

Textbook Non-Malleable Commitments

Vipul Goyal*

Omkant Pandey†

Silas Richelson‡

Abstract

We present a new non-malleable commitment protocol. Our protocol has the following features:

- The protocol has only *three rounds* of interaction. Pass (TCC 2013) showed an impossibility result for a two-round non-malleable commitment scheme w.r.t. a black-box reduction to any “standard” intractability reduction. Thus, this resolves the round complexity of non-malleable commitment at least w.r.t. black-box security reductions. Our construction is secure as per the standard notion of non-malleability w.r.t. commitment.
- Our protocol is *truly efficient*. In our basic protocol, the entire computation of the committer is dominated by just three invocations of a non-interactive statically binding commitment scheme, while, the receiver computation (in the commitment stage) is limited to just sampling a random string. Unlike many previous works, we directly construct a protocol for large tags and hence avoid any non-malleability amplification steps.
- Our protocol is based on a black-box use of any non-interactive statistically binding commitment scheme. Such schemes, in turn, can be based on any one-to-one one-way function (or any one-way function at the cost of an extra initialization round). Previously, the best known black-box construction of non-malleable commitments required a larger (constant) number of rounds.
- Our construction is public-coin and makes use of only black-box simulation. Prior to our work, no public-coin constant round non-malleable commitment schemes were known based on black-box simulation.

Our techniques depart *significantly* from the techniques used previously to construct non-malleable commitment schemes. As a main technical tool, we rely on non-malleable codes in the split state model. Our proofs of security are purely combinatorial in nature.

In addition, we also present a simple construction of constant round non-malleable commitments from any one-way function. While this result is not new, the main feature is its simplicity compared to *any* previous construction of non-malleable commitments (in any number of rounds). We believe the construction is simple enough to be covered in a graduate level course on cryptography. The construction uses non-malleable codes in the split state model in a black-box way.

*Microsoft Research, India. Email: vipul@microsoft.com.

†University of California, Berkeley. Email: omkant@berkeley.edu.

‡MIT. Email: SiRichel@csail.mit.edu. Work done in part while visiting Microsoft Research, India.

1 Introduction

Man-in-the-middle (MIM) attacks are one of the most basic attacks in cryptography. The notion of non-malleable commitments was introduced in a seminal work of Dolev, Dwork and Naor [DDN91] as a countermeasure against such adversaries. Since their introduction, non-malleable commitments have proven to be capable of preventing MIM attacks in a variety of settings. Non-malleability lies at the heart of secure protocol composition, it allows for round-efficient secure multi-party computation and gives applications to areas as diverse as position based cryptography [CGMO09].

A *commitment scheme* is a useful two party protocol which allows a committer, \mathcal{C} , to send a representation of his message v , $\text{Com}(v; r)$ to a receiver, \mathcal{R} , in such a way so that 1) \mathcal{R} learns nothing about v until \mathcal{C} chooses to open his commitment and 2) \mathcal{C} is bound to v ; he cannot open $\text{Com}(v)$ to any value $v' \neq v$. A commitment scheme is *non-malleable* if for every message v , no MIM adversary, intercepting a commitment $\text{Com}(v; r)$ and modifying it at will, is able to efficiently generate a commitment $\text{Com}(\tilde{v}; \tilde{r})$ to a related message \tilde{v} . Interest in non-malleable commitments is motivated both by the central role that they play in securing protocols under composition [CLOS02, LPV09] and by the unfortunate reality that many widely used commitment schemes are actually highly malleable. Indeed, man-in-the-middle (MIM) attacks occur quite naturally when multiple concurrent executions of protocols are allowed, and can be quite devastating.

Since their conception, non-malleable commitment has been studied extensively, and, with increasing success in terms of characterizing its round complexity. The original construction of [DDN91] gave a protocol with logarithmically many rounds. Barak [Bar02] gave a constant round construction based on non-black-box simulation (which was further improved by Pass and Rosen [PR05b]). More recently, constant round protocols for non-malleable commitment with black-box proofs of security were given by Goyal [Goy11] and Lin and Pass [LP11]. Other constructions include [PR05a, LP09, LPV08, PPV08, PW10, Wee10, GLOV12]. The current state of art is represented by the work of Goyal, Richelson, Rosen and Vald [GRRV14] whose scheme requires only four rounds of interaction. On the negative side, Pass [Pas13] showed that two-round non-malleable commitments cannot exist w.r.t. black-box proofs of security based on any “standard” intractability assumption. The lower bound of Pass holds even if the construction uses the underlying assumption in a non-black-box way. Thus, the main remaining open problem regarding the round complexity of non-malleable commitment is

Does there exist a protocol for non-malleable commitment with only three rounds of interaction?

In this work we present a three round protocol, thus giving a positive answer to the above question.

Zero-Knowledge: A Barrier to Three-Round Non-Malleable Commitment. Almost all previous schemes invoke some sort of proof of consistency. These are usually critical to the proofs of non-malleability, as without consistency one runs into a host of “selective \perp attacks” (where the MIM plays in such a way so that whether or not his commitment is valid, depends on the value inside the commitment he receives) which are difficult to rule out. Generally, zero-knowledge is used for this purpose. For example, the recent works of [GRRV14, BGR⁺15] use a three round “commit-and-prove” sigma protocol along with a GMW-style zero-knowledge proof of correctness [GMW87]. They then use the Feige-Shamir paradigm [FS90] in order to parallelize their protocol down to four rounds.

Zero-knowledge, however, is known to require 4-rounds [GK96] (at least w.r.t. black-box simulation), so if we hope to get three round non-malleable commitment, we must overcome our dependency on zero-knowledge. The main observation which makes this possible is that using zero-knowledge to prove consistency is actually overkill. We do not need to be certain that \mathcal{M} has played honestly; we only require that whether \mathcal{M} has played honestly or not, cannot depend on \mathcal{C} 's commitment v . Capitalizing on this observation, however, is challenging and requires new ideas and a new protocol. Indeed, if we simply remove the zero-knowledge from [GRRV14], the resulting protocol is subject to easy mauling attacks (see the Appendix).

We also point out that also that zero-knowledge is usually the most computationally expensive component in protocols for non-malleable commitment. Indeed, all previous schemes for non-malleable commitment are considerably slower than their ordinary statically binding counterparts.

Our Contributions. We present a new construction of non-malleable commitment which has the following features:

- The protocol has only *three rounds* of interaction. Pass [Pas13] showed that two-round non-malleable commitments unfortunately cannot exist w.r.t. black-box proofs of security based on any “standard” intractability assumption. The lower bound of Pass holds even if the construction uses the underlying assumption in a non-black-box way. Thus, this resolves the round complexity of non-malleable commitments at least w.r.t. black-box security reductions. Our construction is secure as per the standard notion of non-malleability w.r.t. commitment.
- Our protocol is simple and *truly efficient*. In our basic protocol, the entire computation of the committer is dominated by just three invocations of a non-interactive statically binding commitment scheme, while, the receiver computation (in the commitment stage) is limited to just sampling a random string. The decommitment stage is equally basic: the committer would send the openings of these commitments, while, the receiver would be required to check these openings for correctness and perform some simple computations. The protocol is easy to describe, the main complexity lies in the analysis rather than the construction.

In several previous works (including [GRRV14, BGR⁺15]), first a non-malleable commitment scheme for “small” tags is constructed. Then, a scheme for large tags is obtained using non-malleability amplification [DDN91, LP09]. This adds a significant multiplicative overhead to the computation of each party: the multiplicative overhead is typically related to the number of bits in the large tags. Unlike these previous works, our basic protocol works directly with large tags, and hence, we avoid any expensive amplification steps. Our basic protocol provides security only against synchronizing adversaries. Extension to non-synchronizing adversaries is addressed later, though still with a three round protocol.

- Our protocol is based on black-box use of any non-interactive statistically binding commitment scheme. Such schemes, in turn, can be based on any one-to-one, one-way function, or, at the cost of an extra initialization round, any one-way function. Previously, the best known black-box construction of non-malleable commitments required a larger constant number of rounds [GLOV12, KMO14]. Furthermore, the previous constructions, even though black-box, were significantly less efficient [GLOV12, LP12, KMO14, Kiy14]. For example, the construction of Goyal et al [GLOV12] used “MPC in the head techniques” of Ishai et. al [IKOS07].
- Our construction is public-coin and makes use of only black-box simulation. Prior to our work, no public-coin constant round non-malleable commitment schemes were known based on black-box simulation. The structure of our basic protocol is arguably “as basic as it can be”: \mathcal{C} sends a single commitment to some string, the receiver sends a random challenge, and, in the final round, sender sends another string (but doesn’t send any opening).
- Finally, we present a simple — in fact, almost elementary — construction of a constant-round non-malleable commitment scheme using split state non-malleable codes. This construction can be based on any one-way function. While such a construction is not new, the main feature is its simplicity over any previous construction: it has a simple “commit-and-prove” structure and can be based on any split state non-malleable code (without any additional properties). The reduction makes a standard use of adaptive security, and may be of independent interest. We believe that this protocol is simple enough to be included in a graduate level textbook on cryptography or be taught in a graduate level cryptography course.

Starting Technical Idea. Our key technical tool will be non-malleable codes in the split-state model [DPW10, DKO13]. We will use these codes in conjunction with a techniques introduced by Goyal et. al. [GRRV14]. Non-malleable codes in the split-state model are codes whose codewords are pairs $(L, R) \in \mathcal{L} \times \mathcal{R}$, and the tampering function family is

$$\mathcal{F}_{\text{split}} = \{(f, g) \mid f : \mathcal{L} \rightarrow \mathcal{L}, g : \mathcal{R} \rightarrow \mathcal{R}\}.$$

That is, one encodes a message by breaking it into two states, and is ensured that non-malleability holds as long as the adversary tampers each state separately. Several recent and exciting works [DKO13, ADL14, CZ14, ADKO15, CGL15] establish a connection between split-state non-malleable codes and various types of randomness extractors. Thus, split-state non-malleable codes allow us to harness deep theorems from the extensive randomness extraction literature and direct them towards the seemingly unrelated area of cryptographic non-malleability.

As a first attempt towards constructing non-malleable commitments, consider the protocol where the committer \mathcal{C} simply commits separately to L and R? This does not work as the underlying commitment scheme may have some homomorphic properties allowing the MIM to maul L and R “jointly”. Our starting idea is as follows. Let us focus our attention on synchronizing adversaries¹. \mathcal{C} encodes the message v as L and R, and, in the first round, sends a commitment $\text{Com}(L)$ to L. The receiver responds back with an acknowledgement message, at which point \mathcal{C} sends R *in the clear*. This scheme does seem to have some non-malleability features. In the first round, the MIM must maul $\text{Com}(L)$ into $\text{Com}(\tilde{L})$ without knowledge of R, while in the final round, the MIM receives R and must produce \tilde{R} given only $\text{Com}(L)$ (rather than L itself). While this is indeed our starting point, this intuition turns out to be not sound (see Appendix for an “explicit” attack). Our basic protocol is quite simple and is given below.

- **Committer’s Input:** A value v to commit to.
- **1. $\mathcal{C} \rightarrow \mathcal{R}$:** \mathcal{C} chooses $(L, R) \leftarrow \text{Enc}(v)$ where L is viewed as a field element in \mathbb{Z}_q ; \mathcal{C} also draws $r \leftarrow \mathbb{Z}_q$ at random and sends $\text{Com}(L \circ r)$ to R where Com is a non-interactive, statistically binding commitment scheme.
- **2. $\mathcal{R} \rightarrow \mathcal{C}$:** \mathcal{R} chooses a random $\alpha \leftarrow \mathbb{Z}_q^*$ and sends it to \mathcal{C} .
- **3. $\mathcal{C} \rightarrow \mathcal{R}$:** \mathcal{C} sends $a = r\alpha + L$ and R to \mathcal{R} .
- **Decommitment:** To decommit, \mathcal{C} decommits to the commitment in **1**.

Intuitively, \mathcal{C} commits to a polynomial-based 2-out-of-2 secret sharing of L in the first round, and in the third round sends R along with one share. The same polynomial based commit-and-reveal mechanism is used in [GRRV14], but no non-malleable codes are used and so a zero-knowledge proof of consistency is needed.

To prove security of the above protocol, we must reduce any “successful” mauling attack to a mauling attack on the underlying non-malleable code. The adversary for the non-malleable codes (*i.e.*, a split-state tampering function pair (f, g)) would have to run the left execution (with the MIM) using the given L and R and extract \tilde{L} and \tilde{R} from the right execution. If MIM successfully mauls the commitment scheme, then the extracted \tilde{L} and \tilde{R} would decode to \tilde{v} which is related to \mathcal{C} ’s committed value v represented by L and R. This would presumably contradict the security of the non-malleable code. However one must keep in mind that (f, g) must be split-state and so neither function is allowed to see L and R at once. Hence, it cannot simply run our non-malleable commitment protocol, and extract the tampered \tilde{L} and \tilde{R} .

To complete the proof of security, we need to construct split-state (f, g) (which can use the MIM and the distinguisher for our protocol internally) so that f outputs \tilde{L} using only L and likewise g outputs \tilde{R} using only R. Therefore f and g cannot complete the protocol execution with MIM to extract the tampered code (since each

¹Roughly, this means that the MIM sends the i -th round message on the right immediately after getting the i -th round message in the left interaction.

will be missing one of L and R)! Thus, the idea of reducing the security of our construction to the security of non-malleable codes (in the split-state model) seems like a non-starter. *This is the key technical challenge we encounter in our proof of non-malleability.*

Our proof strategy, at a very high level is to have $f(L)$ and $g(R)$ execute independent interactions with MIM, and output \tilde{L} and \tilde{R} , respectively. This, however, leads to \tilde{L} and \tilde{R} being extracted from two different protocol transcripts, and so there is no clear way to relate the extracted values back to MIM's mauled commitment \tilde{v} . We then show that there exists a single protocol transcript (from the correct distribution) such that the left execution in that transcript is completed using L and R, and the right uses \tilde{L} and \tilde{R} . Thus, if MIM is successful in mauling the commitment scheme, then (f, g) is split-state and succeeds in mauling the non-malleable code. A more precise technical overview is given later in this section.

We are not able to make the above argument go through based on standard split-state non-malleable codes. We need the following additional properties described informally below. See Section 3 for a more precise description.

1. The code must be an *augmented* split-state non-malleable code [AAG⁺16]. This means that the distinguisher for the non-malleable code is given R as input, in addition to the tampered decoded message.
2. We need the code to be *conditionally* non-malleable. Intuitively, this means that non-malleability holds when L is chosen randomly along with (and independently from) the tampering functions (f, g) , and R is chosen randomly in the tampering experiment. This is in contrast to the usual security game where (f, g) are fixed and then (L, R) are both drawn during the tampering experiment. We define this new property formally in Section 3.1.
3. The code must satisfy what we call the *simulatable right state* property. Roughly speaking, this means that given a random L, the sets $\{R : \text{Dec}(L, R) = v\}$ and $\{R : \text{Dec}(L, R) = v'\}$ should be indistinguishable for all v, v' . See Section 3.2 for more details.

Our construction is based on the recent split state non-malleable code of Aggarwal et. al [ADL14]. It turns out that the code in [ADL14] already satisfies the first two of the above properties. We then present a modification to add the simulatable right state property. Note that the code of [ADL14] is purely information theoretic, and, in comparison to cryptographic objects (such as commitments or one-way functions), very efficient. Encoding and decoding simply requires sampling random vectors and taking their inner product, etc. To add the strong hiding property, we add a commitment and a symmetric encryption to the encoding and decoding procedures. Thus, our overall basic protocol has computation which is dominated by three invocations of a statistically binding commitment scheme (or rather two invocation of a statistically binding commitment and one symmetric encryption).




Figure 1: Protocol with Man-in-the-Middle

Overview of the Proof of Non-Malleability. Our protocol (under a man-in-the-middle attack) is given in Figure 1. For the time being, think of $(L, R) \leftarrow \text{Enc}(v)$ where $L, R \in \mathbb{Z}_q$, v is the value to which \mathcal{C} wishes to commit and (Enc, Dec) is a generic split-state non-malleable code. We will point out where each of the three above additional properties on (Enc, Dec) is needed. The first is easy to see. The distinguisher for non-malleable commitment gets the transcript (which contains R in the clear) as well as the decommitment \tilde{v} , so we need (Enc, Dec) to be augmented non-malleable.

As mentioned above, we prove non-malleability by reducing any mauling attack to a tampering function (f, g) which mauls the underlying code. This must be done carefully as on the one hand, we need to ensure that f and g tamper v correctly to M 's commitment \tilde{v} , on the other hand, (f, g) must be split-state, and so cannot naively run MIM using both L and R .

Both functions will obtain their output using the third message. They share a partial transcript consisting of the first two messages and the value a . Though this information contains L and so it shouldn't be given to g , L is computationally hidden so we will be able to hybrid it away later. $f(L)$ extracts \tilde{L} by choosing a random value R_\S and sending (a, R_\S) , receiving (\tilde{a}_\S, \cdot) ; then it rewinds M and asks another random challenge $\tilde{\beta}$ on the right, it receives β on the left and sends (b, R_\S) where $b = (a - L)(\beta/\alpha) + L$, and receives (\tilde{b}, \cdot) . It outputs \tilde{L} , the constant term on the line spanned by $\{(\tilde{\alpha}, \tilde{a}_\S), (\tilde{\beta}, \tilde{b})\}$. $g(R)$ also shares the random value R_\S and so can compute \tilde{a}_\S . $g(R)$ rewinds M and sends (a, R) on the left and receives (\tilde{a}, \tilde{R}) on the right. If $\tilde{a} = \tilde{a}_\S$, $g(R)$ outputs \tilde{R} , otherwise it outputs \perp .

Note that for (f, g) to succeed in extracting (\tilde{L}, \tilde{R}) , it must be that the answer \tilde{a}_\S M provides when given the random R_\S is equal to the \tilde{a} he provides given R . This will follow from the right state simulatability of (Enc, Dec) . Given this property, we can show that the chance that M answers correctly (*i.e.*, consistently with the linear map he committed to in the first round) given R_\S is about the same as the chance he answers correctly given R . So either both are incorrect with high probability, in which case M is always committing to $\tilde{v} = \perp$ and so cannot be mauling; or else both \tilde{a}_\S and \tilde{a} are correct with non-negligible probability. In this case, we can show that (f, g) succeed in extracting correct (\tilde{L}, \tilde{R}) with non-negligible probability. One subtle point is that (f, g) are defined using L . This means that even after we hybrid this dependence away, the mauling experiment for the resulting tampering functions will have L fixed. This is why we need (Enc, Dec) to be conditionally non-malleable.

Extension to non-synchronizing adversaries. While in some applications, security against synchronizing adversaries is all one needs (e.g., constructing round efficient multi-party computation), in others, non-malleability against arbitrary schedulings is required. Our basic protocol only provides security against synchronizing adversaries, and actually is susceptible to selective bot attacks against general schedulings. Since our protocol has only three rounds, we can enumerate over all possible schedulings and check that the only other potentially problematic scheduling is the sequential one where the left execution finishes entirely even before the right execution starts.

To extend to non-synchronizing adversaries, we make our protocol extractable by running a 3-round extractable (malleable) commitment scheme in parallel to our basic protocol (we do not need any "proofs of consistency" between the two parallel executions). Extraction immediately yields non-malleability against a sequential adversary as we may rewind M and extract his commitment without having to rewind the honest committer. The main technical challenge for this portion is proving non-malleability against a synchronizing adversary; *i.e.*, that the extractable commitment doesn't "interfere" with the basic protocol's synchronizing non-malleability. To achieve this, we will use an extractable commitment scheme such that extraction requires two rewinds instead of just one. This is inspired by a technique from the constant round protocol of Lin and Pass [LP11]. Our final protocol requires significantly more invocations of the underlying commitment scheme, however we stress that even so, it is significantly more efficient than any of the prior ones schemes in the literature (on top of requiring only 3-rounds of interaction). This protocol is described in Section 7.

A simple constant-round non-malleable commitment scheme from any one-way function. We now briefly discuss the ideas behind our simplest protocol. This protocol, and its proof, use only elementary techniques and are simple enough to be included in a graduate level course in cryptography.

The protocol has a simple “commit-and-prove” structure where commitments are executed using Naor’s commitment [Nao91]. Recall that Naor’s commitment uses a random string ρ sent by the receiver. If ρ has a special form, then the commitment is “equivocal” and can be opened to both 0 and 1; otherwise it is statistically binding. Let us denote this commitment by com_ρ . In our protocol ρ is obtained by simulatable coin-flipping.

Informally, the protocol between the committer \mathcal{C} and receiver \mathcal{R} is as follows:

1. **Sample ρ with coin-tossing:** \mathcal{C} commits to a random string ρ' and sends it to the receiver — this can be done using standard 2-round Naor’s commitment. \mathcal{R} responds with a random string ρ'' . \mathcal{C} sends $\rho = \rho' \oplus \rho''$ to \mathcal{R} and proves in (constant round) zero-knowledge that ρ is indeed correct. Parse $\rho = \rho_1 \parallel \rho_2 \parallel \rho_3$.
2. **Commit the first state:** \mathcal{C} encodes the message m (and identity id) using a standard split-state non-malleable code to obtain two states (L, R) . It then commits to L using ρ_1 as the first message of Naor: $c_1 = \text{com}_{\rho_1}(L)$.
3. **Start a proof of consistency.** Before sending the second state, \mathcal{C} and \mathcal{R} start a zero-knowledge proof-of-knowledge of “consistency” in which the statement to be proven is decided in the last step.

This is done by using the (standard) Feige-Shamir ZK protocol [FS89] which has constant rounds. Parties exchange *all but the last* message of this ZK protocol. Each commitment of this ZK protocol is implemented using a unique part of ρ_2 (as first message of Naor).

4. **Send second state and finish the proof.** \mathcal{C} now just need to send the second state R and prove that it is consistent with the value in c_1 . We do this slightly differently: \mathcal{C} commits to R using the third part ρ_3 , i.e., $c_2 = \text{com}_{\rho_3}(R)$ and proves that (c_1, c_2) are commitments to consistent states by sending the last message of the proof. It then *opens* c_2 to R by sending appropriate decommitment. \mathcal{R} simply checks that the opening of c_2 to R is valid and the proof verifies before accepting the commitment.

The proof crucially relies on the fact that Naor’s commitment is equivocal if ρ can be set appropriately. Specifically, we first “cheat” in the first step where ρ is set to be the special string which allows equivocation. This is done by using the simulator of zero-knowledge proof in the first step. For now, assume that the adversary is synchronous. This step ensures that there is no information about the states in the commitments or the proof transcripts.

We now construct split state functions by relying on the equivocality of Naor’s commitment. Specifically, our functions will share the same random tape ϕ and will be denoted by (f_ϕ, g_ϕ) for every ϕ . They will both sample a transcript using random values for the states (instead of actual L and R) up to the point where last message is to be sent. At this point, they will output corrupted states differently, as follows.

The function f_ϕ , upon receiving L , will use the equivocality of com w.r.t. strings ρ_1, ρ_2, ρ_3 to find appropriate randomness that is consistent with L and the transcript and construct a “honest committer” algorithm. It will then extract \tilde{L} from this machine by rewinding in the proof-of-knowledge part. During rewidings, it suffices to use random (incorrect) values for R ; it can be shown that as long as correct L is used in the extraction, we will obtain an appropriately distributed value for \tilde{L} on right.

The function g_ϕ , upon receiving R , will compute the last message of the proof just like function f_ϕ . However, it will open c_2 to the correct value R and feed it to MIM. If it receives a valid last message on the right with a state \tilde{R} , it will output that state (and \perp) otherwise.

We remark that even though the transcripts and the commitments generated in the description of (f_ϕ, g_ϕ) are not binding, the proof ensures that the values output by them are distributed appropriately as long as correct states are given to them as input. The asynchronous adversary is handled using a simple scheduling argument (without relying on non-malleable code). We present the full protocol and proof in Section 8.

2 Preliminaries

We use λ for the security parameter and $\text{negl}(\lambda)$ or negl for a function which tends to zero faster than λ^{-k} for any constant k . We say that two probability distributions X and Y are computationally indistinguishable, and write $X \approx_c Y$ if for all probabilistic polynomial time (PPT) distinguishers D ,

$$\left| \Pr_{x \leftarrow X}(D(x) = 1) - \Pr_{y \leftarrow Y}(D(y) = 1) \right| = \text{negl}.$$

2.1 Cryptographic Building Blocks

Symmetric Key Encryption. A symmetric key encryption scheme is a tuple of algorithms (G, E, D) : the key generation algorithm G takes input 1^λ and outputs a private key k ; the encryption algorithm E takes a key k , a plaintext message msg and randomness r as input and produces a ciphertext $\text{ct} = E_k(\text{msg}; r)$; and decryption takes a key k and a ciphertext ct and outputs $\text{msg} = D_k(\text{ct})$. We require

- **Correctness:** For any message msg , the process: $k \leftarrow G(1^\lambda)$, $\text{ct} \leftarrow E_k(\text{msg})$, $\text{msg}' = D_k(\text{ct})$ outputs $\text{msg}' = \text{msg}$ with probability 1.
- **Semantic Security:** For any messages msg, msg' , we have

$$\{\text{ct} : k \leftarrow G(1^\lambda), \text{ct} \leftarrow E_k(\text{msg})\} \approx_c \{\text{ct}' : k \leftarrow G(1^\lambda), \text{ct}' \leftarrow E_k(\text{msg}')\}.$$

It is known how to construct symmetric key encryption schemes from any one-way function.

Commitment Schemes. A commitment scheme, $\langle \mathcal{C}, \mathcal{R} \rangle$ is a two-phase, two party protocol between a committer \mathcal{C} and a receiver \mathcal{R} . In the commit phase, \mathcal{C} uses secret input v and interacts with \mathcal{R} who uses no input. Let $z = \text{Com}(v; r)$ denote \mathcal{R} 's view after the commit phase. Let $(w, v) = \text{Decom}(z, v, r)$ denote \mathcal{R} 's view after the decommit phase, which \mathcal{R} either accepts or rejects. We say that $\langle \mathcal{C}, \mathcal{R} \rangle$ is a statistically binding commitment scheme if the following properties hold:

- **Correctness:** If parties follow the protocol, then $R(z, w, v) = 1$;
- **Binding:** With high probability over \mathcal{R} 's randomness, there does not exist a (w', v') with $v' \neq v$ such that $R(z, w', v') = 1$;
- **Hiding:** For all $v, v', \{\text{Com}(v; r)\}_r \approx_c \{\text{Com}(v'; r')\}_{r'}$.

It is known how to construct a non-interactive, perfectly binding commitment scheme from any one-way permutation [Blu81]. Alternatively, any one-way function can be used to construct a two-round, statistically binding commitment scheme, where the binding property holds with high probability over \mathcal{R} 's randomness [Nao91].

Id-based Commitment Scheme. Following [PR05b, DDN91], we consider id-based commitment schemes where, in addition to the security parameter, the committer and the receiver also receive an identity $\text{id} \in \{0, 1\}^\lambda$ as common input.

2.2 Non-malleable commitments

Non-malleable commitment is defined using the real/ideal paradigm. In the real interaction, there is a man-in-the-middle adversary M interacting with a committer, \mathcal{C} , in the left session a receiver \mathcal{R} in the right. We denote the various quantities associated with the right interaction as “tilde’d” versions of their left counterparts. So for example, \mathcal{C} commits to v in the left interaction while M commits to \tilde{v} in the right. Let MIM_v denote a random

variable that describes (VIEW, \tilde{v}) , consisting of M 's view in the experiment and the value M commits to in the right interaction, given that C has committed to v in the left interaction. The ideal interaction is the same, except that C commits to an arbitrary fixed value, say 0, on the left. Let MIM_0 be the random variable describing (VIEW, \tilde{v}) in this interaction. We will use id-based commitment schemes and we force M to use an identity id on the right which is distinct from id used on the left. We enforce this by stipulating that MIM_v and MIM_0 output the special symbol \perp_{id} when M has used the same identity on the right which he has received on the left. This is analogous to the uninteresting case when M is simply acting as a channel, forwarding messages from C to \mathcal{R} and back. We let $\text{MIM}_v(y)$ and $\text{MIM}_0(y)$ be the distributions where M gets a string y as auxiliary input.

Definition 1 (Non-Malleable Commitments). *A commitment scheme $\langle C, \mathcal{R} \rangle$ is non-malleable if for every PPT man-in-the-middle adversary M , and for all v , we have $\{\text{MIM}_v(y)\}_{y \in \{0,1\}^*} \approx_c \{\text{MIM}_0(y)\}_{y \in \{0,1\}^*}$.*

On the Presence of Identities. In the real world, nothing can be done to prevent a MIM from simply acting as a channel between C and \mathcal{R} , forwarding messages back and forth. Therefore, the most basic definition of non-malleable commitment requires that the MIM cannot distinguish whether he is in the real or ideal world, even when he is given his committed value on the right, as long as he is doing something (anything) other than copying. [DDN91] noticed that this is equivalent to requiring every committer to have a public identity and forcing the MIM to use an identity which is distinct from C 's. This observation has persisted ever since, so non-malleable commitment is usually defined with respect to identities. The equivalence is based on the following observation: C can choose his identity to be the public verification key for a signature and send a signature of the entire transcript in the last round. The security of the signature scheme then ensures that either M has copied C 's identity, in which case he cannot produce a valid signature unless he has copied every message as well, or he has used a distinct identity.

2.3 Non-Malleable Codes

A coding scheme is a pair of functions (Enc, Dec) where $\text{Enc} : \mathcal{M} \rightarrow \mathcal{C}$ and $\text{Dec} : \mathcal{C} \rightarrow \mathcal{M}$ for a message space \mathcal{M} and codeword space \mathcal{C} . It should be the case that $\text{Dec} \circ \text{Enc}(m) = m$ for all $m \in \mathcal{M}$ with high probability over the randomness of Enc (it needn't be the case that Enc is randomized at all, in which case correctness requires $\text{Dec} \circ \text{Enc}(m) = m$ with probability 1). Historically, coding schemes are usually designed in order to be resilient to some form of tampering. In their important 2010 paper Dziembowski, Pietrzak and Wichs [DPW10] introduced non-malleable codes which are codes with strong security in the presence of tampering. Informally, for a family $\mathcal{F} \subset \{f : \mathcal{C} \rightarrow \mathcal{C}\}$, we say that (Enc, Dec) is non-malleable with respect to \mathcal{F} if for all $f \in \mathcal{F}$, the tamper distribution $(\text{Dec} \circ f \circ \text{Enc})(m)$ (over the randomness of Enc) outputs \tilde{m} which is either equal to m if copying or else is independent of m . In this work we are interested in split-state non-malleable codes.

Let (Enc, Dec) be a split state coding scheme so $\text{Enc} : \mathcal{M} \rightarrow \mathcal{L} \times \mathcal{R}$ and let

$$\mathcal{F}_{\text{split}} = \{(f, g) \mid f : \mathcal{L} \rightarrow \mathcal{L}, g : \mathcal{R} \rightarrow \mathcal{R}\}$$

be the set of split-state tampering functions.

Definition 2 (Tampering Distribution). *Fix $m \in \mathcal{M}$, $(f, g) \in \mathcal{F}_{\text{split}}$. The tampering distribution, denoted $\mathbb{T}_{m,f,g}$ is: draw $(L, R) \leftarrow \text{Enc}(m)$, set $(\tilde{L}, \tilde{R}) = (f(L), g(R))$ and output $\tilde{m} = \text{Dec}(\tilde{L}, \tilde{R})$.*

Definition 3 (Simulatable Distribution). *Let $\{D_m\}_{m \in \mathcal{M}}$ be a family of distributions on \mathcal{M} indexed by m . We say that $\{D_m\}$ is ε -simulatable if there exists a distribution S on $\mathcal{M} \cup \{\text{same}\}$ such that $\Delta(D_m, S_m) < \varepsilon$ for all m , where S_m is the distribution on \mathcal{M} induced by drawing $\tilde{m} \leftarrow S$ and outputting m if $\tilde{m} = \text{same}$, \tilde{m} if not.*

Definition 4 (Split-State Non-Malleable Code). *We say that (Enc, Dec) is ε -non-malleable against $\mathcal{F}_{\text{split}}$ if for all $(f, g) \in \mathcal{F}_{\text{split}}$, $\{\mathbb{T}_{m,f,g}\}_m$ is ε -simulatable.*

The Non-Malleable Code of [ADL14]. The code of [ADL14] maps $m \in \mathcal{M}$ into (L, R) where $L, R \in \mathbb{Z}_p^n$ are random subject to the condition that $\langle L, R \rangle \in H_m \subset \mathbb{Z}_p$ (p is a prime much larger than $|\mathcal{M}|$, and the $\{H_m\}_{m \in \mathcal{M}}$ are carefully chosen disjoint subsets of \mathbb{Z}_p). Non-malleability follows from an extensive analysis of the inner product function which makes heavy use of its properties as a randomness extractor. For any $(f, g) \in \mathcal{F}_{\text{split}}$ and $x \in \mathbb{Z}_p$, the following random process is considered: choose $L, R \in \mathbb{Z}_p^n$ randomly such that $\langle L, R \rangle = x$, set $(\tilde{L}, \tilde{R}) = (f(L), g(R))$, and output $\tilde{x} = \langle \tilde{L}, \tilde{R} \rangle$. The main lemma of [ADL14] says that \tilde{x} is either 1) independent of x , or 2) of the form $\tilde{x} = ax + b$ for some $a, b \in \mathbb{Z}_p$ which depend only on (f, g) . Non-malleability in [ADL14] then follows from the design of *affine evasive sets* as the $\{H_m\}_m$. This aspect of their construction is very elegant but as we will not need to change their $\{H_m\}_m$, we do not discuss this portion further. The interested reader should see [ADL14] for more information. We note that the earlier work of [DKO13] used essentially the same outline in order to give a non-malleable code for one bit messages. Their construction is also very elegant and is much simpler: they use $H_0 = \{0\}$ and $H_1 = \mathbb{Z}_p - \{0\}$.

2.4 Augmented Non-Malleable Codes

Very recently Aggarwal et al. [AAG⁺16] proved that the [ADL14] construction is non-malleable even when the tamper distribution outputs \tilde{m} along with one of the states. Their proof looks at the randomized process: choose $L, R \in \mathbb{Z}_p^n$ randomly such that $\langle L, R \rangle = x$, set $(\tilde{L}, \tilde{R}) = (f(L), g(R))$ and output (R, \tilde{x}) where $\tilde{x} = \langle \tilde{L}, \tilde{R} \rangle$. The same randomness extraction properties of the inner product function used in [ADL14] show that even conditioned on R , \tilde{x} is either independent of x or else $\tilde{x} = ax + b$ for $a, b \in \mathbb{Z}_p$ which depend only on (f, g) . They call this stronger notion augmented non-malleability.

Definition 5 (Augmented Tampering Distribution). Fix $m \in \mathcal{M}$ and $(f, g) \in \mathcal{F}_{\text{split}}$. The augmented tampering distribution, denoted $\mathbb{V}_{m, f, g}$ is: draw $(L, R) \leftarrow \text{Enc}(m)$, set $(\tilde{L}, \tilde{R}) = (f(L), g(R))$ and output (R, \tilde{m}) where $\tilde{m} = \text{Dec}(\tilde{L}, \tilde{R})$.

We use the letter \mathbb{V} for “view”: the output of the augmented tampering distribution will be basically what the MIM sees during a mauling attack on our non-malleable commitment scheme.

Definition 6 (Augmented Simulatable Distribution). Let $\{D_m\}_{m \in \mathcal{M}}$ be a family of distributions on $\mathcal{R} \times \mathcal{M}$ indexed by m , where \mathcal{R} is an arbitrary set. We say that $\{D_m\}$ is ε -augmented simulatable if there exists a distribution S on $\mathcal{R} \times (\mathcal{M} \cup \{\text{same}\})$ such that $\Delta(D_m, S_m) < \varepsilon$ for all m , where S_m is the distribution on \mathcal{M} induced by drawing $(R, \tilde{m}) \leftarrow S$ and outputting (R, m) if $\tilde{m} = \text{same}$, (R, \tilde{m}) if not.

Definition 7 (Augmented Non-Malleable Code). We say that (Enc, Dec) is ε -augmented non-malleable against $\mathcal{F}_{\text{split}}$ if for all $(f, g) \in \mathcal{F}_{\text{split}}$, $\{\mathbb{V}_{m, f, g}\}_m$ is ε -augmented simulatable.

Before moving on, we remark that the proof in [AAG⁺16] actually shows something slightly stronger. Let us think of the randomized process above as drawing $L \leftarrow \mathbb{Z}_p^n$ at random and then drawing $R \leftarrow \{\mathbf{v} \in \mathbb{Z}_p^n : \langle L, \mathbf{v} \rangle = x\}$, computing (\tilde{L}, \tilde{R}) and outputting (R, \tilde{x}) . The analysis of [AAG⁺16] does not actually require L to be uniform in \mathbb{Z}_p^n and works whenever L has sufficient min-entropy. In particular, given $n, p, \mathcal{M}, \{H_m\}_m$ as in [ADL14] and any sufficiently large $\mathcal{L} \subset \mathbb{Z}_p^n$, define the coding scheme:

- $\text{Enc}^{\mathcal{L}}(m)$: choose $L \leftarrow \mathcal{L}$ and $R \leftarrow \mathbb{Z}_p^n$ randomly such that $\langle L, R \rangle \in H_m$.
- $\text{Dec}(L, R)$: if $\langle L, R \rangle \in H_m$, output m , otherwise output \perp .

Claim 1. For all \mathcal{M} , there exist $n, p = \text{poly}(|\mathcal{M}|, \lambda)$ such that for all $\mathcal{L} \subset \mathbb{Z}_p^n$ of size at least $|\mathcal{L}| \geq p^{\eta n}$, $(\text{Enc}^{\mathcal{L}}, \text{Dec})$ is $2^{-\Omega(\lambda)}$ -augmented non-malleable, where $1 - \eta = c \log^{-6} p$ for an absolute constant $c > 0$.

Moreover, the simulator for $(\text{Enc}^{\mathcal{L}}, \text{Dec})$ is identical to the simulator for (Enc, Dec) except that it draws $L \leftarrow \mathcal{L}$ instead of $L \leftarrow \mathbb{Z}_p^n$. Claim 1 follows from the proof of the main theorem in [AAG⁺16]; as pointed out to us by Aggarwal in a personal communication [Agg].

3 New Constructions of Non-Malleable Codes

3.1 Conditional Augmented Non-Malleable Codes

Let (Enc, Dec) be a split-state code with codeword space $\mathcal{L} \times \mathcal{R}$. In proving that our commitment scheme is non-malleable, we will need to choose a random $L \in \mathcal{L}$ and be ensured that the augmented tampering distribution is independent of m even conditioned on L . We define information theoretic and computational variants of these codes.

Definition 8 (Conditional Augmented Tampering Distribution). Fix $m \in \mathcal{M}$, $(f, g) \in \mathcal{F}_{\text{split}}$, and $L \in \mathcal{L}$. The conditional augmented tampering distribution, $\mathcal{V}_{m,f,g}^L$ is: draw $R \leftarrow \text{Enc}(m|L) = \{R' \in \mathcal{R} : \text{Dec}(L, R') = m\}$, set $(\tilde{L}, \tilde{R}) = (f(L), g(R))$ and output (R, \tilde{m}) where $\tilde{m} = \text{Dec}(\tilde{L}, \tilde{R})$.

Definition 9 (Conditional Augmented Simulatable Distribution). Let $\{\mathcal{D}_m^L\}_{m,L}$ be a family of distributions on $\mathcal{R} \times \mathcal{M}$ indexed by $m \in \mathcal{M}$ and $L \in \mathcal{L}$, where \mathcal{L} and \mathcal{R} are arbitrary sets. We say that $\{\mathcal{D}_m^L\}_{m,L}$ is ε -conditionally augmented simulatable if there exists a family of distributions $\{\mathcal{S}^L\}_L$ on $\mathcal{R} \times (\mathcal{M} \cup \{\text{same}\})$ which such that

$$\Pr_L \left[\Delta(\mathcal{D}_m^L, \mathcal{S}_m^L) < \varepsilon \forall m \in \mathcal{M} \right] \geq 1 - \varepsilon,$$

where the probability is over $L \leftarrow \mathcal{L}$ drawn uniformly and where \mathcal{S}_m^L draws $(R, \tilde{m}) \leftarrow \mathcal{S}^L$ and outputs (R, m) if $\tilde{m} = \text{same}$, (R, \tilde{m}) if not. We say $\{\mathcal{D}_m^L\}_{m,L}$ is computationally conditionally augmented simulatable if for all PPT distinguishers D and non-negligible $\delta > 0$, there exist simulators $\{\mathcal{S}^L\}_L$ such that

$$\Pr_L \left[\exists m \in \mathcal{M} \text{ st } \left| \Pr_{(R, \tilde{m}) \leftarrow \mathcal{D}_m^L} (D(R, \tilde{m}) = 1) - \Pr_{(R, \tilde{m}) \leftarrow \mathcal{S}_m^L} (D(R, \tilde{m}) = 1) \right| > \delta \right] = \text{negl}.$$

Definition 10 (Conditional Augmented Non-Malleable Code). We say that (Enc, Dec) is ε -conditionally augmented non-malleable (resp. computationally conditionally augmented non-malleable) against $\mathcal{F}_{\text{split}}$ if for all $(f, g) \in \mathcal{F}_{\text{split}}$, $\{\mathcal{V}_{m,f,g}^L\}_{m,L}$ is ε -conditionally augmented simulatable (resp. computationally conditionally augmented simulatable).

We use a simple probability argument to show that the code from [ADL14] is conditionally augmented non-malleable. This is reminiscent of the way in which one argues that a sufficiently good two-source extractor is also a strong two-source extractor.

Claim 2. The code (Enc, Dec) of [ADL14] is ε' -conditionally augmented non-malleable for some negligible quantity $\varepsilon' > 0$.

Proof. Given $(f, g) \in \mathcal{F}_{\text{split}}$, the simulator $\mathcal{S}_{f,g}$ guaranteed by the augmented non-malleability of (Enc, Dec) behaves as follows: draw $L, R \leftarrow \mathbb{Z}_p^n$, set $(\tilde{L}, \tilde{R}) = (f(L), g(R))$ and output $(R, \langle \tilde{L}, \tilde{R} \rangle)$ unless $\langle \tilde{L}, \tilde{R} \rangle = \langle L, R \rangle$, in which case output (R, same) . We define the family $\{\mathcal{S}_{f,g}^L\}_L$ of simulators similarly: $\mathcal{S}_{f,g}^L$ draws $R \leftarrow \mathbb{Z}_p^n$ at random and outputs $(R, \langle \tilde{L}, \tilde{R} \rangle)$ or (R, same) according to whether $\langle \tilde{L}, \tilde{R} \rangle$ is distinct from or equal to $\langle L, R \rangle$. Define

$$\mathcal{L}_{\text{bad}} = \left\{ L \in \mathbb{Z}_p^n : \exists m \in \mathcal{M} \text{ st } \Delta(\mathcal{V}_{m,f,g}^L, \mathcal{S}_{m,f,g}^L) > \varepsilon' \right\},$$

where $\varepsilon' > 0$ will be specified later. Set $\zeta = p^{-(1-\eta)n}$ for $1 - \eta = c \log^{-6} p$ for an absolute constant $c' > 0$. If $|\mathcal{L}_{\text{bad}}| < \zeta p^n$ then we are done so assume $|\mathcal{L}_{\text{bad}}| \geq \zeta p^n$. By Claim 1, the restricted code $(\text{Enc}^{\mathcal{L}_{\text{bad}}}, \text{Dec})$ is $2^{-\Omega(\lambda)}$ -augmented non-malleable with simulator $\mathcal{S}_{f,g}^{\mathcal{L}_{\text{bad}}}$ identical to that for (Enc, Dec) except that the initial choices of L, R are $L \leftarrow \mathcal{L}_{\text{bad}}, R \leftarrow \mathbb{Z}_p^n$. However, this cannot be: by definition of \mathcal{L}_{bad} , there exists some $m \in \mathcal{M}$ such that $\Delta(\mathcal{V}_{m,f,g}^L, \mathcal{S}_{m,f,g}^L) = 2^{-o(\lambda)}$; so $\mathcal{S}_{f,g}^L$ does not simulate $\{\mathcal{V}_{m,f,g}^L\}_m$. Therefore, it must be that $\Pr_L [L \in \mathcal{L}_{\text{bad}}] < \zeta < \varepsilon'$ and so (Enc, Dec) is ε' -conditionally augmented non-malleable (we have chosen ε' larger than ζ but still negligible). \square

3.2 Adding the Hiding Property

We will also need our non-malleable code to have a computational hiding property resembling semantic security in order to rule out selective \perp attacks. We formalize the property we need using a game between a challenger \mathcal{C} and a PPT adversary \mathcal{A} , parametrized by $N = \text{poly}(\lambda)$, a message $m \in \mathcal{M}$ and a distribution \mathcal{D}_{hid} on \mathcal{R} .

- \mathcal{C} draws $L \leftarrow \mathcal{L}$, $b \leftarrow \{0, 1\}$, sends R_1, \dots, R_N to \mathcal{A} where $R_i \leftarrow \text{Enc}(m|L)$ if $b = 0$; $R_i \leftarrow \mathcal{D}_{\text{hid}}$ if $b = 1$.
- \mathcal{A} outputs b' and wins if $b' = b$.

Definition 11 (Codes with Simulatable State). We say that a split-state code (Enc, Dec) has simulatable right state if there is a distribution \mathcal{D}_{hid} on \mathcal{R} such that for all PPT \mathcal{A} , $N = \text{poly}(\lambda)$, and $m \in \mathcal{M}$, the probability that \mathcal{A} wins the above game is at most $1/2 + \text{negl}$.

Hiding Game Variant. We will use another game parametrized by polynomials $N, N' = \text{poly}(\lambda)$. In this game \mathcal{A} sends \mathcal{C} two messages $m, m' \in \mathcal{M}$, \mathcal{C} draws a secret $L \leftarrow \mathcal{L}$ and sends \mathcal{A} the tuple $(R, \{R_1\}, \dots, \{R_N\})$ where $R \leftarrow \text{Enc}(m|L)$ and each set has N' elements. Moreover, $R_i \leftarrow \text{Enc}(m'|L)$ for all $R_i \in \{R_i\}$ and all $i = 1, \dots, N$ except for one random i^* for which $R_{i^*} \leftarrow \text{Enc}(m|L)$ for all $R_{i^*} \in \{R_{i^*}\}$; \mathcal{A} tries to guess i^* . If (Enc, Dec) has simulatable right state then a PPT adversary \mathcal{A} can guess i^* with probability at most $1/N + \text{negl}$.

Construction. Let $(\text{Enc}_0, \text{Dec}_0)$ be an ε -conditional augmented non-malleable code. Let (G, E, D) be a symmetric key encryption scheme, and let $(\text{Com}, \text{Decom})$ be a non-interactive, perfectly binding commitment scheme. The new coding scheme, (Enc, Dec) is defined as follows.

- $\text{Enc}(m)$: Draw $(L_0, R_0) \leftarrow \text{Enc}_0(m)$, $k \leftarrow G(1^\lambda)$, $\sigma \leftarrow \$$, and $c \leftarrow E_k(R_0)$. Set $z = \text{Com}(k, \sigma)$ and output (L, R) where $L = (L_0, (k, \sigma))$, $R = (c, z)$.
- $\text{Dec}(L, R)$: If either $L, R = \perp_{\text{com}}$ output \perp_{com} . Otherwise, parse $L = (L_0, (k, \sigma))$ and $R = (c, z)$, check that $\text{Decom}(z) = (k, \sigma)$. If so set $R_0 = D_k(c)$, output $\text{Dec}_0(L_0, R_0)$; if not output \perp_{com} .

Claim 3. (Enc, Dec) has simulatable right state.

Proof. Define three challengers $\mathcal{C}_0, \mathcal{C}_1, \mathcal{C}_2$ for the above hiding game, played with some fixed $m \in \mathcal{M}$. Each challenger draws $L = (L_0, (k, \sigma)) \leftarrow \mathcal{L}$ and for $i = 1, \dots, N$

- \mathcal{C}_0 sets $R_i = (c_i, z)$ where $z = \text{Com}(k, \sigma)$ and $c_i \leftarrow E_k(R_0)$ for some R_0 with $\text{Dec}_0(L_0, R_0) = m$.
- \mathcal{C}_1 sets $R_i = (c_i, z)$ where $z = \text{Com}(0)$ and $c_i \leftarrow E_k(R_0)$.
- \mathcal{C}_2 sets $R_i = (c_i, z)$ where $z = \text{Com}(0)$ and $c_i \leftarrow E_{k'}(0)$, where $k' \leftarrow G(1^\lambda)$.

Note that \mathcal{C}_0 draws each R_i from $\text{Enc}(m|L)$, whereas \mathcal{C}_2 draws R_i from a distribution which is independent of m and L ; call this distribution \mathcal{D}_{hid} . Moreover, \mathcal{A} cannot distinguish between his interaction with \mathcal{C}_0 and \mathcal{C}_1 by the hiding of Com . Likewise, he cannot distinguish between his interaction with \mathcal{C}_1 and \mathcal{C}_2 by semantic security. \square

Lemma 1. (Enc, Dec) is computationally conditionally augmented non-malleable against

$$\mathcal{F}_{\text{split}}^{\text{poly}} = \{(f, g) \in \mathcal{F}_{\text{split}} : f \text{ and } g \text{ polytime}\}.$$

Lemma 1 is proven in full in Section 6. In the remainder of this section we give a high-level overview.

Proof Idea. Fix $(f, g) \in \mathcal{F}_{\text{split}}^{\text{poly}}$. We describe a family of simulators for $\{V_{m,f,g}^L\}_{m,L}$. $S_{f,g}^L$ will be one of two distributions depending on L . The first simply draws $R \leftarrow \mathcal{D}_{\text{hid}}$ and outputs (R, \perp_{com}) . This will simulate $\{V_{m,f,g}^L\}_m$ whenever $p_{L,m}$ is small for all $m \in \mathcal{M}$ where $p_{L,m}$ is shorthand for $\Pr_{(R,\tilde{m}) \leftarrow V_{m,f,g}^L}(\tilde{m} \neq \perp_{\text{com}})$. In other words, if (f, g) always tampers to \perp_{com} , then non-malleability follows from right state simulatability.

On the other hand, if $p_{L,m}$ is large for all $m \in \mathcal{M}$ then $V_{m,f,g}^L$ is identical to $V_{m,f_0,g_0}^{L_0}$, a conditional augmented tampering distribution of $(\text{Enc}_0, \text{Dec}_0)$, for some tampering functions $(f_0, g_0) \in \mathcal{F}_{\text{split}}$ which are related to (f, g) . In this case, we can use $S_{f_0,g_0}^{L_0}$ to construct $S_{f,g}^L$.

The final piece of the proof involves ruling out selective bot attacks. We show that if the likelihood of (f, g) tampering to \perp_{com} depends on the encoded message (*i.e.*, if $p_{L,m}$ is significantly larger than $p_{L,m'}$ for some m, m'), then (f, g) can be used to win the above hiding game with non-negligible advantage. Since f and g are polytime, either $p_{L,m}$ is small for all m , or $p_{L,m}$ is large for all m and so we are in one of the cases above.

The first point (when $p_{L,m}$ is small for all $m \in \mathcal{M}$) is relatively straightforward, but we expand a bit on the second and third points. The functions (f_0, g_0) are defined as follows.

- **Random Choices:** Draw $k \leftarrow G(1^\lambda)$, $\sigma \leftarrow \mathcal{S}$, $c^\mathcal{S} \leftarrow E_k(0)$, set $z = \text{Com}(k, \sigma)$, and $(\cdot, \tilde{z}^\mathcal{S}) = g(c^\mathcal{S}, z)$.
- $f_0(L_0)$: Compute $(\tilde{L}_0, (\tilde{k}, \tilde{\sigma})) = f(L_0, (k, \sigma))$. Output \tilde{L}_0 .
- $g_0(R_0)$: Draw $c \leftarrow E_k(R_0)$ and set $(\tilde{c}, \tilde{z}) = g(c, z)$. If $\tilde{z} \neq \tilde{z}^\mathcal{S}$, output \perp_{com} . Otherwise, use superpolynomial time to break open \tilde{z} and recover the pair $(\tilde{k}', \tilde{\sigma}')$. Output $\tilde{R}_0 = D_{\tilde{k}'}(\tilde{c})$.

Remark. Note that g_0 above does not run in polynomial time. It is possible to change the construction and get the proof to work using polynomial time (f_0, g_0) , however the proof would become longer and more difficult.

The main observation is that whenever $(L, c^\mathcal{S})$ are such that $\text{Decom}(\tilde{z}^\mathcal{S}) = (\tilde{k}, \tilde{\sigma})$, the distributions $V_{m,f,g}^L$ and $V_{m,f_0,g_0}^{L_0}$ are the same (up to encrypting the right state output by $V_{m,f_0,g_0}^{L_0}$). In this case, we can simulate $V_{m,f,g}^L$ using $S_{f_0,g_0}^{L_0}$, so it suffices to show that $c^\mathcal{S}$ exists such that $\text{Decom}(\tilde{z}^\mathcal{S}) = (\tilde{k}, \tilde{\sigma})$. This follows from the semantic security of (G, E, D) as $(c^\mathcal{S}, z)$ and $R = (c, z) \leftarrow \text{Enc}(m|L)$ differ only in their encrypted values.

We prove the claim of the third point above, that $p_{L,m}$ and $p_{L,m'}$ cannot differ very much, as it exhibits a type of argument which will be very useful to us in the rest of the paper.

The Hiding Machine. We mention here one technique which will be very useful to us moving forward. The cryptography we have inserted into the right side of our code gives us a way to rule out certain tampering behavior via reductions to computational security. We refer to such arguments as “the hiding machine”. At its core, the hiding machine is just a reduction to one of the hiding games above but there are many moving parts and so things often get quite complicated. Claim 4 below is a particularly simple example of the hiding machine in action and makes for a good first encounter. In general, the hiding machine is useful in ruling out various selective bot attacks.

Claim 4. For any $m, m' \in \mathcal{M}$, $(f, g) \in \mathcal{F}_{\text{split}}^{\text{poly}}$ and non-negligible $\xi > 0$ we have $\Pr_L \left[|p_{L,m} - p_{L,m'}| > \xi \right] = \text{negl}$.

Proof. Fix $m, m' \in \mathcal{M}$, $(f, g) \in \mathcal{F}_{\text{split}}^{\text{poly}}$ and non-negligible $\xi = \xi(\lambda) > 0$. Let BAD be the set of $L \in \mathcal{L}$ such that $p_{L,m} > \xi + p_{L,m'}$, and suppose for contradiction that there is a non-negligible $\xi' = \xi'(\lambda) > 0$ such that $\Pr_L [L \in \text{BAD}] \geq \xi'$. Fix $N = 3/(\xi\xi')$, $N' = \Omega(\lambda\xi^{-2})$ and consider the PPT adversary \mathcal{A} who plays the (N, N') -way hiding game for (Enc, Dec) against a challenger \mathcal{C} as follows.

- \mathcal{A} sends m, m' to \mathcal{C} and receives $(R, \{R_1\}, \dots, \{R_N\})$.
- \mathcal{A} computes $(\tilde{c}, \tilde{z}) = g(R)$ and $(\tilde{c}_i, \tilde{z}_i) = g(R_i)$ for each $i = 1, \dots, N$ and $R_i \in \{R_i\}$.

- For $i = 1, \dots, N$, \mathcal{A} sets $p_i = \Pr_{R_i \in \{\mathcal{R}_i\}}(\tilde{z}_i = \tilde{z})$ and outputs i^* such that p_{i^*} is maximal.

Note that if the random secret $L \in \mathcal{L}$ chosen by \mathcal{C} is in BAD then $\text{Decom}(\tilde{z}) = (\tilde{k}, \tilde{\sigma})$ with probability at least ξ , where $f(L) = (\tilde{L}_0, (\tilde{k}, \tilde{\sigma}))$. If $\text{Decom}(\tilde{z}) = (\tilde{k}, \tilde{\sigma})$, then for each i and $R_i \in \{\mathcal{R}_i\}$, $\text{Dec}(\tilde{L}, \tilde{R}_i) \neq \perp_{\text{com}}$ if and only if $\tilde{z}_i = \tilde{z}$. In this case, p_i approximates p_{L, m_i} where $m_i = m'$ if $i \neq i^*$ and $m_{i^*} = m$. Therefore,

$$p_{i^*} \geq p_{L, m} - \frac{\xi}{3} > p_{L, m'} + \xi - \frac{\xi}{3} \geq p_i + \frac{\xi}{3},$$

for all $i \neq i^*$. We have used the Chernoff-Hoeffding bound, facilitated by our choice of large N' . So we see that

$$\Pr(\mathcal{A} \text{ wins}) \geq \Pr(L \in \text{BAD}) \Pr(\text{Decom}(\tilde{z}) = (\tilde{k}, \tilde{\sigma}) | L \in \text{BAD}) (1 - \text{negl}) = \xi \xi' - \text{negl} > \frac{2}{N},$$

and so \mathcal{A} breaks the right state simulatability of (Enc, Dec) . \square

4 The Basic Protocol

The protocol is shown in Figure 2.

Setup: Let Com be a non-interactive, perfectly binding commitment scheme. Let (Enc, Dec) be a computational, conditional, augmented non-malleable code. Fix a large prime q . Let $\text{id} \in \{0, 1\}^\lambda$ be \mathcal{C} 's identity.

Committer's Private Input: $v \in \mathcal{M}_{\langle \mathcal{C}, \mathcal{R} \rangle_{\text{synch}}}$ to be committed to.

Commit Phase:

1. $\mathcal{C} \rightarrow \mathcal{R}$: Set $m = v \circ \text{id}$ and draw $(L, R) \leftarrow \text{Enc}(m)$, where $L \in \mathcal{L} \subset \mathbb{Z}_q$. Choose random $r \in \mathbb{Z}_q$ and send $\text{Com}(L \circ r)$ to \mathcal{R} .
2. $\mathcal{R} \rightarrow \mathcal{C}$: Send random challenge $\alpha \in \mathbb{Z}_q^*$.
3. $\mathcal{C} \rightarrow \mathcal{R}$: Send response $a = r\alpha + L \in \mathbb{Z}_q$ and also send R .

Decommit Phase:

1. $\mathcal{C} \rightarrow \mathcal{R}$: Open the commitment sent in step 1. Let $L' \circ r' \in \mathbb{Z}_q$ be the decommitted value.

Receiver's Output: If L' and r' do not satisfy $r'\alpha + L' = a$ then output the special symbol \perp_{inc} . Otherwise, compute $m' = \text{Dec}(L', R)$ and parse $m' = v' \circ \text{id}'$. Output v' if $\text{id}' = \text{id}$, \perp_{id} if not.

Figure 2: Non-malleable commitment scheme $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{synch}}$.

Claim 5. $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{synch}}$ is a perfectly binding commitment scheme.

Proof Sketch. Perfect binding follows immediately from the perfect binding of Com. Computational hiding follows in a straightforward fashion from the hiding of Com and the well known fact that any split-state non-malleable code is also a 2-out-of-2 secret sharing scheme. \square

Theorem 1. $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{synch}}$ is non-malleable against a synchronizing adversary.

A Hiding Game For $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{synch}}$. Before proving Theorem 1, we specify a hiding game for $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{synch}}$, analogous to the hiding game for (Enc, Dec). Consider the following interaction between a challenger \mathcal{C} and a PPT adversary \mathcal{A} .

- **Partial Transcript:** \mathcal{C} chooses $L \leftarrow \mathcal{L}$, $r \leftarrow \mathbb{Z}_q$ and sends $\mathbf{Com}_L = (\text{Com}(L), \text{Com}(r))$ to \mathcal{A} , \mathcal{A} returns $\alpha \in \mathbb{Z}_q$ and receives $a = r\alpha + L$ from \mathcal{C} .
- **Message Choice:** \mathcal{A} chooses $m, m' \in \mathcal{M}$ and sends them to \mathcal{C} .
- **Challenge Message:** \mathcal{C} chooses $b \leftarrow \{0, 1\}$ and sends the pair (R_0, R_1) to \mathcal{A} where $R_b \leftarrow \text{Enc}(m|L)$ and $R_{1-b} \leftarrow \text{Enc}(m'|L)$.
- **Guess:** \mathcal{A} outputs a guess $b' \in \{0, 1\}$ and wins if $b' = b$.

Just as for the hiding game of the code (Enc, Dec), we will usually use an (N, N') -way variant of the above game, where the challenge message is $(R, \{R_1\}, \dots, \{R_N\})$ where $R \leftarrow \text{Enc}(m|L)$, each $\#\{R_i\} = N'$, and also $R_i \leftarrow \text{Enc}(m'|L)$ for all $R_i \in \{R_i\}$ and all i except for a random i^* , for which $R_{i^*} \leftarrow \text{Enc}(m|L)$ for all $R_{i^*} \in \{R_{i^*}\}$. In the (N, N') -way variant, \mathcal{A} wins if he guesses i^* .

Claim 6. *If Com is computationally hiding and (Enc, Dec) has the hiding property then for all PPT adversaries \mathcal{A} , the probability that \mathcal{A} wins the above game (resp. its (N, N') -way variant) is at most $1/2 + \text{negl}$ (resp. $1/N + \text{negl}$).*

5 Proof of Non-Malleability (Theorem 1)

5.1 Notation

Transcripts. Suppose a PPT man-in-the-middle, M , participates in two protocol executions. We denote the transcript of M 's view with the letter \mathbb{T} . So

$$\mathbb{T} = (\text{id}, \tilde{\text{id}}, \text{Com}(L), \text{Com}(r), \text{Com}(\tilde{L}), \text{Com}(\tilde{r}), \tilde{\alpha}, \alpha, a, R, \tilde{a}, \tilde{R}).$$

We write \mathbf{Com}_L for $\text{Com}(L)$ and $\text{Com}(r)$. Note \mathbf{Com}_L specifies a linear polynomial $\varphi(x) = rx + L$ which \mathcal{C} uses to answer M 's query α . We will usually write a transcript \mathbb{T} more concisely as $\mathbb{T} = (\mathbf{Com}_L, \tilde{\alpha}, a, R)$, suppressing \mathcal{C} 's identity id and the quantities which are outputs of M . Since without loss of generality M is deterministic, these values uniquely define a full transcript.

Partial Transcripts. We also will find it useful to speak of partial transcripts, as this will let us isolate certain random choices made during the execution of \mathbb{T} . We use τ to denote the partial transcript where \mathcal{C} 's value R remains unspecified. We write τ concisely as $\tau = (\mathbf{Com}_L, \tilde{\alpha}, a)$. Note that M 's third message is not specified given τ , however τ extends to a full transcript \mathbb{T} once R is chosen. We write this full transcript $\mathbb{T}(\tau, R)$.

The Distribution M_m^τ and Distinguisher D^τ . Our goal is to use a MIM who breaks the non-malleability of $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{synch}}$ to violate the security of the code (Enc, Dec). In order to do this we make some notational changes which syntactically relate M to the code's non-malleability game. By definition, if M breaks the non-malleability of $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{synch}}$ then there exists $v \in \mathcal{M}_{\langle \mathcal{C}, \mathcal{R} \rangle_{\text{synch}}}$, a PPT distinguisher D , and non-negligible $\delta = \delta(\lambda) > 0$ such that

$$\left| \Pr_{(\mathbb{T}, \tilde{v}) \leftarrow \text{MIM}_v} (D(\mathbb{T}, \tilde{v}) = 1) - \Pr_{(\mathbb{T}, \tilde{v}) \leftarrow \text{MIM}_0} (D(\mathbb{T}, \tilde{v}) = 1) \right| = \delta. \quad (1)$$

Let $m = v \circ \text{id}$, $m' = 0 \circ \text{id}$ and $\tilde{m} = \tilde{v} \circ \tilde{\text{id}}$. So $m, m' \in \mathcal{M}$ are the messages encoded during the left executions of $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{synch}}$ in the real/ideal world, and $\tilde{m} \in \mathcal{M}$ is the message encoded on the right. For a given partial transcript

$\tau = (\mathbf{Com}_L, \tilde{\alpha}, a)$, let M_m^τ be the distribution which draws $R \leftarrow \text{Enc}(m|L)$ and outputs (R, \tilde{m}) , where \tilde{m} is M 's encoded message in $\mathbb{T}(\tau, R)$. Let D^τ be the PPT distinguisher which on input (R, \tilde{m}) , sets $\mathbb{T} = \mathbb{T}(\tau, R)$, parses $\tilde{m} = \tilde{v} \circ \text{id}'$ and outputs $D(\mathbb{T}, \tilde{v})$. With these notational changes in place, (1) gives

$$\Pr_\tau \left[\left| \Pr_{(R, \tilde{m}) \leftarrow M_m^\tau} (D^\tau(R, \tilde{m}) = 1) - \Pr_{(R, \tilde{m}) \leftarrow M_{m'}^\tau} (D^\tau(R, \tilde{m}) = 1) \right| \geq \frac{\delta}{2} \right] \geq \frac{\delta}{2}. \quad (2)$$

Note that since M is required to produce a commitment using a tag $\tilde{\text{id}} \neq \text{id}$, when (R, \tilde{m}) is drawn from M_m^τ or $M_{m'}^\tau$, we will always have $\tilde{m} \notin \{m, m'\}$.

Definition 12 (Malleable Partial Transcripts). For $m, m' \in \mathcal{M}$, write $\tau \in \text{MAUL}_{m, m'}$ if

$$\left| \Pr_{(R, \tilde{m}) \leftarrow M_m^\tau} (D^\tau(R, \tilde{m}) = 1) - \Pr_{(R, \tilde{m}) \leftarrow M_{m'}^\tau} (D^\tau(R, \tilde{m}) = 1) \right| \geq \delta/2.$$

So if M breaks the non-malleability of $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{synch}}$, there exist $m, m' \in \mathcal{M}$ such that $\Pr_\tau [\tau \in \text{MAUL}_{m, m'}] \geq \delta/2$.

5.2 Proof Overview

At this point, we go through the proof at a high level, in order to highlight the key ideas. As mentioned in the intro, our approach is to use an M who mauls $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{synch}}$ to construct polynomial time split state tampering functions $(f, g) \in \mathcal{F}_{\text{split}}^{\text{poly}}$ which maul the code (Enc, Dec) . In fact, we will construct an efficiently samplable distribution $\mathcal{D}_{\text{split}}$ on $\mathcal{F}_{\text{split}}^{\text{poly}}$ such that with non-negligible probability over $(f, g) \leftarrow \mathcal{D}_{\text{split}}$, (f, g) mauls (Enc, Dec) . We proceed in two steps. First, we use M to define an efficiently samplable distribution \mathcal{D}_M which outputs poly-time (f, g) which maul (Enc, Dec) but are not split-state. Then we show that \mathcal{D}_M is indistinguishable from a distribution $\mathcal{D}_{\text{split}}$ on $\mathcal{F}_{\text{split}}^{\text{poly}}$.

The functions (f, g) output by \mathcal{D}_M share a partial transcript $\tau = (\mathbf{Com}_L, \tilde{\alpha}, a)$ and $R_\S \leftarrow \text{Enc}(m^*|L)$ for some arbitrary fixed $m^* \in \mathcal{M}$. This defines a full transcript $\mathbb{T} = \mathbb{T}(\tau, R_\S)$, let \tilde{a}_\S be M 's response in the third message of \mathbb{T} . Given L , f extracts a candidate \tilde{L} from M by rewinding M and asking a new challenge $\tilde{\beta}$, answering with (b, R_\S) on the left, where b is computed using L and the point (α, a) from τ . Defining $g(R)$ is simpler: it sends (a, R) to M receiving (\tilde{a}, \tilde{R}) , if $\tilde{a} = \tilde{a}_\S$ it outputs \tilde{R} , if not it outputs \perp_{inc} .

Clearly, (f, g) are not split-state as the randomness $(\tau, R_\S) = (\mathbf{Com}_L, \tilde{\alpha}, a, R_\S)$ shared by both functions depends on L . However, (τ, R_\S) only depends on information which computationally hides L , and we show in Section 5.5 that indeed $\mathcal{D}_M \approx_c \mathcal{D}_{\text{split}}$ where (f, g) output by $\mathcal{D}_{\text{split}}$ are defined using bogus randomness (τ, R_\S) drawn from some distribution independent of L .

It is much harder to show that with non-negligible probability, $(f, g) \leftarrow \mathcal{D}_M$ mauls (Enc, Dec) . There are two main issues, which correspond exactly to the two additional properties we need from our non-malleable code. The first is that the tampered value output by (f, g) can only be equal to M 's committed value if M 's response \tilde{a}_\S in $\mathbb{T}(\tau, R_\S)$ is correct. For this, we use the right state simulatability of (Enc, Dec) . Since the only difference between R_\S and the real right state R is that (L, R_\S) is an encoding of m^* while (L, R) is an encoding of m , we can show that \tilde{a}_\S is correct with roughly the same probability that \tilde{a} is. This portion of our proof is in Section 5.3. We then find ourselves in one of two cases. Either the probability of correctness is large for both, or it is small for both in which case M is always committing to \perp_{inc} . In the second case, he is not mauling. In the first case we show that (f, g) output M 's committed value with non-negligible probability that is independent of whether M is mauling or not. It follows that if M mauls $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{synch}}$ with non-negligible probability, then (f, g) mauls (Enc, Dec) with related non-negligible probability. This part of our proof is in Section 5.4.

The second issue is that (f, g) are defined using τ which contains L . This means that the mauling experiment for (f, g) given a message m is conditional: draw $R \leftarrow \text{Enc}(m|L)$ where L is as in τ , compute

$(\tilde{L}, \tilde{R}) = (f(L), g(R))$ and output $\tilde{m} = \text{Dec}(\tilde{L}, \tilde{R})$. This is in contrast to the usual security game for non-malleable codes where (f, g) are fixed and then $(L, R) \leftarrow \text{Enc}(m)$ are both drawn during the tampering experiment, and it explains our need for starting with (Enc, Dec) which are conditionally augmented non-malleable.

5.3 Ruling out Selective \perp_{inc} Attacks

Recall that $\tilde{v} = \perp_{\text{inc}}$ when M 's response \tilde{a} is incorrect. In this section we use the shorthand

$$p_{\tau, m} = \Pr_{(R, \tilde{m}) \leftarrow M_m^\tau} (\tilde{v} \neq \perp_{\text{inc}})$$

and we will prove that if M is mauling then he is doing so by answering correctly. The main lemma of this section is the following.

Lemma 2. *Suppose M breaks the non-malleability of $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{synch}}$ with non-negligible probability δ . Then there exist $m, m' \in \mathcal{M}$ and non-negligible $\delta' > 0$ such that for all $m^* \in \mathcal{M}$*

$$\Pr_\tau \left[\tau \in \text{MAUL}_{m, m'} \ \& \ p_{\tau, m^*} \geq \delta' \right] \geq \frac{\delta}{4} - \text{negl}.$$

Lemma 2 follows immediately from (2) and Claims 7 and 8 which rule out two separate types of mauling behavior.

Claim 7. *For all $m, m' \in \mathcal{M}$ and non-negligible $\xi > 0$ we have:*

$$\Pr_\tau \left[\left| p_{\tau, m} - p_{\tau, m'} \right| > \xi \right] = \text{negl}.$$

Proof. We utilize the hiding machine. Fix $m, m' \in \mathcal{M}$ and non-negligible $\xi > 0$. Let BAD be the set of partial transcripts τ for which $p_{\tau, m} > \xi + p_{\tau, m'}$ and suppose for contradiction that there is some non-negligible $\xi' > 0$ such that $\Pr(\tau \in \text{BAD}) \geq \xi'$. Set $N = 4/(\xi\xi')$, $N' = \Omega(\lambda\xi^{-2})$ and consider the PPT adversary \mathcal{A} who interacts with \mathcal{C} in the (N, N') -way hiding game for $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{synch}}$ as follows.

- \mathcal{A} instantiates M . Upon receiving Com_L from \mathcal{C} , it plays the first two rounds of $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{synch}}$ with M , giving input Com_L and uniform $\tilde{\alpha}$ and receiving α in the second round of the left interaction. \mathcal{A} sends α to \mathcal{C} and receives a . This defines a partial transcript $\tau = (\text{Com}_L, \tilde{\alpha}, a)$.
- \mathcal{A} sends (m, m') to \mathcal{C} and receives challenge $(R, \{R_1\}, \dots, \{R_N\})$. \mathcal{A} forwards (a, R) to M and receives (\tilde{a}, \tilde{R}) . This defines a full transcript $\mathbb{T} = \mathbb{T}(\tau, R)$. Moreover, for all $i = 1, \dots, N$ and $R_i \in \{R_i\}$, \mathcal{A} sends (a, R_i) to M and receives $(\tilde{a}_i, \tilde{R}_i)$, defining transcripts $\{\mathbb{T}_i\}$ for $i = 1, \dots, N$, where $\mathbb{T}_i = \mathbb{T}(\tau, R_i)$.
- \mathcal{A} computes $p_i = \Pr_{R_i \in \{R_i\}}(\tilde{a}_i = \tilde{a})$, and outputs i^* such that p_{i^*} is maximal.

Note that if $\tau \in \text{BAD}$, then M 's response \tilde{a} in \mathbb{T} is correct with probability at least ξ . Moreover, if \tilde{a} is correct then $\tilde{v} \neq \perp_{\text{inc}}$ in \mathbb{T}_i if and only if $\tilde{a}_i = \tilde{a}$. Therefore, conditioned on $\tau \in \text{BAD}$ and \tilde{a} being correct, we see that

$$p_{i^*} \geq p_{\tau, m} - \frac{\xi}{3} > p_{\tau, m'} + \xi - \frac{\xi}{3} \geq p_i + \frac{\xi}{3},$$

for all $i \neq i^*$ with probability at least $1 - 2^{-\Omega(\lambda)}$. We have used the Chernoff-Hoeffding bound, made possible by our choice of large enough N' . We conclude that

$$\Pr(\mathcal{A} \text{ wins}) \geq \xi\xi'(1 - \text{negl}) > \frac{2}{N},$$

which violates the security of the hiding game for $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{synch}}$. □

Claim 8. Let $\delta = \delta(\lambda) > 0$ be as the statement of Lemma 2. For all $m, m' \in \mathcal{M}$, we have:

$$\Pr_{\tau} \left[\tau \in \text{MAUL}_{m,m'} \ \& \ \mathfrak{p}_{\tau,m} < \lambda^{-2}\delta^3 \right] \leq \frac{\delta}{4}.$$

Proof. This is another invocation of the hiding machine. Fix $m, m' \in \mathcal{M}$, let BAD' be the partial transcripts $\tau \in \text{MAUL}_{m,m'}$ such that $\mathfrak{p}_{\tau,m} < \lambda^{-2}\delta^3$, and suppose for contradiction that $\Pr_{\tau} [\tau \in \text{BAD}'] > \delta/4$. For almost all $\tau \in \text{BAD}'$ (in particular, for all $\tau \in \text{BAD}'$ except for those for which Claim 7 does not hold), we have:

$$\begin{aligned} \frac{\delta}{2} &\leq \left| \Pr_{(R, \tilde{m}) \leftarrow M_m^{\tau}} \left(D^{\tau}(R, \tilde{m}) = 1 \right) - \Pr_{(R, \tilde{m}) \leftarrow M_{m'}^{\tau}} \left(D^{\tau}(R, \tilde{m}) = 1 \right) \right| \\ &\leq \left| \Pr_{M_m^{\tau}} \left(D^{\tau}(R, \tilde{m}) = 1 \ \& \ \tilde{v} = \perp_{\text{inc}} \right) - \Pr_{M_{m'}^{\tau}} \left(D^{\tau}(R, \tilde{m}) = 1 \ \& \ \tilde{v} = \perp_{\text{inc}} \right) \right| + \mathfrak{p}_{\tau,m} + \mathfrak{p}_{\tau,m'} \\ &< \left| \Pr_{M_m^{\tau}} \left(D^{\tau}(R, \perp_{\text{inc}}) = 1 \mid \tilde{v} = \perp_{\text{inc}} \right) - \Pr_{M_{m'}^{\tau}} \left(D^{\tau}(R, \perp_{\text{inc}}) = 1 \mid \tilde{v} = \perp_{\text{inc}} \right) \right| + 3\lambda^{-2}\delta^3 + \text{negl}. \end{aligned}$$

Now, set $N = 17/\delta$, $N' = \Omega(\lambda\delta^{-2})$ and consider the \mathcal{A} who plays the hiding game against \mathcal{C} as follows.

- \mathcal{A} instantiates M and obtains partial transcript τ . \mathcal{A} sends m, m' , receives $(R, \{R_1\}, \dots, \{R_N\})$, and sets $\mathbb{T}_i = \mathbb{T}(\tau, R_i)$, for all $R_i \in \{R_i\}$.
- \mathcal{A} computes $p_i = \Pr_{R_i \in \{R_i\}} (D^{\tau}(R_i, \perp_{\text{inc}}) = 1)$ and outputs i^* such that p_{i^*} is maximal.

If $\tau \in \text{BAD}'$ then with probability at least $1/2$, M will answer incorrectly in every $\mathbb{T}_i \in \{\mathbb{T}_i\}$ for all i , and so M 's commitment in every \mathbb{T}_i is to \perp_{inc} . Conditioned on every commitment being to \perp_{inc} , p_i approximates $\Pr_{(R, \tilde{m}) \leftarrow M_{m_i}^{\tau}} (D^{\tau}(R, \perp_{\text{inc}}) = 1 \mid \tilde{v} = \perp_{\text{inc}})$ where $m_{i^*} = m$ and $m_i = m'$ when $i \neq i^*$. In this case, we have by Chernoff-Hoeffding,

$$p_{i^*} \geq \Pr_{(R, \tilde{m}) \leftarrow M_m^{\tau}} (D^{\tau}(R, \perp_{\text{inc}}) = 1 \mid \tilde{v} = \perp_{\text{inc}}) - \frac{\delta}{9} > \Pr_{(R, \tilde{m}) \leftarrow M_{m'}^{\tau}} (D^{\tau}(R, \perp_{\text{inc}}) = 1 \mid \tilde{v} = \perp_{\text{inc}}) + \frac{2\delta}{9} \geq p_i + \frac{\delta}{9},$$

for all $i \neq i^*$ with probability $1 - 2^{-\Omega(\lambda)}$. And so $\Pr(\mathcal{A} \text{ wins}) \geq \frac{\delta}{8} - \text{negl} > \frac{2}{N}$, which violates the security of the hiding game for $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{synch}}$. \square

5.4 The Distribution \mathcal{D}_M

We now use M to define a polynomial time sampleable distribution, \mathcal{D}_M , which outputs a tampering function pair (f, g) , as follows.

- **Random Choices:** Instantiate M and play the first two rounds of $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{synch}}$, obtaining a partial transcript $\tau = (\text{Com}_L, \tilde{\alpha}, a)$ where $a = \varphi(\alpha)$ and $\varphi(x)$ is the linear map specified by Com_L . Draw $R_{\S} \leftarrow \text{Enc}(m^* | L)$ for some arbitrary fixed $m^* \in \mathcal{M}$ and let \tilde{a}_{\S} be M 's response in the full transcript $\mathbb{T}(\tau, R_{\S})$. Finally draw $\tilde{\beta} \leftarrow \mathbb{Z}_q$.
- $f_{\tau, R_{\S}, \tilde{\beta}}(L)$: Let $\varphi(x)$ be the unique linear function with constant term L and $\varphi(\alpha) = a$.
 - rewind M back to the second message of the right interaction and ask $\tilde{\beta}$, receive β on the left;
 - send (b, R_{\S}) where $b = \varphi(\beta)$ and receive (\tilde{b}, \cdot) on the right;
 - output \tilde{L} , the constant term of the line spanned by $\{(\tilde{\alpha}, \tilde{a}_{\S}), (\tilde{\beta}, \tilde{b})\}$.
- $g_{\tau, R_{\S}}(R)$: Let (\tilde{a}, \tilde{R}) be M 's final message in $\mathbb{T}(\tau, R)$. If $\tilde{a} = \tilde{a}_{\S}$ output \tilde{R} , otherwise \perp_{inc} .

- **Output:** $(f, g) = (f_{\tau, R_{\S}, \tilde{\beta}}, g_{\tau, R_{\S}})$.

Notice (f, g) output by \mathcal{D}_M are not split-state as the randomness (τ, R_{\S}) shared by both f and g depends on L . Nonetheless, we show in Section 5.5 below that $\mathcal{D}_M \approx_c \mathcal{D}_{\text{split}}$, a distribution on $\mathcal{F}_{\text{split}}^{\text{poly}}$. Combined with the next lemma, this shows that $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{synch}}$ is non-malleable: an M which breaks the non-malleability of $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{synch}}$ can be used to construct a distribution $\mathcal{D}_{\text{split}}$ on $\mathcal{F}_{\text{split}}^{\text{poly}}$ which breaks the security of the code (Enc, Dec).

Lemma 3. *Let $\delta, \delta' > 0$ be as in the statement of Lemma 2. If M breaks the non-malleability of $\langle \mathcal{C}, \mathcal{R} \rangle$ then there exist $m, m' \in \mathcal{M}$ such that*

$$\Pr_{L, (f, g)} \left[\left| \Pr_{(R, \tilde{m}) \leftarrow V_{m, f, g}^L} \left(D^\tau(R, \tilde{m}) = 1 \right) - \Pr_{(R, \tilde{m}) \leftarrow V_{m', f, g}^L} \left(D^\tau(R, \tilde{m}) = 1 \right) \right| > \frac{\delta}{2} \right] > \frac{(\delta\delta')^3}{256} - \text{negl},$$

where the outer probability is over $L \leftarrow \mathcal{L}$ and $(f, g) = (f_{\tau, R_{\S}, \tilde{\beta}}, g_{\tau, R_{\S}}) \leftarrow \mathcal{D}_M$, where $\tau = (\text{Com}_L, \tilde{\alpha}, a)$

Proof Idea. The randomness needed in order to draw $(f, g) \leftarrow \mathcal{D}_M$ consists of a random partial transcript $\tau = (\text{Com}_L, \tilde{\alpha}, a)$, $R_{\S} \leftarrow \text{Enc}(m^*|L)$ and $\tilde{\beta} \leftarrow \mathbb{Z}_q$. Given these choices, the distributions $V_{m, f, g}^L$ and M_m^τ are very similar: both draw $R \leftarrow \text{Enc}(m|L)$ and output (R, \tilde{m}) . We prove Lemma 3 by showing that whenever $(\tau, R_{\S}, \tilde{\beta})$ is such that M 's response \tilde{a}_{\S} is correct, $V_{m, f, g}^L$ and M_m^τ are actually identical for all $m \in \mathcal{M}$. The proof then follows from the observation that \tilde{a}_{\S} is correct with non-negligible probability, which uses Lemma 2 and the right state simulatability of (Enc, Dec).

Proof of Lemma 3. For randomness $(\tau, R_{\S}) = (\text{Com}_L, \tilde{\alpha}, r\alpha + L, R_{\S})$, say the ‘‘extraction event’’, denoted EXT, occurs whenever \tilde{a}_{\S} is correct in $\mathbb{T}(\tau, R_{\S})$, and

$$\Pr_{\tilde{\beta}} \left(\tilde{b} \text{ correct in } \mathbb{T}(\text{Com}_L, \tilde{\beta}, r\beta + L, R_{\S}) \right) \geq \frac{(\delta\delta')^2}{32}.$$

If M mauls $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{synch}}$ then there exist messages $m, m' \in \mathcal{M}$ such that

$$\Pr_{\tau, R_{\S}} \left(\tau \in \text{MAUL}_{m, m'} \ \& \ \text{EXT} \right) \geq \frac{\delta\delta'}{4} - \frac{\delta\delta'}{8} - \text{negl} = \frac{\delta\delta'}{8} - \text{negl},$$

using Lemma 2 and Bayes' theorem. If \tilde{a}_{\S} is correct then $g(R)$ identifies when M is committing to \perp_{inc} on the right (\tilde{a} is correct if and only if $\tilde{a} = \tilde{a}_{\S}$), and so outputs the correct \tilde{R} . Moreover, if \tilde{a}_{\S} and \tilde{b} are correct then $f(L)$ outputs the correct \tilde{L} . So we see that when $\tau \in \text{MAUL}_{m, m'}$ and EXT occurs and \tilde{b} is correct (all of which happen with probability at least $(\delta\delta')^3/256 - \text{negl}$), then $V_{m, f, g}^L \equiv M_m^\tau$ for all $m \in \mathcal{M}$ and so since $\tau \in \text{MAUL}_{m, m'}$,

$$\left| \Pr(D^\tau(V_{m, f, g}^L) = 1) - \Pr(D^\tau(V_{m', f, g}^L) = 1) \right| > \frac{\delta}{2}.$$

□

5.5 A Hybrid Argument to Prove $\mathcal{D}_M \approx_c \mathcal{D}_{\text{split}}$

Note that the distribution \mathcal{D}_M from the previous section does not output split state (f, g) as the randomness $(\tau, R_{\S}) = (\text{Com}(L), \text{Com}(r), \tilde{\alpha}, a, R_{\S})$ shared between f and g depends in three ways on L : 1) τ contains a commitment to L , 2) $a = r\alpha + L$, and 3) $R_{\S} \leftarrow \text{Enc}(m^*|L)$. We show in this section, however, that \mathcal{D}_M is computationally indistinguishable from a polynomial time sampleable distribution $\mathcal{D}_{\text{split}}$ on $\mathcal{F}_{\text{split}}^{\text{poly}}$. This, together with Lemma 3, completes the proof that $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{synch}}$ is non-malleable.

$\mathcal{D}_0 = \mathcal{D}_M$ – This is the distribution defined above. It draws $L \leftarrow \mathcal{L}$, a and also random partial transcript $\tau = (\text{Com}(L), \text{Com}(r), \tilde{\alpha}, a)$ where $a = r\alpha + L$, $R_{\S} \leftarrow \text{Enc}(m^*|L)$ and outputs $(f_0, g_0) = (f_{\tau, R_{\S}}, g_{\tau, R_{\S}})$.

\mathcal{D}_1 – This distribution outputs functions (f_1, g_1) which behave exactly like (f_0, g_0) except that they are seeded with $(\tau, R_\S) = (\text{Com}(L), \text{Com}(r), \tilde{\alpha}, a, R_\S)$, where $a \leftarrow \mathbb{Z}_q$ is random instead of equal to $r\alpha + L$.

\mathcal{D}_2 – This outputs (f_2, g_2) which, again, differ from (f_1, g_1) only in their shared randomness. This time $(\tau, R_\S) = (\text{Com}(0), \text{Com}(r), \tilde{\alpha}, a, R_\S)$, where $a \in \mathbb{Z}_q$ is random. Now the only dependence on L is that $R_\S \leftarrow \text{Enc}(m^*|L)$.

$\mathcal{D}_3 = \mathcal{D}_{\text{split}}$ – This outputs (f_3, g_3) which are the same as (f_2, g_2) except that $R_\S \leftarrow \mathcal{D}_{\text{hid}}$, the distribution on \mathcal{R} guaranteed by the right state simulatability of (Enc, Dec) . Since (τ, R_\S) no longer depends on L , $(f_3, g_3) \in \mathcal{F}_{\text{split}}^{\text{poly}}$.

Claim 9. $\mathcal{D}_0 \approx_c \mathcal{D}_1 \approx_c \mathcal{D}_2 \approx_c \mathcal{D}_3$.

Proof. The first two indistinguishabilities follow from the hiding of Com . For the first, consider an adversary \mathcal{A} who interacts with a challenger \mathcal{C} in the hiding game by choosing $r_0, r_1 \in \mathbb{Z}_q$ at random and sending (r_0, r_1) to \mathcal{C} , receiving a commitment $z = \text{Com}(r_b)$ for a random $b \in \{0, 1\}$. Then \mathcal{A} draws $L \leftarrow \mathcal{L}$ and $\tilde{\alpha} \leftarrow \mathbb{Z}_q$ at random and sends $(\text{Com}(L), z)$ and $\tilde{\alpha}$ to \mathcal{M} (corresponding to the first message of the left interaction and the second message of the right interaction), receiving α as the second message on the left. \mathcal{A} sets $a = r_0\alpha + L$, draws $R_\S \leftarrow \text{Enc}(m_\S|L)$ and outputs $(\text{Com}(L), z, \tilde{\alpha}, a, R_\S)$. If $b = 0$ then \mathcal{A} 's output is distributed according to \mathcal{D}_0 , while if $b = 1$, \mathcal{A} 's output is distributed like \mathcal{D}_1 . This proves that $\mathcal{D}_0 \approx_c \mathcal{D}_1$; $\mathcal{D}_1 \approx_c \mathcal{D}_2$ follows even more readily. Finally, $\mathcal{D}_2 \approx_c \mathcal{D}_3$ follows from the hiding property of (Enc, Dec) . \square

It follows from Lemma 3 and Claim 9 that if \mathcal{M} breaks the non-malleability of $(\mathcal{C}, \mathcal{R})$ then there exist $m, m' \in \mathcal{M}$ and non-negligible $\delta, \delta' > 0$ such that

$$\Pr_{L, (f, g)} \left[\left| \Pr_{(R, \tilde{m}) \leftarrow \mathcal{V}_{m, f, g}^L} \left(D^\tau(R, \tilde{m}) = 1 \right) - \Pr_{(R, \tilde{m}) \leftarrow \mathcal{V}_{m', f, g}^L} \left(D^\tau(R, \tilde{m}) = 1 \right) \right| > \frac{\delta}{2} \right] > \frac{(\delta\delta')^3}{256} - \text{negl},$$

where the outer probability is over $L \leftarrow \mathcal{L}$ and $(f, g) \leftarrow \mathcal{D}_{\text{split}}$, drawn independently. This breaks the computational conditional augmented non-malleability of (Enc, Dec) , thus completing our proof of Theorem 1.

6 Constructing Codes with Simulatable Right State (Proof of Lemma 1)

Lemma 1 (Restated). (Enc, Dec) is computationally conditionally augmented non-malleable against

$$\mathcal{F}_{\text{split}}^{\text{poly}} = \{(f, g) \in \mathcal{F}_{\text{split}} : f \text{ and } g \text{ polytime}\}.$$

Fix $(f, g) \in \mathcal{F}_{\text{split}}^{\text{poly}}$ and non-negligible $\delta = \delta(\lambda) > 0$ (every δ in this section refers to this choice). We must describe simulators $\{S_{f, g}^L\}_L$ such that for all PPT distinguishers D , $\Pr_L[E_{f, g}] < 2\delta + \lambda^{-2}\delta^4 + \text{negl}$, where $E_{f, g}$ is the event

$$\exists m \in \mathcal{M} \text{ st } \left| \Pr_{(R, \tilde{m}) \leftarrow \mathcal{V}_{m, f, g}^L} \left(D(R, \tilde{m}) = 1 \right) - \Pr_{(R, \tilde{m}) \leftarrow S_{m, f, g}^L} \left(D(R, \tilde{m}) = 1 \right) \right| > \delta.$$

We address the three points of the high level proof in Section 3.2 in reverse order. Recall, we already proved Claim 4. Just as in Section 3.2 we use the shorthand $\rho_{L, m}$ for $\Pr_{(R, \tilde{m}) \leftarrow \mathcal{V}_{m, f, g}^L} (\tilde{m} \neq \perp_{\text{com}})$.

Claim 4 (Restated). For any $m, m' \in \mathcal{M}$, $(f, g) \in \mathcal{F}_{\text{split}}^{\text{poly}}$ and non-negligible $\xi = \xi(\lambda) > 0$ we have

$$\Pr_{\mathcal{L}} \left[\left| \rho_{\mathcal{L}, m} - \rho_{\mathcal{L}, m'} \right| > \xi \right] = \text{negl}.$$

Now, fix any $m^* \in \mathcal{M}$ and define $\mathcal{L}_{\text{VALID}} = \{L \in \mathcal{L} : \rho_{\mathcal{L}, m^*} \geq \lambda^{-2}\delta^3\}$, and $\mathcal{L}_{\text{BOT}} = \mathcal{L} \setminus \mathcal{L}_{\text{VALID}}$. We treat the cases $L \in \mathcal{L}_{\text{VALID}}$ and $L \in \mathcal{L}_{\text{BOT}}$ separately, beginning with $L \in \mathcal{L}_{\text{VALID}}$. We use (f, g) to construct (f_0, g_0) which tampers $(\text{Enc}_0, \text{Dec}_0)$. We will then use $S_{f_0, g_0}^{\mathcal{L}_0}$ whose existence is guaranteed by the conditional non-malleability of $(\text{Enc}_0, \text{Dec}_0)$, to construct $S_{f, g}^{\mathcal{L}}$. Whenever $L \in \mathcal{L}_{\text{VALID}}$, $S_{f, g}^{\mathcal{L}}$ will simulate $\{V_{m, f, g}^{\mathcal{L}}\}_m$. The construction of (f_0, g_0) is as follows.

- **Random Choices:** Draw $k \leftarrow G(1^\lambda)$, $\sigma \leftarrow \mathcal{S}$, $c^\mathcal{S} \leftarrow E_k(0)$, set $z = \text{Com}(k, \sigma)$, and $(\cdot, \tilde{z}^\mathcal{S}) = g(c^\mathcal{S}, z)$.
- $f_0(L_0)$: Compute $(\tilde{L}_0, (\tilde{k}, \tilde{\sigma})) = f(L_0, (k, \sigma))$. Output \tilde{L}_0 .
- $g_0(R_0)$: Draw $c \leftarrow E_k(R_0)$ and set $(\tilde{c}, \tilde{z}) = g(c, z)$. If $\tilde{z} \neq \tilde{z}^\mathcal{S}$, output \perp_{com} . Otherwise, use superpolynomial time to break open \tilde{z} and recover the pair $(\tilde{k}', \tilde{\sigma}')$. Output $\tilde{R}_0 = D_{\tilde{k}'}(\tilde{c})$.

Define $S_{f, g}^{\mathcal{L}}$ as follows: parse $L = (L_0, (k, \sigma))$ and draw $c^\mathcal{S} \leftarrow E_k(0)$ such that $\text{Decom}(\tilde{z}^\mathcal{S}) = (\tilde{k}, \tilde{\sigma})$, where $(\cdot, \tilde{z}^\mathcal{S}) = g(c^\mathcal{S}, \text{Com}(k, \sigma))$ and $(\tilde{L}_0, (\tilde{k}, \tilde{\sigma})) = f(L)$ (we prove in Claim 10 below that such $c^\mathcal{S}$ exist with high probability when $L \in \mathcal{L}_{\text{VALID}}$); this defines $(f_0, g_0) \in \mathcal{F}_{\text{split}}$. Draw $(R_0, \tilde{m}) \leftarrow S_{f_0, g_0}^{\mathcal{L}_0}$, $c \leftarrow E_k(R_0)$ and output (R, \tilde{m}) where $R = (c, \text{Com}(k, \sigma))$.

Claim 10. For any PPT distinguisher D , $\Pr_{\mathcal{L}} [L \in \mathcal{L}_{\text{VALID}} \ \& \ E_{f, g}] < \delta + \text{negl}$.

Proof. Assume $\mathcal{L}_{\text{VALID}}$ comprises at least a δ -fraction of \mathcal{L} ; if not we are done. Note that as $c^\mathcal{S}$ is such that $\text{Decom}(\tilde{z}^\mathcal{S}) = (\tilde{k}, \tilde{\sigma})$, tampering with (f_0, g_0) is the same as tampering with (f, g) . Indeed, $(R, \tilde{m}) \leftarrow V_{m, f, g}^{\mathcal{L}}$ has $\tilde{m} = \perp_{\text{com}}$ if and only if $\text{Decom}(\tilde{z}) \neq (\tilde{k}, \tilde{\sigma})$, or equivalently, since Com is perfectly binding, if $\tilde{z} \neq \tilde{z}^\mathcal{S}$. Moreover, if $\tilde{m} \neq \perp_{\text{com}}$ then $g(R) = (E_{\tilde{k}}(g(R_0)), \tilde{z})$. It follows that $V_{m, f, g}^{\mathcal{L}}$ is identical to the distribution: draw $(R_0, \tilde{m}) \leftarrow V_{m, f_0, g_0}^{\mathcal{L}_0}$, set $R = (E_k(R_0), \text{Com}(k, \sigma))$ and output (R, \tilde{m}) . So whenever there exists a $c^\mathcal{S}$ such that $\text{Decom}(\tilde{z}^\mathcal{S}) = (\tilde{k}, \tilde{\sigma})$ we have,

$$\Pr_{\mathcal{L}} [E_{f, g}] \leq \Pr_{\mathcal{L}} [\exists m \text{ st } \Delta(V_{m, f_0, g_0}^{\mathcal{L}_0}, S_{m, f_0, g_0}^{\mathcal{L}_0}) > \delta] < \text{negl},$$

as $(\text{Enc}_0, \text{Dec}_0)$ is ε -conditionally augmented non-malleable for negligible ε . Finally, we show that for almost all $L \in \mathcal{L}_{\text{VALID}}$, there exists $c^\mathcal{S}$ st $\text{Decom}(\tilde{z}^\mathcal{S}) = (\tilde{k}, \tilde{\sigma})$. In fact, $\text{Decom}(\tilde{z}^\mathcal{S}) = (\tilde{k}, \tilde{\sigma})$ holds for a non-negligible fraction of $c^\mathcal{S}$: $\Pr_{c^\mathcal{S}} [\text{Decom}(\tilde{z}^\mathcal{S}) = (\tilde{k}, \tilde{\sigma}) | L \in \mathcal{L}_{\text{VALID}}] \geq \lambda^{-2}\delta^4 - \text{negl}$. This follows immediately from the definition of $\mathcal{L}_{\text{VALID}}$ and semantic security. Indeed, conditioned on $L \in \mathcal{L}_{\text{VALID}}$, we have that $\Pr_{R \leftarrow \text{Enc}(m^* | L)} (\text{Decom}(\tilde{z}) = (\tilde{k}, \tilde{\sigma})) \geq \lambda^{-2}\delta^3$, and the only difference between $R = (c, z) \leftarrow \text{Enc}(m^* | L)$ and $(c^\mathcal{S}, z)$ is their encrypted values. This intuition can be formalized easily using the hiding machine. We complete the proof of Claim 10 by collecting our observations:

$$\begin{aligned} \Pr_{\mathcal{L}} [L \in \mathcal{L}_{\text{VALID}} \ \& \ E_{f, g}] &\leq \Pr_{\mathcal{L}} [L \in \mathcal{L}_{\text{VALID}} | |\mathcal{L}_{\text{VALID}}| < \delta |\mathcal{L}|] + \Pr_{\mathcal{L}} [E_{f, g} | \exists c^\mathcal{S} : \text{Decom}(\tilde{z}^\mathcal{S}) = (\tilde{k}, \tilde{\sigma})] \\ &+ \Pr_{\mathcal{L}} [\text{Decom}(\tilde{z}^\mathcal{S}) \neq (\tilde{k}, \tilde{\sigma}) \ \forall c^\mathcal{S} | L \in \mathcal{L}_{\text{VALID}} \ \& \ |\mathcal{L}_{\text{VALID}}| \geq \delta |\mathcal{L}|] \\ &< \delta + \text{negl}. \end{aligned}$$

□

When $L \in \mathcal{L}_{\text{BOT}}$, (f, g) is almost always tampering to \perp_{com} . In this case, computational conditional augmented non-malleability follows from right state simulatability. We define $\{S_{f, g}^{\mathcal{L}}\}_L$ simply: $S_{f, g}^{\mathcal{L}}$ draws $R \leftarrow \mathcal{D}_{\text{hid}}$ (the right state simulator) and outputs (R, \perp_{com}) .

Claim 11. For any PPT distinguisher D , $\Pr_{\mathcal{L}}[\mathbf{L} \in \mathcal{L}_{\text{BOT}} \ \& \ E_{f,g}] < \delta$.

Proof. We use the hiding machine. Set BAD to be the set of $\mathbf{L} \in \mathcal{L}_{\text{BOT}}$ for which $E_{f,g}$ occurs. Unless \mathbf{L} is in the negligible fraction of \mathcal{L} for which Claim 4 does not hold, $\rho_{\mathbf{L},m} < 2\lambda^{-2}\delta^3$ (using Claim 4 with $\xi = \lambda^{-2}\delta^3$). For such \mathbf{L} , we have

$$\begin{aligned} \Pr_{(\mathbf{R}, \tilde{m}) \leftarrow \mathcal{V}_{m,f,g}^{\mathbf{L}}} (D(\mathbf{R}, \tilde{m}) = 1 \mid \tilde{m} = \perp_{\text{com}}) &\geq \Pr_{(\mathbf{R}, \tilde{m}) \leftarrow \mathcal{V}_{m,f,g}^{\mathbf{L}}} (D(\mathbf{R}, \tilde{m}) = 1) - 2\lambda^{-2}\delta^3 \\ &> \Pr_{(\mathbf{R}, \tilde{m}) \leftarrow \mathcal{S}_{m,f,g}^{\mathbf{L}}} (D(\mathbf{R}, \tilde{m}) = 1) + \delta - 2\lambda^{-2}\delta^3 \\ &> \Pr_{(\mathbf{R}, \tilde{m}) \leftarrow \mathcal{S}_{m,f,g}^{\mathbf{L}}} (D(\mathbf{R}, \tilde{m}) = 1) + \frac{\delta}{2}, \end{aligned}$$

for some $m \in \mathcal{M}$. Assume for contradiction that $\Pr_{\mathcal{L}}[\mathbf{L} \in \text{BAD}] \geq \delta$. Fix $N = 5/\delta$, $N' = \Omega(\lambda\delta^{-2})$ and consider the PPT adversary \mathcal{A} who plays with challenger \mathcal{C} as follows.

- \mathcal{C} chooses $\mathbf{L} \leftarrow \mathcal{L}$, $i^* \leftarrow \{1, \dots, N\}$, sends $(\{\mathbf{R}_1\}, \dots, \{\mathbf{R}_N\})$ to \mathcal{A} where $\mathbf{R}_i \leftarrow \mathcal{D}_{\text{hid}}$ for all $i \neq i^*$, $\mathbf{R}_i \in \{\mathbf{R}_i\}$; and $\mathbf{R}_{i^*} \leftarrow \text{Enc}(m \mid \mathbf{L})$ for all $\mathbf{R}_{i^*} \in \{\mathbf{R}_{i^*}\}$.
- \mathcal{A} receives $(\{\mathbf{R}_1\}, \dots, \{\mathbf{R}_N\})$ and for $i = 1, \dots, N$, sets $p_i = \Pr_{\mathbf{R}_i \in \{\mathbf{R}_i\}} (D(\mathbf{R}, \perp_{\text{com}}) = 1)$ and outputs i^* such that p_{i^*} is maximal.

If $\mathbf{L} \in \mathcal{L}$ chosen by \mathcal{C} is in BAD and such that Claim 4 holds then the expected number of (i, \mathbf{R}_i) such that $i \in \{1, \dots, N\}$, $\mathbf{R}_i \in \{\mathbf{R}_i\}$, and $\tilde{m}_i \neq \perp_{\text{com}}$ is at most $2\lambda^{-2}\delta^3 N N' < 1/2$, where $\tilde{m}_i = \text{Dec}(\tilde{\mathbf{L}}, \tilde{\mathbf{R}}_i)$. With probability at least $1/2$ there exist no such (i, \mathbf{R}_i) . In this case we have

$$p_{i^*} \geq \Pr_{(\mathbf{R}, \tilde{m}) \leftarrow \mathcal{V}_{m,f,g}^{\mathbf{L}}} (D(\mathbf{R}, \tilde{m}) = 1 \mid \tilde{m} = \perp_{\text{com}}) - \frac{\delta}{6} > \Pr_{(\mathbf{R}, \tilde{m}) \leftarrow \mathcal{S}_{f,g}^{\mathbf{L}}} (D(\mathbf{R}, \tilde{m}) = 1) + \frac{\delta}{2} - \frac{\delta}{6} \geq p_i + \frac{\delta}{6},$$

for all $i \neq i^*$. So we see that

$$\Pr(\mathcal{A} \text{ wins}) \geq \Pr(\mathbf{L} \in \text{BAD}) \Pr(\tilde{m}_i = \perp_{\text{com}} \ \forall (i, \mathbf{R}_i) \mid \mathbf{L} \in \text{BAD}) - \text{negl} \geq \frac{\delta}{2} - \text{negl} > \frac{2}{N},$$

and so \mathcal{A} breaks the right state simulatability of (Enc, Dec) . \square

7 The Extended Protocol

In this section we modify the protocol of Section 4 so it remains non-malleable against a non-synchronizing adversary. The only non-synchronizing scheduling available to the adversary which is not trivially dealt with is the sequential scheduling where he lets the left interaction complete before beginning the right. Note this scheduling could not yeild a mauling attack against an extractable commitment, since the sequential scheduling allows one to extract M 's commitment without rewinding \mathcal{C} . However, our protocol in Section 4 is not extractable. In this section we make it so, and thus also non-malleable against a sequential adversary, while still maintaining its non-malleability against a synchronizing adversary.

Let $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{synch}}$ be the commitment of Section 4 and let $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{ext}}$ be a (malleable) three round extractable commitment scheme. Our commitment scheme in this section, $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{nm}}$ commits to v as follows: it uses $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{synch}}$ to commit to v while, in parallel, using $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{ext}}$ to commit to the decommitment information of the first commitment. We prove that this composition enjoys the best of both of its building blocks: it is extractable (and so non-malleable against a sequential adversary) while still being non-malleable against a synchronizing adversary. One technical point is that in the proof of synchronizing non-malleability, we need to rewind the protocol one time, therefore to make our proof go through, we need extraction from $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{ext}}$ to require two rewinds. The protocol is shown in Figure 3.

Setup: Let Com be a non-interactive, perfectly binding commitment scheme. Let (Enc, Dec) be a computational, conditional, augmented non-malleable code. Fix a large prime q . Let $\text{id} \in \{0, 1\}^\lambda$ be \mathcal{C} 's identity.

Committer's Private Input: $v \in \mathcal{M}_{\langle \mathcal{C}, \mathcal{R} \rangle_{\text{synch}}}$ to be committed to.

Commit Phase:

1. $\mathcal{C} \rightarrow \mathcal{R}$:

- Set $m = v \circ \text{id}$ and draw $(L, R) \leftarrow \text{Enc}(m)$, where $L \in \mathcal{L} \subset \mathbb{Z}_q$. Choose random $r \leftarrow \mathbb{Z}_q, \omega \leftarrow \mathbb{Z}_q$ and send $\text{Com}(L \circ r; \omega)$ to \mathcal{R} .
- Set $X = (L \circ r, \omega) \in \{0, 1\}^{\text{poly}(\lambda)}$, viewed as a string and for $i = 1, \dots, \lambda$ choose $X_i^0, X_i^1, X_i^2 \in \{0, 1\}^{\text{poly}(\lambda)}$ randomly such that $X_i^0 \oplus X_i^1 \oplus X_i^2 = X$. Send $Y_i^b = \text{Com}(X_i^b)$ to \mathcal{R} for $b \in \{0, 1, 2\}$.

2. $\mathcal{R} \rightarrow \mathcal{C}$: Send random challenge $\alpha \in \mathbb{Z}_q^*$ and $c \in \{0, 1, 2\}^\lambda$ to \mathcal{C} .

3. $\mathcal{C} \rightarrow \mathcal{R}$: Send $a = r\alpha + L, R$ and $\text{Decom}(Y_i^{c_i})$ for $i = 1, \dots, \lambda$ to \mathcal{C} .

Decommit Phase:

1. $\mathcal{C} \rightarrow \mathcal{R}$: Open the commitments sent in step 1. Let $L', r' \in \mathbb{Z}_q$ and $X_i^b \in \{0, 1\}^{\text{poly}(\lambda)}$ be the decommitted values.

Receiver's Output: If the strings $X_i^0 \oplus X_i^1 \oplus X_i^2$ are not equal for all i or if they are all equal to $X \in \{0, 1\}^{\text{poly}(\lambda)}$ but X is not a valid decommitment to $\text{Com}(L \circ r; \omega)$, sent in step one output \perp_{fail} . If L and r do not satisfy $r\alpha + L = a$ then output \perp_{inc} . Otherwise, compute $m = \text{Dec}(L, R)$ and parse $m = v \circ \text{id}'$. Output v if $\text{id}' = \text{id}$, \perp_{id} if not.

Figure 3: Non-malleable commitment scheme $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{nm}}$.

Claim 12. $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{nm}}$ is a perfectly binding extractable commitment scheme.

Proof Sketch. Perfect binding follows from the perfect binding of Com. Computational hiding follows from the hiding of Com and the fact that any split-state non-malleable code is also a 2-out-of-2 secret sharing scheme. Finally, extractability follows from the extractability of $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{ext}}$ and the fact that the value committed to using $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{ext}}$ is the decommitment of the commitment sent using $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{synch}}$. This means that whether the extracted value is correct can be efficiently checked. \square

Theorem 2. $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{nm}}$ is non-malleable.

Proof. It suffices to prove that it is non-malleable against a synchronizing adversary as the only other non-trivial scheduling is the sequential one and non-malleability against a sequential adversary follows from extractability. We follow the same outline and use the same notation as in the proof of Theorem 1. Recall that if M breaks the non-malleability of $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{nm}}$ then there exist $m, m' \in \mathcal{M}$ and non-negligible $\delta > 0$ st $\Pr_\tau[\tau \in \text{MAUL}_{m, m'}] \geq \delta/2$, where $\tau \in \text{MAUL}_{m, m'}$ if

$$\left| \Pr_{(R, \tilde{m}) \leftarrow M_m^\tau} \left(D^\tau(R, \tilde{m}) = 1 \right) - \Pr_{(R, \tilde{m}) \leftarrow M_{m'}^\tau} \left(D^\tau(R, \tilde{m}) = 1 \right) \right| \geq \frac{\delta}{2}.$$

The partial transcript τ includes the first two rounds of the right execution and the entire left execution except for R (so in particular, τ contains messages from $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{ext}}$ as well as $\langle \mathcal{C}, \mathcal{R} \rangle_{\text{synch}}$). As before, τ is completed to a

full transcript once R is specified. Note that if M 's commitment \tilde{v} is to \perp_{fail} , then M 's first message is bad, either because the strings $\tilde{X}_i^0 \oplus \tilde{X}_i^1 \oplus \tilde{X}_i^2$ are not all equal, or because they are not all equal to a valid decommitment of $\text{Com}(\tilde{L} \circ \tilde{r}; \omega)$. In any case, M cannot be mauling if $\tilde{v} = \perp_{\text{fail}}$ as C 's commitment on the left is not even defined after the first message. Also, the proof of Lemma 2 goes through unchanged for this new protocol and so it follows that if M is mauling with non-negligible probability then he is doing so while also sending the correct value for \tilde{a} with non-negligible probability.

This allows us to build a distribution \mathcal{D}_M on tampering functions which we will use to break the security of (Enc, Dec) . As before $(f, g) \leftarrow \mathcal{D}_M$ share a random partial transcript τ and a random R_{\S} , let \tilde{a}_{\S} be M 's response on the right in the transcript $\mathbb{T}(\tau, R_{\S})$. $f(L)$ extracts L by rewinding M and asking a new challenge $\tilde{\beta}$ on the right, using L and τ to answer on the left. We provide f with the decommitments $\text{Decom}(Y_i^b)$ so he can answer this part of M 's query on the left honestly. $g(R)$ completes τ to $\mathbb{T}(\tau, R)$ and checks whether M 's answer \tilde{a} is equal to \tilde{a}_{\S} or not. If so $g(R) = \tilde{R}$, where \tilde{R} is from M 's final message of $\mathbb{T}(\tau, R)$. If not $g(R) = \perp_{\text{inc}}$. The same proof of Lemma 3 shows that if M breaks the non-malleability of $\langle \mathcal{C}, \mathcal{R} \rangle_{nm}$ then there exist $m, m' \in \mathcal{M}$ and non-negligible $\delta, \delta' > 0$ such that

$$\Pr_{L, (f, g)} \left[\left| \Pr_{(R, \tilde{m}) \leftarrow \mathcal{V}_{m, f, g}^L} \left(D^\tau(R, \tilde{m}) = 1 \right) - \Pr_{(R, \tilde{m}) \leftarrow \mathcal{V}_{m', f, g}^L} \left(D^\tau(R, \tilde{m}) = 1 \right) \right| > \frac{\delta}{2} \right] > \delta', \quad (3)$$

where the outer probability is over $L \leftarrow \mathcal{L}$ and $(f, g) \leftarrow \mathcal{D}_M$ using randomness which depends on L . In order to complete the reduction to the security of (Enc, Dec) we need to exhibit a distribution $\mathcal{D}_{\text{split}}$ on split state tampering function pairs such that

$$\Pr_{L, (f, g)} \left[\left| \Pr_{(R, \tilde{m}) \leftarrow \mathcal{V}_{m, f, g}^L} \left(D^\tau(R, \tilde{m}) = 1 \right) - \Pr_{(R, \tilde{m}) \leftarrow \mathcal{V}_{m', f, g}^L} \left(D^\tau(R, \tilde{m}) = 1 \right) \right| > \frac{\delta}{2} \right] > \delta', \quad (4)$$

where the outer probability is over $L \leftarrow \mathcal{L}$ and $(f, g) \leftarrow \mathcal{D}_{\text{split}}$ drawn independently. In the proof of Theorem 1, we argued that $\mathcal{D}_M \approx_c \mathcal{D}_{\text{split}}$ and so (4) followed straight from (3). Here, this won't quite work because f gets every $\text{Decom}(Y_i^b)$ so it can answer M 's query in the rewind. This means f knows L and so we will not be able to hybrid away f 's dependence on L without f noticing. However, the following observation lets us deduce (4) anyway: as f only rewinds once its output contains no information about $X = X_i^0 \oplus X_i^1 \oplus X_i^2$. Therefore, D^τ cannot distinguish between the output of (f, g) seeded with correct decommitments $\{X_i^b\}$ and the output of (f, g) seeded with random strings. The result follows. \square

8 Simple Constant-Round Non-malleable Commitments

In this section, we describe a very simple non-malleable commitment scheme based on standard split-state non-malleable codes. The scheme has constant rounds and an almost elementary proof. Let us quickly recall the notation for Naor's statistically binding commitment scheme [Nao91] based on a pseudorandom generator $\text{prg} : \{0, 1\}^\lambda \rightarrow \{0, 1\}^{3\lambda}$ (which can be constructed from any one-way function [HILL99]) and Feige-Shamir zero-knowledge proof-of-knowledge [FS89].

Naor's commitment consists of two rounds: first the receiver sends a random string τ of length 3λ ; to commit to a bit b , the sender sends $c = \text{com}_\tau(b)$ as follows: select a uniform seed $r \in \{0, 1\}^\lambda$ and set $c = \text{prg}(r)$ if $b = 0$, and $c = \tau \oplus \text{prg}(r)$ if $b = 1$. To decommit, or "open" c , the sender simply provides r . The receiver accepts 0 as the decommitted value if $c = \text{prg}(r)$, 1 if $c = \tau \oplus \text{prg}(r)$, and rejects otherwise. If τ is not of the form $\text{prg}(r) \oplus \text{prg}(r')$, which happens with 2^{-n} probability, the scheme is perfectly binding. When we want to be explicit about the randomness used for commitment we write $c = \text{com}_\tau(b; r)$.

The protocol has the following **equivocality** property: if τ is of the special form $\tau = \text{prg}(r) \oplus \text{prg}(r')$ then by setting $c = \text{prg}(r)$, the committer can open c to both 0 and 1 by sending either r or r' respectively. A *string* $s = (s_1, \dots, s_n)$ can be committed bit by bit: the receiver selects n strings τ_1, \dots, τ_n , and the sender

sends $\text{com}_{\tau_i}(s_i; r_i)$ for every $i \in [n]$. We abuse the notation and denote this commitment to s by $\text{com}_\tau(s; r) := (\text{com}_{\tau_1}(s_1; r_1), \dots, \text{com}_{\tau_n}(s_n; r_n))$ where $\tau = (\tau_1, \dots, \tau_n)$ and $r = (r_1, \dots, r_n)$. Equivocality for s is achieved through equivocality for each s_i when every τ_i is of special form. When equivocality is not desired, the same string, e.g., τ_1 , can be used for all bits. It will be clear from the context if we want to use a single string for all bits, or a different one for each bit.

Feige-Shamir zero-knowledge proof-of-knowledge [FS89] will be denoted by Π_{FS} . For concreteness, we use the four round variant of [GRRV14], and denote its messages by $(\text{zk}_1, \text{zk}_2, \text{zk}_3, \text{zk}_4)$. This protocol has the property the statement to be proven can be decided in the last round. It is based on the three round protocols of [Blu86, LS90] and inherits their “special soundness” and “proof-of-knowledge” properties; furthermore, the second round of the protocol consists of λn^2 bit-commitments where n is the number of nodes in the graph corresponding to the statement being proven. For concreteness, we denote the “knowledge” extractor of Π_{FS} by Ext and the zero-knowledge simulator by ZKSim . Without loss of generality, we assume that Ext^{P^*} , w.r.t. any cheating prover P^* , works by rewinding P^* only in last two messages: it rewinds P^* until it obtains two accepting transcripts from which a suitable witness can be obtained; the expected running time of Ext^{P^*} is $\text{poly}(\lambda)/p$ where p is the convincing probability of P^* w.r.t. an appropriate zk_3 .

Our commitment protocol $\langle \mathcal{C}, \mathcal{R} \rangle$. We now describe our protocol for committing messages of length n and supporting tags of length t where n, t can be arbitrary polynomials in the security parameter λ . Let (Enc, Dec) be an efficient split-state non-malleable code for messages of length $n + t$ against the class $\mathcal{F}_{\text{split}}$. For simplicity, we assume that the states are of length ℓ , and tampering functions map from $\{0, 1\}^\ell$ to $\{0, 1\}^\ell$. Let com and Π_{FS} be Naor’s commitment and Feige-Shamir protocols respectively as defined above.

The input to the commitment algorithm, \mathcal{C} , is a message $m \in \{0, 1\}^n$ and an identity $\text{id} \in \{0, 1\}^t$. The protocol proceeds in following stages:

Stage 1. \mathcal{C} and \mathcal{R} perform a coin-tossing protocol.

- (a) \mathcal{R} sends a random $\tau \leftarrow \{0, 1\}^{3\lambda}$ (the first message of Naor’s commitment)
- (b) \mathcal{C} commits to a random string ρ' i.e., $c = \text{com}_\tau(\rho'; r)$
- (c) \mathcal{R} sends a random string ρ''
- (d) \mathcal{C} sends ρ'
- (e) \mathcal{C} proves that: “ $\exists r : c = \text{com}_\tau(\rho'; r)$ ” using protocol Π_{FS} .
It uses com_τ as the commitment scheme for *every* commitment in Π_{FS} .
The transcript of this execution of Π_{FS} is denoted by $(\text{zk}_1, \text{zk}_2, \text{zk}_3, \text{zk}_4)$.
- (f) Define $\rho = \rho' \oplus \rho''$, and parse $\rho := \rho_1 \parallel \rho_{\text{FS}} \parallel \rho_2$

Stage 2. \mathcal{C} encodes the message $m \parallel \text{id}$, and commits to the first state *using* ρ_1 , i.e.,

- (a) $(L, R) \leftarrow \text{Enc}(m \parallel \text{id}; v)$
- (b) Send $c_1 \leftarrow \text{com}_{\rho_1}(L; v_1)$

Stage 3. \mathcal{C} and \mathcal{R} start a fresh execution of Π_{FS} where the statement x to be proven is decided in the last round.

Let G_x denote the Hamiltonian graph corresponding to x , and q be the number of nodes in G_x . \mathcal{C} and \mathcal{R} complete the first three rounds of Π_{FS} , denoted $(\text{zk}'_1, \text{zk}'_2, \text{zk}'_3)$ where every commitment in α'_2 is performed using a *unique part* of ρ_{FS} .² In the final round, \mathcal{C} sends the following:

- (a) Commitment to the second state: $c_2 \leftarrow \text{com}_{\rho_2}(R; v_2)$, and

²In more detail, recall that Π_{FS} consists of λ parallel repetitions of [Blu86, LS90], each of which requires committing a cycle as a $q \times q$ matrix. We view $\rho_{\text{FS}} = \{\rho_{i,j,k}\}$ where $\rho_{i,j,k}$ is of length 3λ and used to commit to (i, j) -th entry of the the matrix in k -th repetition using the function $\text{com}_{\rho_{i,j,k}}$, for every $i, j \in [q], k \in [\lambda]$. Also note that the size of x , and hence the value of q is a fixed and a-priori known polynomial in (λ, n, t, ℓ) .

(b) zk'_4 which is the last message of Π_{FS} proving the following statement x :

$$\exists(L, R, v_1, v_2, m) : c_1 = \text{com}_{\rho_1}(L; v_1) \wedge c_2 = \text{com}_{\rho_2}(R; v_2) \wedge \text{Dec}(L, R) = m \parallel \text{id}."$$

Stage 4. Com “opens” c_2 to R by sending v_2 .

To open the commitment, the committer simply reveals v_1 (and hence L). The message is recovered by decoding L and R (available as part of the commitment transcript).

This concludes the description of our commitment scheme. For concreteness, let us note that the size of ρ_1, ρ_2 is $3\lambda\ell$ each, size of statement x in 3(b) is $6\ell\lambda + t$ and its witness has length $(6\lambda + 2)\ell + n$, $q = q(\lambda, n, t, \ell)$ is a fixed polynomial defined by the NP-reduction from x to G_x , and ρ_{FS} has length $3\lambda \cdot \lambda q^2$. Strings ρ, ρ', ρ'' thus have length $3\lambda(2\ell + \lambda q^2)$ each.

Theorem 3. *Protocol $\langle \mathcal{C}, \mathcal{R} \rangle$ is a tag-based statistically-binding commitment scheme that is non-malleable with respect to commitments for messages and tags of arbitrary polynomial length.*

Proof. The scheme is statistically-binding for every efficient committers due to the fact that com_τ is statistically binding and protocol Π_{FS} is sound. It is computationally hiding because com_{ρ_1} and com_{ρ_2} are computationally hiding and the protocol Π_{FS} in stage 4 is zero-knowledge.

We now prove that the scheme is non-malleable w.r.t. commitments. We divide the proof in two parts: in Section 8.1, lemma 4, we prove that the scheme is non-malleable w.r.t. synchronous strategies; in Section 8.2, lemma 5, we prove that it is also non-malleable w.r.t. every M who employs any scheduling except synchronous. It follows that the scheme is non-malleable w.r.t. all adversaries since any successful M must also succeed, with at least half the advantage, by following one of these strategies. \square

8.1 Non-malleability against Synchronous Strategies

Proof overview. As discussed earlier, the proof will consider several hybrid experiments and create a situation where the commitments of man-in-the-middle to the two states can be separated into two tampering functions which tamper the states independently. This will be done by eventually obtaining a transcript that is statistically independent of the states, yet extraction of the states can still be performed “faithfully.”

Lemma 4. *$\langle \mathcal{C}, \mathcal{R} \rangle$ is non-malleable w.r.t. every synchronous adversary.*

Proof. Suppose that the lemma is not true, and there exists a synchronous man-in-the-middle adversary M , an efficient distinguisher D , and a polynomial p such that for infinitely many values of λ , there exist strings (m, m') and a tag id , such that $|\delta_\lambda - \delta'_\lambda| \geq 1/p(n)$ where $\delta_\lambda := \Pr[D(\text{MIM}_{\lambda, \text{id}, z}(m)) = 1]$, $\delta'_\lambda := \Pr[D(\text{MIM}_{\lambda, \text{id}, z}(m')) = 1]$ where variable $\text{MIM}_{\lambda, \text{id}, z}$ is defined w.r.t. M (and (λ, id, z) plays the role of auxiliary input, specified here for emphasis).

Fix any such $\lambda, m, m', \text{id}, z$. Recall that $\text{MIM}_{\lambda, \text{id}, z}$ outputs a pair of the form $(\tilde{\alpha}, \text{VIEW})$. Define ϵ_λ to be the probability that distribution $\text{MIM}_{\lambda, \text{id}, z}(m)$ outputs a VIEW that is accepting on right, i.e., M makes a successful commitment in the right session of VIEW. Define ϵ'_λ analogously w.r.t. m' . By computational hiding of our scheme, $\epsilon_\lambda, \epsilon'_\lambda$ are negligibly close.

We now design a sequence of hybrid experiments where the first hybrid \mathcal{H}_0 corresponds to sampling from the distribution $\text{MIM}_{\lambda, \text{id}, z}(m)$ and the last hybrid employs properties of non-malleable codes to sample a computationally indistinguishable output without knowing m .

\mathcal{H}_0 : This hybrid is identical to the experiment $\text{MIM}_{\lambda, \text{id}, z}(m)$. Specifically, it incorporates machines M and D internally, and interacts with M by committing m with identity id and simultaneously receiving a commitment from M w.r.t. some identity $\tilde{\text{id}}$. Let VIEW denote the view of M at the end of these executions. If VIEW is accepting on right, compute the (unique) value \tilde{m} committed by M (by running in exponential time) and return the output of $D(\tilde{m}, \text{VIEW})$.

By construction, VIEW accepts on right with probability ϵ_λ , and \mathcal{H}_0 outputs 1 with probability δ_λ .

\mathcal{H}_1 : This hybrid is identical to \mathcal{H}_0 except that it does not extract \tilde{m} in exponential time; instead it applies the “proof-of-knowledge” extractor to stage-3. We provide the (straightforward) details for future reference. \mathcal{H}_1 proceeds as follows:

1. Proceed identically to \mathcal{H}_0 until the second messages of stage-3 on both sides (zk'_3, \tilde{zk}'_3) have been sent.³ Let st denote the state of \mathcal{H}_1 up to this point.
2. Continue the execution from st identically to \mathcal{H}_0 to finish the last round. Let $VIEW$ be M 's view. If $VIEW$ is rejecting on right, or $\tilde{id} = id$, output $D(\perp, VIEW)$. Otherwise, let \tilde{x} be the statement proven by M in stage-3 on right. Also let \tilde{R} be the “second state” in the last message on right.
3. Define a “prover” machine $P_{L,v_1,m,\zeta}^{st}$ for proof system Π_{FS} as follows. $P_{L,v_1,m,\zeta}^{st}$ is the machine \mathcal{H}_0 initialized to state st with hardwired values (L, v_1, m) and sufficient randomness ζ ; it proceeds identically to \mathcal{H}_0 from state st but forwards the stage-3 messages on right to an outside verifier. It simulates the left execution for M internally, by computing the last message as follows: sample a random R so that it is consistent with L and $m||id$, i.e., $Dec(L, R) = m||id$, and then use (L, v_1, R, v_2, m) to send the last message where v_2 is chosen randomly.
4. Run the knowledge-extractor $Ext^{P_{L,v_1,m,\zeta}^{st}}$ to get a witness \tilde{w} which includes opening of commitment \tilde{c}_1 to a state \tilde{L} . Ignore the rest of the witness, compute $\tilde{m} = Dec(\tilde{L}, \tilde{R})$, and output $D(\tilde{m}, VIEW)$.

Observe that the distribution of $VIEW$ is identical in $\mathcal{H}_1, \mathcal{H}_0$, and Ext extracts a valid witness, and hence correct decommitment of \tilde{c}_1 , with probability $1 - \text{negl}(\lambda)$. The outputs of $\mathcal{H}_1, \mathcal{H}_0$ are thus statistically close.

The running time of \mathcal{H}_1 is expected polynomial as follows: for a state st , let ϵ_{st} denote the probability that $P_{L,v_1,m,\zeta}^{st}$ proves a statement; then, the expected running time of \mathcal{H}_1 is given by

$$\sum_{st} (\epsilon_{st} \cdot E \left[\text{Run time of } Ext^{P_{L,v_1,m,\zeta}^{st}} \right]) \cdot \Pr[st] = \sum_{st} \epsilon_{st} \cdot \frac{\text{poly}(\lambda)}{\epsilon_{st}} \cdot \Pr[st] = \text{poly}(\lambda).$$

\mathcal{H}_2 . This hybrid is identical to the previous hybrid except that in the extraction step, it extracts from the prover $P_{L,v_1,m^*,\zeta}^{st}$ where m^* is a random message. Note that this means that during extraction, machine $P_{L,v_1,m^*,\zeta}^{st}$ will now sample state R^* that is consistent with L and message $m^*||id$ (fresh every time), so that $Dec(L, R^*) = m^*||id$.

We prove that \mathcal{H}_1 and \mathcal{H}_2 are statistically close as follows. The extraction procedure $Ext^{P_{L,v_1,m^*,\zeta}^{st}}$ finds a witness which includes a valid opening of \tilde{c}_1 with high probability. However, since \tilde{c}_1 is statistically binding, and distributed identically in both hybrids, so is the opened state \tilde{L} . It follows that the two hybrids are statistically close. Next, we prove that:

Claim 13. \mathcal{H}_2 runs in expected polynomial time.

As noted earlier, the proof of this claim does not follow from standard averaging argument. This is because the success probabilities of $P_{L,v_1,m^*,\zeta}^{st}$ and $P_{L,v_1,m,\zeta}^{st}$ could be different. We therefore adopt a different strategy by considering several hybrids to move to a statistically, in fact perfectly, hiding transcript/view and then revert back. For better readability, we defer the proof of this claim towards the end.

Let us also note that at this point, we have already started to “split off” the extraction procedure for \tilde{L} without the help of R . We will now continue to make small changes until we the extraction of two states can be completely separated from each other.

\mathcal{H}_3 . This hybrid is identical to the previous hybrid, except that it uses the simulator $ZKSim$ to complete stage 1(e), instead of the real prover algorithm. All other steps are performed as in \mathcal{H}_2 , using honest values:

³Since M is synchronous, these messages are sent before the last rounds of both protocols start.

i.e., the hybrid defines st after stage-2 and completes stages 3 and 4 honestly. If the view is accepting, it performs extraction as before (using $P_{L,v_1,m^*,\zeta}^{\text{st}}$).

It is straightforward to see that \mathcal{H}_2 runs in expected polynomial time. This is because M is *synchronous* and hence the simulation and extraction steps are not intertwined. Furthermore, due to the zero-knowledge of Π_{FS} , the outputs of $\mathcal{H}_2, \mathcal{H}_1$ are computationally indistinguishable.

\mathcal{H}_4 : This hybrid is identical to \mathcal{H}_3 except that it sets ρ to be of the special form which allows equivocation. More specifically, it selects $N = 2\ell + \lambda q^2$ pairs of seeds of $\{(r_0^i, r_1^i)\}_{i=1}^N$ and sets $\rho := R_1 \| \dots \| R_N$ where $R_i = \text{prg}(r_0^i) \oplus \text{prg}(r_1^i)$ for all $i \in [N]$. \mathcal{H}_3 performs all other steps exactly as in \mathcal{H}_3 .

Clearly, \mathcal{H}_3 runs in expected polynomial time. Its output is computationally indistinguishable from that of \mathcal{H}_2 due to the pseudorandomness of prg .

Remark. At this stage, the commitments used in \mathcal{H}_4 (for stages 2, 3, and 4 of the protocol on left) are in fact perfectly hiding and independent of the underlying message.

\mathcal{H}_5 : This hybrid is actually identical to \mathcal{H}_3 , but we describe it differently using the equivocality property of com . Roughly speaking, the hybrid will simply sample a transcript randomly without knowing the message or the states, but then rely on the equivocality to later find randomness consistent with actual states, to compute the last message. We describe this hybrid in detail, since it will serve as the basis for describing our tampering functions in the next hybrid.

\mathcal{H}_5 proceeds as follows:

1. Complete stage-1 exactly as in \mathcal{H}_3 using ZKSim to set up $\rho = R_1 \| \dots \| R_N$ where $R_i = \text{prg}(r_0^i) \oplus \text{prg}(r_1^i)$ for all $i \in [N]$.
2. Complete stage-2 and first two messages of stage-3 by sending *equivocal* commitment strings. Specifically, set $c_1 := \text{prg}(r_0^1) \| \dots \| \text{prg}(r_0^\ell)$, $\text{zk}'_2 := \text{prg}(r_0^{\ell+1}) \| \dots \| \text{prg}(r_0^{\ell+\lambda q^2})$, and $c_2 := \text{prg}(r_0^{\ell+\lambda q^2+1}) \| \dots \| \text{prg}(r_0^N)$. Let st denote the state of the hybrid at this point. Continue the execution to also complete the third message of stage-3.
3. Complete the last message zk'_4 by using appropriate seeds, i.e., obtain $(L, R) \leftarrow \text{Enc}(m \| \text{id})$, and define randomness v_1, v_2 based on the bits of L, R , namely, $v_1 := (r_0^{L[1]} \| \dots \| r_0^{L[\ell]})$, and v_2 likewise. This defines the witness $w = (L, R, v_1, v_2, m)$ for statement x . Next, choose λ random cycles $C := (C_1, \dots, C_\lambda)$ and define randomness u so that zk'_2 is a commitment to C using u w.r.t. string ρ .⁴ Complete the last message, zk'_4 , using (w, C, u) .
4. Complete stage-4 by sending v_2 as computed above. Let VIEW be M 's view at this point. If VIEW rejects on right, output $D(\perp, \text{VIEW})$.
5. Otherwise, proceed as before: extract a witness from $P_{L,v_1,m^*,(C,u,\zeta')}^{\text{st}}$ where $\zeta = (C, u, \zeta')$ acts as the randomness. The extracted witness contains \tilde{L} and VIEW contains \tilde{R} in the last message; output $D(\tilde{m}, \text{VIEW})$ where m is obtained by decoding (\tilde{L}, \tilde{R}) .

Using step by step verification, it is easy to see that (\tilde{m}, VIEW) is distributed identically in both $\mathcal{H}_5, \mathcal{H}_4$. It is also easy to check that \mathcal{H}_5 runs in expected polynomial time.

Remark: Note that only step 3 above needs access to the message m for completing the last zero-knowledge message zk'_4 .

\mathcal{H}_6 : In this hybrid, instead of using m to compute the last zero-knowledge message zk'_4 , we use states of a random message. All other steps are still performed using correct states (L, R) . Formally, \mathcal{H}_6 starts by sampling two encodings: $(L, R) \leftarrow \text{Enc}(m \| \text{id})$ and $(\hat{L}, \hat{R}) \leftarrow \text{Enc}(\hat{m} \| \text{id})$ where \hat{m} is chosen randomly. It then proceeds as follows:

⁴Informally, for every bit $b_{i,j,k} = C_k[i, j]$, define $u_{i,j,k} = r_{b_{i,j,k}}^{i,j,k}$ so that u is the collection of all $u_{i,j,k}$.

1. Proceed identically to \mathcal{H}_4 up to the first three messages of stage-3
2. Obtain randomness $\widehat{v}_1, \widehat{v}_2$ w.r.t. \widehat{L}, \widehat{R} as before to define the witness $\widehat{w} = (\widehat{L}, \widehat{v}_1, \widehat{R}, \widehat{v}_2, \widehat{m})$. Compute message zk'_4 using (\widehat{w}, C, u) as in the previous hybrid and finish stage-3.
3. Finish stage-4 using the *correct* states as in the previous hybrid. I.e., find randomness v_1, v_2 w.r.t. L, R and send v_2 in stage-4. Let VIEW^* denote M 's view at this point.
4. If VIEW^* is accepting on right, extract \widetilde{L} from $P_{L, v_1, m^*, (C, u, \zeta')}^{\text{st}}$ as before and use D to return the output.

The only difference between two hybrids is in the computation of zk'_4 . However, zk'_4 is independent of the underlying “witness” due to the perfect-hiding property of underlying equivocal commitments. Therefore $\mathcal{H}_6, \mathcal{H}_5$ are statistically close.

\mathcal{H}_7 : In this hybrid we first define a family of tampering functions which use \mathcal{H}_6 internally. Observe that \mathcal{H}_6 does not use (L, R) until the last message (which consists of zk'_4 , and v_2). Furthermore, it computes \widetilde{R} without the knowledge of L and, likewise, extracts \widetilde{L} without the knowledge of R . Therefore, we already have the ability to perform extraction (or “tampering”) individually. With this observation, we define the following family $\mathcal{F}_{\text{split}}$ of tampering functions.

Let Φ denote the set of all possible random tapes for hybrid \mathcal{H}_6 . Define $F = \{(f_\phi, g_\phi)\}_{\phi \in \Phi}$ where f_ϕ, g_ϕ are defined below (we define g_ϕ first):

1. **Function** $g_\phi : \{0, 1\}^\ell \rightarrow \{0, 1\}^\ell$ takes as input a state R . It proceeds exactly as \mathcal{H}_6 with randomness ϕ up to the completion of stage-4 where VIEW^* is obtained. Let \widetilde{R} denote the state opened by M in stage-4 on right. Output \widetilde{R} .
2. **Function** $f_\phi : \{0, 1\}^\ell \rightarrow \{0, 1\}^\ell$ takes as input a state s_1 . It proceeds exactly as \mathcal{H}_6 with randomness ϕ up to the completion of stage-3. It does not complete stage-4; instead, it directly runs the extractor exactly as \mathcal{H}_6 , and outputs the first component of the extracted witness, denoted L .

We now describe \mathcal{H}_7 using functions in F . Specifically, \mathcal{H}_7 , samples states $(L, R) \leftarrow \text{Enc}(m \parallel \text{id})$, and proceeds identically to \mathcal{H}_6 with fresh randomness ϕ and state R to obtain VIEW^* . However, \mathcal{H}_6 does not run the extractor directly. Instead, if VIEW^* is accepting, it computes

$$\widetilde{L} = f_\phi(L), \quad \widetilde{R} = g_\phi(R), \quad \widetilde{m} \parallel \text{id} = \text{Dec}(\widetilde{L}, \widetilde{R}),$$

and outputs $D(\widetilde{m}, \text{VIEW}^*)$ if $\widetilde{\text{id}} \neq \text{id}$. Otherwise it outputs $D(\perp, \text{VIEW}^*)$.

By construction, $\mathcal{H}_6, \mathcal{H}_7$ are identical. Let $\delta_\lambda^{(7)}$ denote the probability that \mathcal{H}_7 outputs 1. We have that $|\delta_\lambda^{(7)} - \delta_\lambda| \leq \text{negl}(\lambda)$. We claim that $\delta_\lambda^{(7)}$ is independent of m . We prove this by obtaining an explicit expression for $\delta_\lambda^{(7)}$ using the non-malleability of the code.

Let us fix some notation. Observe that fixing a random tape ϕ , completely fixes variable VIEW^* sampled above. Therefore, we will write VIEW_ϕ^* to refer to VIEW^* corresponding to ϕ . Let \mathbb{A} denote the set of VIEW^* that are “accepting” on right. Note that we have already argued that VIEW^* is independent of the message m . Therefore, $\Pr[\phi \in \mathbb{A}]$ is also independent of m . Also observe that sampling of $\widetilde{m} \parallel \text{id}$ in \mathcal{H}_7 w.r.t. tape ϕ is identical to sampling from the “tampered distribution” $\text{T}_{m \parallel \text{id}, f_\phi, g_\phi}$ and this distribution is simulatable, i.e., there exists negligible function negl_ϕ and a simulator $\text{Sim}_{f_\phi, g_\phi}$ such that $\Delta(\text{T}_{m \parallel \text{id}, f_\phi, g_\phi}, \text{Copy}_{\text{Sim}_{f_\phi, g_\phi}}^{(m \parallel \text{id})}) = \text{negl}_\phi(\lambda)$ where $\text{Copy}_{\text{Sim}_{f_\phi, g_\phi}}^{(m \parallel \text{id})}$ samples $\widetilde{m} \leftarrow \text{Sim}_{f_\phi, g_\phi}$ and outputs $m \parallel \text{id}$ if $\widetilde{m} = \text{same}$ and \widetilde{m} otherwise.

We have,

$$\begin{aligned}
\delta_\lambda^{(7)} &= \sum_{\phi \in \mathbb{A}} \Pr \left[D(\tilde{m}, \text{VIEW}_\phi^*) = 1 \mid \begin{array}{l} \tilde{m} \parallel \tilde{\text{id}} \leftarrow \text{T}_{m \parallel \text{id}, f_\phi, g_\phi} \\ \text{if } \tilde{\text{id}} = \text{id set } \tilde{m} := \perp \end{array} \right] \cdot \Pr[\phi] + \underbrace{\sum_{\phi \notin \mathbb{A}} \Pr [D(\perp, \text{VIEW}_\phi^*) = 1] \cdot \Pr[\phi]}_{=p \text{ (independent of } m)} \\
&= \sum_{\phi \in \mathbb{A}} \Pr \left[D(\tilde{m}, \text{VIEW}_\phi^*) = 1 \mid \begin{array}{l} \tilde{m} \parallel \tilde{\text{id}} \leftarrow \text{Copy}_{\text{Sim}_{f_\phi, g_\phi}}^{(m \parallel \text{id})} \\ \text{if } \tilde{\text{id}} = \text{id set } \tilde{m} := \perp \end{array} \right] \cdot \Pr[\phi] + \sum_{\phi \in \mathbb{A}} \pm \text{negl}_\phi \cdot \Pr[\phi] + p \\
&= \sum_{\phi \in \mathbb{A}} \Pr \left[D(\tilde{m}, \text{VIEW}_\phi^*) = 1 \mid \begin{array}{l} \beta \leftarrow \text{Sim}_{f_\phi, g_\phi} \\ \text{if } \beta = \text{same set } \tilde{m} := \perp \\ \text{else } \tilde{m} := \beta[1 \dots n] \end{array} \right] \cdot \Pr[\phi] + \underbrace{\sum_{\phi \in \mathbb{A}} \pm \text{negl}_\phi \cdot \Pr[\phi] + p}_{=\text{negligible}} \\
&= \text{value independent of } m.
\end{aligned}$$

Now consider the case of m' and the variable δ'_λ w.r.t. $\text{MIM}_{\lambda, \text{id}, z}(m')$. Using the same argument as above, we can conclude that $\text{MIM}_{\lambda, \text{id}, z}(m')$ is also computationally indistinguishable from hybrid \mathcal{H}_7 . Therefore, $|\delta'_\lambda - \delta_\lambda^{(7)}| \leq \text{negl}(\lambda)$ and hence $|\delta_\lambda - \delta'_\lambda| \leq 4\text{negl}(\lambda)$ which is a contradiction. \square

Proof of claim 13. To prove that \mathcal{H}_2 runs in expected polynomial time, we have to consider a series of hybrid experiments to obtain statistically independent transcripts. At this point, we can calculate the running time in a standard manner since there would be no difference in success probabilities when we switch m to m^* in the prover machine. We now provide a sketch of the hybrids (since they are very similar to our previous hybrids).

Consider the following sequence of hybrids:

- $\mathcal{H}_{1:1}$. Identical to \mathcal{H}_1 except that it uses ZKSim to complete stage 1(e). Note that $\mathcal{H}_{1:1}$ is expected polynomial time; furthermore, the extracted witness and in particular its first component $\tilde{\text{L}}$ is computationally indistinguishable from \mathcal{H}_1 . (Note that M is still synchronous)
- $\mathcal{H}_{1:2}$. Identical to $\mathcal{H}_{1:1}$ except that we switch the strings ρ to be of the special form. (This is done in the same manner as hybrid \mathcal{H}_4 and we rely on the security of prg). Note that at this point the transcripts are independent of the underlying witness. (See \mathcal{H}_5 for a detailed explanation)
- $\mathcal{H}_{1:3}$. Identical to $\mathcal{H}_{1:2}$ except that the extraction is now performed using the prover machine $P_{L, v_1, m^*, \zeta}^{\text{st}}$ where m^* is a random message.

We claim that $\mathcal{H}_{1:3}$ runs in expected polynomial time. Let ϵ_{st} denote the success probability of prover $P_{L, v_1, m, \zeta}^{\text{st}}$ (in the previous hybrid). Observe that the two hybrids define the same distribution on states st . The only difference between $P_{L, v_1, m, \zeta}^{\text{st}}$ and $P_{L, v_1, m^*, \zeta}^{\text{st}}$ is that they compute the last zero-knowledge message zk'_4 (of stage-3) using a different witness. However, since the protocol transcript hides the witness perfectly, we have that for every st , $\epsilon_{\text{st}} = \epsilon_{\text{st}}^*$ where ϵ_{st}^* is the success probability of $P_{L, v_1, m^*, \zeta}^{\text{st}}$. The running time of $\mathcal{H}_{1:3}$ is thus given by $\sum_{\text{st}} \left(\epsilon_{\text{st}} \cdot \mathbb{E} \left[\text{Run time of Ext}_{L, v_1, m^*, \zeta}^{\text{st}} \right] \right) \cdot \Pr[\text{st}] = \sum_{\text{st}} \epsilon_{\text{st}} \cdot \frac{\text{poly}(\lambda)}{\epsilon_{\text{st}}^*} \cdot \Pr[\text{st}] = \text{poly}(\lambda)$.

Note that we have already switched to the prover $P_{L, v_1, m^*, \zeta}^{\text{st}}$. Now we simply reverse the hybrids: in $\mathcal{H}_{1:4}$ we switch back to using a normal string ρ , and in $\mathcal{H}_{1:5}$ we start using the real prover algorithm instead of ZKSim (details are straightforward). Observe that $\mathcal{H}_{1:4} = \mathcal{H}_2$, and hence the claim. \square

8.2 Non-malleability against Asynchronous Strategies

Lemma 5. $\langle \mathcal{C}, \mathcal{R} \rangle$ is non-malleable w.r.t. every asynchronous adversary.

Proof. We now consider an asynchronous adversary who follows any strategy different from the synchronous. We perform a case by case analysis, depending upon the schedule of the adversary, and show that each case either reduces to the synchronous case, or it is trivially non-malleable due to the hiding and extractability properties of the scheme. Specifically, there are three representative cases, as described below (and depicted pictorially in figure 4):⁵

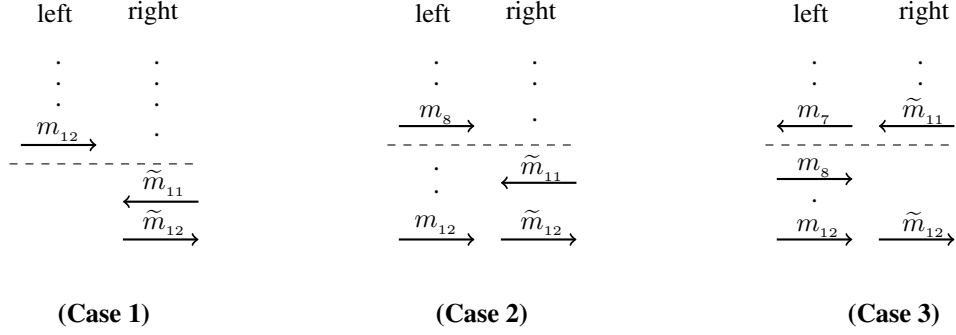


Figure 4: Representative schedules for an asynchronous adversary for proof of lemma 5.

Case 1: The last two messages on right appear after all messages on left.

This case is trivially non-malleable due to the hiding of left side commitment since \tilde{s}_1 on right can be extracted without rewinding the left protocol to obtain the committed message \tilde{m} .

Note: If the last message on right, \tilde{m}_{12} , appears after the left protocol ends, then also the protocol is non-malleable since M has committed to his message before the message on left has been fixed. Therefore, in the following cases, we only need to consider the case when \tilde{m}_{12} is synchronous with m_{12} .

Case 2: The last two messages on right appear after stage-1 Π_{FS} on left.

This case is handled the same way as the synchronous case. Specifically, since the extraction on right still does not interfere with simulation on left, the proof for this case goes through exactly as for the synchronous case.

Case 3: The non-malleability for this case also follows from hiding of left side commitment, because the last two message on right “completely contain” the stage-3 and 4 of left protocol. Therefore, rewinding on right will completely rewind this stage (past the first message) and hence computational hiding of left protocol is maintained. We provide a proof sketch.

As in proof of lemma 4, fix λ, M, D, z and values $\epsilon_\lambda, \delta_\lambda, \delta'_\lambda$ for one of infinitely many λ such that ϵ_λ (the probability of a successful commitment on right) is a inverse polynomial so is the difference between $\delta_\lambda, \delta'_\lambda$. We perform extraction for this λ and violate hiding of stage-4 commitment (which is actually a ZK protocol but can be viewed as a commitment to m (or m') through commitment to its first state and the second state given in plain) for this λ .

Consider a machine D' which internally incorporates M and D and has been initialized up to the point where it awaits \tilde{m}_{11} on right (see fig. 4). It gets a commitment from outside (to either m or m') and forwards it internally to M . If M fails on right, D' outputs a random bit as the guess. Otherwise, D' proceeds for extraction (from the last two messages)—while completing the left side execution using a *random* message whenever needed.

⁵Before reading the cases/figure, let us note that the messages of the left protocol are denoted by m_1, \dots, m_{12} where m_8 denotes the end of stage-1 as well as 2 (since both messages are sent together), and m_7 is the challenge for stage-1 Π_{FS} on left. The corresponding messages on right are denoted by $\tilde{m}_1, \dots, \tilde{m}_{12}$.

Let ϵ_λ^* denote the probability that M makes a successful commitment on right when the left side is completed with a random message m^* . Clearly, $\epsilon_\lambda, \epsilon_\lambda^*$ are negligibly close. Note that by our assumption, ϵ_λ is an inverse polynomial, and therefore so must be ϵ_λ^* . After extraction, D' outputs whatever D would.

The expected running time of D' is dominated by the term $\epsilon_\lambda \cdot \frac{\text{poly}(\lambda)}{\epsilon_\lambda^*}$; which is bounded by a polynomial due to our assumptions on $\epsilon_\lambda, \epsilon_\lambda^*$. It is clear that in this case the value is extracted with probability $1 - \text{negl}(\lambda)$ and the output is the same as that of D in the real execution with m (resp. m'). Therefore D' breaks hiding if D breaks non-malleability.

□

9 Acknowledgements

We would like to thank Divesh Aggarwal, Anand Degwekar, Yevgeniy Dodis and Antigoni Polychroniadou for helpful discussions and insights.

References

- [AAG⁺16] Divesh Aggarwal, Shashank Agrawal, Divya Gupta, Hemanta Maji, Omkant Pandey, and Manoj Prabhakaran. Optimal computational split-state non-malleable codes. In *TCC*, 2016.
- [ADKO15] Divesh Aggarwal, Yevgeniy Dodis, Tomasz Kazana, and Maciej Obremski. Non-malleable reductions and applications. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pages 459–468, 2015.
- [ADL14] Divesh Aggarwal, Yevgeniy Dodis, and Shachar Lovett. Non-malleable codes from additive combinatorics. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 774–783, 2014.
- [Agg] Divesh Aggarwal. Personal communication, 10/30/2015.
- [Bar02] Boaz Barak. Constant-Round Coin-Tossing with a Man in the Middle or Realizing the Shared Random String Model. In *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science, FOCS '02*, pages 345–355, 2002.
- [BGR⁺15] Hai Brenner, Vipul Goyal, Silas Richelson, Alon Rosen, and Margarita Vald. Fast non-malleable commitments. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, Denver, CO, USA, October 12-6, 2015*, pages 1048–1057, 2015.
- [Blu81] Manuel Blum. Coin flipping by telephone. In *Advances in Cryptology: A Report on CRYPTO 81, CRYPTO 81, IEEE Workshop on Communications Security, Santa Barbara, California, USA, August 24-26, 1981.*, pages 11–15, 1981.
- [Blu86] Manuel Blum. How to prove a theorem so no one else can claim it. In *Proceedings of the International Congress of Mathematicians, Berkeley, CA*, pages 1444–1451, 1986.
- [CGL15] Eshan Chattopadhyay, Vipul Goyal, and Xin Li. Non-malleable extractors and codes, with their many tampered extensions. *CoRR*, abs/1505.00107, 2015.
- [CGMO09] Nishanth Chandran, Vipul Goyal, Ryan Moriarty, and Rafail Ostrovsky. Position based cryptography. In Shai Halevi, editor, *CRYPTO*, volume 5677 of *Lecture Notes in Computer Science*, pages 391–407. Springer, 2009.

- [CLOS02] Ran Canetti, Yehuda Lindell, Rafail Ostrovsky, and Amit Sahai. Universally composable two-party and multi-party secure computation. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, STOC '02, pages 494–503, 2002.
- [CZ14] Eshan Chattopadhyay and David Zuckerman. Non-malleable codes against constant split-state tampering. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 306–315, 2014.
- [DDN91] Danny Dolev, Cynthia Dwork, and Moni Naor. Non-Malleable Cryptography (Extended Abstract). In *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing*, STOC '91, pages 542–552, 1991.
- [DKO13] Stefan Dziembowski, Tomasz Kazana, and Maciej Obremski. Non-malleable codes from two-source extractors. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part II*, volume 8043 of *Lecture Notes in Computer Science*, pages 239–257. Springer, 2013.
- [DPW10] Stefan Dziembowski, Krzysztof Pietrzak, and Daniel Wichs. Non-malleable codes. In *Innovations in Computer Science - ICS 2010, Tsinghua University, Beijing, China, January 5-7, 2010. Proceedings*, pages 434–452, 2010.
- [FS89] Uriel Feige and Adi Shamir. Zero knowledge proofs of knowledge in two rounds. In *Advances in Cryptology - CRYPTO '89, 9th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 1989, Proceedings*, pages 526–544, 1989.
- [FS90] Uriel Feige and Adi Shamir. Witness indistinguishable and witness hiding protocols. In *STOC*, pages 416–426. ACM, 1990.
- [GK96] Oded Goldreich and Hugo Krawczyk. On the composition of zero-knowledge proof systems. *SIAM J. Comput.*, 25(1):169–192, 1996.
- [GLOV12] Vipul Goyal, Chen-Kuei Lee, Rafail Ostrovsky, and Ivan Visconti. Constructing non-malleable commitments: A black-box approach. In *FOCS*, pages 51–60. IEEE Computer Society, 2012.
- [GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In *Proceedings of the 19th Annual ACM Symposium on Theory of Computing, 1987, New York, New York, USA*, pages 218–229, 1987.
- [Goy11] Vipul Goyal. Constant Round Non-malleable Protocols Using One-way Functions. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing*, STOC '11, pages 695–704. ACM, 2011.
- [GRRV14] Vipul Goyal, Silas Richelson, Alon Rosen, and Margarita Vald. An algebraic approach to non-malleability. In *FOCS*, 2014.
- [HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A Pseudorandom Generator from any One-way Function. *SIAM J. Comput.*, 28(4):1364–1396, 1999.
- [IKOS07] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Zero-knowledge from Secure Multiparty Computation. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing*, STOC '07, pages 21–30, 2007.

- [Kiy14] Susumu Kiyoshima. Round-efficient black-box construction of composable multi-party computation. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part II*, volume 8617 of *Lecture Notes in Computer Science*, pages 351–368. Springer, 2014.
- [KMO14] Susumu Kiyoshima, Yoshifumi Manabe, and Tatsuaki Okamoto. Constant-round black-box construction of composable multi-party computation protocol. In Yehuda Lindell, editor, *Theory of Cryptography - 11th Theory of Cryptography Conference, TCC 2014, San Diego, CA, USA, February 24-26, 2014. Proceedings*, volume 8349 of *Lecture Notes in Computer Science*, pages 343–367. Springer, 2014.
- [LP09] Huijia Lin and Rafael Pass. Non-malleability Amplification. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing*, STOC '09, pages 189–198, 2009.
- [LP11] Huijia Lin and Rafael Pass. Constant-round Non-malleable Commitments from Any One-way Function. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing*, STOC '11, pages 705–714, 2011.
- [LP12] Huijia Lin and Rafael Pass. Black-box constructions of composable protocols without set-up. In Reihaneh Safavi-Naini and Ran Canetti, editors, *Advances in Cryptology - CRYPTO 2012 - 32nd Annual Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2012. Proceedings*, volume 7417 of *Lecture Notes in Computer Science*, pages 461–478. Springer, 2012.
- [LPV08] Huijia Lin, Rafael Pass, and Muthuramakrishnan Venkitasubramaniam. Concurrent Non-malleable Commitments from Any One-Way Function. In *Theory of Cryptography, 5th Theory of Cryptography Conference, TCC 2008*, 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 *Proceedings of the 41st Annual ACM Symposium on Theory of Computing*, STOC '09, pages 179–188, 2009.
- [LS90] Dror Lapidot and Adi Shamir. Publicly verifiable non-interactive zero-knowledge proofs. In *CRYPTO*, pages 353–365, 1990.
- [Nao91] Moni Naor. Bit Commitment Using Pseudorandomness. *J. Cryptology*, 4(2):151–158, 1991.
- [Pas13] Rafael Pass. Unprovable security of perfect NIZK and non-interactive non-malleable commitments. In *TCC*, pages 334–354, 2013.
- [PPV08] Omkant Pandey, Rafael Pass, and Vinod Vaikuntanathan. Adaptive One-Way Functions and Applications. In *Advances in Cryptology — CRYPTO '08*, pages 57–74, 2008.
- [PR05a] Rafael Pass and Alon Rosen. Concurrent Non-Malleable Commitments. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '05, pages 563–572, 2005.
- [PR05b] Rafael Pass and Alon Rosen. New and improved constructions of non-malleable cryptographic protocols. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, STOC '05, pages 533–542, 2005.
- [PW10] Rafael Pass and Hoeteck Wee. Constant-Round Non-malleable Commitments from Sub-exponential One-Way Functions. In *Advances in Cryptology — EUROCRYPT '10*, pages 638–655, 2010.
- [Wee10] Hoeteck Wee. Black-Box, Round-Efficient Secure Computation via Non-malleability Amplification. In *Proceedings of the 51th Annual IEEE Symposium on Foundations of Computer Science*, pages 531–540, 2010.

A Failed Attempts: Selective \perp Attacks

We describe a few “false starts” which exemplify the difficulty in attaining three round non-malleable commitment. We restrict our attention to mauling attacks which could be ruled out with a zero-knowledge proof of correctness as other attacks have already been addressed in prior work. We leave out the identities for this discussion. To start with (and disregarding the lower bound of [Pas13]), consider the non-interactive “encode-then-encrypt” protocol where \mathcal{C} commits to v by drawing $(L, R) \leftarrow \text{Enc}(v)$ and sending $\text{Com}(L), \text{Com}(R)$ to \mathcal{R} . One might hope that committing separately to the states would allow us to argue that in any mauling attack, the PPT M must treat the states individually and so non-malleability would follow from that of the code. This intuition does not work as Com possesses no non-malleability guarantees, and so we cannot rule out the possibility that M will maul (L, R) jointly to (\tilde{L}, \tilde{R}) , an encoding of say $v + 1$.

Consider instead the scheme where \mathcal{C} sends $\text{Com}(L)$ in the first round, \mathcal{R} checks back with an acknowledgement message, then \mathcal{C} sends R . This protocol seems like it should be non-malleable: M is forced to maul L to \tilde{L} before he sees R , and the hiding of Com should ensure that M is mauling R without knowledge of L . However, this intuition turns out not to be sound. Given any non-malleable code $(\text{Enc}_0, \text{Dec}_0)$, consider the pathological non-malleable code where $\text{Enc}(m)$ samples $(L_0, R_0) \leftarrow \text{Enc}_0(m)$ and outputs $L = (L_0, \sigma)$ and $R = (R_0, z)$ where $\sigma, z \leftarrow \$$; $\text{Dec}(L, R) = \text{Dec}_0(L_0, R_0)$ unless z is a commitment to the message (L_0, σ) using randomness σ , in which case $\text{Dec}(L, R) = \text{Dec}_0(L_0, R_0) + 1$. It is easy to show that (Enc, Dec) is non-malleable if $(\text{Enc}_0, \text{Dec}_0)$ is. But now if the above protocol is instantiated on top of such a code, M can maul $\text{Com}((L_0, \sigma); \tau)$ to $\tilde{z} = \text{Com}((L_0, \tilde{\sigma}); \tilde{\sigma})$ (this mauling is happening under the Com so M needn’t know $\tilde{\sigma}$ to perform such an attack), then M can maul (R_0, z) to (R_0, \tilde{z}) . M has successfully committed to $v + 1$. It is important to force M to demonstrate knowledge of L in the third round.

Finally, consider a version of our protocol but without any non-malleable codes. In this scheme \mathcal{C} sends $\text{Com}(v)$ and $\text{Com}(r)$ for a random $r \leftarrow \mathbb{Z}_q$, \mathcal{R} sends a random $\alpha \leftarrow \mathbb{Z}_q$ and \mathcal{C} answers with $r\alpha + v$. This is the main subprotocol from the [GRRV14] scheme but with the zero-knowledge proof of consistency removed. Also forget about the copying attack available to M in this discussion as this can be prevented using non-malleability amplification. Still, M has the following selective \perp attack available. Upon receiving $\text{Com}(v), \text{Com}(r)$ M produces $\text{Com}(\tilde{v}), \text{Com}(v)$ on the right, where \tilde{v} is the (unrelated) value to which he wants to commit. He disregards the remainder of the left protocol. Upon receiving $\tilde{\alpha}$, M answers with $\tilde{\alpha} + \tilde{v}$. If $v = 1$, then M has successfully committed to \tilde{v} ; if $v \neq 1$, he has committed to \perp . Such selective \perp attacks seem to be available to M in any protocol in which \mathcal{C} ’s message is fully specified in the first round.