58th Annual IEEE Symposium on Foundations of Computer Science

How to Achieve Non-Malleability in One or Two Rounds

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*—Non-malleable commitments, introduced by Dolev, Dwork and Naor (STOC 1991), are a fundamental cryptographic primitive, and their round complexity has been a subject of great interest. And yet, the goal of achieving non-malleable commitments with only *one or two rounds* has been elusive. Pass (TCC 2013) captured this difficulty by proving important impossibility results regarding two-round non-malleable commitments. This led to the widespread belief that achieving two-round non-malleable commitments was impossible from standard assumptions. We show that this belief was false. Indeed, we obtain the following positive results:**

- **We construct two-message non-malleable commitments satisfying non-malleability with respect to commitment, based on standard sub-exponential assumptions, namely: sub-exponential one-way permutations, sub-exponential ZAPs, and sub-exponential DDH. Furthermore, our protocol is *public-coin*.**
- **We obtain two-message *private-coin* non-malleable commitments with respect to commitment, assuming only sub-exponential DDH or QR or N^{th}-residuosity.**
- **We bootstrap the above protocols (under the same assumptions) to obtain two round constant bounded-concurrent non-malleable commitments. In the simultaneous message model, we obtain unbounded concurrent non-malleability in two rounds.**
- **In the simultaneous messages model, we obtain *one-round* non-malleable commitments, with unbounded concurrent security with respect to opening, under standard sub-exponential assumptions.**
  - **This implies non-interactive non-malleable commitments with respect to opening, in a restricted model with a broadcast channel, and a-priori bounded polynomially many parties such that every party is aware of every other party in the system. To the best of our knowledge, this is the first protocol to achieve completely non-interactive non-malleability in any plain model setting from standard assumptions.**
  - **As an application of this result, in the simultaneous exchange model, we obtain two-round multi-party pseudorandom coin-flipping.**
- **We construct two-message zero-knowledge arguments with super-polynomial *strong* simulation (SPSS-ZK), which also serve as an important tool for our constructions of non-malleable commitments.**
- **In order to obtain our results, we develop several techniques that may be of independent interest.**
  - **We give the first two-round black-box rewinding strategy based on standard sub-exponential assumptions, in the plain model.**
  - **We also give a two-round tag amplification technique for non-malleable commitments, that amplifies a 4-tag scheme to a scheme for all tags, while relying on sub-exponential DDH. This includes a more efficient alternative to the DDN encoding.**

The full version of this paper is available online at: https://eprint.iacr.org/2017/291.pdf.

I. Introduction

The notion of non-malleability was introduced by Dolev, Dwork and Naor [1] in 1991, to counter the ubiquitous problem of man-in-the-middle (MIM) attacks on cryptographic protocols. An MIM adversary participates in two or more instantiations of a protocol, trying to use information obtained in one execution to breach security in the other protocol execution. A non-malleable protocol should ensure that such an adversary gains no advantage. Let's call any interactive protocol between two parties, where both parties send at least one message to each other, a conversation. In this paper, we ask if we can provably embed non-malleability into two-party conversations. We focus on a core non-malleable cryptographic primitive: non-malleable commitments. Thus, the main question we consider in this work is,

Can we get two-message non-malleable commitments from standard sub-exponential assumptions?

A commitment scheme is a two-party protocol between a committer and a receiver. The committer has input message m, while the receiver obtains no input. The two parties engage in a probabilistic interactive commitment protocol, and the receiver's view at the end of this protocol is denoted by $\mathsf{com}(m)$. Later, in the opening phase, the committer sends an opening message to the receiver, allowing the receiver to verify that the message m was really the message committed during the commitment phase.

In a (statistically) binding commitment, the transcript $\mathsf{com}(m)$ should be binding in the sense that with high

0272-5428/17 $31.00 © 2017 IEEE
DOI 10.1109/FOCS.2017.58

probability, there should not exist an opening message that would convince the receiver that the committer 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 hiding; that is, for any pair of messages (m, m') the distributions $\mathsf{com}(m)$ and $\mathsf{com}(m')$ should be computationally indistinguishable. Finally, such a scheme is said to be *non-malleable* with respect to commitment, 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 $\mathsf{com}(m')$ such that message m' is related to the original message m.

In the standard model, we call each message sent by any party a *round*. We will also consider the simultaneous-message model, wherein a round consists of both (or all) parties sending a single message simultaneously. Non-malleable commitments are among the core building blocks of (and therefore have a direct impact on the round complexity of) various cryptographic protocols such as coin-flipping, secure auctions, electronic voting, non-malleable proof systems and multi-party computation protocols.

The goal of achieving non-malleable commitment protocols with only *two messages* has been particularly elusive. Notably, Pass [2] proved that two-message non-malleable commitments (satisfying non-malleability with respect to commitment) are impossible to construct with a black-box reduction to any polynomial falsifiable assumption. However, another claim from [2] stated that two-message non-malleable commitments are impossible to construct with a black-box reduction to any *sub-exponentially* hard falsifiable assumptions, seemingly cutting off hope of achieving two-message non-malleable commitments from standard assumptions.

a) On the impossibility result of [2]: Let us examine the impossibility result of [2]: it considers the setting where there are only two identities/tags in the system, and discusses how one cannot achieve non-malleability even in this restricted setting via black-box reductions to falsifiable hardness. The impossibility builds as a counter-example, a MIM that *runs the reduction* in order to break hiding of an honest commitment and carry out a successful mauling attack. If the assumption is with regard to any polynomial-time attacker with inverse polynomial advantage, then this proof works, and the impossibility holds. It might appear that this argument should also extend to assumptions that require security against sub-exponential attackers with inverse sub-exponential advantage. However, we observe that an actual MIM only participates in at most a polynomial number of interactions and is required to break non-malleability in one of them[1], whereas a (sub-exponential) time reduction has oracle access to an adversary – and can therefore participate in sub-exponentially many interactions.

This gap between the number of sessions that the reduction can participate in, and the number of sessions in which participation is possible for any adversary that wants to "run the reduction," precludes the impossibility claim. Therefore, Theorem 5.11 as stated in [3], is incorrect[2]. Indeed, we show how to contradict this statement by achieving several positive results from standard sub-exponential assumptions.

We stress that when considering a reduction that can run in sub-exponential time, a reduction that participates in sub-exponentially many sessions is no worse asymptotically than a reduction that participates in only polynomially many sessions. For example, let $\delta < \epsilon$, and suppose that we consider a reduction $\mathcal{R}$ that runs in time 2^{n^ϵ}, and participates in m sessions with an adversary MIM that runs in time 2^{n^δ}. Then observe:

- If $\mathcal{R}$ participates in $\mathsf{poly}(n)$ sessions, then the total security loss is $2^{n^\epsilon} + \mathsf{poly}(n) \cdot 2^{n^\delta} = O(2^{n^\epsilon})$.
- If $\mathcal{R}$ participates in 2^{n^δ} sessions, the security loss is $2^{n^\epsilon} + 2^{n^\delta} \cdot 2^{n^\delta} = 2^{n^\epsilon} + 2^{2n^\delta} = O(2^{n^\epsilon})$.

Thus, it makes sense asymptotically to consider reductions that can participate in sub-exponentially many sessions.

b) The state of the art before our work.: There has been a long line of work on constructing non-malleable commitments with respect to commitment, in the plain model in as few rounds as possible (e.g.[1], [4], [5], [6], [7], [8], [9], [10], [11], [12], [13], [14]). In a major advance, [12] showed how to construct three-message non-malleable commitments, and subsequently [13], [14] obtained concurrent three-message non-malleable commitments. These results relied on super-polynomial or sub-exponential injective one-way functions to achieve general notions of non-malleability in three rounds. Thus, to the best of our knowledge, current constructions of even 3-message non-malleable commitments (with respect to commitment) require super-polynomial assumptions. In contrast, in this paper, we will construct the first *2-message non-malleable commitments* with respect to commitment, from standard sub-exponential assumptions.

In our work, we will also consider a weaker notion of malleable commitments called non-malleability with

[1]Alternately, an MIM is required to maul with some inverse polynomial probability in a single interaction.

[2]We contacted the author via personal communication, and he explicitly agreed that the impossibility result as stated in [3] is incorrect. As we note above, however, the only case not ruled out by Pass is a reduction that makes super-polynomially many queries to the adversary.

respect to opening (see below for a discussion of this definition), where our goal will be to construct *one-round non-malleable commitments* in the simultaneous-message model. Prior to our work, no one-round non-malleable commitment with respect to opening was known, for any flavor of the definition, in any communication model, without setup and based on standard assumptions. Before our work, the work of [15] had the fewest rounds of interaction for non-malleable commitment with respect to opening from standard assumptions. That work showed how to construct *two-round unidirectional* non-malleable commitments achieving a form of non-malleability with respect to opening, from polynomial hardness of injective one-way functions. The model and definition in [15] were carefully chosen to avoid the impossibility of [2] for two rounds, even in the polynomial hardness regime. As a result, [15] achieve a weaker definition of non-malleability with respect to opening than ours, achieve non-malleability only with respect to synchronizing adversaries, and require two rounds in the commit phase.

A. Our Results

As mentioned above, broadly speaking, there are two flavors of definitions for non-malleable commitment that have been considered in the literature, called non-malleability with respect to commitment, and non-malleability with respect to opening. We will obtain different positive results for each of these definitions.

a) Non-Malleable Commitments with respect to Commitment.: We first consider the standard model, where each round consists of a single message from one party to another. In the standard model, we work with the stronger of the two standard definitions of non-malleability, namely non-malleability with respect to commitment (against both synchronous and asynchronous adversaries). Informally, this definition requires that non-malleability hold with respect to the underlying message as soon as the commitment phase completes. Thus, even if an adversary MIM never actually opens its commitment, nevertheless we can be assured that the message underlying his commitment did not depend on the message committed to by the honest party.

In the standard model, we obtain the first positive results from standard sub-exponential assumptions, for two-round non-malleable commitments with respect to to commitment.

- We construct two-message *public-coin* non-malleable commitments with respect to commitment, assuming sub-exponentially hard one-way permutations, sub-exponential ZAPs, and sub-exponentially hard DDH.
- We obtain two-message *private-coin* non-malleable commitments with respect to commitment, assuming only sub-exponentially hard DDH or QR or N^{th}-residuosity.
- We bootstrap the above (under the same assumptions) to obtain constant[3] bounded-concurrent non-malleability while preserving round complexity.

b) Another viewpoint: Non-interactive non-malleability with a tamperable CRS.: If we were willing to rely on a trusted setup that generates a common random string (CRS) for all parties, constructions of non-interactive non-malleable commitments become much simpler [16]. However, a major design goal of all of theoretical cryptography is to reduce global trust as much as possible. A trusted CRS is a straightforward example of the kind of global trust that we would like to avoid.

Indeed, we can interpret our result above through the lens of an *untrusted* CRS: what if the man-in-the-middle attacker can arbitrarily tamper with a CRS, and convince an honest committer to generate his commitment with respect to this tampered CRS? For all prior constructions, in this situation, all bets would be off. On the other hand, our work shows the first solution to this problem: we obtain non-interactive non-malleable commitment with respect to commitment, where the honest committer must use a *tampered CRS*.

c) Non-Malleable Commitments with respect to Opening.: We next consider the simultaneous-message model, where a round consists of both (or all) parties sending a single message to all other parties. We consider the standard asynchronous model with rushing adversaries.

We achieve the first one-round non-malleable commitment protocols in this model under standard sub-exponential assumptions. To achieve one-round protocols, we work with the other definition of non-malleable commitments, called non-malleability with respect to opening. Roughly speaking, this definition requires that the adversary cannot *open* his commitment to a value related to the honest party's opened value. There are several ways to formulate the definition of non-malleability with respect to opening [17], [5], [18], [15]. We formulate a simulation-based definition that is both simpler and more powerful than the recent indistinguishability-based definition of [15] (in particular, our definition implies the definition of [15]). Furthermore, we require and obtain security against asynchronous adversaries, whereas the work of [15] required an additional round and only obtained non-malleable commitments with respect to opening against synchronous adversaries.

[3]Our actual construction imposes a trade-off between the concurrent non-malleability and the tag space. Please see the full version for a discussion of this tradeoff, and the actual bounds that we have in different settings.

In particular, in the simultaneous-message model, we obtain the following results from standard sub-exponential assumptions:

- We compile the previously described two-round protocols in the standard model to obtain *one-round* non-malleable commitments with respect to opening, in the simultaneous-message model. The opening phase of this protocol remains non-interactive.
- We further show how to transform this protocol to achieve fully concurrent non-malleable commitments with respect to opening, in the simultaneous-message exchange model, still using only one round. The opening phase of this transformed protocol remains non-interactive.
- We show that this implies concurrent *completely non-interactive non-malleable commitments with respect to opening*, in a model with a broadcast channel, and an a-priori fixed polynomial number of parties such that every party is aware of every other party in the system. To the best of our knowledge, this is the first protocol to achieve completely non-interactive non-malleability in *any* plain model setting from standard assumptions.

d) Applicability.: The general applicability of non-malleable commitments within cryptography is well known; a classic simple example is conducting sealed-bid auctions online. As stated above, in a setting where there are a fixed polynomial number of participants and a broadcast channel, our results give the first completely non-interactive method of conducting sealed-bid auctions based on standard sub-exponential assumptions.

Can we break round-complexity barriers in other settings as well? Indeed, consider the classic question of secure coin flipping [19] in a multi-party setting, where parties wish to agree on a shared random string. Note that the standard model of interaction in this setting is the simultaneous-message model. The work of [20] establish a lower bound of 4 rounds for secure multi-party coin-flipping with black-box security from polynomial hardness assumptions (with polynomial simulation). We show that by moving to the super-polynomial regime (with super-polynomial simulation), we can cut this lower bound in half! We give the first two-round bounded multiparty secure coin flipping protocol (with super-polynomial simulation) from standard super-polynomial assumptions. Note that super-polynomial simulation also implies two-round pseudorandom coin-flipping, where the output of the coin flipping protocol is indistinguishable from random even to super-polynomial time distinguishers.

B. Related Work

In less than three messages, the only prior method of achieving 2-message non-malleable commitments with respect to commitment was via the assumption of adaptive one-way functions [21], which essentially assumes the existence of a one-way function that already exhibits strong non-malleability properties. Such assumptions are very different in spirit from traditional hardness assumptions, and are both non-falsifiable [22] and not complexity assumptions in the sense of [23]. We also note that constructions of non-malleable commitments in two rounds were previously not known even based on indistinguishability obfuscation.

C. Concurrent and Independent Work

In a fascinating concurrent and independent work, Lin, Pass, and Soni (LPS) [24] construct two-message concurrent non-malleable commitments, and non-interactive non-malleable commitments with respect to commitment against uniform adversaries. Their work is substantially different from ours in terms of techniques as well as assumptions.

The constructions of LPS require many assumptions, most notably a novel sub-exponential variant of the Rivest-Shamir-Wagner (RSW) assumption first proposed for constructing time-lock puzzles by [25]. Roughly speaking, the RSW assumption considers the Repeated Squaring Algorithm for computing $h = g^{2^n}$, and requires that the natural algorithm for computing h in time n cannot be sped up by parallel computation. The novel variant of the RSW assumption considered by [24] is essentially a "two-dimensional" family of assumptions: there is a security parameter n and another parameter t, and it is required that computing $h = g^{2^{2^t}}$ cannot be done by circuits of overall size 2^{n^ϵ} and depth 2^{t^δ}, for constants ϵ and δ.

In contrast, standard subexponential assumptions in cryptography – including the assumptions that we make in our work – require only security against circuits of subexponential size, regardless of the depth of these circuits. In this way, the assumption of [24] falls outside the definition of falsifiable assumptions ruled out by Pass [2]. The authors in [24] note that assumptions of this type were previously used only in time-release cryptography. On the other hand, the assumptions that we use in our work have been considered by many previous works constructing cryptographic protocols, including secure computation protocols.

Finally, on a quantitative level, we only require $O(\log^* n)$ levels of complexity leveraging, thereby only requiring sub-subexponential hardness assumptions as per the new definition of [24].

In terms of techniques, the novel assumption on parallel complexity allows LPS to[4] construct a pair of commitment schemes Com_1 and Com_2 that are simultaneously harder than the other, in different axes. In

[4]The following text is largely copied directly from [24].

particular, Com_2 is harder in the axis of circuit-size, in the sense that Com_1 admits an extractor of size S while Com_2 is secure against all circuits of size S; on the other hand, Com_1 is harder in the axis of circuit-depth, in the sense that it admits an extractor of depth D (and some size S) while Com_1 is hiding against all circuits with depth D (and size S). This scheme already achieves a flavor of non-malleability for two tags.

In contrast, we develop new techniques to work by assuming only a single axis of hardness, in order to rely on standard subexponential hardness. Indeed, a lot of work in our paper goes into constructing extractable commitments that help us obtain a non-malleable commitment scheme for just two tags (please refer to Section II for more details).

II. Overview of Techniques

As we already discussed, we would like to build protocol that admits a security reduction that can access the (adversarial) committer a super-polynomial number of times, while an actual adversary can only interact with the honest committer in polynomially many executions. Any hope of obtaining a positive result requires us to exploit this disparity between the MIM and the reduction, otherwise our approach would succumb to the impossibility result of [3].

a) Main Tool: Extractable Commitments.: The crux of this question boils down to building a special kind of extractable commitment with just two messages. In such a commitment scheme, informally speaking, there is a black-box extractor algorithm that runs in time T', that extracts the values committed to by any malicious polynomial-time committer. Popular intuition so far has been that rewinding with only two rounds is useless: whatever the extractor can do, a malicious receiver can also do.

However, in our new extractable commitment, we will require that the hiding property of the commitment scheme holds with respect to any malicious receiver that runs in time T that exceeds T'. This seemingly contradictory requirement means that a malicious receiver should not be able to run the extractor on his own.

This is the point at which we will use the disparity in the number of interactions that a malicious receiver can participate in, versus those that an extractor can participate in. Our techniques will be centered around the following question for cryptographic protocols between parties Alice and Bob:

Can extractor E with black-box access to Alice*, gain an advantage in just 2 messages, over (malicious)* Bob *interacting with* Alice *in the actual protocol?*

As we have already discussed, we do not want to restrict the running time of Bob to be less than that of the extractor. Prior to our work, achieving black-box extraction in just 2 rounds from standard assumptions eluded all attempts at analysis.

A. Our Approach: Extractable Commitments

We devise a completely new simulation strategy that allows the reduction to gain an advantage over a malicious receiver potentially running in more time than the reduction itself. We begin by giving a high-level overview of the properties that this strategy should satisfy, after which we describe how it is implemented.

We will think of every execution of the committer as being analogous to taking *one* random walk. The receiver is also allowed *one* random walk. The receiver is given the ability to "steal" the committed value, without the committer's realization, *if and only if* the receiver's path ends up being the same as the path chosen by the committer. We set parameters up so that this event occurs with probability exactly $\frac{1}{T'}$, even if one of the parties is malicious. On the other hand, with probability $1 - \frac{1}{T'}$, the committer is "safe" in any single execution and the committed value remains well-hidden. In fact, the parameters are set so that the committed value remains well-hidden *even against* a receiver that runs in time T that is much larger than T', and interacts with the committer in polynomially many executions (we note that T' and T are set to be super-polynomial).

At the same time, an extractor that runs in time slightly larger than T' can keep rewinding a malicious committer T' times, using honest receiver strategy with fresh randomness each time. With overwhelming probability, such an extractor will succeed in crossing the committer's path in at least one execution – thereby extracting the value committed in this interaction. It is important that the committer be unable to tell whether the extractor was able to extract the committed value from a particular execution, to ensure that the distribution of extracted values is not skewed.

a) Implementing extractable commitments.: We now turn to describing the construction of extractable commitments. The commitments will be hiding against T-time receivers, and yet will be extractable by T'-time extractors where T' is much smaller than T. Formally, we will write $T' \ll T$ to mean that T' is smaller than T multiplied by any polynomial in the security parameter. At this point in the technical overview, it will be useful to assume that we have two idealized technical tools. We will in fact make do with less ideal tools, as we discuss later[5]. For now, assume that we have the following two

[5]It turns out that two round secure two-party computation with indistinguishability-based security, together with two-round zero-knowledge with super-polynomial simulation(SPS), will suffice. If uniform reductions are required, the two-round SPS ZK can be replaced with two-round strong WI [26] at the cost of requiring private coins.

primitives that can be leveraged to be secure against T-time adversaries:

- Two-message two-party computation, against semi-honest senders and malicious receivers.
- Two-message "ideal" ZK arguments.

The leprechauns described above will be implemented using secure two-party computation for the following functionality: $\mathcal{F}\big((x, M), y\big) = \left\{ \begin{array}{ll} \perp & \text{if } x \neq y \\ \mathrm{M} & \text{if } x = y \end{array} \right\}$

Intuitively, this functionality denotes the committer choosing path x and the leprechaun choosing path y, such that the leprechaun steals the committed message M if and only if $x = y$.

More formally, the receiver will sample a random challenge $\mathsf{ch} \xleftarrow{\$} \{0,1\}^m$ and the committer will sample another challenge $r \xleftarrow{\$} \{0,1\}^m$ independently. In order to commit to message M, the committer and receiver run secure two-party computation for $\mathcal{F}\big((r, M), \mathsf{ch}\big)$. The committer will also prove, via the ZK argument, that he correctly computed the output of the functionality.

Note that a malicious receiver, running in time T and participating in only a single execution, will have probability at most 2^{-m} of guessing the committer's challenge r. Thus, the commitment will still be computationally hiding against such a receiver.

On the other hand, an extractor that interacts with the committer super-polynomially many times, will have a good probability of obtaining at least one "extracting" transcript where $\mathsf{ch} = r$, and will thus find M after only slightly more than $T' = 2^m$ attempts. We must also ensure that the distribution over messages M output by the extractor is indistinguishable from the actual distribution of committed messages. We will exploit the security of two-party computation protocol against semi-honest senders, and additional complexity leveraging to ensure that the distribution of values committed by the committer cannot change between extracting transcripts and transcripts that don't allow extraction.

Finally, note that in this construction, the honest receiver is only required to verify the ZK argument (which will be public coin) – and doesn't actually need to observe the output of the two-party computation protocol. Thus, such a receiver can sample uniformly random coins to compute his message for the two-party computation protocol.

This completes an informal description of our extractable commitment, and we have the following (informal) theorem:

Informal Theorem 1. *Assume sub-exponential security of DDH, together with sub-exponentially hard one-way permutations and sub-exponential ZAPs. Then there exists a statistically binding two-round public-coin extractable commitment scheme, that is hiding against malicious receivers running in time T and extractable in time $T' \ll T$.*

For technical reasons, our actual construction of extractable commitments is a slight variant of the scheme outlined above, and can be found in the full version. In fact, this type of extractable commitment is the main technical tool that we will use to obtain our results on non-malleable commitments.

B. Two-Message Non-Malleable Commitments

1) Model: Our main result is the construction of a public-coin bounded-concurrent two-message non-malleable commitment scheme with respect to commitment, assuming sub-exponentially hard ZAPs, sub-exponential one-way permutations, and sub-exponential hardness of DDH. We also get a private coin construction assuming only sub-exponential DDH or QR or N^{th} residuosity.

Very roughly, non-malleability requires that a man-in-the-middle adversary participating in two executions, acting as a receiver interacting with an honest committer in a "left" execution, and acting as committer interacting with an honest receiver in a "right" execution, is unable to commit to a message $\widetilde{m}$ on the right, that is nontrivially related to the message m committed by the honest committer on the left.

We require non-malleability against both synchronous and asynchronous adversaries. A synchronous MIM adversary observes an honest receiver message on the right, and then generate its own (malicious) receiver message for the left execution. Then, on obtaining an honestly generated left commitment, it generates a (malicious) right commitment. An asynchronous adversary is one that completes the entire left commitment, before generating its own right commitment. Typically (and this will especially be true in our situation), it is more difficult to prove security against synchronous adversaries than against asynchronous adversaries.

In this paper, we consider a setting where parties have identities or tags, typically in $[2^n]$ and only require non-malleability to hold when the tag used by the adversary is different from the tag used by an honest party. We note that this can be compiled in a standard way (using one-time signatures) to a notion without tags that requires the MIM's committed message to be independent from that of the honest committer, unless the MIM copies the entire left transcript [1].

We now discuss a basic scheme, secure in a restricted setting where there are only two tags in the system, and the MIM's tag is guaranteed to be different from the honest committer's tag.

2) A basic scheme for two tags: The impossibility in [2] is stated for the setting of just two tags, therefore overcoming it using sub-exponential assumptions is

already non-trivial. As stated in the introduction, this will require us to exploit the gap between the number of executions available to the MIM versus those available to the reduction.

Recall that we achieved two-round extractable commitments that are secure against malicious receivers running in time T_{hid}, while extractable by extractors running in time $T_{\mathsf{Ext}} \ll T_{\mathsf{hid}}$. Having achieved such an extractable commitment scheme, we obtain a non-malleable commitment scheme for just two tags in the following way.

Let us first consider a one-sided non-malleable commitment: Suppose there are two tags 0 and 1. Then a one-sided non-malleable commitment would guarantee that the commitment with tag 1 cannot depend on a commitment with tag 0, but it would potentially enable arbitrary malleability in the other direction. Pass and Wee [27] demonstrated how to obtain a *one-sided* non-malleable commitment in this setting, based on sub-exponential assumptions.

We now illustrate how the gap between extraction and hiding of our two-round extractable commitment scheme can be used to enable two-sided non-malleable commitments, by appropriately leveraging hardness to exploit this gap. We use a two-round extractable commitment ext-com with security parameter n, that is extractable in time T_{Ext} and hiding against adversaries running in time $T_{\mathsf{hid}} \gg T_{\mathsf{Ext}}$. We also make use of a non-interactive commitment com leveraged so that it is hiding against adversaries running in time T_{Ext}, and trivially breakable in time T_{com}. We set parameters such that $T_{\mathsf{hid}} \gg T_{\mathsf{com}} \gg T_{\mathsf{Ext}}$. Then consider the following protocol:

- If $\mathsf{tag} = 0$, commit to the message m using the non-interactive commitment scheme com.
- If $\mathsf{tag} = 1$, commit to the message m using the extractable commitment scheme ext-com.

This scheme is represented in Figure 1. We consider two representative settings, one where the man-in-the-middle (MIM) is the receiver on the left, and the committer on the right (thus, $R_1 = C_2$), and second, where the MIM is the receiver on the right, and committer on the left (thus, $R_2 = C_1$).

First, we consider the case where an honest committer uses tag 0 to commit to message m, while the MIM uses tag 1. A challenger against the hiding of the non-interactive commitment com, can obtain $\mathsf{com}(0)$ or $\mathsf{com}(m)$ externally, and then exploit the extractability of ext-com that is being used by the MIM, to extract the value committed by the MIM, in time T_{Ext}.

However, the non-interactive commitment is hiding against adversaries running in time T_{Ext}. Thus, if the MIM's commitment is related to m, such a challenger can break hiding of com, by extracting the value committed by the MIM, which contradicts the T_{Ext}-hiding of the non-interactive commitment.

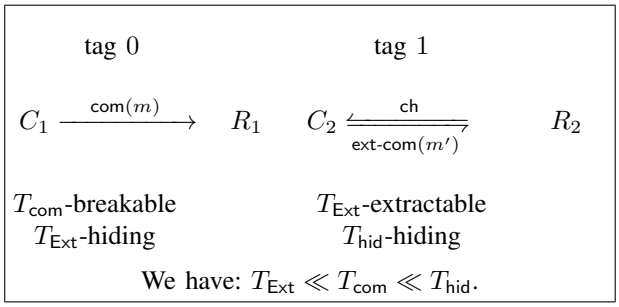

Fig. 1: A scheme for two tags

Next, let us consider the complementary case where an honest committer uses tag 1 to commit to message m, while the MIM uses 0. A challenger against the hiding of the extractable commitment ext-com, can obtain $\mathsf{ext\text{-}com}(0)$ or $\mathsf{ext\text{-}com}(m)$ externally, and then break com that is being used by the MIM, via brute-force to extract the value committed by the MIM, in time T_{com}.

However, ext-com is hiding against adversaries running in time $T_{\mathsf{hid}} \gg T_{\mathsf{com}}$. Thus, if the MIM's commitment is related to m, such a challenger can break hiding of ext-com, by breaking the commitment of the MIM using com, and extracting the value committed, in time only $T_{\mathsf{com}} \ll T_{\mathsf{hid}}$, contradicting the hiding of ext-com. We must now extend the above construction for two tags, all the way to tags in $[2^n]$. Pass and Wee [27] noted that assuming sub-exponential hardness, it is possible to obtain $O(\frac{\log n}{\log \log n})$ levels of hardness. Thus, simple complexity leveraging, even if it could be used in some way, would not help us directly go beyond $O(\frac{\log n}{\log \log n})$ tags. As a first step, we describe how the construction above can be extended to a constant number of tags.

3) A construction for constant number of tags: Note that the 2-tag construction relied on extractability of ext-com to achieve non-malleability when the adversary uses $\mathsf{tag} = 1$. Implicit in the description above, was a crucial reliance on the *non-interactivity* of the other (non-extractable) commitment.

Indeed, a problem arises when using ext-com on both sides: the extractor that extracts from the MIM on the right, naturally needs to rewind the MIM. This may result in the MIM implicitly rewinding the honest committer, possibly causing extraction even from the honest committer. If the honest commitment is non-interactive, this is not a problem because it is possible to send the *same externally obtained string* to the MIM, every time the honest committer interaction is rewound. In other words, there is no rewinding allowed in the left interaction. However, if the honest interaction consists of two rounds, then the initial challenge of the MIM to

the honest committer may change, and require a new response on the left from the honest committer. How should we simulate this response?

Let us illustrate this issue more concretely: A natural way of extending our 2-tag construction to a constant number of tags is illustrated in Figure 2, with parameters of various extractable commitment schemes adjusted (via leveraging, like in Figure 1) to ensure that:

1) For every pair of tags $\mathsf{tag} > \mathsf{tag}'$, the commitment for tag is hiding with respect to the time it takes to brute-force break the commitment for tag'.
2) The commitment associated with each tag is extractable in time less than the time with respect to which hiding is guaranteed all the tags: thus when $\mathsf{tag} < \mathsf{tag}'$ we will extract the commitment for tag' while trying to rely on the hiding of tag.

In the figure, by T-breakable, we always mean that the underlying commitment in ext-com is breakable using brute-force in time T.

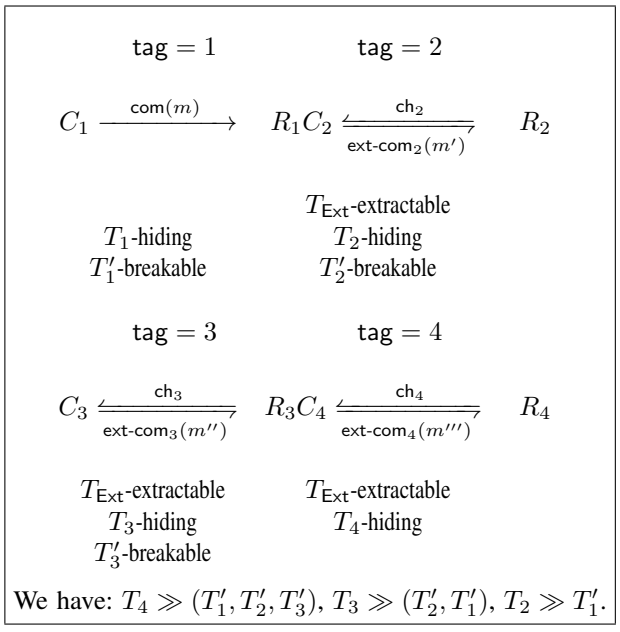

Fig. 2: An illustrative natural (but incomplete) extension to four tags

We recall (from the two-tag case) that an extractor has two possible strategies, depending on whether the honest tag is larger or smaller than the MIM's tag. If the MIM's tag is smaller than the honest tag, then it is possible to argue non-malleability by breaking (via brute-force) the commitment generated by the MIM. This part of the argument goes through exactly as in the two-tag case.

However, the proof runs into the subtle issue mentioned above when the MIM's tag is larger than the honest tag. In this case, the reduction must run the extractor on the commitment generated by the MIM. However, every time the MIM is rewound by the extractor (using different challenges for the ext-com), the MIM may generate its own fresh challenges for the honest commitment. Therefore, while extracting from the MIM, we may end up inadvertently also be extracting from the honest commitment – which would not let us achieve any contradictions. Recall that the entire point of this experiment was to extract from the man-in-the-middle while preserving hiding in the commitment generated by the honest committer.

a) Our Solution.: Our main idea to solve this problem is as follows: We set our parameters in such a way that we can "modulate" the extractability of the commitment scheme. In other words, when the MIM's tag is larger than the honest tag, the MIM's commitment will be extractable in time $T_{\mathsf{Ext},\mathsf{tag}'}$ that is *much smaller than* the time taken to extract from the honest commitment $T_{\mathsf{Ext},\mathsf{tag}}$.

In a nutshell, we will set challenge spaces (for extraction) so that, when the MIM's tag is larger than the honest tag, the MIM's challenge space is also exponentially larger than the honest challenge space. This is accomplished, in particular, just by setting the length of ch corresponding to tag, to be $(\mathsf{tag} \times p(n))$, where $p(n)$ is some fixed (small) polynomial in the security parameter n.

Not only this, we will in fact require that the honest commitment corresponding to tag be *hiding* even under $T_{\mathsf{Ext},\mathsf{tag}'}$ attempts to extract from it. This will be achieved by leveraging the advantage of the adversaries in SPS ZK and secure two-party computation appropriately. We will still be careful so that time taken for *any extraction* will be much smaller than the time required to break hiding of any of the commitments. The flexibility of our construction of extractable commitments ensures that we can set parameters appropriately.

b) Bounded-Concurrent Security.: We also prove a stronger security guarantee about the scheme outlined above, that is, we consider a setting where the MIM participates in $\ell(n)$ sessions with honest receiver(s) in which he acts as malicious committer, while obtaining a single commitment from an honest committer. We require that the *joint distribution* of the view and value committed by the MIM is unrelated to the message committed by the honest committer[6].

We prove $\ell(n)$-bounded-concurrent non-malleability of the scheme described above for polynomial $\ell(n) \ll m$, where m denotes the length of the challenge string ch for extraction. We need to set parameters appropriately for bounds $\ell(n)$. To ensure $\ell(n)$-bounded non-

[6]This notion is called one-many non-malleability (with a bounded number of right executions), and implies many-many non-malleability [5], [28].

malleability, in the sessions where MIM commits to messages, we require an an extractor that extracts the *joint distribution* of messages committed by the MIM committer.

However, upon careful observation, our extraction strategy turns out to have the following problem: The extractor extracts the value from some "rare" transcripts, and when the MIM generates multiple transcripts simultaneously, the rare transcripts that the extractor is able to extract from, may not occur simultaneously at all. We therefore need to modify the extraction strategy to keep running until it succeeds in simultaneously extracting from all of the MIM's transcripts. Note that this only happens when the extractor is able to guess *all* the challenges generated by the MIM in all its commitment sessions.

In order for such an extractor to contradict non-malleability, we need to set the parameters large enough so that hiding of the challenge commitment holds even against adversaries running in time $\mathcal{T}$, where $\mathcal{T}$ is the time taken to extract from *all* the MIM's sessions simultaneously. This helps prove bounded-concurrent non-malleability.

Our techniques for handling a constant number of tags as well as bounded-concurrent non-malleability are novel and very specific to our construction. Next, we bootstrap a (sub-exponentially secure) non-malleable commitment scheme for just 4 tags into a scheme for all tags, in a way that only requires two rounds, and preserves bounded-concurrent non-malleability. Before this, we will review a new technical tool that will help in our two-round tag amplification scheme.

4) Two-round ZK with Strong Superpolynomial Simulation: Standard constructions of two round zero-knowledge arguments with superpolynomial simulation can be described as follows: the verifier generates a challenge that is hard to invert by adversaries running in time T, then the prover proves (via a ZAP) that either the statement being proven is in the language, or that he knows the inverse of the challenge used by the verifier. This ZAP is such that the witness used by the prover can be extracted (via brute-force) in time $T' \ll T$. Naturally, this restricts the argument to be zero-knowledge against verifiers that run in time $T_{\mathsf{zk}} \ll T' \ll T$.

Thus, if a prover generates an accepting proof for a false statement, the ZAP can be broken in time T' to invert the challenge, leading to a contradiction. On the other hand, there exists a simulator that runs in time $T_{\mathsf{Sim}} \gg T$ to invert the receiver's challenge and simulate the proof (alternatively, such a simulator can non-uniformly obtain the inverse of the receiver's challenge). Thus, we have $T_{\mathsf{Sim}} \gg T_{\mathsf{zk}}$.

The notion of Zero-Knowledge with Strong Superpolynomial Simulation (SPSS-ZK) was defined by [29] as ZK with super-polynomial simulation, such that $T_{\mathsf{Sim}} \ll T_{\mathsf{zk}}$. At first glance such a primitive may seem impossible to realize[7], but let us revisit the construction of SPS ZK described above, through the lens of non-malleability.

In order to ensure soundness, what we actually require is that a cheating prover, be unable to "maul" the challenge sent by the verifier, into a witness for his own ZAP. A simple way to do this is to use complexity leveraging to get one-sided non-malleability, which is what the construction described above achieves.

However, this *constrains* $T' \ll T$, which in turn constrains $T_{\mathsf{Sim}} \gg T_{\mathsf{zk}}$. We would like to look for a different way of achieving non-malleability, which potentially allows $T' \gg T$. In other words, we would like a more efficient way of extracting the witness from the NIWI than directly breaking it via brute force. This is *exactly* the kind of non-malleability that is supported our basic construction of two-sided non-malleable commitments for two tags.

Specifically, we will just let the verifier use a non-interactive non-malleable commitment corresponding to $\mathsf{tag} = 0$, whereas the prover will use a two-message non-malleable (extractable) commitment corresponding to $\mathsf{tag} = 1$. We can now set parameters such that $T \ll T'$, which allows $T_{\mathsf{Sim}} \ll T_{\mathsf{zk}}$. On the other hand, in order to ensure soundness, we rely on the *extractability* of the prover's commitment in time $T_{\mathsf{Ext}} \ll T$.

We will use this primitive in the next subsection, while performing tag amplification while preserving bounded-concurrent non-malleability. We also believe that this primitive may be of independent interest.

5) Two-round tag amplification from 4 tags: While tag amplification has been extensively studied in the non-malleability literature (e.g. [1], [31], [6], [9]), no previous work applied to the case of 2-round protocols. We give the first tag amplification technique, for non-malleable commitments with respect to commitment, that requires just two rounds and only 4 tags to begin with, and only makes standard sub-exponential assumptions. In fact, we are able to amplify tags in by bootstrapping from a bounded-concurrent non-malleable commitment scheme for 4 tags to a bounded-concurrent non-malleable commitment scheme for all tags. Apart from being an important ingredient of our construction, this result may be of independent interest.

To build our tag amplification mechanism for 2-round protocols, we use some ideas from previous constructions [1], [31], [6], [9], while introducing new ideas to

[7]We thank Rafael Pass for pointing out that in fact, this primitive was proved impossible to realize via black-box reductions to sub-exponential assumptions in [30]. However, just like the impossibility in [3], the impossibility in [30] also no longer holds when the reduction is allowed to interact with the adversary in superpolynomially many sessions

keep the protocol at two rounds, and to minimize the number of tags that we bootstrap from.

Let us begin by recalling some ideas from previous work. Suppose we had a non-malleable commitment scheme for tags in $[2n]$. The popular DDN [1] encoding suggests a method of breaking a large tag T^j (say, in $[2^n]$) into n small tags $t_1^j, t_2^j, \ldots t_n^j$, such that for two different large tags $T^1 \neq T^2$, there exists at least one index i such that $t_i^2 \notin \{t_1^1, t_2^1, \ldots t_n^1\}$. As in other tag amplification schemes [6], [8], we will recursively apply an encoding with the property specified above. However, we would also like to be able to begin with as few tags as possible. To accomplish this, we first observe that a different encoding also achieves the same effect as DDN, but with better efficiency.

Suppose we had a scheme for tags in $[n]$. We will directly obtain a scheme for tags in $\left[\binom{n}{n/2}\right]$. Let the tag $T \in \left[\binom{n}{n/2}\right]$ itself denote a subset of $[n]$ of size $n/2$. Let $t_1, t_2, \ldots t_{n/2}$ denote the elements in T. These will now represent the small tags using which the committer must generate commitments. Note that this also satisfies the property that, given $T \neq T'$, at least one of the small tags in the set generated by T', differs from every single tag in the set generated by T, since no two sets of $n/2$ elements in $[n]$ can dominate the other. This property is sufficient for the rest of our proof to go through. Furthermore, this allows us to begin with just 4 tags and obtain $\binom{4}{2} = 6$ tags, and keep amplifying repeatedly thereafter.

Given the property of the encoding scheme, we consider the following construction: To commit to a value with large tag T, commit to the value multiple times with small tags $t_1, t_2, \ldots t_n$ that correspond to an appropriate encoding of T. Simultaneously, provide a 2-round ZK argument that all commitments are to the same value. We require the proof to be ZK against adversaries running in time T, where T is the time required to brute-force break (all components of) the underlying non-malleable commitment scheme for small tags.

In order to prove $\ell(n)$-bounded-concurrent non-malleability of the resulting scheme, we will focus on the index i_j in the MIM's j^{th} commitment, for $j \in \ell(n)$, such that the tag $\tilde{t}_{i_j} \notin \{t_1^1, t_2^1, \ldots t_n^1\}$. In the real interaction, by soundness of the ZK argument, the value committed by the MIM is identical to the value committed using $\tilde{t}_{i_j}$. Thus, it suffices to argue that this value is generated independent of the honest committer's value. Because the argument is ZK against adversaries running in time T (that is, $T_{\mathsf{zk}} \gg T$), where T is the time required to brute-force break (all components of) the non-malleable commitment with $\tilde{t}_{i,j}$, the value committed remains indistinguishable even when a challenger generates the honest commitment by simulating the ZK argument.

Next, it is possible to switch commitments using tags $t_1^1, t_2^1, \ldots t_n^1$ one by one, while the joint distribution of the values committed using tag $\tilde{t}_{i_j}$ does not change, because of $\ell(n)$-bounded concurrent non-malleability of the underlying commitment scheme. Note that here we are running in super-polynomial time T_{Sim}, so we require non-malleability to hold even against T_{Sim}-time adversaries. By our constraint on the ZK property of the argument, we will end up requiring that $T_{\mathsf{Sim}} \ll T_{\mathsf{zk}}$. This is exactly where our two-round SPSS ZK helps.

We note that this amplification can be applied recursively, several times, until non-malleability is obtained for all tags in $[2^n]$. The resulting protocol for tags in $[2^n]$ still only uses $\mathsf{poly}(n)$ commitments with small tags. Furthermore, at each recursion, the ZK argument we use will require stronger parameters. However, since the tag space grows exponentially, starting with a constant number of tags, recursion only needs to be applied $O(\log^* n)$ times. Thus, we only require $O(\log^* n)$ levels of security for the ZK and for the non-malleable commitments, which can be obtained based on sub-exponential hardness, as was also shown by Pass and Wee [27]. Apart from minor technical modifications to ensure that the resulting protocol remains efficient, this is essentially how we construct non-malleable commitments for larger tags.

Informal Theorem 2. *Assume sub-exponential security of DDH, together with sub-exponentially hard one-way permutations and sub-exponential ZAPs. Then there exists a constant bounded-concurrent statistically binding two-round public-coin non-malleable commitment scheme with respect to commitment.*

6) Instantiating the primitives: Throughout the discussion above, we assumed some idealized 2-round primitives, most notably a 2-round ZK argument, and 2-round secure two party computation. We note that almost everywhere above (except when SPSS ZK is explicitly stated), the 2-round ZK argument can be instantiated with the work of Pass [29] that builds 2-round public coin super-polynomial simulation ZK arguments. At the same time, however, it turns out that some of our proofs only need a distinguisher-dependent notion of simulation called weak ZK. Recently, a construction of such weak ZK arguments (albeit with private coins) was given in [26], [32], and by using this recent construction we also enjoy the ability to instantiate this 2-round weak ZK argument from any of the subexponential assumptions given in the set $\mathcal{Y}$ referenced in our informal theorem statements above. We note that the same construction also satisfies ZK with super-polynomial simulation.

Obtaining 2-round secure two-party computation is simpler [32]: We can use 2-round OT, secure against malicious receivers, together with garbled circuits to implement this; OT security guarantees hiding of the receiver input against semi-honest senders. We additionally rely on leveraging to ensure that the sender input is chosen independently of the receiver input. To argue sender input-indistinguishability, we extract the OT receiver's choice bits, and then can invoke the security of the garbled circuit scheme. Another option is to adapt the proof strategy in [26] to provide distinguisher-based polynomial extraction of the OT choice bits that suffice in the circumstances where we need sender input-indistinguishability.

C. One Round Non-Malleable Commitments w.r.t. Opening, with Simultaneous Messages

We define non-malleability with respect to opening by requiring that the joint distribution of the view (including both the commit and opening phase) and the value *committed* by the MIM remain indistinguishable between real and simulated executions. Of course, in the real experiment, the MIM obtains the honest committer's opening once the commit phase is over, and therefore, the simulator is also given the honest committer's opening. This definition is similar to several previously considered definitions, with the main exception being that it allows super-polynomial simulation (this restriction is because of the two-round setting). In particular, our definition implies the recent indistinguishability-based definition in [15].

a) Reordering Non-Malleable Commitments.: We observe that the extractable commitments described above can be deconstructed into two sub-protocols that occur in parallel: one sub-protocol (which we will call the commitment sub-protocol) is used to generate the actual commitment, and the other sub-protocol (which we will call the extraction sub-protocol) consists of the two-party computation together with proof of correct computation. The extraction sub-protocol is carried out purely to assist the extractor. Furthermore, the sub-protocol that generates the commitment can be made completely non-interactive by using a non-interactive statistically binding commitment based on injective one-way functions.

Moreover, the relative ordering of messages between these sub-protocols can be arbitrarily altered without affecting security. More specifically, we can reorder the extractable commitment, into the following different (still, two-round) extractable commitment in the simultaneous exchange model: In the first round of simultaneous exchange, the committer sends the commitment sub-protocol, whereas the receiver sends the first message of the extraction sub-protocol. In the second round of simultaneous exchange, the committer responds to the receiver's message for the extraction sub-protocol. This reordered scheme satisfies the same extraction properties as the previously considered scheme. In fact, in the simultaneous exchange model, this reordered scheme has an additional property: the committer is bound to his message by the end of the first round.

The non-malleable commitment scheme described previously can be similarly reordered, as we illustrate in more detail in the full version of the paper. At this point, we have a two round non-malleable commitment scheme $\mathsf{NM}-\mathsf{Com}$, with respect to commitment, in the simultaneous exchange model, that is binding in the first round.

b) Non-Malleability with respect to Opening.: The natural next step, after obtaining a non-malleable commitment scheme in the simultaneous message model, that is binding in the first round, is to try and push the second message of the non-malleable commitment into the opening phase, and send this message together with an opening.

We accomplish this by setting up the opening phase in a specific way, additionally making use of an SPSS ZK argument, with low T_{Sim} (lower than other parameters of the $\mathsf{NM}-\mathsf{Com}$) and high T_{zk} (higher than other parameters of the $\mathsf{NM}-\mathsf{Com}$).

We also obtain *fully concurrent two round non-malleable commitments with respect to commitment* in the simultaneous message setting (where the MIM can participate as malicious committer and malicious receiver in an unbounded number of sessions), details of which are provided in the full version. We use these to obtain *fully concurrent one-round non-malleable commitments with respect to opening* in the simultaneous exchange setting. These protocols make a more central use of SPSS ZK, in fact they work by first modifying the SPSS ZK to obtain a variant of simulation soundness, and then using techniques similar to those of [33] to obtain concurrent non-malleability. We believe that our round-optimal non-malleable protocols will find several other interesting applications, to low-round secure computation.

Informal Theorem 3. *Assume sub-exponential security of DDH, together with sub-exponentially hard one-way permutations and sub-exponential ZAPs. Then there exists a fully concurrent statistically binding two-round public-coin non-malleable commitment scheme with respect to commitment, in the simultaneous exchange model. Furthermore, there exists a one round fully concurrent statistically binding public-coin non-malleable commitment scheme with respect to opening, in the simultaneous exchange model.*

ACKNOWLEDGMENT

We thank Rafael Pass and Rachel Lin for valuable discussions regarding [34], [3], [24].

Research supported in part from a DARPA/ARL SAFEWARE award, NSF Frontier Award 1413955, NSF grants 1619348, 1228984, 1136174, and 1065276, BSF grant 2012378, 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 authors and do not reflect the official policy or position of the Department of Defense, the National Science Foundation, or the U.S. Government.

REFERENCES

[1] D. Dolev, C. Dwork, and M. Naor, "Non-Malleable Cryptography (Extended Abstract)," in *STOC 1991*, 1991. 1, 2, 6, 9, 10

[2] R. Pass, "Unprovable security of perfect NIZK and non-interactive non-malleable commitments," in *TCC*, 2013, pp. 334–354. 2, 3, 4, 6

[3] ——, "Unprovable security of perfect NIZK and non-interactive non-malleable commitments," *Computational Complexity*, vol. 25, no. 3, pp. 607–666, 2016. [Online]. Available: http://dx.doi.org/10.1007/s00037-016-0122-2 2, 5, 9, 12

[4] B. Barak, "Constant-Round Coin-Tossing with a Man in the Middle or Realizing the Shared Random String Model," in *FOCS 2002*, 2002, pp. 345–355. 2

[5] R. Pass and A. Rosen, "New and improved constructions of non-malleable cryptographic protocols," in *STOC 2005*, 2005, pp. 533–542. 2, 3, 8

[6] H. Wee, "Black-box, round-efficient secure computation via non-malleability amplification," in *FOCS 2010*, 2010, pp. 531–540. 2, 9, 10

[7] R. Pass and H. Wee, "Constant-round non-malleable commitments from sub-exponential one-way functions," in *EUROCRYPT 2010*, 2010, pp. 638–655. 2

[8] H. Lin and R. Pass, "Constant-round Non-malleable Commitments from Any One-way Function," in *STOC 2011*, pp. 705–714. 2, 10

[9] V. Goyal, "Constant Round Non-malleable Protocols Using One-way Functions," in *STOC 2011*. ACM, 2011, pp. 695–704. 2, 9

[10] V. Goyal, C.-K. Lee, R. Ostrovsky, and I. Visconti, "Constructing non-malleable commitments: A black-box approach," in *FOCS*, 2012. 2

[11] V. Goyal, S. Richelson, A. Rosen, and M. Vald, "An algebraic approach to non-malleability," in *FOCS 2014*, 2014, pp. 41–50. 2

[12] V. Goyal, O. Pandey, and S. Richelson, "Textbook non-malleable commitments," in *STOC*. New York, NY, USA: ACM, 2016, pp. 1128–1141. [Online]. Available: http://doi.acm.org/10.1145/2897518.2897657 2

[13] M. Ciampi, R. Ostrovsky, L. Siniscalchi, and I. Visconti, "Concurrent non-malleable commitments (and more) in 3 rounds," in *Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part III*, ser. Lecture Notes in Computer Science, M. Robshaw and J. Katz, Eds., vol. 9816. Springer, 2016, pp. 270–299. [Online]. Available: https://doi.org/10.1007/978-3-662-53015-3_10 2

[14] ——, "Four-round concurrent non-malleable commitments from one-way functions," in *Annual International Cryptology Conference*. Springer, 2017, pp. 127–157. 2

[15] V. Goyal, D. Khurana, and A. Sahai, "Breaking the three round barrier for non-malleable commitments," in *FOCS*, 2016. 3, 11

[16] G. D. Crescenzo, Y. Ishai, and R. Ostrovsky, "Non-interactive and non-malleable commitment," in *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998*, J. S. Vitter, Ed. ACM, 1998, pp. 141–150. [Online]. Available: http://doi.acm.org/10.1145/276698.276722 3

[17] R. Pass and A. Rosen, "Concurrent Non-Malleable Commitments," in *Proceedings of the 46th Annual IEEE Symposium on Foundations of ComputerScience*, ser. FOCS '05, 2005, pp. 563–572. 3

[18] R. Ostrovsky, G. Persiano, and I. Visconti, *Simulation-Based Concurrent Non-malleable Commitments and Decommitments*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 91–108. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-00457-5_7 3

[19] M. 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.*, 1981, pp. 11–15. 4

[20] S. Garg, P. Mukherjee, O. Pandey, and A. Polychroniadou, "The exact round complexity of secure computation," in *EUROCRYPT 2016*, 2016, pp. 448–476. [Online]. Available: http://dx.doi.org/10.1007/978-3-662-49896-5_16 4

[21] O. Pandey, R. Pass, and V. Vaikuntanathan, "Adaptive One-Way Functions and Applications," in *Advances in Cryptology — CRYPTO '08*, 2008, pp. 57–74. 4

[22] M. Naor, "On cryptographic assumptions and challenges," in *CRYPTO 2003*, 2003, pp. 96–109. 4

[23] O. Goldreich and H. Krawczyk, "On the composition of zero-knowledge proof systems," *SIAM Journal on Computing*, vol. 25, pp. 169–192, 1990. 4

[24] H. Lin, R. Pass, and P. Soni, "Two-round and non-interactive concurrent non-malleable commitments from time-lock puzzles," Cryptology ePrint Archive, Report 2017/273, 2017, http://eprint.iacr.org/2017/273. 4, 12

[25] R. L. Rivest, A. Shamir, and D. A. Wagner, "Time-lock puzzles and timed-release crypto," 1996. 4

[26] A. Jain, Y. T. Kalai, D. Khurana, and R. Rothblum, *Distinguisher-Dependent Simulation in Two Rounds and its Applications*. Cham: Springer International Publishing, 2017, pp. 158–189. [Online]. Available: https://doi.org/10.1007/978-3-319-63715-0_6 5, 10, 11

[27] R. Pass and H. Wee, "Black-Box Constructions of Two-Party Protocols from One-Way Functions," in *TCC 2009*. 7, 10

[28] H. Lin, R. Pass, and M. Venkitasubramaniam, "Concurrent Non-malleable Commitments from Any One-Way Function," in *TCC 2008*, pp. 571–588. 8

[29] R. Pass, "Simulation in quasi-polynomial time, and its application to protocol composition," in *EUROCRYPT 2003*, 2003, pp. 160–176. 9, 10

[30] K.-M. Chung, E. Lui, M. Mahmoody, and R. Pass, "Unprovable security of two-message zero knowledge," *IACR Cryptology ePrint Archive*, vol. 2012, p. 711, 2012. 9

[31] H. Lin and R. Pass, "Non-malleability Amplification," in *Proceedings of the 41st Annual ACM Symposium on Theory of Computing*, ser. STOC '09, 2009, pp. 189–198. 9

[32] S. Badrinarayanan, S. Garg, Y. Ishai, A. Sahai, and A. Wadia, "Two-message witness indistinguishability and secure computation in the plain model from new assumptions," *IACR Cryptology ePrint Archive*, vol. 2017, p. 433, 2017. [Online]. Available: http://eprint.iacr.org/2017/433 10, 11

[33] H. Lin, R. Pass, and M. 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*, ser. STOC '09, 2009, pp. 179–188. 11

[34] R. Pass, "Simulation in Quasi-Polynomial Time, and Its Application to Protocol Composition," in *Advances in Cryptology — EUROCRYPT '03*, 2003, pp. 160–176. 12