the interaction between $\text{Sen}^*/\text{Sen}_a^*$ and \mathcal{O}_{\perp}^1 perfectly, and whenever the commitment is non-aborting, Sim_a outputs fail as Sen_a^* does. Therefore, Sim_a perfectly simulates the output state of Sen_a^* .

$$\text{Sim}_a(\text{Sen}^*, \rho) \equiv \text{Sen}_a^{*\mathcal{O}_{\perp}^1}(\rho)$$

Simulating the non-aborting execution: Consider a modified adversary $\text{Sen}_{na}^{*\mathcal{O}_{\perp}^1}(\rho)$ that proceeds identically to $\text{Sen}^{*\mathcal{O}_{\perp}^1}(\rho)$, except that after sending a commitment of $\langle \text{ExCom.Sen}, \text{ExCom.Rec} \rangle$ to \mathcal{O}_{\perp}^1 , it aborts and outputs fail if the commitment is aborting. If the commitment is non-aborting, it proceeds as Sen^* does – possibly requests break to obtain the committed value, and outputs ϕ . Since Sen_{na}^* simply aborts if the commitment is its output is identically distributed when it has access to \mathcal{O}_{\perp}^1 or \mathcal{O}^1 , that is, $\text{Sen}_{na}^{*\mathcal{O}_{\perp}^1}(\rho) \equiv \text{Sen}_{na}^{*\mathcal{O}}(\rho)$.

Given that $\langle \text{ExCom.Sen}, \text{ExCom.Rec} \rangle$ is extractable, there is a simulator Sim^1 that on input $(\text{Sen}_{na}^*, \rho)$ simulates the output state of $\text{Sen}_{na}^{*\mathcal{O}^1}(\rho)$, and hence also $\text{Sen}_{na}^{*\mathcal{O}_{\perp}^1}(\rho)$. Therefore, we simply let the simulator $\text{Sim}_{na}(\text{Sen}^*, \rho)$ output $\phi \leftarrow \text{Sim}(\text{Sen}_{na}^*, \rho)$, and have

$$\text{Sim}_{na}(\text{Sen}^*, \rho) \approx \text{Sen}_{na}^{*\mathcal{O}_{\perp}^1}(\rho).$$

Next, we combine the above two simulators into Sim_{comb} :

Combined Simulator: $\text{Sim}_{comb}(\text{Sen}^*, \rho)$ samples $b \leftarrow \{0, 1\}$, and outputs $\phi \leftarrow \text{Sim}_{na}(\text{Sen}^*, \rho)$ if $b = 0$ and $\phi \leftarrow \text{Sim}_a(\text{Sen}^*, \rho)$ if $b = 1$.

It holds that *i*) the probability that Sim_{comb} outputs fail is negligibly close to $1/2$ for every (Sen^*, ρ) , and *ii*) the output state of Sim_{comb} conditioned on not outputting fail is indistinguishable to the output state of $\text{Sen}^{*\mathcal{O}_{\perp}^1}(\rho)$.

Finally, we observe that Sim_{comb} satisfies the required conditions for applying Watrous's quantum rewinding lemma, Lemma 3.1, in order to amplify the success probability from negligibly close to $1/2$ to negligibly close to 1. Formally, let R be the algorithm from Lemma 3.1. Our final simulator is thus:

Final Simulator: $\text{Sim}_{\perp}^1(\text{Sen}^*, \rho)$ outputs $\leftarrow R(\text{Sim}_{comb}, (\text{Sen}^*, \rho), \lambda)$.

More precisely, Sim_{comb} satisfies that for every input (Sen^*, ρ) , the probability that it does not output fail is negligibly close to $1/2$. Thus it satisfies the premise of Lemma 3.1 w.r.t. $q = 1/2$, $p_0 = 1/3$, and $\epsilon = 1/p(\lambda)$ for every polynomial p . Moreover, $t = \lambda \geq \frac{\log 1/\epsilon}{4p_0(1-p_0)}$ for every inverse polynomial ϵ . Therefore, the trace distance between the output state of Sim_\perp^1 and that of Sim_{comb} conditioned on not outputting fail is bounded by every inverse polynomial, and hence Sim_\perp^1 simulates the output state of $\text{Sen}^{*\mathcal{O}_\perp^1}(\rho)$. Finally, Sim_\perp^1 runs in polynomial time $\lambda \text{poly}(|\text{Sim}_{comb}|)$. This concludes the proof of the claim. \square

4.3 Non-Interleaving Extraction

In this section, we leverage two-sided extraction to prove a *non-interleaving extraction lemma*. Roughly speaking, the lemma allows simulating any man-in-the-middle while extracting the committed values of sessions on the right that do not interleave with the interaction on the left.

Specifically, consider a man-in-the-middle adversary A that participates in an *arbitrary* left interaction with an interactive PPT machine B , while acting as a sender to multiple sequential commitments on the right. We say that a right session is *free* if during its execution, no message was exchanged between A and B . Consider a restricted commitment oracle $\mathcal{O}_{free}^\infty$ that behaves like the standard commitment oracle \mathcal{O}^∞ , with the exception that it responds with \perp to any break request in any session that is *not* free. We denote this experiment by $\langle B(z, 1^\lambda), A_\lambda^{\mathcal{O}_{free}^\infty}(\rho_\lambda) \rangle$.

Lemma 4.1. *Let $\langle \text{ExCom.Sen}, \text{ExCom.Rec} \rangle$ be an extractable commitment scheme. There exists a PPT simulator Sim_{free}^∞ , such that, for every quantum polynomial-size adversary $A = \{A_\lambda, \rho_\lambda\}_{\lambda \in \mathbb{N}}$, the following two ensembles are computationally indistinguishable,*

- $\left\{ \text{OUT}_{A_\lambda} \langle B(z, 1^\lambda), A_\lambda^{\mathcal{O}_{free}^\infty}(\rho_\lambda) \rangle \right\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*}$
- $\left\{ \text{OUT}_{\text{Sim}_{free}^\infty} \langle B(z, 1^\lambda), \text{Sim}_{free}^\infty(A_\lambda, \rho_\lambda) \rangle \right\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*}$

Proof. Using the extractor Sim_\perp^1 given by Claim 4.1, we construct a sequential robust extractor Sim_{free}^∞ next. Fix a security parameter λ , auxiliary input z , and adversary A with auxiliary state ρ . The goal of Sim_{free}^∞ is simulating the output state of $A^{\mathcal{O}_{free}^\infty}(\rho)$ in interaction with $B(z, 1^\lambda)$, without oracle $\mathcal{O}_{free}^\infty$. That is,

$$\text{OUT}_A \langle B(z, 1^\lambda), A^{\mathcal{O}_{free}^\infty}(\rho) \rangle \approx \text{OUT}_{\text{Sim}_{free}^\infty} \langle B(z, 1^\lambda), \text{Sim}_{free}^\infty(A, \rho) \rangle$$

Recall that A interacts with $\mathcal{O}_{free}^\infty$ in many sessions. In each session i , A first interacts with $\mathcal{O}_{free}^\infty$ in the commit phase of $\langle \text{ExCom.Sen}, \text{ExCom.Rec} \rangle$, and produce a commitment c_i . After the commit phase, A can request $\mathcal{O}_{free}^\infty$ to return the value committed in c_i , if A was non-aborting and did not interact with B (i.e., session i is free). We denote by ϕ_{i-1} and ϕ_i the quantum states of A at the beginning and end of session i ; at the very beginning of the execution $\phi_0 = \rho$.

We construct $\text{Sim}_{free}^\infty(A, \phi_0 = \rho)$ to simulate the evolution of the state of A , $\{\phi_i\}$, in sequence as follows: In every session i ,

1. Sim_{free}^∞ first constructs a stand-alone sender C_i that on input state ϕ_i and with access to the enhanced oracle \mathcal{O}_\perp^1 , emulates the execution of A in session i till A sends a message to B , or reaches the end of session i . More precisely, C_i runs $A(\phi_i)$ internally and forwards messages in the commitment phase between A and \mathcal{O}_\perp^1 .

Case 1: A communicates with B in session i . Whenever A sends a message q to B and expects a reply, C_i aborts in the commitment phase, and obtains from \mathcal{O}_\perp^1 the current state $\phi_{R,i}$ of the honest receiver ExCom.Rec . It then outputs $(b = 1, \phi'_i, q, \phi_{R,i})$, where ϕ'_i is the state A right after sending the message q .

Case 2: A does not communicate with B in session i . Otherwise, if no message is sent to B till the end of the commit phase and the commit phase is non-aborting, C_i continues to forward A 's break request to \mathcal{O}_\perp^1 and returns the committed value m from \mathcal{O}_\perp^1 to A . It then outputs $(b = 2, \phi'_i, \perp, \perp)$, where ϕ'_i is the state of A at the end of session i .

Observe that $C_i^{\mathcal{O}_\perp^1}(\phi_i)$ emulates the execution with $A(\phi_i)$ in session i perfectly till it sends a message to B or session i ends.

2. $\text{Sim}_{\text{free}}^\infty$ uses the extractor Sim_\perp^1 guaranteed by Claim 4.1 to simulate the output state of $C_i^{\mathcal{O}_\perp^1}(\phi_i)$, that is, $(\tilde{b}, \tilde{\phi}_i, q, \tilde{\phi}_{R,i}) \leftarrow \text{Sim}_\perp^1(C_i, \phi_i)$. It then simulates the state ϕ_{i+1} of A at the end of session i as follows:

Case 1: If $\tilde{b} = 1$, $\tilde{\phi}_i$ is the simulated state of A after sending q to B . $\text{Sim}_{\text{free}}^\infty$ sends q to B and receives a reply a , and completes simulating session i as follows: i) it sends a to $A(\tilde{\phi}_i)$, ii) emulates an honest receiver ExCom.Rec starting with state $\tilde{\phi}_{R,i}$ for A in the commit phase, and iii) since this session is not free, A cannot request for the committed value after the commit phase. $\text{Sim}_{\text{free}}^\infty$ then obtains a simulated state ϕ_{i+1} of A at the end of session i .

Case 2: If $\tilde{b} = 2$, $\tilde{\phi}_i$ is already the simulated state of A at the end of session i . $\text{Sim}_{\text{free}}^\infty$ simply sets $\phi_{i+1} = \tilde{\phi}_i$.

3. After the last iteration I , $\text{Sim}_{\text{free}}^\infty$ outputs the final state ϕ_I .

It follows from Claim 4.1 that Sim_\perp^1 simulates the output state of A at the end of session i for all i in an indistinguishable way. Therefore, by a hybrid argument, the final state ϕ_I is indistinguishable to the real output state of $A^{\mathcal{O}_{\text{free}}^\infty}(\rho)$ interacting with B . Finally, observe that $\text{Sim}_{\text{free}}^\infty$ does not have access to any oracle and interacts with B in a straight-line fashion. The running time of $\text{Sim}_{\text{free}}^\infty$ is polynomial as for every i , $\text{Sim}_\perp^1(C_i, \phi_i)$ runs in polynomial time, and the rest of the steps all have fixed polynomial time proportional to the size/time of A . Therefore, we conclude the lemma. \square

Remark 4.1. We note that Lemma 4.1 in particular implies Lemma 3.2, where there is no interaction on the left (i.e., B is empty).

5 Post-quantum Non-malleable Commitment For Few Tags

In this section, we present our construction of a classical post-quantum non-malleable commitment protocol with at most a logarithmic number of tags τ .

Ingredients and notation:

- A quantumly-extractable classical commitment scheme ($\text{ExCom.Sen}, \text{ExCom.Rec}$) with first-message binding.
- A post-quantum classical zero-knowledge argument (P, V).

We describe the protocol in Figure 4.

Using post-quantum extractable commitments with k rounds one can obtain post-quantum zero-knowledge arguments with $k + O(1)$ rounds [Ros04, BS20]. It follows that the number of rounds in Protocol 4 is $k^{O(\tau)}$. Statistical binding of the commitment scheme follows readily from the statistical binding of the extractable commitment scheme. Hiding of any commitment scheme follows directly from non-malleability, so it remains for us to show that our commitment protocol is non-malleable. Later, we also show that our commitment scheme satisfies r -robustness (Definition 3.6.1), a property of the commitment protocol which we use in our tag amplification scheme in Section 6.

Proposition 5.1. *The protocol in Figure 4 is non malleable.*

Protocol 4

Parameters:

- λ is the security parameter.
- r is the robustness parameter.
- k is the total number of messages in the extractable commitment protocol.
- $\tau \leq O(\log_k(\lambda))$ is the number of tags.
- ℓ is the maximum between (1) the robustness parameter r , and (2) the total number of messages in the zero knowledge argument system.

Common input: Security parameter $\lambda \in \mathbb{N}$, robustness parameter $r \leq \text{poly}(\lambda)$, an identification tag $\text{tg} \in [\tau]$ for the sender.

Sender private input: A value $v \in \{0, 1\}^*$ to commit to.

Phase 1: Commitments to Secret Shares of Value:

- Let $n := (k + 1)^{\text{tg}}$ and $\bar{n} := (k + 1)^{2\tau - \text{tg}}$.
- Sen secret-shares the value v twice, first into n shares and second into \bar{n} shares: $\mathbf{u} = (u_1, \dots, u_n)$ and $\bar{\mathbf{u}} = (\bar{u}_1, \dots, \bar{u}_{\bar{n}})$, respectively.
- Sen provides extractable commitments to the two sequences of shares:
 1. An extractable commitment to u_i , for every $i \in [n]$, sequentially, one after the other.
 2. An extractable commitment to \bar{u}_i , for every $i \in [\bar{n}]$, sequentially, one after the other. This sequential commitment to $\bar{\mathbf{u}}$ is repeated $\ell + 1$ times, sequentially.

Phase 2: Zero-knowledge Argument of Consistency: The protocol ends with Sen giving a ZK argument that its generated transcript is consistent; namely, there exists private input and randomness for the honest sender inducing the transcript.

Decommitment. If the interaction ends in an accepting proof, the decommitment information includes the shares u_1, \dots, u_n along with the decommitment information for each of their corresponding extractable commitments. The decommitment verification algorithm checks that the shares yield the value v and then runs the decommitment verification algorithm of the extractable commitment on each of the shares and its decommitment information. If the ZK argument is not accepting, or the sender prematurely aborts, the verification algorithm rejects, regardless of the decommitment information given.

Figure 4: A τ -tag post-quantum non-malleable commitment (Sen, Rec).

5.1 Ideally-Scheduled Block Commitments

Before turning to prove Proposition 5.1, we state and prove a combinatorial claim regarding the structure of executions. We first fix relevant terminology for addressing different parts of the protocol.

Block Commitments. For $m, N \in \mathbb{N}$, a block commitment of length N and sub-block length m for a string $s = s_1, s_2, \dots, s_N \in \{0, 1\}^{m \times N}$ (such that $\forall i \in [N] \ s_i \in \{0, 1\}^m$) consists of N sequential

extractable commitment to each of the strings s_1, \dots, s_N in their respective order. In particular, note that in Phase 1 of our Protocol 4, the sender gives one block commitment of length n with sub-block length $|v|$ to $\mathbf{u} = (u_1, \dots, u_n)$ and $\ell + 1$ block commitments to $\bar{\mathbf{u}} = (\bar{u}_1, \dots, \bar{u}_{\bar{n}})$, each of length \bar{n} and sub-block length $|v|$.

Ideally Scheduled Block Commitments. Consider a two-sided MIM execution of Protocol 4; that is, the MIM adversary \mathcal{A} interacts with Sen on the left and Rec on the right.

We call an execution of a block commitment on the left *free* on index i with respect to a given block commitment on the right, if interaction during the i -th extractable commitment in that block commitment does not interleave with the interaction during the given right block commitment. We call an execution of a block commitment on the right free if it does not interleave with the interaction during Phase 2 of the protocol on the left.

An execution I of a block commitment on the right is *ideally scheduled* if all of the above hold:

- It is free (with respect to the second phase on the left).
- There is some index $i \in [n]$ such that the block commitment to \mathbf{u} on the left is free on index i with respect to I .
- There is some index $j \in [\bar{n}]$ such that *all* $\ell + 1$ block commitments to $\bar{\mathbf{u}}$ on the left are free on *the same* index j with respect to I .

In case, the MIM adversary aborts, we assume w.l.o.g it keeps sending messages \perp according to some schedule, so that the above notion is always defined.

Claim 5.1. *In every MIM execution of Protocol 4 with tag tg on the left and tag $\tilde{\text{tg}}$ on the right, if $\text{tg} \neq \tilde{\text{tg}}$, there is an ideally scheduled execution of a block commitment on the right.*

Proof. Throughout, we refer to the first block commitment as the *top* block commitment, and to the subsequent $\ell + 1$ block commitments as the *bottom* block commitments. We also refer to the commitments given from Sen to the MIM adversary as the *left* commitments, and accordingly, to the commitments given from the adversary to Rec as the *right* commitments. Recall that for k the number of messages in the extractable commitment protocol we have:

- The block length of the top left block commitment is $n = (k + 1)^{\text{tg}}$.
- The block length of each of the $\ell + 1$ bottom left block commitments is $\bar{n} = (k + 1)^{2\tau - \text{tg}}$.
- The block length of the top right block commitment is $n' = (k + 1)^{\tilde{\text{tg}}}$.
- The block length of each of the $\ell + 1$ bottom right block commitments is $\bar{n}' = (k + 1)^{2\tau - \tilde{\text{tg}}}$.

We first note that since there are at most ℓ messages in the second phase of the protocol and in particular in the left execution of the protocol, and since there are $\ell + 1$ block commitments given for the bottom block on the right, it follows that in every execution of the protocol there is always a free block commitment on the right (with respect to the second phase on the left).

We now consider several cases.

- $\text{tg} > \tilde{\text{tg}}$: We further divide this case into sub-cases.
 - **Top right block commitment ends after beginning of Phase 2 on the left:** It follows that all bottom right block commitments are non-interleaving with all of the block commitments on the left. Now, since one of the bottom right block commitments is always free, it follows that one of the bottom right block commitments is ideally scheduled.

- **Top right block commitment ends before beginning of Phase 2 on the left:** First, note that the top right block commitment is free (with respect to the second phase on the left). Second, the top right block commitment has kn' messages, and accordingly interleaves with at most kn' extractable commitments on the left. Since $\text{tg} > \tilde{\text{tg}}$, the number of shares n in the top block commitment and the number of shares \bar{n} in the $\ell + 1$ bottom block commitments satisfy

$$\bar{n} \geq n > kn' .$$

This means that there exist shares u_i, \bar{u}_j on the left, such that the extractable commitment to u_i , and all $\ell + 1$ extractable commitments to \bar{u}_j , are free (with respect to the top right block commitment).

It follows that the top right block commitment is ideally scheduled.

- $\text{tg} < \tilde{\text{tg}}$: We further divide this case into sub-cases.

- **Top extractable commitment in top left block ends before beginning of bottom right block commitments:** In this case, all $\ell + 1$ bottom right blocks do not interleave with the first extractable commitment from the top left block. Recall that it is always the case that one of the $\ell + 1$ bottom right block commitments is free with respect to Phase 2 on the left - let's denote such right block with blk . Furthermore, blk has $k\bar{n}'$ messages, and since $\text{tg} < \tilde{\text{tg}}$, the number of shares \bar{n} on the left satisfies

$$\bar{n} > k\bar{n}' .$$

By the same counting argument as in the second sub-case above, there exists a share \bar{u}_j on the bottom left such that all $\ell + 1$ extractable commitments to \bar{u}_j are free (with respect to blk).

It follows that blk is ideally scheduled.

- **Top extractable commitment in top left block ends after beginning of bottom right block commitments:** It follows that the top right block can interleave with at most the top extractable commitment from the top left block. The second extractable commitment on the left (note that $n = (k + 1)^{\text{tg}} \geq 2$), all subsequent extractable commitments on the left, as well as Phase 2 messages on the left, do not interleave with the top right block.

It follows that the top right block is ideally scheduled.

□

5.2 Adversaries with Predetermined Ideal Schedule

Before proving Proposition 5.1, we prove a lemma that basically says that we can restrict attention to MIM adversaries that always announce ahead of time the structure of the ideal schedule. This lemma will later simplify our proof of Proposition 5.1.

In what follows, let N be a bound on the size of $n := (k + 1)^{\text{tg}}$, $\bar{n} := (k + 1)^{2\tau - \text{tg}}$, for every possible tg . We consider *configurations* of the form

$$C = (i, c, \bar{c}, w) \in [\ell + 2] \times [N] \times [N] \times \{\text{IP}_2, \text{P}_2\text{I}\} .$$

We say that a given MIM execution is consistent with such a configuration C if:

- The i -th block commitment on the right is the first ideally scheduled block.

- The commitment to u_c (in the first block) on the left is free with respect to the ideally scheduled block i .
- The commitment to $\bar{u}_{\bar{c}}$ in every one of the blocks $2, \dots, \ell + 2$ on the left is free with respect to the ideally scheduled block i .
- If the first ideally scheduled block i on the right ends before Phase 2 on the left begins, $w = \text{IP}_2$. Otherwise (Phase 2 on the left begins before the first ideally scheduled block i has ended), $w = \text{P}_2\text{I}$. Note that in case $w = \text{P}_2\text{I}$, due to the fact that block i on the right is ideally scheduled and in particular is continuous with respect to Phase 2 on the left, we can also say that block i on the right *begins* after Phase 2 on the left has started (rather than say that it only ends after the beginning of Phase 2 on the left).

Note that the number of possible configurations is bounded by

$$\Delta := (\ell + 2) \times N \times N \times 2 = \text{poly}(\lambda) .$$

Definition 5.1 (MIM with predetermined ideal schedule). *A MIM QPT adversary $\mathcal{A} = \{\mathcal{A}_\lambda, \rho_\lambda\}_\lambda$ has a predetermined ideal schedule $C = \{C_\lambda\}_\lambda$, if any execution in which \mathcal{A}_λ participates is consistent with configuration C_λ .*

Lemma 5.1. *If the protocol in Figure 4 is secure against MIM QPT adversaries with predetermined ideal schedule, then it is also secure against arbitrary MIM QPT adversaries.*

Proof. Given an arbitrary MIM QPT \mathcal{A} and QPT distinguisher D that break non-malleability for some values v, v' with advantage ε , we construct an MIM QPT adversary with predetermined schedule, which breaks the scheme with probability ε/Δ .

Consider an adversary \mathcal{A}' that first samples uniformly at random a configuration $C \leftarrow [\ell + 2] \times [N] \times [N] \times \{\text{IP}_2, \text{P}_2\text{I}\}$. It then emulates \mathcal{A} , and if at any point the execution is about to become inconsistent with C , \mathcal{A}' stops emulating \mathcal{A} , completes the execution consistently with C , and eventually outputs \perp . If the emulation of \mathcal{A} is completed consistently with C , \mathcal{A}' outputs the same as \mathcal{A} .

Then, since every execution has an ideally scheduled block (Claim 5.1), \mathcal{A}' breaks non-malleability with probability exactly ε/Δ (with respect to the same distinguisher D and v, v'). Finally, by an averaging argument, we fix the choice of \mathcal{A}' for a configuration to be the configuration C that maximizes D 's distinguishing advantage. We obtain a corresponding MIM with predetermined ideal schedule with the same advantage ε/Δ . \square

5.3 Proof of Proposition 5.1

We prove the Proposition by a hybrid argument, specifically, we show that the MIM experiment output distribution for any value v on the left is indistinguishable from an experiment independent of v . Following Lemma 5.1, we restrict attention to a MIM adversary with a predetermined ideal schedule $C = (i, c, \bar{c}, w)$.

\mathcal{H}_0 : The original MIM experiment output. This includes the output of the MIM adversary in the experiment and the committed value on the right.

\mathcal{H}_1 : Inefficient extraction from ideally-scheduled block. In this hybrid, instead of the committed value \tilde{v} on the right, we consider the value \tilde{v}_1 reconstructed from the shares of the ideally scheduled block i on the right. If the value of any of the commitments to these shares is \perp , we set $\tilde{v}_1 = \perp$.

Claim 5.2. $\mathcal{H}_0 \approx_s \mathcal{H}_1$.

Proof. This indistinguishability follows from the soundness of the ZK argument that \mathcal{A} gives to the receiver on the right in Phase 2. Specifically, given the correctness of the proven statement, the value \tilde{v}_1 reconstructed from the set of shares of the block commitments all yield the same value, and thus $\tilde{v}_1 = \tilde{v}$, where the latter is the value reconstructed from the first block commitment. \square

\mathcal{H}_2 : Alternative description via oracle extraction. In this hybrid we consider an augmented adversary $\mathcal{A}_2^{\mathcal{O}^\infty}$, which is given access to the sequential committed-value oracle $\mathcal{O}^\infty = \mathcal{O}^\infty[\text{ExCom.Sen}, \text{ExCom.Rec}]$ and acts as follows:

- \mathcal{A}_2 emulates \mathcal{A} .
- On the left, \mathcal{A}_2 relays all messages between \mathcal{A} and the sender.
- On the right,
 - During the ideally scheduled block i , \mathcal{A}_2 interacts with its oracle \mathcal{O}^∞ , in every extractable commitment. Recall that \mathcal{O}^∞ acts as the honest receiver, and answers break requests with the corresponding committed value. \mathcal{A}_2 submits such a break request after each of the commitments and stores the received share value.
 - In any other (than i) block in Phase 1, \mathcal{A}_2 internally emulates the the receiver on the right.
 - In Phase 2, \mathcal{A}_2 internally emulates the the zero knowledge verifier on the right.
- Eventually, \mathcal{A}_2 outputs the output of \mathcal{A} as well as the value \tilde{v}_1 reconstructed from the ideal block shares obtained from the oracle \mathcal{O}^∞ .

The output of this hybrid is the output of \mathcal{A}_2 .

Claim 5.3. $\mathcal{H}_1 \equiv \mathcal{H}_2$.

Proof. This follows directly from the construction of $\mathcal{A}_2^{\mathcal{O}^\infty}$ and the definition of \mathcal{O}^∞ . \square

\mathcal{H}_3 : Efficient extraction on the right when $w = \text{P}_2\text{I}$. This hybrid, differs from the previous hybrid only if $w = \text{P}_2\text{I}$; namely, Phase 2 on the left begins before the ideally scheduled block commitment i on the right had started. In such executions, for the ideally scheduled block commitment i , we perform sequential extraction to obtain the corresponding shares.

In more detail, let ψ be the (quantum) state of \mathcal{A}_2 when it initiates the ideally scheduled block i on the right, and let $\bar{\mathcal{A}}_2^{\mathcal{O}^\infty}$ be the adversary that starting from ψ , emulates $\mathcal{A}_2^{\mathcal{O}^\infty}$ during block i and outputs its state at the end (Note that since block i is ideally scheduled and also starts after Phase 2 on the left, it follows that $\bar{\mathcal{A}}_2$ does not perform *any* interaction on the left during the right block i).

In \mathcal{H}_3 , we consider another augmented adversary \mathcal{A}_3 that acts like \mathcal{A}_2 , only that instead of executing $\bar{\mathcal{A}}_2^{\mathcal{O}^\infty}$ during block i , it invokes the sequentially-extracting simulator $\text{Sim}^\infty(\bar{\mathcal{A}}_2, \psi)$, given by Lemma 3.2, which eliminates the use of the commitment oracle \mathcal{O}^∞ .

Claim 5.4. $\mathcal{H}_2 \approx_c \mathcal{H}_3$.

Proof. Computational indistinguishability of \mathcal{H}_2 and \mathcal{H}_3 follows directly from the sequential extraction guarantee (Lemma 3.2). \square

\mathcal{H}_4 : Simulating the ZK argument on the left. In this hybrid, the ZK argument on the left is generated by the zero knowledge simulator.

Specifically, let ψ be the state of \mathcal{A}_3 when the zero-knowledge argument is initiated on the left. We consider the zero-knowledge verifier V^* that starting from ψ emulates \mathcal{A}_3 in the rest of the interaction while forwarding its messages on the left to the zero-knowledge prover, and eventually outputs the same.

In particular, if $w = P_2I$ then the code of V^* includes the code of the simulator Sim^∞ , which is applied to $(\bar{\mathcal{A}}_2, \psi)$ as part of the execution of \mathcal{A}_3 . Note that in both cases $w = IP_2$ and $w = P_2I$, once Phase 2 on the left starts, \mathcal{A}_3 no longer makes oracle calls to \mathcal{O}^∞ , so the code of V^* is fully specified and executes in polynomial time.

In \mathcal{H}_4 , we consider an augmented adversary \mathcal{A}_4 that acts as \mathcal{A}_3 , only that when Phase 2 starts on the left, instead of executing V^* and interacting on the left with the zero knowledge prover, \mathcal{A}_4 runs the zero knowledge simulator $\text{Sim}(V^*, \psi)$, and outputs the same.

Claim 5.5. $\mathcal{H}_3 \approx_c \mathcal{H}_4$.

Proof. First, note that by construction, the output of V^* is identically distributed to the output of \mathcal{H}_3 . Computational indistinguishability of \mathcal{H}_3 and \mathcal{H}_4 now follows from the zero knowledge simulation guarantee (we note that any use of the inefficient oracle \mathcal{O}^∞ , in case $w = IP_2$, occurs before Phase 2 on the left, and can thus be non-uniformly fixed). \square

\mathcal{H}_5 and \mathcal{H}_6 : Interchangeably, changing left committed values and efficient extraction threshold.

As a preliminary high-level explanation to the next step, at this point in our hybrid distributions, we consider the $1 + (\ell + 1)$ block commitments given to the MIM adversary on the left, and in each block, we'll switch a commitment for a secret share (of v), to a commitment for a string of zeros. For this, we will need to use the computational hiding property of the extractable commitments. The point, however, is to be able to use the hiding of the extractable commitments while still being able to *efficiently* extract the value \tilde{v}_1 from the right interaction with the MIM adversary⁵.

Formally, we next define two sequences of hybrids $\mathcal{H}_{5,j}$ and $\mathcal{H}_{6,j}$ (for $j \in [\ell + 3]$) that augment one another interchangeably:

$$\mathcal{H}_4 = \mathcal{H}_{5,\ell+3} \rightarrow \mathcal{H}_{6,\ell+2} \rightarrow \mathcal{H}_{5,\ell+2} \rightarrow \cdots \rightarrow \mathcal{H}_{5,2} \rightarrow \mathcal{H}_{6,1} \rightarrow \mathcal{H}_{5,1} .$$

In what follows, recall that \mathcal{A}_4 in \mathcal{H}_4 is following a predetermined ideal schedule $C = (i, c, \bar{c}, w)$.

$\mathcal{H}_{5,j}$, for $j = \ell + 3, \dots, 1$: Swapping one more free commitment to zeros. In this hybrid, we simulate the most bottom free commitment on the left. Formally:

- $\mathcal{H}_{5,\ell+3}$ is defined as \mathcal{H}_4 .
- For $j \leq \ell + 2$, $\mathcal{H}_{5,j}$ is defined exactly as $\mathcal{H}_{6,j}$, except that the left extractable commitment c_j (to share u_c or $\bar{u}_{\bar{c}}$) in the left block j is replaced with a commitment to $0^{|v|}$.

$\mathcal{H}_{6,j}$, for $j = \ell + 2, \dots, 1$: Raising the threshold for efficient extraction. Recall \mathcal{A}_4 in \mathcal{H}_4 interacts with the sender in Phase 1 on the left and in case $w = IP_2$, namely, the ideally scheduled block on the right ends before Phase 2 on the left begins, \mathcal{A}_4 interacts with the sequential commitment oracle \mathcal{O}^∞ on the right during block i . For a left block index $j \in [\ell + 2]$, we denote by c_j the corresponding free extractable commitment; namely, $c_j = c$ if $j = 1$, and $c_j = \bar{c}$ if $j \geq 2$.

Informally, in hybrid $\mathcal{H}_{6,j}$, we move to simulating the oracle \mathcal{O}^∞ in any right extractable commitment that starts after the free left commitment c_j . Formally, $\mathcal{H}_{6,j}$ is different from $\mathcal{H}_{5,j+1}$ only if $w = IP_2$. In this case, we consider an augmented adversary $\mathcal{A}_{6,j}$ defined as follows for $j \in [\ell + 2]$:

- $\mathcal{A}_{6,j}$ acts as $\mathcal{A}_{6,j+1}^{\mathcal{O}^\infty}$ until the first right extractable commitment (in the ideally scheduled right block i) in which the first sender message is sent after the free left commitment c_j .
- $\mathcal{A}_{6,j}$ simulates the remaining calls to \mathcal{O}^∞ as follows:

⁵Recall that currently, if $w = IP_2$, we extract \tilde{v}_1 inefficiently using the sequential committed-value oracle $\mathcal{O}^\infty = \mathcal{O}^\infty[\text{ExCom.Sen}, \text{ExCom.Rec}]$. If $w = P_2I$ we don't have this problem, as the ideally scheduled right block commitment i starts after the beginning of Phase 2 on the left.

- Let ψ be the state of $\mathcal{A}_{6,j+1}^{\mathcal{O}^\infty}$ at the abovementioned point, just before the right extractable commitment begins.
- Let $\bar{\mathcal{A}}_{6,j+1}^{\mathcal{O}^\infty}$ be the adversary that starting from ψ emulates $\mathcal{A}_{6,j+1}^{\mathcal{O}^\infty}$ in the following right extractable commitments, up to those that are already simulated, while internally emulating the sender in any left commitment.
- $\mathcal{A}_{6,j}$ invokes the sequentially-extracting simulator $\text{Sim}^\infty(\bar{\mathcal{A}}_{6,j+1}, \psi)$ to remove the use \mathcal{O}^∞ .
- $\mathcal{A}_{6,j}$ then completes the execution as $\mathcal{A}_{6,j+1}$ and outputs the same.

To complete the proof of Proposition 5.1, we prove the following claim.

Claim 5.6.

- The output of $\mathcal{H}_{5,1}$ is independent of the committed value v .
- $\forall j \in [\ell + 2] : \mathcal{H}_{5,j+1} \approx_c \mathcal{H}_{6,j}$.
- $\forall j \in [\ell + 2] : \mathcal{H}_{6,j} \approx_c \mathcal{H}_{5,j}$.

Proof. We now prove each of the three parts of the claim.

Independence. In $\mathcal{H}_{5,1}$, every left block commitment j is independent of the share (u_c or $\bar{u}_{\bar{c}}$) in the free commitment c_j . By the secret sharing guarantee, all blocks are information theoretically independent of v .

$\mathcal{H}_{5,j+1} \approx_c \mathcal{H}_{6,j}$. This indistinguishability follows from the sequential extraction guarantee (Lemma 3.2). Here too, we use the fact that the inefficient oracle \mathcal{O}^∞ is only ever invoked before the simulator is applied, and thus it can be non-uniformly fixed. In case the extractable commitment starts with a receiver message, and the left free commitment c_j splits a right commitment right after the first receiver message and before the first sender message, we also rely on the fact that the sequential extractor Sim^∞ simulates the first receiver message honestly and independently of the sender circuit (Lemma 3.3).

$\mathcal{H}_{6,j} \approx_c \mathcal{H}_{5,j}$. This indistinguishability follows from the hiding of the left free commitment c_j in block j . Here, we need to address two types of calls to the inefficient oracle \mathcal{O}^∞ . First, are calls that occur strictly before the commitment c_j has started — these can be non-uniformly fixed. The second type corresponds to a right extractable commitment that is split by the left commitment c_j . Here we rely on first-message binding (Definition 3.12), which says that the value of the commitment is fixed by the first sender message. Since this first message, occurs before the commitment c_j has started, we can non-uniformly fix it, and thus also efficiently simulate a corresponding call to oracle \mathcal{O}^∞ . \square

This concludes our proof for Proposition 5.1. \square

5.4 Robustness

We show that the commitment is also r -robust.

Proposition 5.2. *The commitment scheme from Protocol 4 is extractable with r -robustness.*

Proof sketch. We describe an r -robust extractor $\text{Sim}_r(\mathcal{A}, \rho)$ that given a sender \mathcal{A} , interacting in a left r -message protocol B and with oracle access to $\mathcal{O}^\infty = \mathcal{O}^\infty[\text{Sen}, \text{Rec}]$ on the right, removes the use of the oracle.

$\text{Sim}_r(\mathcal{A}, \rho)$ invokes the non-interleaving extractor $\text{Sim}_{\text{free}}^\infty(\bar{\mathcal{A}}, \rho)$ with augmented adversary $\bar{\mathcal{A}}$. Adversary $\bar{\mathcal{A}}$ interacts with B on the left and on the right has oracle access to the sequential oracle $\mathcal{O}_{\text{free}}^\infty = \mathcal{O}_{\text{free}}^\infty[\text{ExCom.Sen}, \text{ExCom.Rec}]$ of the extractable commitment underlying $\langle \text{Sen}, \text{Rec} \rangle$. Recall that the latter oracle only answers break requests in sessions that do not interleave with the left protocol

B . $\bar{\mathcal{A}}^{\mathcal{O}_{\text{free}}^\infty}$ emulates $\mathcal{A}^{\mathcal{O}^\infty}$ and emulates interaction with \mathcal{O}^∞ using its own oracle $\mathcal{O}_{\text{free}}^\infty$. $\bar{\mathcal{A}}$ submits a break request after each free session on the right, and uses the obtained shares to answer \mathcal{A} 's break requests to \mathcal{O}^∞ .

Recall that in $\langle \text{Sen}, \text{Rec} \rangle$ there are $\ell + 2$ block commitments to the shares of the committed value \tilde{v} . By the soundness of the zero knowledge argument in Phase 2 of $\langle \text{Sen}, \text{Rec} \rangle$, the value \tilde{v} can be extracted from any one of these blocks. Since the number of block commitments is $\ell + 2 > r$, every execution contains a block on the right, which is free of interaction with B on the left. In that block, all the break requests by $\bar{\mathcal{A}}$ will succeed, allowing it to recover the shares of the committed value \tilde{v} and to properly answer the break requests made by \mathcal{A} . Thus $\bar{\mathcal{A}}^{\mathcal{O}_{\text{free}}^\infty}$ statistically simulates $\mathcal{A}^{\mathcal{O}^\infty}$. The validity of Sim_r now follows from that of the non-interleaving simulator $\text{Sim}_{\text{free}}^\infty$.

A note on using extractable commitments with quantum communication.

Remark 5.1. We remark that our construction of non-malleability commitments for a small number of tags can also work with a quantumly-extractable commitment scheme that uses quantum communication, as long as it has the property that in the protocol, the sender sends a classical statistically binding commitment which defines the committed value. The resulting non-malleable commitments also uses quantum communication, whose non-malleability is defined as in Definition 3.4 and statistical binding property states that the *classical transcript* of a commitment defines the committed values. More formally, in an interaction between an unbounded malicious committer Sen^* and the honest receiver Rec , resulting in a transcript of classical messages τ (does not contain quantum messages), the probability that Sen can output two decommitments (τ, v_0, d_0) and (τ, v_1, d_1) and Rec accepts both is negligible.

Recall that starting with an arbitrary extractable commitment $(\text{ExCom.Sen}', \text{ExCom.Rec}')$, we first transform it into another extractable commitment $(\text{ExCom.Sen}, \text{ExCom.Rec})$ that is first-message binding (Definition 3.12) using the transformation in Lemma 3.3, and then plug it into the construction in Figure 4 to obtain a non-malleable commitments for a small number of tags. The key observation is that both the transformation in 3.3 and the construction in Figure 4 essentially makes black-box use of their underlying extractable commitments, $(\text{ExCom.Sen}', \text{ExCom.Rec}')$ and $(\text{ExCom.Sen}, \text{ExCom.Rec})$ respectively, except for using ZK arguments to prove about the consistency of their committed values. By definition, the scheme $(\text{ExCom.Sen}', \text{ExCom.Rec}')$ we start with has the property that the committed value is defined by a classical statistically binding commitment. Hence the ZK proof of consistency only needs to refer to that classical statistically binding commitment (and ignore the other messages). Furthermore, the transformation in Lemma 3.3 outputs a scheme $(\text{ExCom.Sen}, \text{ExCom.Rec})$ where the sender starts by sending a Naor's 2-round statistically binding commitment to the value to be committed. Hence for the same reason, the proof of consistency in Figure 4 can simply refer to this part and ignore the rest of the commitment, which may use quantum communication.

All security analyses stay the same, since they generically reduce to the hiding and binding property of $(\text{ExCom.Sen}, \text{ExCom.Rec})$ and $(\text{ExCom.Sen}', \text{ExCom.Rec}')$.

6 Tag Amplification

In this section, we present a tag amplification transformation that converts a non-malleable commitment scheme $\langle \text{Sen}, \text{Rec} \rangle$ for $t \in [3, O(\log(\lambda))]$ bit tags into a non-malleable commitment scheme $\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle$ for $T = 2^{t-1}$ bit tags. The transformation requires $\langle \text{Sen}, \text{Rec} \rangle$ to additionally satisfy k -robustness, where k is the number of rounds of a witness-indistinguishable (WI) argument. The transformation preserves the k -robustness and incurs an additive polynomial overhead in the complexity and an additive negligible security loss. As such, the transformation can be applied iteratively to amplify the number of tags from constant to exponential in the security parameter λ ,

The transformation uses the following ingredients:

- A post-quantum secure one-way function f .

- Naor’s 2-message statistically binding commitment [Nao91] instantiated with a post-quantum secure pseudo-random generator, which in turn can be based on post-quantum one-way functions. The receiver of Naor’s protocol is public coin and sends a random string a as the first message, the sender then responds with $c = \text{Com}_a(m; d)$ depending on a ; the decommitment is simply sender’s private random coins. The receiver can reuse a across many commitments sent to it, and we can effectively use the second message of Naor’s commitments as a non-interactive commitment.
- A post-quantum secure extractable commitment scheme ECom. Let k_1 be the number of rounds in this commitment scheme.
- A post-quantum secure WI protocol which can be based on any post-quantum one-way functions. Let k_2 be the number of rounds of WI.
- A non-malleable commitment scheme $\langle \text{Sen}, \text{Rec} \rangle$ for $t \geq 3$ bit tags that is also r -robust (Definition 3.11) for $r = k_1 + k_2$. Let n be the length of messages the scheme can commit to.

The transformed non-malleable commitment $\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle$ for $T = 2^{t-1}$ tags is presented in Figure 5.

We show that $\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle$ is statistically binding, r -robust and post-quantum non-malleable. The full proof, as well as the detailed analysis of the complexity growth and security loss, is deferred to section 6.1.

6.1 Analysis of The Tag Amplification

We show the following properties of $\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle$ (Protocol 5).

Proposition 6.1. *The protocol $\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle$ is statistically binding.*

Proposition 6.2. *The protocol $\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle$ is k -robust.*

Theorem 6.1. *The protocol $\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle$ is post-quantum non-malleable.*

The statistical binding property of $\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle$ follows immediately from that the fact that the committed value is defined by the Naor’s commitments in the initial commitment step, which is statistically binding. We prove that $\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle$ is k -robust and non-malleable below.

Proof of r -Robustness, Proposition 6.2 We want to construct a PPT simulator $\widehat{\text{Sim}}_r$, such that, for every quantum polynomial-size adversary $\hat{A} = \{ \hat{A}_\lambda, \hat{\rho}_\lambda \}_{\lambda \in \mathbb{N}}$, and every every PPT r -message PPT B , the output of \hat{A} with access to the committed value oracle $\mathcal{O}[\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle]$, and in interaction with B , can be simulated by $\widehat{\text{Sim}}_r$, without access to $\mathcal{O}[\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle]$ and only interacting with B . For simplicity of notation, we suppressed subscript λ below. Formally,

$$\begin{aligned} & \left\{ \text{OUT}_{\hat{A}} \langle B(z, 1^\lambda), \hat{A}^{\mathcal{O}[\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle]}(\hat{\rho}) \rangle \right\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*} \\ & \approx \left\{ \text{OUT}_{\widehat{\text{Sim}}_r} \langle B(y), \widehat{\text{Sim}}_r(\hat{A}, \hat{\rho}) \rangle \right\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*} \end{aligned}$$

To show this, we will reduce to the r -robustness of the underlying non-malleable commitments scheme $\langle \text{Sen}, \text{Rec} \rangle$. Let Sim_r be the simulator for $\langle \text{Sen}, \text{Rec} \rangle$, such that, for every polynomial-sized adversary $A = \{ \mathcal{A}, \rho \}_{\lambda \in \mathbb{N}}$ and every PPT r -round B , we have

$$\begin{aligned} & \left\{ \text{OUT}_A \langle B(z, 1^\lambda), A^{\mathcal{O}[\langle \text{Sen}, \text{Rec} \rangle]}(\rho) \rangle \right\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*} \\ & \approx \left\{ \text{OUT}_{\text{Sim}_r} \langle B(z, 1^\lambda), \text{Sim}_r(A, \rho) \rangle \right\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*} \end{aligned}$$

Protocol 5

Common Input: Security parameter $\lambda \in \mathbb{N}$ and a tag $\hat{\text{tg}} \in \{0, 1\}^T$ for the sender, where $T = 2^{t-1}$.

$\widehat{\text{Sen}}$'s private input: A message $m \in \{0, 1\}^n$ to commit to.

1. **Trapdoor Setup:** $\widehat{\text{Rec}}$ sends two random images $y_1 = f(u_1)$ and $y_2 = f(u_2)$ of the one-way function f , where $u_1 \leftarrow \{0, 1\}^\lambda$, $u_2 \leftarrow \{0, 1\}^\lambda$. $\widehat{\text{Rec}}$ proves using WI that either y_1 or y_2 is in the image of f for λ -bit inputs. We refer to u_1 and u_2 as the trapdoors.
2. **Initial Commitment:** $\widehat{\text{Rec}}$ sends the first message a of Naor's commitment. $\widehat{\text{Sen}}$ commits to m using Naor's commitment $c = \text{Com}_a(m; d)$ w.r.t. receiver's message a , using random coins d .
3. **$\langle \text{Sen}, \text{Rec} \rangle$ commitments:** For every bit $\hat{\text{tg}}_i$ in the $T = 2^{t-1}$ bit tag $\hat{\text{tg}}$, define tag $\text{tg}_i = (i, \hat{\text{tg}}_i)$, which has exactly t bits.

For every $i \in [T]$, $\widehat{\text{Sen}}$ commits to m using $\langle \text{Sen}, \text{Rec} \rangle$ and tag tg_i ; let c_i denote the produced commitment and d_i the decommitment. All commitments are sent in parallel.

4. **Proof of Consistency:** $\widehat{\text{Sen}}$ first commits to 0^λ using the extractable commitment scheme ECom. Let c_e denote the produced commitment.

$\widehat{\text{Sen}}$ proves using WI that either $c, c_1 \dots, c_T$ are all valid commitments to m , or c_e commits to a preimage of y_1 or y_2 . Formally, it proves that the statement $X = (a, c, c_1, \dots, c_T, c_e, y_1, y_2)$ belongs to the language \mathcal{L} defined by the following witness relation:

$$\begin{aligned} \mathcal{R}_{\mathcal{L}}(X, W = (m, d, d_1, \dots, d_T, d_e, u)) &= 1 \text{ iff} \\ \text{Either } c &= \text{Com}_a(m; d) \wedge \forall i \in [T], \text{open}_{\langle \text{Sen}, \text{Rec} \rangle}(c_i, m, d_i) = 1, \\ \text{Or } \text{open}_{\text{ECom}}(c_e, u, d_e) &= 1 \text{ and } (y_1 = f(u) \text{ or } y_2 = f(u)) \end{aligned}$$

5. **Receiver's Decision:** $\widehat{\text{Rec}}$ accepts the commitment iff the proof of consistency is accepting.
6. **Decommitment:** $\widehat{\text{Sen}}$ outputs decommitment d . The decommitment is accepted if $c = \text{Com}_a(m; d)$.

Figure 5: Post-quantum tag amplification.

We construct $\widehat{\text{Sim}}_r$ for $\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle$ using Sim_r for $\langle \text{Sen}, \text{Rec} \rangle$ and show its correctness via a sequence of hybrids, \mathcal{H}_0 to \mathcal{H}_x , where \mathcal{H}_3 defines the simulator $\widehat{\text{Sim}}_r$.

\mathcal{H}_0 consists of an honest interaction $\langle B(z, 1^\lambda), \hat{A}^{\mathcal{O}[\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle]}(\hat{\rho}) \rangle$ and outputs the output of \hat{A} .

\mathcal{H}_1 proceeds identically to \mathcal{H}_0 except that the oracle $\mathcal{O}[\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle]$ is replaced with by the following procedure:

Special Extraction Procedure: For every commitment \hat{c} of $\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle$, emulate the honest receiver $\widehat{\text{Rec}}$ for \hat{A} , but forward the first commitment c_1 of $\langle \text{Sen}, \text{Rec} \rangle$ in it (using tag tg_1) to the committed value oracle $\mathcal{O}[\langle \text{Sen}, \text{Rec} \rangle]$ of $\langle \text{Sen}, \text{Rec} \rangle$. If \hat{A} is non-aborting (which implies that the final proof of consistency is accepting and c_1 is non-aborting) and queries break, call the

committed value oracle \mathcal{O} to break c_1 and obtain $m_1 \in \{0, 1\}^n \cup \{\perp\}$, and return m_1 to \hat{A} as the value committed in c .

\mathcal{H}_1 and \mathcal{H}_0 are statistically close since interaction with the special extraction procedure and $\mathcal{O}[\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle]$ are statistically close. To see this, first observe that they both emulate the honest receiver $\widehat{\text{Rec}}$ perfectly. Second, we claim that $\mathcal{O}[\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle]$ and the special extraction procedure return the same committed values with overwhelming probability. The former returns the committed value m in Naor's commitment c , while the latter returns that m_1 in c_1 , only if the proof of consistency is accepting. If $m \neq m_1$, by the soundness of WI we can efficiently extract from the ECom commitment a fake witness which is a preimage r of one of the two random images y_1, y_2 of f sent in the trapdoor setup step. It then follows from the one-wayness of f and the witness indistinguishability of WI that the probability of successful inversion is negligible. Therefore, the special extraction procedure returns the same value as $\mathcal{O}[\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle]$ with overwhelming probability.

\mathcal{H}_2 proceeds identically to \mathcal{H}_1 except that we view it as an interaction $\langle B(z, 1^\lambda), A^{\mathcal{O}[\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle]}(\hat{\rho}) \rangle$ between B and a wrapper adversary A with access to $\mathcal{O}[\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle]$. The wrapper adversary A emulates an execution of \mathcal{H}_1 with $\hat{A}(\hat{\rho})$ internally, by forwarding messages between B and \hat{A} and running the special extraction procedure using its oracle $\mathcal{O}[\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle]$ as in \mathcal{H}_1 . By definition \mathcal{H}_2 and \mathcal{H}_1 are distributed identically.

\mathcal{H}_3 consists of the simulation $\langle B(z, 1^\lambda), \text{Sim}_r(A, \hat{\rho}) \rangle$ using the simulator Sim_r of \mathcal{O} . It follows directly from the r -robustness of $\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle$ that \mathcal{H}_2 and \mathcal{H}_3 are indistinguishable.

\mathcal{H}_3 defines the following simulator $\widehat{\text{Sim}}_r$ for $\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle$: $\widehat{\text{Sim}}_r$ given $(\hat{A}, \hat{\rho})$ simply constructs the wrapper adversary A internally and runs Sim_r with the same common input and private input $(A, \hat{\rho})$.

Proof of Non-Malleability, Theorem 6.1 We want to show the following indistinguishability that for every PPT man-in-the-middle adversary $\{\hat{A}_\lambda, \rho_\lambda\}_\lambda$ (we'll suppress the subscript λ for simplicity of notation),

$$\begin{aligned} & \left\{ \text{mim}_{\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle}(\hat{A}, \rho, m_0) \right\}_{\lambda \in \mathbb{N}, m_0 \in \{0, 1\}^\lambda, m_1 \in \{0, 1\}^\lambda} \\ & \approx \left\{ \text{mim}_{\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle}(\hat{A}, \rho, m_1) \right\}_{\lambda \in \mathbb{N}, m_0 \in \{0, 1\}^\lambda, m_1 \in \{0, 1\}^\lambda} \end{aligned}$$

We show this via a sequence of hybrids.

\mathcal{H}_0^b outputs $\text{mim}_{\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle}(\hat{A}, \rho, m_b)$ – the output $\text{OUT}_{\hat{A}}$ of \hat{A} in a man-in-the-middle experiment receiving a left commitment to m_b using tag tg , and the value \tilde{m} it commits to in the right commitment (which is \perp if the right commitment is aborting or the right tag $\tilde{\text{tg}}$ is the same as the left tag tg).

\mathcal{H}_1^b proceeds identically to \mathcal{H}_0^b except for that the right committed value \tilde{m} is computed differently. In both \mathcal{H}_0^b and \mathcal{H}_1^b , if the right commitment is not accepting or the left and right tags are the same $\text{tg} = \tilde{\text{tg}}$, the right committed value \tilde{m} is set to \perp . Otherwise, consider two cases:

Case 1 : If Step 3 of the right commitment starts before the first message y_1, y_2 in the left commitment is sent, then in both \mathcal{H}_0^b and \mathcal{H}_1^b , \tilde{m} is set to the value committed in the Naor's commitment \tilde{c} in Step 2.

Case 2 : Else if Step 3 of the right commitment starts after the first message y_1, y_2 in the left commitment is sent, in \mathcal{H}_0^b , \tilde{m} is still set to the value committed in the Naor's commitment \tilde{c} in Step 2, whereas in \mathcal{H}_1^b , \tilde{m} is set to \tilde{m}_i committed in the i 'th $\langle \text{Sen}, \text{Rec} \rangle$ commitment \tilde{c}_i in Step 3 of the right commitment, for an i satisfying that $\text{tg}_i \neq \tilde{\text{tg}}_i$.

By the same argument as in the proof of r -robustness (Proposition 6.2), by the soundness and witness indistinguishability of WI, the extractability of ECom, and the one-wayness of f , with overwhelming probability, the value committed in Naor's commitment in Step 2 and \tilde{m}_i committed in the i 'th $\langle \text{Sen}, \text{Rec} \rangle$ commitment in Step 3 are the same. Therefore, \mathcal{H}_0^b and \mathcal{H}_1^b are statistically close.

\mathcal{H}_2^b proceeds identically to \mathcal{H}_1^b except that the Step 4 WI of consistency in the left commitment is generated using a fake witness. More specifically, in Step 1 of the left commitment, \hat{A} sends y_1, y_2 and gives a WI argument that one of them is in the image of f . If this argument is accepting, with overwhelming probability, there is a fake witness u that is the preimage of y_1 or y_2 (or else, the honest committer $\widehat{\text{Sen}}$ of the left commitment would abort). In this case, in Step 4, instead of sending an ECom commitment of 0^λ and using the honest witness, send an ECom commitment of u and generate the proof of consistency using the fake witness u .

We show that by the soundness and the witness indistinguishability of WI, the hiding of ECom, and the r -robustness of the input non-malleable commitment scheme $\langle \text{Sen}, \text{Rec} \rangle$, \mathcal{H}_1^b and \mathcal{H}_2^b are indistinguishable.

Claim 6.1. For every $b \in \{0, 1\}$, $\mathcal{H}_1^b \approx \mathcal{H}_2^b$.

Proof. Suppose for contradiction that the view of \hat{A} and the right committed value \tilde{m} in \mathcal{H}_1^b and \mathcal{H}_2^b are distinguishable with advantage $1/p(\lambda)$ for some polynomial p . Observe that \mathcal{H}_1^b and \mathcal{H}_2^b proceed identically before Step 4 in the left commitment starts, and thus conditioned on the Step 1 WI from \hat{A} not convincing, these two hybrids are identically distributed. Therefore, the contradiction hypothesis implies that conditioned on the WI in Step 1 of the left commitment being accepting, the output of \hat{A} and the right committed value \tilde{m} in \mathcal{H}_1^b and \mathcal{H}_2^b are distinguishable with advantage $1/p(\lambda)$. In this case, by the argument of knowledge property of WI in Step 1, with overwhelming probability there exists a trapdoor u that is either the preimage of y_1 or y_2 . Thus, there must exist a prefix τ of execution of $\mathcal{H}_1^b/\mathcal{H}_2^b$ up to the point that y_1, y_2 are sent, such that, i) there exists a fake witness u , and ii) conditioned on τ occurring, the output of \hat{A} and the right committed value \tilde{m} in \mathcal{H}_1^b and \mathcal{H}_2^b is distinguishable with advantage $1/p(\lambda)$. Furthermore, if Case 1 occurs, that is, Step 3 of the right commitment starts inside τ , let \tilde{m}' be the value committed in Naor's commitment in Step 2. Otherwise, $\tilde{m}' = \perp$.

Using τ, \tilde{m}' and \hat{A} we construct a wrapper adversary A that distinguishes an ECom commitment to 0^λ and WI argument generated using an honest witness, from an ECom commitment to u and WI argument generated using a fake witness, while having access to the committed value oracle $\mathcal{O}[\langle \text{Sen}, \text{Rec} \rangle]$ on the right. More specifically, $A(\tau, \tilde{m}')$ internally emulates $\mathcal{H}_1^b/\mathcal{H}_2^b$ with \hat{A} starting from τ , by emulating the honest committer $\widehat{\text{Sen}}(m_b)$ on the left and receiver $\widehat{\text{Rec}}$ on the right, except that, i) it forwards the ECom commitment and WI argument it receives as the Step 4 argument in the left commitment, and ii) if Case 2 occurs, that is, Step 3 of the right commitment starts outside τ , it forwards the i 'th $\langle \text{Sen}, \text{Rec} \rangle$ commitment \tilde{c}_i in Step 3 for i s.t. $\text{tg}_i \neq \tilde{\text{tg}}_i$ to $\mathcal{O}[\langle \text{Sen}, \text{Rec} \rangle]$ and obtains the value \tilde{m}_i committed in \tilde{c}_i . Finally, it outputs the output of \hat{A} in the emulation and right committed value $\tilde{m} = \tilde{m}'$ if Case 1, and $\tilde{m} = \tilde{m}_i$ if Case 2. If the right commitment is not accepting or $\text{tg} = \text{tg}_i$, \tilde{m} is overwritten by \perp . Observe that depending on the value committed in the ECom commitment and the witness used in the WI that A receives,

it emulates \mathcal{H}_1^b or \mathcal{H}_2^b perfectly for \hat{A} , and A 's output is exactly the output of \hat{A} and the right committed value output by \mathcal{H}_1^b or \mathcal{H}_2^b .

Finally, by the r -robustness of $\langle \text{Sen}, \text{Rec} \rangle$, we can simulate the output of $A^{\mathcal{O}[\langle \text{Sen}, \text{Rec} \rangle]}(\tau, \tilde{m}')$ when receiving an ECom commitment and WI argument of total r messages, by the output of the r -robust simulator $\text{Sim}_r(\tau, \tilde{m}')$ of $\langle \text{Sen}, \text{Rec} \rangle$ when receiving the same r -round WI argument. Therefore, by the fact that the output of A is distinguishable with advantage $1/p(\lambda)$, we have that the output of Sim_r is also distinguishable, which contradicts either the hiding of ECom or the witness indistinguishability of WI. \square

\mathcal{H}_3^b proceeds identically to \mathcal{H}_2^b except that in the initial commitment step, it commits to 0 using Naor's commitment scheme instead of committing to m_b . By a similar argument as that for Claim 6.1, it follows from the computational hiding property of Naor's commitment and the r -robustness of $\langle \text{Sen}, \text{Rec} \rangle$ that the output of \hat{A} and the right committed value \tilde{m} in \mathcal{H}_2^b and \mathcal{H}_3^b are indistinguishable.

$\mathcal{H}_{4:i}^b$ for $i \in \{0, 1, \dots, T\}$ proceeds identically to \mathcal{H}_3^b except that in the $(\widehat{\text{Sen}}, \widehat{\text{Rec}})$ commitments step, the first i commitments of $\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle$ commit to 0 instead of m_b . It follows from the non-malleability of $\langle \text{Sen}, \text{Rec} \rangle$ that if there exists i such that $\text{tg}_i \neq \tilde{\text{tg}}_i$, the output of \hat{A} and the value \tilde{m}_i committed in the i 'th $\langle \text{Sen}, \text{Rec} \rangle$ commitment \tilde{c}_i on the right is indistinguishable in $\mathcal{H}_{4:0}^b$ and \mathcal{H}_3^b , as well as in $\mathcal{H}_{4:i}^b$ and $\mathcal{H}_{4:i+1}^b$. As in \mathcal{H}_2^b , the right committed value \tilde{m} is set to \tilde{m}_i in Case 2. Otherwise, \tilde{m} is set to the value committed in the Naor's commitment in Step 2 in Case 1, which occurs before the left commitment starts, and hence are the same in neighboring hybrids. Therefore, we have that the output of \hat{A} and \tilde{m} is indistinguishable in $\mathcal{H}_{4:0}^b$ and \mathcal{H}_3^b , as well as in $\mathcal{H}_{4:i}^b$ and $\mathcal{H}_{4:i+1}^b$.

Finally, observe that $\mathcal{H}_{4:T}^b$ is independent of b . Hence by a hybrid argument we have that \mathcal{H}_0^0 and \mathcal{H}_0^1 , which shows that $\langle \widehat{\text{Sen}}, \widehat{\text{Rec}} \rangle$ is non-malleable.

Iteratively applying the Transformation – Complexity Growth and Security Loss

We now analyze the complexity growth and security loss when applying the transformation iteratively to transform a non-malleable commitment for a constant number of tags into one for exponential number of tags.

First observe that the transformation increases the round complexity additively by $2k + O(1)$. We show below that if the complexity of the protocol $(\widehat{\text{Sen}}_1, \widehat{\text{Rec}}_1)$ produced after l iterations is D_l , the complexity of that $(\widehat{\text{Sen}}_{l+1}, \widehat{\text{Rec}}_{l+1})$ produced after $l + 1$ iterations is $D_1 T_{l+1} + \text{poly}(\lambda)$, where T_{l+1} is the length of tags of $(\widehat{\text{Sen}}_{l+1}, \widehat{\text{Rec}}_{l+1})$. Therefore, iteratively applying the transformation for $L = O(\log^* \lambda)$ times on a non-malleable commitment with constant-bit tags produces one for polynomial-bit tags, while increasing the round complexity additively by $(\log^* \lambda)(2k + O(1))$, and the computational complexity of the final scheme is $\text{poly}(\lambda)$. To see the latter note that all but the final transformation produces a scheme with logarithmic length tags, $T_l = O(\log \lambda)$ for every $l < L$. Thus the complexity of the second last scheme $(\widehat{\text{Sen}}_{L-1}, \widehat{\text{Rec}}_{L-1})$ is bounded by $\text{poly}(\lambda)O(\log \lambda)^L = \text{poly}(\lambda)$. Then, the complexity of the last scheme $(\widehat{\text{Sen}}_L, \widehat{\text{Rec}}_L)$ is also $\text{poly}(\lambda)$ as $T_L = \text{poly}(\lambda)$.

It remains to argue that the complexity of $(\widehat{\text{Sen}}_{l+1}, \widehat{\text{Rec}}_{l+1})$ is $D_1 T_{l+1} + \text{poly}(\lambda)$. It is easy to see that Step 1 and 2 has fixed $\text{poly}(\lambda)$ complexity, and Step 3 has complexity T_{l+1} times the complexity D_l of $(\widehat{\text{Sen}}_l, \widehat{\text{Rec}}_l)$. At a first glance, Step 4 would have complexity polynomial in that of Step 2 and 3, since it proves using WI a statement involving the Naor's commitment in Step 2, the T_{l+1} $\langle \text{Sen}, \text{Rec} \rangle$ commitments in Step 3, and the ECom commitment in Step 4. However, a more careful examination shows that the value committed by $\langle \text{Sen}, \text{Rec} \rangle$ (except for the initial non-malleable commitment scheme in 5) is always defined by a Naor's commitment, if the $\langle \text{Sen}, \text{Rec} \rangle$ commitment is accepting. More

precisely, for every $l \geq 1$, $(\widehat{\text{Sen}}_l, \widehat{\text{Rec}}_l)$ has form as in Figure 5, whose committed value is defined by the Naor’s commitment in Step 2. Therefore, the complexity of Step 4 is polynomial in the complexity of $T_{l+1} + 1$ Naor’s commitments, which is bounded by a fixed polynomial $\text{poly}(\lambda)$. Therefore, the complexity of $(\widehat{\text{Sen}}_{l+1}, \widehat{\text{Rec}}_{l+1})$ is $D_1 T_{l+1} + \text{poly}(\lambda)$.

Finally, we analyze the security loss. First observe that the proofs of r -robustness and non-malleability for each transformation go through hybrids, where the indistinguishability of neighboring hybrids either reduces to the security properties of some basic primitives, including Naor’s commitment, WI, ECom, and one-way function f , or reduces to the r -robustness or non-malleability of the input non-malleable commitment. For all reductions, the size of the reduction is larger than the size of the adversary by an additive polynomial factor. Therefore, after $O(\log^* \lambda)$ iterations, the security loss in the size of the adversary is bounded by an additive polynomial factor. Next observe that the proof of r -robustness essentially reduces the r -robustness of the output non-malleable commitment scheme to that of the input scheme, while incurring an additive negligible increase in the advantage (of violating r -robustness). The proof of non-malleability, on the other hand, reduces the non-malleability of the output scheme to that of the input scheme, while incurring both an additive negligible increase and a multiplicative T fold increase in the advantage, where T is the length of tags of the output scheme. Therefore, by a similar analysis as above for complexity growth, after $O(\log^* \lambda)$ iterations, the advantage of the adversary in violating either r -robustness and non-malleability increases by at most a multiplicative and an additive polynomial factor. Overall, the security loss is polynomial.

A note on using extractable commitments with quantum communication.

Remark 6.1. Recall that in remark 5.1 we argued that if we started with a quantumly-extractable commitment protocol (ExCom.Sen, ExCom.Rec) that uses quantum communication, satisfying the property that the sender sends a classical statistically binding commitment in the protocol to commit to the value, we would obtain a non-malleable commitments for a small number of tags also using quantum communication satisfying that the *classical transcript* of a commitment defines the committed values.

By similar argument as in remark 5.1, our tag-amplification in Figure 5 can also be applied to such a non-malleable commitment scheme and the extractable commitment scheme (ExCom.Sen, ExCom.Rec). This is because it essentially makes black-box use of the underlying non-malleable commitment and extractable commitment schemes, except for using ZK arguments to prove about the consistency of their committed values. Therefore, since the values committed to by both schemes are defined by the classical transcript, the ZK proof of consistency only needs to refer to the classical transcripts of these commitments (and ignore other messages). Furthermore, observe that the output non-malleable commitment scheme commits to values using Naor’s commitment scheme. Hence the invariant that values are defined by the classical transcript is kept and the tag-amplification can be iteratively applied. All security analyses stay the same, since they generically reduce to the non-malleability, hiding, and binding property of the the component commitment schemes.

References

- [ABG⁺20] Amit Agarwal, James Bartusek, Vipul Goyal, Dakshita Khurana, and Giulio Malavolta. Post-quantum multi-party computation in constant rounds. *CoRR*, abs/2005.12904, 2020.
- [AP19] Prabhanjan Ananth and Rolando L. La Placa. Secure quantum extraction protocols. *IACR Cryptol. ePrint Arch.*, 2019:1323, 2019.
- [Bar02] Boaz Barak. Constant-round coin-tossing with a man in the middle or realizing the shared random string model. pages 345–355, 2002.

- [BG08] Boaz Barak and Oded Goldreich. Universal arguments and their applications. *SIAM J. Comput.*, 38(5):1661–1694, 2008.
- [BL18] Nir Bitansky and Huijia Lin. One-message zero knowledge and non-malleable commitments. In *Theory of Cryptography - 16th International Conference, TCC 2018, Panaji, India, November 11-14, 2018, Proceedings, Part I*, pages 209–234, 2018.
- [Blu87] Manuel Blum. How to prove a theorem so no one else can claim it. In *In: Proceedings of the International Congress of Mathematicians*, pages 1444–1451, 1987.
- [Bra18] Zvika Brakerski. Quantum FHE (almost) as secure as classical. pages 67–95, 2018.
- [BS20] Nir Bitansky and Omri Shmueli. Post-quantum zero knowledge in constant rounds. In Konstantin Makarychev, Yury Makarychev, Madhur Tulsiani, Gautam Kamath, and Julia Chuzhoy, editors, *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020*, pages 269–279. ACM, 2020.
- [CCY20] Nai-Hui Chia, Kai-Min Chung, and Takashi Yamakawa. Black-box approach to post-quantum zero-knowledge in constant round. *arXiv preprint arXiv:2011.02670*, 2020.
- [CLP16] Ran Canetti, Huijia Lin, and Rafael Pass. Adaptive hardness and composable security in the plain model from standard assumptions. *SIAM J. Comput.*, 45(5):1793–1834, 2016.
- [CMSZ21] Alessandro Chiesa, Fermi Ma, Nicholas Spooner, and Mark Zhandry. Post-quantum succinct arguments. *arXiv preprint arXiv:2103.08140*, 2021.
- [COSV16] Michele Ciampi, Rafail Ostrovsky, Luisa Siniscalchi, and Ivan Visconti. Concurrent non-malleable commitments (and more) in 3 rounds. pages 270–299, 2016.
- [COSV17] Michele Ciampi, Rafail Ostrovsky, Luisa Siniscalchi, and Ivan Visconti. Four-round concurrent non-malleable commitments from one-way functions. pages 127–157, 2017.
- [CR87] Benny Chor and Michael O. Rabin. Achieving independence in logarithmic number of rounds. In *Proceedings of the Sixth Annual ACM Symposium on Principles of Distributed Computing, Vancouver, British Columbia, Canada, August 10-12, 1987*, pages 260–268, 1987.
- [DDN03] Danny Dolev, Cynthia Dwork, and Moni Naor. Nonmalleable cryptography. *SIAM Review*, 45(4):727–784, 2003.
- [FLS99] Uriel Feige, Dror Lapidot, and Adi Shamir. Multiple noninteractive zero knowledge proofs under general assumptions. *SIAM J. Comput.*, 29(1):1–28, 1999.
- [FP96] Christopher A. Fuchs and Asher Peres. Quantum-state disturbance versus information gain: Uncertainty relations for quantum information. *Phys. Rev. A*, 53:2038–2045, Apr 1996.
- [GKLW20] Rachit Garg, Dakshita Khurana, George Lu, and Brent Waters. Black-box non-interactive non-malleable commitments. Cryptology ePrint Archive, Report 2020/1197, 2020. <https://eprint.iacr.org/2020/1197>.
- [GKS16] Vipul Goyal, Dakshita Khurana, and Amit Sahai. Breaking the three round barrier for non-malleable commitments. pages 21–30, 2016.
- [GLOV12] Vipul Goyal, Chen-Kuei Lee, Rafail Ostrovsky, and Ivan Visconti. Constructing non-malleable commitments: A black-box approach. pages 51–60, 2012.

- [GLSV21] Alex B. Grilo, Huijia Lin, Fang Song, and Vinod Vaikuntanathan. Oblivious transfer is in minicrypt. In *Eurocrypt*, 2021. <https://eprint.iacr.org/2020/1500>.
- [GMW91] Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. *Journal of the ACM*, 38(3):690–728, 1991.
- [Goy11] Vipul Goyal. Constant round non-malleable protocols using one way functions. pages 695–704, 2011.
- [GPR16a] Vipul Goyal, Omkant Pandey, and Silas Richelson. Textbook non-malleable commitments. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 1128–1141, 2016.
- [GPR16b] Vipul Goyal, Omkant Pandey, and Silas Richelson. Textbook non-malleable commitments. pages 1128–1141, 2016.
- [GR19] Vipul Goyal and Silas Richelson. Non-malleable commitments using goldreich-levin list decoding. In *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019*, pages 686–699, 2019.
- [HSS15] Sean Hallgren, Adam Smith, and Fang Song. Classical cryptographic protocols in a quantum world. *International Journal of Quantum Information*, 13(04):1550028, 2015. Preliminary version appeared in IACR Crypto 2011.
- [Khu17] Dakshita Khurana. Round optimal concurrent non-malleability from polynomial hardness. In *Theory of Cryptography - 15th International Conference, TCC 2017, Baltimore, MD, USA, November 12-15, 2017, Proceedings, Part II*, pages 139–171, 2017.
- [KK19] Yael Tauman Kalai and Dakshita Khurana. Non-interactive non-malleability from quantum supremacy. pages 552–582, 2019.
- [KS17] Dakshita Khurana and Amit Sahai. How to achieve non-malleability in one or two rounds. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 564–575, 2017.
- [LN11] Carolin Lunemann and Jesper Buus Nielsen. Fully simulatable quantum-secure coin-flipping and applications. pages 21–40, 2011.
- [LP09] Huijia Lin and Rafael Pass. Non-malleability amplification. pages 189–198, 2009.
- [LP11] Huijia Lin and Rafael Pass. Constant-round non-malleable commitments from any one-way function. pages 705–714, 2011.
- [LP12] Huijia Lin and Rafael Pass. Black-box constructions of composable protocols without set-up. pages 461–478, 2012.
- [LPS17] Huijia Lin, Rafael Pass, and Pratik Soni. Two-round and non-interactive concurrent non-malleable commitments from time-lock puzzles. pages 576–587, 2017.
- [LPV08a] Huijia Lin, Rafael Pass, and Muthuramakrishnan Venkatasubramanian. Concurrent non-malleable commitments from any one-way function. In *Theory of Cryptography, Fifth Theory of Cryptography Conference, TCC 2008, New York, USA, March 19-21, 2008.*, pages 571–588, 2008.

- [LPV08b] Huijia Lin, Rafael Pass, and Muthuramakrishnan Venkitasubramaniam. Concurrent non-malleable commitments from any one-way function. pages 571–588, 2008.
- [LPV09] Huijia Lin, Rafael Pass, and Muthuramakrishnan Venkitasubramaniam. A unified framework for concurrent security: universal composability from stand-alone non-malleability. In Michael Mitzenmacher, editor, *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 179–188. ACM, 2009.
- [Mah18] Urmila Mahadev. Classical homomorphic encryption for quantum circuits. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 332–338. IEEE, 2018.
- [Nao91] Moni Naor. Bit commitment using pseudorandomness. *J. Cryptology*, 4(2):151–158, 1991.
- [PPV08] Omkant Pandey, Rafael Pass, and Vinod Vaikuntanathan. Adaptive one-way functions and applications. pages 57–74, 2008.
- [PR05a] Rafael Pass and Alon Rosen. Concurrent non-malleable commitments. pages 563–572, 2005.
- [PR05b] Rafael Pass and Alon Rosen. New and improved constructions of non-malleable cryptographic protocols. pages 533–542, 2005.
- [PW10] Rafael Pass and Hoeteck Wee. Constant-round non-malleable commitments from sub-exponential one-way functions. pages 638–655, 2010.
- [Ros04] Alon Rosen. A note on constant-round zero-knowledge proofs for np. In *Theory of Cryptography Conference*, pages 191–202. Springer, 2004.
- [Unr12] Dominique Unruh. Quantum proofs of knowledge. In *Annual international conference on the theory and applications of cryptographic techniques*, pages 135–152. Springer, 2012.
- [Wat09] John Watrous. Zero-knowledge against quantum attacks. *SIAM Journal on Computing*, 39(1):25–58, 2009.
- [Wee10] Hoeteck Wee. Efficient chosen-ciphertext security via extractable hash proofs. pages 314–332, 2010.
- [WZ82] W. K. Wootters and W. H. Zurek. A single quantum cannot be cloned. *Nature*, 299:802–803, 1982.