2016 IEEE 57th Annual Symposium on Foundations of Computer Science

Breaking the Three Round Barrier for Non-malleable Commitments

Vipul Goyal
Microsoft Research
Bangalore, India
Email: vipul@microsoft.com

Dakshita Khurana
Department of Computer Science
University of California, Los Angeles
Los Angeles, California 90024
Email: dakshita@cs.ucla.edu

Amit Sahai
Department of Computer Science
University of California, Los Angeles
Los Angeles, California 90024
Email: sahai@cs.ucla.edu

***Abstract*—We construct two-message non-malleable commitments with respect to opening in the standard model, assuming only one-to-one one-way functions. Our protocol consists of two unidirectional messages by the committer (with no message from the receiver), and is secure against all polynomial-time adversaries in the standard synchronous setting.**

Pass (TCC 2013) proved that any commitment scheme with non-malleability with respect to commitment, using only 2 rounds of communication, cannot be proved secure via a black-box reduction to any "standard" intractability assumption. We extend this by showing a similar impossibility result for commitments with non-malleability with respect to *opening*, another standard notion of non-malleability for commitments, for any 2-message challenge-response **protocol, as well.**

However, somewhat surprisingly, we show that this barrier breaks down in the setting of two unidirectional messages by the committer (with no message from the receiver), for non-malleability with respect to opening.

- **Our protocol makes only *black-box* use of any non-interactive statistically binding commitment scheme. Such a scheme can be based on any one-to-one one-way function.**
- **Our techniques depart significantly from the** commit-challenge-response **structure followed by nearly all prior works on non-malleable protocols in the standard model. Our methods are combinatorial in nature.**
- **Our protocol resolves the round complexity of commitments with non-malleability with respect to opening via natural (non-embedding) black-box security reductions. We show that completely non-interactive non-malleable commitments w.r.t. opening cannot be proved secure via most natural black-box reductions. This result extends to also rule out bi-directional two-message non-malleable commitments w.r.t. opening in the synchronous or asynchronous setting.**
- **Our protocol, together with our impossibility result, also resolves the round complexity of block-wise non-malleable codes (Chandran et al) w.r.t. natural black-box reductions.**

I. Introduction

Man-in-the-middle (MIM) adversaries participate in two or more instantiations of a protocol, and try to use information obtained in one execution to breach security in the other protocol execution. Non-malleable commitments were introduced in the seminal work of Dolev, Dwork and Naor [1] as countermeasures against such attacks. They have proved to be useful and versatile building blocks in the construction of non-malleable cryptographic protocols such as coin-flipping, non-malleable proof systems (zero-knowledge, witness indistinguishable and even multi-prover interactive proofs), and multiparty computation protocols.

A commitment scheme is a two-party protocol between a committer and a receiver. The committer has a message m as input, while the receiver is inputless. The two parties engage in a probabilistic interactive commitment protocol, and we denote the receiver's view of this protocol by $\mathsf{com}(m)$. Then, later, the committer sends to the receiver an opening message, that allows the receiver to confirm that the message m was really the message committed to during the commitment protocol. In a statistically binding commitment, the receiver's view $\mathsf{com}(m)$ should be binding in the sense that with high probability, there should not exist an opening message that would convince the receiver that the committer had used any string $m' \neq m$. In short, we say that the commitment cannot be later opened to any message $m' \neq m$. A commitment should also be computationally hiding; that is, for any pair of messages (m, m') the distributions $\mathsf{com}(m)$ and $\mathsf{com}(m')$ should be computationally indistinguishable. Informally speaking, such a scheme is said to be *non-malleable* if for every message m, no MIM adversary, intercepting a commitment protocol $\mathsf{com}(m)$ and modifying every message sent during this protocol arbitrarily, is able to efficiently *generate* a commitment $\mathsf{com}(\tilde{m})$ that can be opened to a message $\tilde{m}$ related to the original message m.

The original construction of non-malleable commitments of [1] was conceptually simple and efficient, because of its instantiability with highly efficient cryptographic sub-protocols. However, it required logarithmically many rounds. Subsequently, Barak [2], Pass [3], and Pass and Rosen [4] constructed constant-round

0272-5428/16 $31.00 © 2016 IEEE
DOI 10.1109/FOCS.2016.12

protocols relying on highly inefficient non-black box techniques. Wee [5] gave a constant-round black-box construction of non-malleable commitments assuming the sub-exponential hardness of one-way functions. Goyal [6] and Lin and Pass [7] constructed constant-round non-malleable commitments under the minimum assumption that one-way functions exist. Goyal, Lee, Ostrovsky and Visconti [8] converted the protocol of Goyal into a fully black-box protocol based on any one-way function.

The last five years have seen significant progress in understanding the necessity for interaction in non-malleable commitments, in terms of the concrete number of messages required. In particular, Goyal, Richelson, Rosen and Vald [9] constructed four round non-malleable commitments in the standard model based on the existence of one-way functions. Finally, Goyal, Pandey and Richelson [10] constructed three round non-malleable commitments via a black-box use of injective one-way functions, by exploiting properties of non-malleable codes.

A. The Three Round Barrier

Pass [11] showed an impossibility for non-malleable commitments using 2 rounds of communication, via a black-box reduction to any "standard" intractability assumption. Pass gave this result for the stronger of two standard notions of non-malleability: This stronger notion is called non-malleability with respect to commitment, and it requires that the adversary cannot even create a valid commitment $\mathsf{com}(\tilde{m})$ to a related string $\tilde{m}$. Very roughly, Pass' impossibility result holds because at least three rounds are required by a simulator to perform extraction via a commit-challenge-response paradigm. Such an extraction is required for proving non-malleability with respect to commitment.

The other standard notion is called non-malleability with respect to *opening* [12], [13], [4]. This notion requires that no matter what commitment com the adversary creates, the adversary will not be able to open it to a string $\tilde{m}$ related to the string committed by the honest party.

In fact, we are able to generalize the result of Pass to rule out 2-round challenge-response commitment protocols also for non-malleability with respect to opening, whose security is to be based on natural (so-called "non-embedding") black-box reductions. Therefore, going below three rounds seems like a natural barrier for non-malleable commitments with security proven via black-box reductions. In this paper, *we ask whether it is possible to circumvent this barrier.*

Surprisingly, we show that it is possible to circumvent this barrier for non-malleability with respect to opening, by considering *uni-directional* commitment protocols where the committer sends two messages, and the receiver sends no message at all. The uni-directional message model can be seen as a relaxation of the non-interactive message model where the protocol proceeds in phases and in each phase, the first party sends a message to the second party. In the context of 2-round non-malleable commitments, the man-in-the-middle sees the committer message on the left in the first phase, and can then decide its own first message in an arbitrary manner. Next, the man-in-the-middle sees the second committer message on the left and then decides its own second message. This marks a significant departure from all previous non-malleable commitment protocols in the plain model, which generally engaged in at least one round of challenge-response in order to establish non-malleability.

To illustrate our uni-directional message model, consider the setting of sealed-bid auctions. After obtaining a commitment to some bid b, we would like to ensure that the adversary is unable to generate *and open* a valid commitment to the related bid $(b+1)$. In our setting, the commitment phase of the auction would proceed in two stages. In the first stage (say on day 1), all the committers would be required to submit their first message (of the non-malleable commitment scheme), and, in the second stage (on day 2), the committers must submit their second message. Then assume that on each day, a (rushing) adversary Bob gets to see an honest bidder Alice's messages (committing to b) and decides its own. Our protocol would ensure that the Bob cannot open his bid to be $(b+1)$, or to any other message related to Alice's bid b. Furthermore in applications like secure multi-party computation (where the round complexity is determined by the round complexity of non-malleable commitments), such a model is a natural fit since we assume that all the parties send their i-th round message before seeing the $(i+1)$-th round message of any party.

B. Our Results

Our main result establishes the existence of a two message commitment scheme that is non-malleable with respect to opening, based only on the assumption that injective one-way functions exist. We now elaborate:

- As discussed above, our protocol is *uni-directional* – that is, both messages are sent from the committer to the receiver. We prove security in the standard synchronous setting for non-malleability, where the adversary can modify each successive message of the protocol arbitrarily (we also discuss the asynchronous setting in our impossibility results below). We have already discussed above that the uni-directionality of our protocol is motivated by the impossibility of constructing such a protocol in the more traditional two-message challenge-response setting. But in fact, uni-directional protocols like ours enjoy other notable features:

- Our construction provides security even in cases where the honest receiver in the right execution is partially corrupted. For example, if an adversary could somehow corrupt the randomness of the honest receiver (or learn it in advance before the protocol starts), the security of existing non-malleable commitment protocols such as [10] would completely fall apart. However, in our protocol, the receiver literally does nothing until receiving the final opening message of the committer. Thus, there is no randomness to corrupt, and security would hold even against such a strengthened adversary. This also means we get a form of resettable security, where the adversary can reset the receiver, for free.
- Furthermore, an interesting theoretical feature of uni-directional protocols like ours is in the scenarios where the receiver may not wish to be online. For instance, in the context of the sealed-bid auction example, our protocol would allow a committer Alice to send her bid by post (sending the first message on the first day and the second message on the second), whereas the receiver would only have to collect all incoming bids and then declare the winner on day 3, when the opening messages are received. In particular, the receiver would not have to send any messages until all the bids are collected and the winner is declared. Similarly, uni-directional protocols are interesting in situations where it may not be feasible to wait for the receiver to respond. This situation could arise when long distances are involved in the communication, or in settings where non-digital means are used for communication.

Apart from non-malleable commitments, we believe it would be interesting to study the round complexity of various other cryptographic tasks in the uni-directional message model.

- ◦ Our protocol makes only *black-box* use of any non-interactive statistically binding commitment scheme. In the standard model, such a scheme can be realized based on any injective one-way function. Furthermore, our protocol itself is quite simple and intuitive, and most of the complexity of our paper lies in the analysis of this protocol.
- ◦ Our scheme can also be viewed as a three block block-wise non-malleable code [14] requiring only injective one-way functions, used in a black-box manner. The only previously known feasibility result for such non-malleable codes required $O(\mathsf{poly}(n))$ blocks, where n is the security parameter, and relied on sub-exponential hardness assumptions. Thus, we significantly improve simultaneously on the hardness assumptions as well as the number of blocks.
- ◦ Finally, we extend the impossibility result of Pass [11] to the bi-directional two message commitments with non-malleability with respect to opening.
  - We show that no natural[1] (so-called "non-embedding") black-box reduction can be used to prove one-sided security of bi-directional two message non-malleable commitments with respect to opening, based on any standard polynomial intractability assumption.
  - We show also that no such natural black-box reduction can be used to prove two-sided security of bi-directional two message non-malleable commitments with respect to opening, based on any standard sub-exponential intractability assumption.
  - This also shows that no such natural black-box reduction can be used to build a block-wise non-malleable code with fewer than three blocks.
  - We note that in the asynchronous setting, any uni-directional protocol can be treated as a one-message (non-interactive) protocol for the purpose of proving non-malleability, and therefore our impossibility result applies. This shows that no two-message commitment protocol of any kind can be non-malleable with respect to opening in the asynchronous adversarial model; i.e. in the asynchronous setting, the three-round barrier is insurmountable with respect to protocols with natural black-box proofs of security. However, we stress that the synchronous setting is the more common one for typical applications of non-malleable commitments within other protocols, such as the secure computation protocols discussed above.

II. Overview of Our Techniques

Before describing the main ideas behind our construction, we recall the essence of techniques used in recent prior works on constant-round non-malleable commitments [6], [9], [10]. Note that our protocol, like all these recent protocols, assume that parties have "tags" (or id's), denoted $\mathtt{tag}$, and we require non-malleability to hold whenever the adversary is trying to commit on behalf of a $\widetilde{\mathtt{tag}}$ that is different from the $\mathtt{tag}$ used by the honest committer. Thus, even copying every message sent by the honest party does not yield a successful

[1] To the best of our knowledge, all reductions in the literature for non-malleable commitments fall into this class and are therefore ruled out.

attack, since this would be a valid commitment only with respect to the tag of the honest party, and not with respect to the adversary's distinct $\widetilde{\mathsf{tag}}$.

All recent protocols use "slots" (or non-malleable codes [10]) to create imbalances between protocols corresponding to different tags. They also use a commit-challenge-response structure so that a simulator can, while simulating the honest execution (without access to the honest receiver's input), *extract* the value committed by the man-in-the-middle.

Goyal, Pandey and Richelson [10] constructed a three round non-malleable commitment scheme with respect to commitment, matching the lower bound of Pass [11]. Indeed, black-box extraction requires at least three messages – and therefore, it would seem that this is perhaps the end of the story on optimizing the round complexity of non-malleable commitments. Surprisingly, we demonstrate that this is not the case.

Indeed, in the case of non-malleable commitments with respect to opening, we show how it is possible to use the opening phase itself for extraction. Crucially, our opening message does not merely consist of all randomness used by the committer during commitment; instead we craft it carefully to enable our proof strategy to work. We highlight the core ideas behind our two round non-malleable commitments with respect to opening, in Section II-A. We then show how to extend and polish these ideas to obtain a full-fledged non-malleable commitment scheme w.r.t. opening, in Section II-B, Section II-C and Section II-D.

A. One-Sided Non-Malleable Commitment Scheme for Small Tags, Against Perfectly Copying Adversaries

a) A Very Simple Scheme for Two Tags.: At the intuitive heart of our paper is the following idea: let's make "the level of commitment" provided by the first message of the protocol depend on the tag of the user making the commitment. This initial idea will only take us so far, but we will find a way to extend the combinatorics upon which this idea rests to obtain our final result.

To explain this intuition more precisely, we begin by constructing a simple commitment scheme, such that any adversary trying to copy a commitment constructed with $\mathsf{tag} = 1$ to a commitment with $\mathsf{tag} = 0$, fails with at least a constant probability[2]. We call such a scheme "one-sided" because we only guarantee security if the adversary uses a $\widetilde{\mathsf{tag}}$ that is *smaller* than the tag used by the honest party.

Let $\mathsf{com}(\cdot)$ denote a non-interactive statistically binding commitment scheme. When $\mathsf{tag} = 1$, in order to commit to a message $m \in \{0,1\}^n$, the committer must pick a random string $s \xleftarrow{\$} \{0,1\}^n$, randomness $r_1, r_2 \xleftarrow{\$} \{0,1\}^{2n}$, and compute $\mathsf{com}(m; r_1), \mathsf{com}(s; r_2)$. It picks a position $p \xleftarrow{\$} \{1,2\}$. If $p = 1$, the committer sends $\mathsf{com}(m; r_1), \mathsf{com}(s; r_2)$ to the receiver. If $p = 2$, the committer sends $\mathsf{com}(s; r_2), \mathsf{com}(m; r_1)$ to the receiver. In the second round, the committer sends p to the receiver. It is easy to see that this commitment is statistically binding (by the end of the second round) and computationally hiding based on the properties of the underlying commitment scheme $\mathsf{com}(\cdot)$. In order to open the commitment, the committer simply decommits to the commitment at the p^{th} position. On the other hand, when $\mathsf{tag} = 0$, in order to commit to a message $\tilde{m} \in \{0,1\}^n$, the committer picks a random string $r \xleftarrow{\$} \{0,1\}^n$ and sends $\mathsf{com}(\tilde{m}; r)$ to the receiver. In the second round, the committer sends position $p = 1$ to the receiver (since there is only one first round commitment when $\mathsf{tag} = 0$). In order to open the commitment, the committer decommits to the only commitment he sent.

Now, consider a man-in-the-middle(MIM) adversary that obtains the first message of the commitment from an honest committer in a left execution, with $\mathsf{tag} = 1$, and then outputs his first message corresponding to $\mathsf{tag} = 0$ in the right execution. Intuitively, any such MIM adversary is already committed to his message *even before* he knows whether the honest committer placed his message at (and will be opening) position 1 or 2. In this case, it is possibly to formally show that an MIM will not succeed at copying the value of the honest committer with probability greater than $1/2$ (which is the probability with which the MIM guesses the position that the honest committer will open, and only mauls the commitment at this position).

b) A Simple Scheme for Small Tags.: We develop the above idea further and extend the setting in two ways: First, we will allow parties to use somewhat larger tags, namely $\mathsf{tag} \in [n]$. Second, we will consider *arbitrary* polynomial-time mauling strategies of the adversary – the adversary doesn't just have to copy a commitment entirely, as we considered above. Nevertheless, we are still considering the one-sided setting, and we are only going to rule out an adversary that tries to *always* succeed in copying the value committed to by the honest (left) committer. The protocol is in Figure 1. This scheme is computationally hiding and statistically binding (by the second round) because of the hiding and binding properties of the underlying commitment scheme.

An interesting property of our scheme is that – unlike most other non-malleable w.r.t. commitment schemes in the literature – the decommitment phase does not include revealing all the coins of the committer. This is actually crucial for our proofs to work.

Despite the fact that the adversary can now perform

[2]Our final commitment scheme allows string commitment, for tags in $[2^n]$, and is secure with overwhelming probability, whenever the left tag tag is not equal to the right tag $\widetilde{\mathsf{tag}}$.

Let $\mathsf{com}(\cdot)$ denote a non-interactive statistically binding commitment scheme.
Tag: Let the tag for the interaction be $\mathsf{tag} \in [n]$.
Input: Committer $\mathcal{C}$ has private input message $m \in \{0,1\}^n$.

1) **Commit Stage:**
   - **First Message.**
     $\mathcal{C}$ picks random position $v \in [\mathsf{tag}]$. It sets $c_v = \mathsf{com}(m; r)$ for randomly chosen r.
     It samples $(s_i; r_i) \xleftarrow{\$} \{0,1\}^{2n}$ and generates $c_i = \mathsf{com}(s_i; r_i)$ for $i \in [\mathsf{tag}] \setminus \{v\}$. $\mathcal{C}$ sends $c_1, c_2, \ldots c_{\mathsf{tag}}$ to $\mathcal{R}$.
   - **Second Message.** $\mathcal{C}$ sends v to $\mathcal{R}$.
2) **Reveal Stage:** $\mathcal{C}$ outputs m and decommits to the v^{th} commitment c_v.
   $\mathcal{R}$ verifies that the decommitment is correct and equal to m.

Fig. 1: One-Sided Non-Malleable Commitment for $\mathsf{tag} \in [n]$, Secure Against Perfect Copying Adversaries

arbitrary mauling, for example by algebraically combining all the commitments on the left to produce a single commitment on the right, we still want to appeal to the pigeonhole principle in our proof. However, this will now be more technical to capture.

Let the tag used in the right execution be $\widetilde{\mathsf{tag}}$ and the tag used in the left execution be tag, such that $\mathsf{tag} > \widetilde{\mathsf{tag}}$. Recall that the simplified MIM adversary that we consider is not allowed to abort and must successfully copy always. Thus, if we fix the first messages of both the honest committer and the adversary, by the pigeonhole principle, there exist at least two positions $v_1, v_2 \in [\mathsf{tag}]$ output by the honest committer on the left in the second message, that correspond to the same position $\tilde{v} \in [\widetilde{\mathsf{tag}}]$ output by the MIM adversary on the right. We want to obtain a contradiction by exploiting this.

To do so, we design a sequence of hybrids that can aid us in carrying out our combinatorial argument. The present scenario that we consider offers a simplified setting for our hybrid argument, and the intuition developed here will be useful later. We assume the existence of a successful adversary, and then consider the following sequence of hybrids between the case when the left execution consists of an honest commitment to m versus an honest commitment to $m' \neq m$, in order to reach a contradiction.

Hybrid_0 : This is similar to the real world experiment, with one abort condition added. Just as in the real world, challenger carries out the entire left execution committing to m, outputting v as its second message, and then decommitting to message m. It records the transcript generated by the MIM along with the decommitted value (which should be m since our MIM always copies successfully). Let the second message of the MIM in this execution be $\tilde{v}$.

The challenger then picks a random position $v' \xleftarrow{\$} [\mathsf{tag}], v' \neq v$, intuitively hoping that this choice of v' will be the one guaranteed by the pigeonhole principle. Keeping the first messages of the transcript on both left and right fixed, it rewinds the protocol execution to the beginning of the second round and sends v'. If the MIM outputs the same value $\tilde{v}$, as we hoped he would, the challenger outputs the original view (using v) and the value opened by the MIM, otherwise it aborts.

Note that in this hybrid, our random choice of v' is going to work with probability at least $1/n$ by the pigeonhole principle, and thus our copying MIM opens his commitment to message m with probability at least $1/n$ and aborts otherwise.

Hybrid_1 : This experiment is identical to Hybrid_0 except that the challenger picks the random position $v' \in [\mathsf{tag}], v' \neq v$ at the beginning of this experiment, and generates a commitment to m' at position v' (instead of committing to a random string $s_{v'}$). Then, just like Hybrid_0, the challenger carries out the entire left execution committing to m, outputting v as its second message and then decommitting to message m. It records the transcript generated by the MIM along with the decommitted value (which is m since our MIM copies always). Let the second message of the MIM in this execution be $\tilde{v}$.

Keeping the first message of the transcript fixed, just as in Hybrid_0, the challenger rewinds the protocol execution to the second round and sends v'. If the MIM outputs $\tilde{v}$, the challenger outputs the original view and the value opened by the MIM, otherwise it aborts.

Crucially, the only difference between Hybrid_1 and Hybrid_0 is that one of the commitments that is never opened by the challenger is either a commitment to m', in Hybrid_1, or a commitment to $s_{v'}$, in Hybrid_0. Thus, these two hybrids are computationally indistinguishable by the hiding property of com. Thus, our copying MIM still opens his commitment to message m with probability at least roughly $1/n$ and aborts otherwise.

Analysis of Hybrid_1: Looking at Hybrid_1, observe that v and v' play completely symmetrical roles: The same first message of the left commitment, if we view v as being chosen first, is a commitment to m. However, if we view v' as being chosen first, then it is a commitment to m'.

Furthermore, in fact, we know that with probability at least $1/n$, when we view the commitment as being a commitment to m', when we send the value v' as the second message, the adversary chooses a value $\tilde{v}$ such that the adversary's commitment in his first message at position $\tilde{v}$ is a *statistically binding* commitment to m. Thus, we know that with probability at least $1/n$, even if the original commitment had been a commitment to m', the adversary will have no choice but to open to m.

This contradicts our assumption that the adversary always succeeds in copying the message committed to by the honest committer on the left, and shows that such an adversary cannot exist. The hybrid argument above illustrates the simplest case where we combine cryptographic computational indistinguishability arguments with combinatorial arguments (here, just a simple application of the pigeonhole principle). Later on, especially as we remove the unrealistic assumption that the adversary always succeeds, these "combinatorial hybrid arguments" become more subtle and complex.

B. One-Sided Non-Malleable Commitment Scheme for Small Tags, Against General Synchronous Adversaries

The protocol we described in the previous section, provides only partial non-malleability against adversaries that do not abort. In particular, it does not guarantee full security against general (possibly aborting) adversaries. Indeed, with probability $\frac{1}{n}$, the MIM can directly guess the position that the honest committer plans to decommit, and copy the commitment at that position. In this case, the MIM succeeds in copying the value committed in the left execution.

In order to handle general adversaries, the protocol in Figure 1 is modified as follows: The committer computes an N-out-of-N additive secret sharing of the message m, for $N = n^3$ (where n denotes the security parameter). Then, the same basic protocol of Figure 1 is repeated in parallel N times with fresh randomness each time, to commit to each of the N shares of the message m. An MIM adversary can no longer "guess" all these positions correctly. Indeed, we show that if an adversary tries to maul in all possible positions, w.h.p. he will fail and abort the protocol.

For $\mathsf{tag} > \widetilde{\mathsf{tag}}$ (where tag is the honest/left execution tag and $\widetilde{\mathsf{tag}}$ is the MIM's tag), we note that the honest committer generates a total of $\mathsf{tag} \cdot N$ commitments in the first message, and the MIM generates much fewer – i.e., a total of $\widetilde{\mathsf{tag}} \cdot N$ commitments – in his first message. The space of possible positions chosen to be opened on the left, is $(\mathsf{tag})^N$, and that on the right is only $(\widetilde{\mathsf{tag}})^N$. Then on an average, there are an exponential number of position tuples on the left mapping to each position tuple on the right. However, unlike the hybrids in the previous protocol, it is no longer possible to efficiently "find" such collisions.

We overcome this issue using ideas inspired from those used by Goyal [6], i.e. by extracting the "commitment dependency graph" created by the adversary. However, our proofs depart significantly from [6] because our setting is inherently uni-directional. In particular, unlike [6], the receiver on the right (or the challenger) no longer has the power to choose which positions the MIM will open.

Informally, given the first message of the protocol: we say that a right commitment Y of the adversary depends on a left commitment X of the honest committer, if Y is chosen to be opened by the MIM with noticeable probability conditioned on a random choice of position tuple on left, but is chosen to be opened by the MIM with close to "negligible" probability if X is not chosen to be opened on the left. Please refer to the full version for a formal definition of dependency. Then, it is possible to show using combinatorial arguments, that for "most transcripts", there exists at least one position I that is chosen to be opened on the left, such that none of the positions that are chosen to be opened on the right, lie in the dependent set of the commitment at position I.

Given this lemma, it is possible to replace the commitment at position I, with an externally generated commitment. Since none of the positions that are chosen to be opened on the right (let these be denoted by the set $\widetilde{\ell}$), lie in the dependent set of the commitment at position I, by the definition of dependency, it is possible to rewind and provide various other openings on the left to "extract" the values committed to by the MIM for each of the commitments at positions in $\widetilde{\ell}$. Note that because the message is shared using an N-out-of-N additive secret sharing scheme, in order to change the left commitment from m to m', it suffices to just change the committed value on the left corresponding to the single position I.

Finally, we note that in our actual protocol (Refer to the full version), we do not actually use tag commitments in each parallel repetition but rather use $n + \mathsf{tag}$ commitments. This is done for technical reasons, to ensure that none of the sets is too "small" in size. Also, in the actual protocol, all commitments are to random values, but in the second message, together with the indices to be opened, a correction factor is sent. This is done for technical reasons as well, to make the dependency graph well-defined for the purpose of our combinatorial analysis.

C. Two-Sided Non-Malleable Commitment Scheme for All Tags, Against General Synchronous Adversaries

The scheme described in the previous section is non-malleable only if the tag $\widetilde{\mathsf{tag}} \in [n]$ of the right execution, is smaller than the tag $\mathsf{tag} \in [n]$ of the left execution. We show how to extend the ideas developed in the previous section to obtain a commitment scheme that is non-malleable whenever $\widetilde{\mathsf{tag}} \neq \mathsf{tag}$.

We first consider the case of a perfectly copying adversary. The basic protocol (Figure 1) is modified as follows: The committer computes a 2-out-of-2 secret sharing of the message m, and let the shares be denoted by m_1 and m_2. Then the protocol from Figure 1 is used to commit to the shares m_1 and m_2 in parallel, using tags (tag) and $(2n - \mathsf{tag})$ respectively. The modified basic protocol, in simplified form for the purpose of this technical overview, is described in Figure 2.

Let $\mathsf{com}(\cdot)$ denote a non-interactive statistically binding commitment scheme.
Tag: Let the tag for the interaction be $\mathsf{tag} \in [n]$.
Input: Committer $\mathcal{C}$ has private input message $m \in \{0,1\}^n$.

1) **Commit Stage:**
   - **First Message.** $\mathcal{C}$ picks random position $v \in [\mathsf{tag}]$ and sets $c_v = \mathsf{com}(m_1; r)$ for random r. $\mathcal{C}$ also picks random position $\bar{v} \in [2n - \mathsf{tag}]$ and sets $c_{\bar{v}} = \mathsf{com}(m_2; \bar{r})$ for random $\bar{r}$.
     It samples $(s_i, r_i) \xleftarrow{\$} \{0,1\}^{2n}$ and generates $c_i = \mathsf{com}(s_i; r_i)$ for $i \in [\mathsf{tag}] \setminus \{v\}$. It samples $(\bar{s}_i, \bar{r}_i) \xleftarrow{\$} \{0,1\}^{2n}$ and generates $\bar{c}_i = \mathsf{com}(\bar{s}_i; \bar{r}_i)$ for $\bar{i} \in [2n - \mathsf{tag}] \setminus \{\bar{v}\}$. $\mathcal{C}$ sends $c_1, c_2, \ldots c_{\mathsf{tag}}, \bar{c}_1, \bar{c}_2, \ldots \bar{c}_{2n-\mathsf{tag}}$ to $\mathcal{R}$.
   - **Second Message.** $\mathcal{C}$ sends $v, \bar{v}$ to $\mathcal{R}$.
2) **Reveal Stage:** $\mathcal{C}$ outputs m and decommits to the v^{th} commitment c_v and the $\bar{v}^{th}$ commitment $c_{\bar{v}}$. $\mathcal{R}$ verifies that the decommitments are correct and XORs to m.

Fig. 2: Simplified Two-Sided Commitment for $\mathsf{tag} \in [n]$, secure against Non-Aborting Adversaries

The hiding and binding properties of this scheme follow directly from the hiding and binding properties of the underlying statistically binding commitment scheme.

Interestingly, in this protocol, the total number of commitments generated is exactly $2n$, irrespective of the tag. This is reminiscent of the approach of [9]: who set up challenge spaces of identical length for all tags (while nearly all prior works used imbalanced spaces to obtain non-malleability). However, they use more algebraic structure on the challenge spaces whereas our techniques are purely combinatorial while still maintaining just two rounds in the scheme.

Now, the honest committer in the left execution, opens 1 out of (tag) commitments to m_1, and 1 out of $(2n - \mathsf{tag})$ commitments to m_2. On the other hand, the MIM must open 1 out of $(\widetilde{\mathsf{tag}})$ commitments to $\tilde{m}_1$, and 1 out of $(2n - \widetilde{\mathsf{tag}})$ commitments to $\tilde{m}_2$. For simplicity, here we consider the case where MIM creates a one-to-one mapping of the $2n$ commitments (although our proof rules out other arbitrary strategies of any MIM).

We also observe that the sets of commitments in our protocol satisfy a specific property: *when* $\mathsf{tag} \neq \widetilde{\mathsf{tag}}$, *there is no perfect embedding of the left sets* $([\mathsf{tag}], [2n - \mathsf{tag}])$ *into the right sets* $([\widetilde{\mathsf{tag}}], [2n - \widetilde{\mathsf{tag}}])$. In particular, when $\mathsf{tag} > \widetilde{\mathsf{tag}}$, then $\widetilde{\mathsf{tag}} < (2n - \mathsf{tag})$ also (since $\mathsf{tag}, \widetilde{\mathsf{tag}} \in [n]$). Therefore, for any mapping $f : [2n] \rightarrow [2n]$ of the MIM, there is significant probability (at least $\frac{1}{n}$), that $f^{-1}(i)$ is not opened on the left for any element i of the set $[\widetilde{\mathsf{tag}}]$. In other words, no matter how the adversary maps commitments in the sets $([\mathsf{tag}], [2n - \mathsf{tag}])$ to commitments in $[\widetilde{\mathsf{tag}}]$, there is a chance that none of the left commitments that were mapped to the set $[\widetilde{\mathsf{tag}}]$ are opened. In this case, the adversary must either abort, or open a right commitment that depends on no left commitment.

Again, when $\mathsf{tag} < \widetilde{\mathsf{tag}}$, then $(2n - \mathsf{tag}) > (2n - \widetilde{\mathsf{tag}})$. Then, for any mapping $f : [2n] \rightarrow [2n]$ of the MIM, there is significant probability (at least $\frac{1}{n}$), that $f^{-1}(i)$ is not opened on the left for any element i of the either the set $[\widetilde{\mathsf{tag}}]$ or the set $[2n - \widetilde{\mathsf{tag}}]$. If the adversary only maps commitments in the set $[2n - \mathsf{tag}]$ to commitments in $[2n - \widetilde{\mathsf{tag}}]$, then it is easy to see that this statement is true because the size of the set $[2n - \widetilde{\mathsf{tag}}]$ is smaller than the size of $[2n - \mathsf{tag}]$. The only case in which the inverse on the left is always opened for all elements in the set $[2n - \widetilde{\mathsf{tag}}]$, is when all commitments in the set $[\mathsf{tag}]$ on the left are mapped to commitments within the set $[2n - \widetilde{\mathsf{tag}}]$ on the right. But in this case, no matter how the adversary maps the remaining commitments in the set $[2n - \mathsf{tag}]$ to commitments in $([\mathsf{tag}], [2n - \widetilde{\mathsf{tag}}])$, there exists a significant probability that none of the left commitments that were mapped to the set $[\widetilde{\mathsf{tag}}]$ will be opened. Again, in this case, the adversary must either abort, or open a right commitment that depends on no left commitment.

a) Moving to General Synchronous Adversaries.: Like in Section II-A, the protocol in Figure 2 is only partially non-malleable. But the protocol can be modified, via N-out-of-N secret sharing, to obtain security against general synchronous adversaries. Again, the proof runs into issues because of the exponential possibilities for the position tuple in the second round. These are handled in the same way as the one-sided setting, with one additional key idea: In most sets of $[\widetilde{\mathsf{tag}}]$ or $[2n - \widetilde{\mathsf{tag}}]$ commitments on the right, there must exist at least one "reserve commitment", which does not *depend* on any left commitment. We just showed that for any set of $([\widetilde{\mathsf{tag}}], [2n - \widetilde{\mathsf{tag}}])$ right commitments which all depend on left commitments, none of the left commitments on which such a set depends, are opened with probability at least $\frac{1}{n}$. Therefore, if there are too many such sets that depend on distinct left commitments and have no "reserve commitments", the adversary will end up aborting in the real execution with overwhelming probability.

The existence of these reserve commitments helps us establish, for any adversary that does not abort nearly all the time, a "gap" between the number of possible positions that can be opened in the left execution (i.e. $2nN$) and the number of possible positions that can depend on them in the right execution (which is less than $2nN$ because of the presence of "reserve commitments" that depend on nothing). This gap can then be exploited, just like in the basic protocol, to prove non-malleability.

Finally, we note that, as before, in our actual protocol (Refer to the full version), we do not actually set the size of these sets to $(\mathsf{tag}, 2n-\mathsf{tag})$, but set it to $(2n+\mathsf{tag}, 4n-\mathsf{tag})$. This is done for technical reasons, to ensure that none of the sets is too "small" in size. Also as before, in the actual protocol, all commitments are to random values, but in the second message, together with the indices to be opened, a correction factor is sent.

D. Full-Fledged Non-Malleability Against All Synchronous Adversaries

Our full-fledged non-malleable commitment scheme (for $\mathsf{tag} = (t_1||t_2||\ldots||t_n) \in [2^n]$), is obtained by computing an n-out-of-n additive secret sharing of the message m to be committed to, and using the two-sided non-malleable commitment scheme discussed in the last subsection, to commit in parallel to each of the n shares, using tags $(1||t_1,\ 2||t_2,\ \ldots n||t_n)$. The protocol resembles the classic DDN [1] tag amplification construction – however, its analysis is significantly different. In particular, Dolev, Dwork and Naor [1] rely on a one-many non-malleable commitment scheme in order to achieve such amplification. However, our two-sided non-malleable scheme is only one-one secure, and therefore we cannot hope to use the [1] proof strategy in our setting. Interestingly, in this scheme (just as in the two-sided scheme) the total number of commitments generated remains fixed and independent of the tag. Furthermore, we can show that our construction is secure by using the set-embedding argument from the previous section.

Suppose the left tag is $(t_1||t_2||\ldots t_n)$ and the right tag is $(\tilde{t}_1|||\tilde{t}_2||\ldots \tilde{t}_n)$. Note that there exists an index i, such that $\tilde{t}_i \neq t_i$; equivalently, $i||\tilde{t}_i \neq j||t_j$ for all $j \in [n]$. In this case, none of the left sets corresponding to indices $(j||t_j)$ for $j \in [n]$ would possibly completely embed in the set corresponding to $(i||\tilde{t}_i)$. We set our parameters to ensure that there do not exist any combinations of left sets that embed in the set corresponding to the differing index $(i||\tilde{t}_i)$. The core of the argument is to show that such embeddings cannot exist in this setting.

Then, a similar argument as in the previous section can be used to show that the adversary must maintain "reserve commitments", thereby creating a gap between the number of possible commitments that can be opened in the left execution (i.e. $2nN$) and the number of commitments positions that can depend on them in the right execution. The gap can then be exploited, just like in the one-sided protocol, to prove non-malleability for all tags.

Once we obtain a non-malleable commitment scheme for all tags, it is possible to use standard techniques from the literature to obtain a full-fledged non-malleable commitment scheme (without tags) by relying on one-way functions. This is done by generating the tag as the verification key of a signature scheme, and using the corresponding signing key to sign each message in the protocol.

E. Impossibility of Non-Interactive/Two-Round Asynchronous Non-Malleable Commitments w.r.t. Opening

In this section, we illustrate the key ideas behind our impossibility result. Assume that Π is a non-interactive non-malleable commitment scheme with respect to opening[3].

Suppose there exists a reduction $\mathcal{R}$ such that $\mathcal{R}^A$ breaks some underlying assumption C, whenever the MIM adversary A breaks non-malleability of Π. We will assume that any possible open queries of $\mathcal{R}$ for a particular session appear immediately after commit queries for the same session (we call such a reduction non-embedding). We will describe the ideas in our impossibility only for non-embedding reductions, but they also easily extend to reductions that nest $O(\log n)$ sessions within each other by relying on standard repeated rewinding techniques.

First consider a computationally unbounded copying A that picks identity 0 on the left and 1 on the right, and upon receiving a left commitment to some value v, recovers v using brute force (we assume the commitment is statistically binding so there exists such a unique v). Then, A commits to v using tag 1. When $\mathcal{R}$ decommits to v on the left, A decommits to v on the right. Since A breaks non-malleability of Π, $\mathcal{R}^A$ must also break the underlying commitment C.

Implementing such an A requires super-polynomial time by the hiding property of Π – however, we construct another algorithm $\tilde{A}$ that can efficiently emulate A. Roughly, $\tilde{A}$ proceeds as follows:

- ◦ $\tilde{A}$ picks the same identities, but simply commits to 0^n in the right interaction (irrespective of the messages it received on the left).
- ◦ If $\mathcal{R}$ does not send an opening and starts a fresh commitment phase – $\tilde{A}$ continues the game by

[3]In this overview, for simplicity, we only consider a non-interactive scheme. It will be easy to see that the same result will extend to a two-round scheme where each party sends a single message.

committing to 0^n each time each time $\mathcal{R}$ generates some commitment. Because of the hiding property of the commitment scheme, $\mathcal{R}$ cannot distinguish A's commitments to v from $\tilde{A}$'s commitments to 0^n, without asking $\tilde{A}$ to open its commitment. Since a non-embedding $\mathcal{R}$ does not nest sessions, once $\mathcal{R}$ starts a new commit phase, it will never ask for an opening for any previous session. Now, if $\mathcal{R}$ does not require an opening to break the underlying assumption, this roughly means that the scheme is non-malleable w.r.t. commitment and the same ideas as Pass' [11] impossibility apply.

- The more interesting situation is when $\mathcal{R}$ does ask $\tilde{A}$ to give an opening on the left. In this situation, $\tilde{A}$ emulates A as follows: it observes the opening (to value v) given by $\mathcal{R}$, and then rewinds $\mathcal{R}$ to just after the corresponding commit stage message (to v) on the left (observe that $\mathcal{R}$ sends a single message in the commit stage). It then commits to value v that it extracted before rewinding. It generates this commitment honestly using randomness r, and then waits for $\mathcal{R}$ to decommit to v on the left. When $\mathcal{R}$ does decommit to v, $\tilde{A}$ opens his commitment to v as well.

Such an adversary $\tilde{A}$ successfully emulates A against any black-box reductions to underlying hardness assumptions, that nest upto $\log(n)$ commit and open queries (which we call $\log(n)$-embedding reductions and define in detail in the full version). Thus, any black-box reduction $\mathcal{R}$ must succeed in breaking the underlying assumption, when interacting with $\tilde{A}$. Since $\tilde{A}$ can be efficiently emulated by $\mathcal{R}$ itself, this gives a polynomial-time algorithm to break the underlying assumption.

We note that to the best of our knowledge, all known reductions in literature are non-embedding, and therefore our impossibility rules out all such reductions.

III. Our Protocols

In this section, we describe our protocols. Our one-sided non-malleable commitment scheme w.r.t. opening against synchronizing adversaries, for small tags is described in Figure 3. Our two-sided non-malleable commitment scheme w.r.t. opening against synchronizing adversaries, for small tags, is described in Figure 4.

In Figure 5, we describe our full-fledged non-malleable commitment scheme with respect to large tags (i.e., all tags in $[2^n]$). While the scheme is based on the DDN encoding [1], our analysis is entirely different because we do not rely on one-many non-malleability of the underlying scheme for small tags ($\mathsf{tag} \in [n]$).

The binding and hiding properties of these schemes follow in a straightforward manner from the statistical binding and computational hiding of the underlying commitment scheme. Our proofs proceed in two phases:

Let $\mathsf{com}(\cdot)$ denote a non-interactive statistically binding commitment. Let n denote the security parameter and set $N = n^5$.
Tag: Let the tag be $\mathsf{tag} \in [n]$. Set $t = n + \mathsf{tag}$.
Input: Committer $\mathcal{C}$ has private input $m \in \{0,1\}^n$.

1) **Commit Stage: First Message.**
   - For $i \in [N]$, $v \in [t]$, $\mathcal{C}$ samples $b_{i,v} \xleftarrow{\$} \{0,1\}^n$ and $r_{i,v} \xleftarrow{\$} \{0,1\}^n$.
   - For $i \in [N]$, $v \in [t]$, $\mathcal{C}$ computes $c_{i,1} = \mathsf{com}(b_{i,1}, r_{i,1}), c_{i,2} = \mathsf{com}(b_{i,2}, r_{i,2}), \ldots c_{i,t} = \mathsf{com}(b_{i,t}, r_{i,t})$; sends $c_{i,1}, c_{i,2}, \ldots c_{i,t}$.

   Second Message. $\mathcal{C}$ picks $v_i \xleftarrow{\$} [t]$ for each $i \in [N]$, and sets correction factor $\mathsf{cf} = (\bigoplus_{i \in [N]} b_{i,v_i}) \oplus m$. $\mathcal{C}$ sends cf along with v_i for each $i \in [N]$.
2) **Reveal Stage:** For each $i \in [N]$, $\mathcal{C}$ decommits to the v_i^{th} commitment c_{i,v_i}.
   $\mathcal{R}$ verifies that the decommitment is correct, and for each $i \in [N]$, the decommitments obtained together with the correction factor are XOR shares of message m.

Fig. 3: One-Sided Non-Malleable Protocol for $\mathsf{tag} \in [n]$

in the first phase, we argue that for most honest (left) transcripts, there exists an index (or a secret share) that is chosen to be opened in the second message of the left execution, such that it is possible to "extract" the MIM's opening in the right execution without having to open this share at all in the left execution. In the second phase, we show that this property can be leveraged to replace the commitment (or the share) at such an index on the left, with an external challenge commitment, such that if the MIM's opening changes based on the externally obtained challenge, then this MIM can be used to break hiding of the challenge commitment. Please refer to the full version for complete proofs.

Acknowledgment

Research of the second and third authors supported in part from a DARPA/ARL SAFEWARE award, NSF Frontier Award 1413955, NSF grants 1228984, 1136174, 1118096, a Xerox Faculty Research Award, a Google Faculty Research Award, an equipment grant from Intel, and an Okawa Foundation Research Grant. This material is based upon work supported by the Defense Advanced Research Projects Agency through the ARL under Contract W911NF-15-C-0205. The views expressed are those of the author and do not reflect the official policy or position of the Department of Defense, the National Science Foundation, or the U.S. Government.

Let $\mathsf{com}(\cdot)$ denote a non-interactive statistically binding commitment, set $N = n^5$.
Tag: Let the tag be $\mathsf{tag} \in [n]$. Let $t = \mathsf{tag} + 2n$.
Input: Committer $\mathcal{C}$ has private input $m \in \{0,1\}^n$.

1) **Commit Stage: First Message.**
   - For each $(i \in [N], v \in [t])$ and each $(i \in [N+1, 2N], v \in [6n-t])$, $\mathcal{C}$ samples $b_{i,v} \xleftarrow{\$} \{0,1\}^n$ and $r_{i,v} \xleftarrow{\$} \{0,1\}^n$. Then, $\mathcal{C}$ computes $c_{i,v} = \mathsf{com}(b_{i,v}, r_{i,v})$.

   Second Message. For each $i \in [N]$, $\mathcal{C}$ picks and sends $v_i \xleftarrow{\$} [t]$. For each $i \in [N+1, 2N]$, $\mathcal{C}$ picks and sends $v_i \xleftarrow{\$} [6n-t]$. Finally, $\mathcal{C}$ also sends $(\bigoplus_{i\in[2N]} \mathsf{cf}_{i,v_i}) \oplus m$.
2) **Reveal Stage:** For each $i \in [2N]$, $\mathcal{C}$ decommits to the v_i^{th} commitment c_{i,v_i}.
   $\mathcal{R}$ verifies that the decommitment is correct, and for each $i \in [2N]$, the decommitments obtained, together with the correction factor, are XOR shares of message m.

Fig. 4: Two-Sided Non-Malleable Protocol for $\mathsf{tag} \in [n]$

Let $\mathsf{com}(\cdot)$ denote a non-interactive statistically binding commitment scheme, and $\mathsf{small-NM}_{\mathsf{tag}}(\cdot)$ denote the commitment messages of the one-one two-sided non-malleable commitment scheme with respect to some $\mathsf{tag} \in [n]$, from Figure 4.
Tag: Let the tag for the interaction be $\mathsf{tag} \in [2^n]$, with bit-representation $\mathsf{tag} = t_1, t_2, \ldots t_n$.
Input: Committer $\mathcal{C}$ has private input $m \in \{0,1\}^n$.

1) **Commit Stage:**
   - $\mathcal{C}$ samples uniformly random shares $m_1, m_2, \ldots m_n$ such that $m = \bigoplus_{i\in[n]} m_i$.
   - $\mathcal{C}$ commits to in parallel to m_i with tag $i||t_i$, using scheme $\mathsf{small-NM}$ as follows: $\mathsf{small-NM}_{1||t_1}(m_1)$, $\mathsf{small-NM}_{2||t_2}(m_2)$ $\ldots \mathsf{small-NM}_{n||t_n}(m_n)$.
2) **Reveal Stage:** C decommits in parallel to each of the $\mathsf{small-NM}$ generated above.
   $\mathcal{R}$ verifies that the decommitments are correct, and for each $i \in [n]$, the decommitments obtained are shares of message m.

Fig. 5: Two-Sided Non-Malleable Protocol for $\mathsf{tag} \in [2^n]$

REFERENCES

[1] D. Dolev, C. Dwork, and M. Naor, "Non-Malleable Cryptography (Extended Abstract)," in *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing*, ser. STOC '91, 1991, pp. 542–552. 1, 8, 9

[2] B. 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*, ser. FOCS '02, 2002, pp. 345–355. 1

[3] R. Pass, "Bounded-Concurrent Secure Multi-Party Computation with a Dishonest Majority," in *Proceedings of the 36th Annual ACM Symposium on Theory of Computing*, ser. STOC '04, 2004, pp. 232–241. 1

[4] R. Pass and A. Rosen, "New and improved constructions of non-malleable cryptographic protocols," in *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, ser. STOC '05, 2005, pp. 533–542. 1, 2

[5] H. Wee, "Black-Box, Round-Efficient Secure Computation via Non-malleability Amplification," in *Proceedings of the 51th Annual IEEE Symposium on Foundations of Computer Science*, 2010, pp. 531–540. 2

[6] V. Goyal, "Constant Round Non-malleable Protocols Using One-way Functions," in *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing*, ser. STOC '11. ACM, 2011, pp. 695–704. 2, 3, 6

[7] H. Lin and R. Pass, "Constant-round Non-malleable Commitments from Any One-way Function," in *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing*, ser. STOC '11, 2011, pp. 705–714. 2

[8] V. Goyal, C.-K. Lee, R. Ostrovsky, and I. Visconti, "Constructing non-malleable commitments: A black-box approach," in *FOCS*. IEEE Computer Society, 2012, pp. 51–60. 2

[9] V. Goyal, S. Richelson, A. Rosen, and M. Vald, "An algebraic approach to non-malleability," in *FOCS*, 2014. 2, 3, 7

[10] V. Goyal, O. Pandey, and S. Richelson, "Textbook non-malleable commitments," in *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, D. Wichs and Y. Mansour, Eds. ACM, 2016, pp. 1128–1141. [Online]. Available: http://doi.acm.org/10.1145/2897518.2897657 2, 3, 4

[11] R. Pass, "Unprovable security of perfect NIZK and non-interactive non-malleable commitments," in *TCC*, 2013, pp. 334–354. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-36594-2_19 2, 3, 4, 9

[12] G. D. Crescenzo, Y. Ishai, and R. Ostrovsky, "Non-interactive and non-malleable commitment," in *STOC*, 1998. 2

[13] M. Fischlin and R. Fischlin, "Efficient non-malleable commitment schemes," in *Advances in CryptologyCRYPTO 2000*. Springer, 2000, pp. 413–431. 2

[14] N. Chandran, V. Goyal, P. Mukherjee, O. Pandey, and J. Upadhyay, "Block-wise non-malleable codes," *IACR Cryptology ePrint Archive*, vol. 2015, p. 129, 2015. [Online]. Available: http://eprint.iacr.org/2015/129 3