58th Annual IEEE Symposium on Foundations of Computer Science

Two-Round and Non-interactive Concurrent Non-Malleable Commitments from Time-Lock Puzzles

Huijia Lin*
UC Santa Barbara

Rafael Pass†
Cornell University

Pratik Soni*
UC Santa Barbara

***Abstract*—Non-malleable commitments are a fundamental cryptographic tool for preventing against (concurrent) man-in-the-middle attacks. Since their invention by Dolev, Dwork, and Naor in 1991, the round-complexity of non-malleable commitments has been extensively studied, leading up to constant-round concurrent non-malleable commitments based only on one-way functions, and even 3-round concurrent non-malleable commitments based on subexponential one-way functions.**

But constructions of *two-round*, or *non-interactive*, non-malleable commitments have so far remained elusive; the only known construction relied on a strong and non-falsifiable assumption with a non-malleability flavor. Additionally, a recent result by Pass shows the impossibility of basing two-round non-malleable commitments on falsifiable assumptions using a polynomial-time black-box security reduction.

In this work, we show how to overcome this impossibility, using super-polynomial-time hardness assumptions. Our main result demonstrates the existence of a two-round concurrent non-malleable commitment based on sub-exponential "standard-type" assumptions—notably, assuming the existence of the following primitives (all with subexponential security): (1) non-interactive commitments, (2) ZAPs (i.e., 2-round witness indistinguishable proofs), (3) collision-resistant hash functions, and (4) a "weak" time-lock puzzle.

Primitives (1),(2),(3) can be based on e.g., the discrete log assumption and the RSA assumption. Time-lock puzzles—puzzles that can be solved by "brute-force" in time 2^t, but cannot be solved significantly faster even using parallel computers—were proposed by Rivest, Shamir, and Wagner in 1996, and have been quite extensively studied since; the most popular instantiation relies on the assumption that 2^t repeated squarings mod $N = pq$ require "roughly" 2^t parallel time. Our notion of a "weak" time-lock puzzle, requires only that the puzzle cannot be solved in parallel time 2^{t^ε} (and thus we only need to rely on the relatively mild assumption that there are no *huge* improvements in the parallel complexity of repeated squaring algorithms).

We additionally show that if replacing assumption (2) for a non-interactive witness indistinguishable proof (NIWI), and (3) for a *uniform* collision-resistant hash function, then a *non-interactive* (i.e., one-message) version of our protocol satisfies concurrent non-malleability w.r.t. uniform attackers.

***Keywords*-Non-malleable commitment; 2-message; non-interactive; time-lock puzzles**

I. Introduction

Commitment schemes are one of the most fundamental cryptographic building blocks. Often described as the "digital" analogue of sealed envelopes, commitment schemes enable a *sender* to commit itself to a value while keeping it secret from the *receiver*. This property is called *hiding*. Furthermore, the commitment is *binding*, and thus in a later stage when the commitment is opened, it is guaranteed that the "opening" can yield only a single value determined in the committing stage.

For many applications, however, the most basic security guarantees of commitments are not sufficient. For instance, the basic definition of commitments does not rule out an attack where an adversary, upon seeing a commitment to a specific value v, is able to commit to a related value (say, $v - 1$), even though it does not know the actual value of v. To address this concern, Dolev, Dwork and Naor (DDN) introduced the concept of *non-malleable commitments* [1]. Loosely speaking, a commitment scheme is said to be non-malleable if it is infeasible for an adversary to "maul" a commitment to a value v into a commitment to a related value $\tilde{v}$. The notion of a *concurrent non-malleable commitment* [1], [2] further requires non-malleability to hold even if the adversary receives many commitments and can itself produce many commitments.

The first non-malleable commitment protocol was constructed in the original work of [1] in 1991, based on the minimal assumption of one-way functions. The

* {rachel.lin,pratik_soni}@cs.ucsb.edu.
Supported in part by NSF grants CNS-1528178, CNS-1514526 and CNS-1652849 (CAREER).
† rafael@cs.cornell.edu. Supported in part by an Alfred P. Sloan Fellowship, a Microsoft New Faculty Fellowship, NSF Awards CNS-1217821 and CCF-1214844, NSF CAREER Award CCF-0746990, AFOSR Award FA9550-08-1-0197, AFOSR YIP Award FA9550-10-1-0093, BSF Grant 2006317, and DARPA and AFRL under contract FA8750-11-2-0211. The views and conclusions contained in this document are those of the authors and should not be interpreted as the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the US Government.

0272-5428/17 $31.00 © 2017 IEEE
DOI 10.1109/FOCS.2017.59

IEEE computer society

first concurrently secure construction was provided by Pass and Rosen in 2005 [2]. Since then, a central question in the study of non-malleability has been to determine the exact number of communication rounds needed for achieving (concurrent) non-malleable commitments. Significant progress has been made over the years [2]–[12]. The current state-of-the-art is that 4-round concurrent non-malleable commitments can constructed based on one-way functions [13] and 3-round concurrent non-malleable commitments can be constructed from subexponentially-secure injective one-way functions [14]–[16].

On the Existence of Two-Round or Non-Interactive Non-malleable Commitments: The situation changes drastically when it comes to two-round or non-interactive (i.e., one-message) protocols: Pandey, Pass and Vaikuntanathan [7] provided a construction of a non-interactive non-malleable commitment based on a new *non-falsifiable* hardness assumption, namely, the existence of an *adaptively-secure injective one-way function*—roughly speaking, a one-way function f that is hard to invert on a random point $y = f(x)$ even if you get access to an inversion oracle that inverts it on every *other* point $y' \neq y$. This assumption is not falsifiable since the inversion oracle cannot be implemented in "real-life"; additionally, note that the assumption also has a strong non-malleability flavor—in particular, the assumption would clearly be false if one could "maul" $y = f(x)$ to e.g., $y' = f(x+1)$. As such, this construction gives us little insight into whether we can obtain two-round "non-malleability" from "pure scratch" (i.e., from "hardness" alone). Indeed, a recent work by Pass [17] showed that there are some inherent limitations to reducing 2-round non-malleability from falsifiable assumptions. More precisely, Pass shows that if a 2-round non-malleable commitment that can be proven secure using a polyomial-time (or even super-polynomial, but security preserving) black-box reduction R, then the reduction R can itself break the assumption. In particular, this rules out basing 2-round non-malleability (using black-box reduction) on falsifiable polynomial-time hardness assumptions.

Towards overcoming this barrier, a recent work by Goyal, Khurana and Sahai [18] presents a two-message protocol in a stronger "synchronous model" of communication (and achieving only a weaker notion of notion of non-malleability "w.r.t. opening"). In this work, we focus on the standard communication model (and the standard notion of non-malleability) and explore whether super-polynomial-time hardness assumptions (and using non-security preserving reductions) can be used to overcome this barrier:

Can we have non-interactive or 2-round non-malleable commitment from super-polynomial "standard-type" assumptions?

A. Our Results

Our main result demonstrates the existence of a two-round concurrent non-malleable commitment scheme based on sub-exponential "standard-type" assumptions—notably, assuming the existence of the following primitives (all with subexponential security): (1) non-interactive commitments, (2) ZAPs (i.e., 2-round witness indistinguishable proofs) [19], (3) collision-resistant hash functions, and (4) a "weak" time-lock puzzle [20].

Primitives (1),(2),(3) are all very commonly used and can be based on e.g., the discrete log assumption and the RSA assumption. Primitive (4) deserves some more discussion: *Time-lock puzzles*—roughly speaking, puzzles that can be solved in "brute-force" in time 2^t, but cannot be solved "significantly faster" even using parallel computers—were proposed by Rivest, Shamir, and Wagner in 1996 [20] (following May's work on time-release cryptography [21]), and have since been quite extensively used in the area of time-release cryptography. A bit more precisely, a $(T(\cdot), B(\cdot))$-*time-lock puzzle* enables a "sender" to efficiently generate a puzzle puzz with a solution s and a designated "level" of hardness $t = t(n)$ where n is the security parameter, so that: (i) the puzzle solution can be found in (uniform) time 2^t, but (ii) the puzzle solution cannot be recovered by any $B(n)$-size attacker with (parallel) running-time (i.e., circuit depth) $T = T(t)$ (where $T(t) << t$ determines the "hardness gap" of the puzzle). Typical applications of time-lock puzzles only require security against polynomial-size attackers, thus it suffices to let $B(\cdot)$ be any slightly super-polynomial function; however, they require the hardness gap to be very small—namely, $T = 2^{\delta t}$ or even $T = \delta 2^t$ (i.e., the problem is inherently "sequential" and the honest puzzle solver is essentially is optimal, even if you have access to parallel computers). In this work, we will need security against subexponential-size attackers, but in contrast, only require the existence of a time-lock puzzle with a relatively "large" hardness gap—we only need the puzzle to be hard to break for time $T = 2^{n^\varepsilon}$ for some constant ε.

Theorem 1 (Main Theorem, Informal)**.** *Let T and B be two arbitrary subexponential functions. Assume the existence of non-interactive commitments, a ZAP,*

a family of collision-resistant hash functions, all with subexponential-security, and the existence of a (T, B)-time-lock puzzle. Then, there exists a 2-round concurrent non-malleable commitment.

The original construction of time-lock puzzles due to Rivest, Shamir, and Wagner [20] is based on the hardness of a very natural strengthening of the factoring problem referred to as the *repeated squaring problem*: given a random RSA-modulus $N = pq$, and a random (or appropriately chosen) element g, compute

$$g^{2^{2^t}} \mod N$$

Clearly, this can be done using 2^t repeated squarings. The RSW assumption is that this task cannot be significantly sped up, even using parallel resources, unless N can be factored. Given the current state-of the art, the repeated squaring problem appears to be hard for *strongly exponential* parallel-time: $T(t) = \delta 2^t$ (that is, basically, no non-trivial speed-up to repeated squaring is possible); indeed, this strong assumption is typically used in the literature on time-release cryptography (in fact, several significantly stronger versions of this assumption, where additional leakage is given, are also typically considered—see e.g., the "generalized Blum-Blum-Schub assumption" of Boneh-Naor [22].)

Since we only need a "weakly"-secure time-lock puzzle where the hardness gap is large, it suffices for us to make a significantly weaker, *subexponential*, repeated squaring assumption, that is,

$$2^t \textit{ repeated squarings (modulo } N = pq) \textit{ cannot be done in parallel-time } 2^{t^\varepsilon}$$

More formally:

Assumption 1 (Subexponential Repeated Squaring Assumption)**.** *There exists subexponential functions T, B and a constant c such that for every function $t(\cdot)$ such that $c \log n < t(n) < B(n)$, the following holds: For every size $B(\cdot)$-attacker A with running-time (i.e,. circuit depth) $T(t(\cdot))$, there exists a negligible function μ such that for every $n \in \mathbb{N}$, the probability that A, given g, N where N is a randomly chosen n-bit RSA-modulus, and g is a randomly chosen (or appropriately fixed) element in Z_N^*, can compute $g^{2^{2^t}} \mod N$ is bounded by $\mu(n)$.*

We remark that, in our eyes, the subexponential repeated squaring assumption is milder than most "standard" subexponential assumptions used in the cryptographic literature (such as e.g., the subexponential DDH assumption, which is a decisional assumption), and has a stronger "win-win" flavor than most cryptographic assumptions: Repeated squaring is a problem that arises naturally in the design of algorithms (e.g., any improvement on repeated squaring would yield improved efficiency for the verification of RSA-based signatures.)

We finally mention that the time-lock puzzle needed for our construction can also be based on the existence of a parallel-time hard language and indistinguishability obfuscation (with subexponential security) by the work of Bitansky *et al.* [23].)

Towards Non-interactive Non-malleable Commitments: We also address the question of whether fully non-interactive (i.e., single-message) non-malleable commitments are possible. We show that if we replace the assumption of the existence of ZAPs (i.e., two-message witness indistinguishability) with non-interactive witness indistinguishable proofs (NIWI) [24]–[26], and the existence of families of collision-resistant hash functions for a *single*, *uniform*, collision-resistant hash function [27], [28], then a slightly modified *non-interactive* version of our protocol satisfies concurrent non-malleability w.r.t. *uniform attackers*: Basically, the first message of our two-round protocol only contains the first message of the ZAP, and the index of the hash function, so by relying on a NIWI and a single hash function (secure against uniform subexponential-time attackers), the first message can be skipped.

Theorem 2 (Informal)**.** *Let T and B be two arbitrary subexponential functions. Assume the existence of non-interactive commitments, a NIWI, a uniform collision-resistant function, all with subexponential-security, and the existence of a (T, B)-time-lock puzzle. Then, there exists a one-message concurrent non-malleable commitment secure w.r.t. uniform polynomial-time adversaries.*

We leave open the question of whether we can get a non-interactive non-malleable commitment w.r.t. also non-uniform attackers.

A Remark on "Sub-subexponential" Security: Let us finally mention that although for the simplicity of notation we rely on subexponential hardness assumption, our actual proof reveals that we only need to rely on "sub-subexponential" hardness assumption for all the primitives we rely on: namely, we only require security to hold w.r.t. attackers of size (and depth) $2^{n^{1/\log\log n}}$ (and in fact, even slightly less).

Why Time-Lock Puzzles? Our Ideas In a Nut Shell: In cryptography, the power, or *resource*, of attackers is usually measured by their running-time when represented as Turing machines, or equivalently by their circuit-size when represented as circuits. Time-lock puzzles, and more generally time-release cryptography [21], [22], [29]–[31], on the other hand, measure the resource

of attackers by their parallel running-time or equivalently by their circuit-depth. Our 2-round non-malleable commitments crucially rely on the synergy of these two types of resources. The key idea is, instead of measuring the hardness of commitment schemes in a single "axis" of resource, measure the hardness in two axes, one refers to circuit-size and the other to circuit-depth. By doing so, we can construct a pair of commitment schemes $\mathsf{Com}_1, \mathsf{Com}_2$ that are simultaneously harder than the other, in different axes. In 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). Such a pair of commitment schemes that are mutually harder than each other already has a weak flavor of non-malleability, which can then be amplified to achieve full-fledged non-malleability. More precisely, we transform the aforementioned commitment schemes, which are non-malleable w.r.t. short "tags" to that for much longer "tags" (explained below), while keeping two rounds. A step in the transformation lifts non-malleability in the stand-alone setting to that in the concurrent setting.

B. Concurrent and Independent Work

A concurrent and independent, beautiful, work by Khurana and Sahai (KS) [32], [33] also presents a construction of 2-round non-malleable commitments from subexponential "standard-type" assumptions. The results, however, are incomparable, both in terms of assumptions, and also in terms of the achieved results (and use significantly different techniques).

In terms of the achieved results, our protocols satisfy *full* concurrent non-malleability, whereas the KS protocol only satisfies "bounded-concurrent" non-malleability—which is a weaker notion of concurrent non-malleability where the number of sessions is an *a-priori bounded* by some pre-determined polynomial in the security parameter; in particular, the communication complexity of their protocol grows super linearly with the bound on the number of sessions, and the complexity assumptions they rely on need to be parametrized by it. Additionally, we also present a fully non-interactive protocol, whereas their technique appears to be inherently limited to two-round protocols.

In terms of assumptions, the key difference is that KS does not rely on time-lock puzzles but rather on the existence of certain 2-round secure two-party computation protocols (with super-polynomial-time simulation security); they also claim that such protocols can be constructed based on the subexponential DDH assumption, or the subexponential QR assumption. These assumptions are incomparable to the subexponential repeated squaring assumption. While DDH and QR are clearly more typical assumptions in the literature on cryptographic protocols, as we mentioned above, the repeated squaring assumption is, in our eyes, perhaps an even more natural computational problem that has been extensively studied over the years. On a qualitative level, it is also a search assumption (and thus our construction of non-malleable commitments can be based on search assumptions), whereas the KS construction (due to the above DDH, or QR, assumption) relies on "decisional assumptions".

C. Organization

In Section II, we give a detailed overview of our approach for constructing 2-round non-malleable commitments. Namely, first we construct commitment schemes for short tags that are both mutually hard but in different axes and show that these already have a weak flavour of non-malleability. Then, we amplify the weak non-malleability to standard notion of non-malleability using a novel round-preserving transformation, a detailed account of which is presented in Section III. We refer the reader to [34] for more details.

II. Overview

Every statistically binding commitment scheme is *hiding* against polynomial-sized circuits, while *extractable* by some exponential-sized circuit (such an extractor is guaranteed to exist since one can always find the committed value by brute force). In this work, we pay special attention to the *gap* between the "resources" of attackers and that of extractors. Moreover, we crucially rely on the synergy between different resources — in particular, *circuit-size* and *circuit-depth*, which are captured by the following two basic types of commitment schemes:

Size-Robust Commitments are parametrized versions of classical commitments: An (S, S')-*size-robust commitment* is hiding against any size-$\mathsf{poly}(S)$ attackers, and extractable by some size-S' extractor, for an $S' = S^{\omega(1)}$ denoted as $S' >> S$. Importantly, the extractor has large size, but *shallow* polynomial depth. Such extractors can be implemented using the naïve brute force strategy of enumerating all possible decommitments, which is a time-consuming but a highly-parallelizable task.

Depth-Robust Commitments are natural analogues of size-robust commitments, but with respect to the resource of circuit-depth. A (D, D')-*depth-robust commitment* is hiding against any depth-$\mathsf{poly}(D)$ circuits with size up to a large upper bound B, and extractable by some *size-D'* extractor for a $D' >> D$ that necessarily has a depth super-polynomially larger than D. In this work, we consider a subexponential size upper bound $B = 2^{n^\varepsilon}$ for some constant $\varepsilon > 0$; for simplicity of exposition, we ignore this upper bound in the rest of this overview (see Section 4 in [34] for more detail).

Size-Robust Commitments from Subexponential Injective OWFs: Size-robust commitments can essentially be instantiated using any off-the-shelf commitment schemes that are subexponential secure, by appropriately scaling the security parameter to control the levels of security and hardness for extraction. Take the standard non-interactive commitment scheme from any injective one-way function f as an example: A commitment to a bit b is of form $f(r), h(r) \oplus b$, consisting of the image $f(r)$ of a random string r of length n, and the committed bit b XORed with the hard-core bit $h(r)$. Assuming that f is subexponentially hard to invert, the commitment is hiding against all size-2^{n^ε} circuits for some constant $\varepsilon > 0$, while extractable in size 2^n (ignoring polynomial factors in n) and polynomial depth. By setting the security parameter n to $(\log S)^{1/\varepsilon}$, we immediately obtain a (S, S')-size robust commitment for $S' = 2^{\log S^{1/\varepsilon}}$.

Depth-Robust Commitments from Time-Lock Puzzles: Depth-robust commitments are naturally connected with cryptographic objects that consider parallel-time complexity, which corresponds to circuit-depth. When replacing subexponentially-hard one-way functions in the above construction with time-lock puzzles, we immediately obtain depth-robust commitments:

- To commit to a bit b, generate a puzzle puzz with a random solution s and a designated level of hardness t, and hide b using the Goldreich-Levin hard-core bit, producing $C = (\mathsf{puzz}, r, \langle r, s\rangle \oplus b)$ as the commitment.
- To decommit, the committer can simply reveal the puzzle solution s together with the random coins ρ used for generating the puzzle. The receiver verifies that the puzzle is honestly generated with solution s, and uses s to recover the committed bit b.

Since the time-lock puzzle solution s is hidden against adversaries in parallel-time $T(t)$ (and overall time $B(n)$), the commitments are hiding against depth-$T(t)$ adversaries (with size up to $B(n)$). Moreover, since the puzzles can be "forcefully" solved in time 2^t, the committed values can be extracted in size 2^t. This gives a $(T, 2^t)$-depth-robust commitment.

Next, we show how to compose the basic size-robust and depth-robust commitment schemes to overcome Pass's impossibility result on 2-round non-malleable commitments.

A. Towards Overcoming the Impossibility Result

In the literature, there are two formulations of non-malleable commitments, depending on whether the commitment scheme uses players' *identities* or not. The formulation with identities, adopted in this work, assumes that the players have identities of certain length ℓ, and that the commitment protocol depends on the identity of the committer, which is also referred to as the *tag* of the interaction. Non-malleability ensures that, as long as the tags of the left and right commitments are different (that is, the man-in-the-middle does not copy the identity of the left committer), no man-in-the-middle attacker can "maul" a commitment it receives *on the left* into a commitment of a related value it gives *on the right*. This is formalized by requiring that for any two values v_1, v_2, the values the man-in-the-middle commits to after receiving left commitments to v_1 or v_2 are indistinguishable.

The length ℓ of the tags can be viewed as a quantitative measure of how non-malleable a scheme is: An ℓ-bit tag non-malleable commitment gives a family of 2^ℓ commitment schemes — each with a hardwired tag — that are "mutually non-malleable" to each other. Therefore, the shorter the tags are, the easier it is to construct such a family. Full-fledged non-malleable commitments have tags of length equal to the security parameter $\ell = n$, and hence corresponds to a exponentially sized family. However, when the number of communication rounds is restricted to 2, Pass [17] showed that even the weakest non-malleable commitment for just *1-bit tags*, corresponding to a size 2 family, cannot be reduced from falsifiable assumptions, via a polynomial-time black-box reduction.

One-Sided Non-Malleability via Complexity Leveraging: It is well known that *one-sided non-malleability* can be achieved easily via complexity leveraging. One-sided non-malleability only prevents mauling attacks when the tag of the left commitment is "larger than" the tag of the right commitment [1]. In the simple case of

[1]The choice that the left tag is smaller than the right tag is not important. One could also require the opposite that the left tag is larger than the right tag. The limitation is that the design of the commitments depends on this arbitrary decision.

1-bit tags, this requires the commitment for tag 1 (on the left) to be non-malleable w.r.t. the commitment for tag 0 (on the right), which holds if the tag-1 commitment is "harder" than the tag-0 commitment. For example, if the tag-1 commitment is (S_1, S_1')-size-robust while the tag-0 commitment is (S_0, S_0')-size-robust for some $S_0 << S_0' << S_1 << S_1'$, then one can extract the right committed value using a size-S_1 extractor, while the left committed value still remain hidden. Therefore, the right committed value must be (computationally) independent of the left. Similarly, we can also achieve one-sided non-malleability using depth-robust commitments, by using a (D_1, D_1')-depth robust commitment scheme for tag 1 and a (D_0, D_0')-depth robust commitment scheme for tag 0, for some $D_0 << D_0' << D_1 << D_1'$.

However, simple complexity leveraging is inherently limited to one-sided non-malleability, since when only one resource is considered, the tag-1 commitment cannot be both harder and easier than the tag-0 commitment.

Two Resources for (Two-Sided) Non-Malleability: Therefore, our key idea is using two resources to create two "axes", such that, the tag-1 commitment and tag-0 commitment are simultaneously "harder" than the other, but, with respect to different resources. This is achieved by combining the basic size-robust and depth-robust commitment schemes in the following simple way.

Basic 1-bit Tag Non-Malleable Commitment:
For some $D_0 << D_0' << D_1 << D_1' << S_0 << S_0' << S_1 << S_1'$,

- a tag-0 commitment to a value v consists of commitments to two random secret shares α, β of v, such that, $v = \alpha + \beta$, where the first share is committed under a (D_0, D_0')-depth-robust commitment scheme and the second under a (S_1, S_1')-size-robust commitment scheme, and
- a tag-1 commitment to v, on the other hand, uses a (D_1, D_1')-depth-robust commitment scheme to commit to the first share and a (S_0, S_0')-size-robust commitment scheme to commit to the second share.

Thus, the tag-1 commitment is harder w.r.t. circuit-depth, while the tag-0 commitment is harder w.r.t. circuit-size. Leveraging this difference, one can extract from a tag-0 commitment (on the right) without violating the hiding property of a tag-1 commitment (on the left), and vice versa — leading to two-sided non-malleability. More specifically, the committed values in a tag-0 commitment can be extracted in depth D_0' and size S_1' by extracting both secret shares from the size- and depth-robust commitments contained in it. Yet, adversaries with such depth and size cannot break the (D_1, D_1')-depth-robust commitment contained in a tag-1 commitment; thus, the value committed to in the tag-1 commitment remains hidden. On the flip side, the committed value in a tag-1 commitment can be extracted in depth D_1' and size S_0', and, similarly, adversaries with such depth and size do not violate the hiding of a tag-0 commitment, due to the fact that the size-robust commitment contained in it is hiding against size-S_1 adversaries.

In summary, combining the two types of commitment schemes gives us depth-and-size robust commitment schemes: A $(D \vee S, D' \wedge S')$-robust commitment is hiding against circuits with depth below D *or* size below S, while extractable by some circuit with depth D *and* size S, as illustrated in Figure 1 (left). In this language, a tag-0 commitment is $(D_0 \vee S_1, D_0' \wedge S_1')$-robust while a tag-1 commitment is $(D_1 \vee S_0, D_1' \wedge S_0')$-robust. They are mutually non-malleable, because the extractor for one falls into the class of adversaries that the other is hiding against.

The Subtle Issue of Over-Extraction: The above argument captures our key idea, but is overly-simplified. It implicitly assumes that the size- and depth-robust commitments are extractable in the perfect manner: 1) Whenever a commitment is valid, in the sense that there exists an accepting decommitment, the extractor outputs exactly the committed value, otherwise, 2) when the commitment is invalid, it outputs $\perp$. Such strong extractability ensures that to show non-malleability that the right *committed* value is independent of the left committed value, it suffices to show that the right *extracted* value is independent of the left committed value, as argued above.

However, our depth-robust commitments from time-lock puzzles do not satisfy such strong extractability. [2] In particular, they do not satisfy the second property above: When commitments are invalid, the extractor can output arbitrary values — this is known as "over-extraction". Over-extraction traces back to the fact that only *honestly generated* time-lock puzzles (*i.e.*, in the domain of the puzzle generation algorithm) are guaranteed to be solvable in certain time. There is no guarantee for ill-generated puzzles, and no efficient procedure for deciding whether a puzzle is honestly generated or not. Observe that this is the case for the time-lock puzzles proposed by Rivest, Shamir, and Wagner [20], since given a puzzle $(s + a^{2^{2^t}} \mod N,\ N)$ one can extract s using 2^t squaring modular N, but cannot obtain a

[2] Our size-robust commitments from injective one-way functions do satisfy such strong extractability.

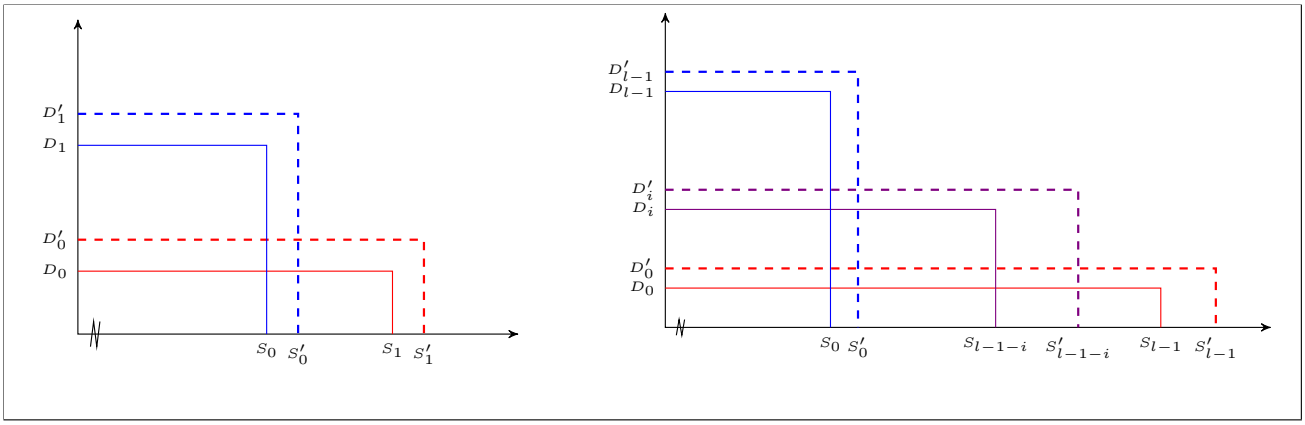

Figure 1: (left) A 1-bit tag based commitment scheme: The tag-0 (resp., tag-1) commitment scheme is hiding for circuits of depth below D_0 (resp., D_1) *OR* size below S_1 (resp., S_0), represented by the solid line joining D_0 (resp., D_1) and S_1 (resp., S_0). The tag-0 (resp., tag-1) commitment scheme admits an extractor of depth at most D'_0 (resp., D'_1) and size at most S'_1 (resp., S'_0). (right) This is a generalization of the 1-bit tag commitment scheme to $\log l$-bits tags, where for tag-i the commitment scheme is hiding for circuits of depth below D_i *OR* size below S_{l-1-i} and exhibits an extractor of depth at most D'_i and size at most S'_{l-1-i}.

proof that N is a valid RSA-modulus; this is also the case for the other puzzle construction [23]. As a result, the extractor of our depth-robust commitments that extracts committed values via solving time-lock puzzles, provides no guarantees when commitments are invalid.

This means that our basic 1-bit tag commitment scheme is over-extractable, and the argument above that reasons about the right extracted value fails to establish non-malleability. Nevertheless, the basic scheme does satisfy a variant of non-malleability that we call *non-malleability w.r.t. extraction*, which ensures that the value *extracted* from the right commitment is independent of the left committed value.[3] When a commitment scheme is perfectly-extractable, this new notion is equivalent to standard non-malleability (w.r.t. commitment), but with over-extraction, it becomes incomparable. The issue of over-extraction has appeared in the literature (*e.g.*, [9], [35]), standard methods for eliminating it requires the committer to additionally prove the validity of the commitment it sends, using for instance zero-knowledge protocols or cut-and-choose techniques. However, these methods take more than 2 rounds of interaction, and do not apply here.

[3]Our notion of non-malleability w.r.t. extraction can be viewed as a special case of the notion of non-malleability w.r.t. replacement defined by Goyal [10], in the sense that the replacer in Goyal's definition is fixed to the over-extractor of the commitment scheme. The benefit of doing so is that we know exactly the complexity of the extractor, which is useful in the rest of the construction.

B. Full-Fledged Non-Malleable Commitments

At this point, we face two challenges towards constructing full-fledged non-malleable commitments:

- *Challenge 1:* We need to go from non-malleability w.r.t. extraction to non-malleability w.r.t. commitment in 2 rounds. Resolving this challenge would give a 2-round 1-bit tag non-malleable commitment scheme.
- *Challenge 2:* The next challenge is going beyond two tags, towards supporting an exponential 2^n number of tags.
  It is easy to generalize our basic 1-bit tag commitment scheme to handle arbitrary l tags, if there exists a "ladder" of l commitment schemes with increasing levels of depth-robustness, and another "ladder" of l schemes with increasing levels of size-robustness. Concretely, the i'th schemes are respectively (D_i, D'_i)-depth robust and (S_i, S'_i)-size robust, for some

$$\cdots << D_i << D'_i << \cdots << D_l << D'_l << \\ S_0 << S'_0 \cdots << S_i << S'_i << \cdots .$$

  A commitment with tag $i \in \{0, \cdots, l-1\}$ combines the i'th (D_i, D'_i)-depth-robust scheme and the $(l-1-i)$'th (S_{l-1-i}, S'_{l-i-1})-size-robust scheme to commit to a pair of secret shares of the committed value. This gives a family of l mutually non-malleable commitment schemes, as illustrated in Figure 1 (right).
  To directly obtain full-fledged non-malleable commitments, we need an exponential number of

levels $l = 2^n$ of depth- and size-robustness, which is, however, impossible from the underlying assumptions. From subexponentially hard injective one-way functions, we can instantiate at most $O(\log n/\log\log n)$ levels of size-robustness, and similarly, from subexponentially parallel-time hard time-lock puzzles, we can instantiate $O(\log n/\log\log n)$ levels of depth-robustness. Therefore, we need to amplify the number of tags.

We address both challenges using the a single transformation.

2-Round Tag Amplification Technique: We present a transformation that converts a 2-round l-tag commitment scheme that is non-malleable w.r.t. extraction, into a 2-round 2^{l-1}-tag commitment scheme that is both non-malleable w.r.t. extraction and w.r.t. commitment. The output protocol can be further transformed to achieve concurrent non-malleability.

With the above transformation, we can now construct full-fledged non-malleable commitment. Start from our basic scheme for a constant $l_0 = O(1)$ number of tags that is non-malleable w.r.t. extraction; apply the tag-amplification technique *iteratively for* $m = O(\log^* n)$ *times* to obtain a scheme for $l_m = 2^n$ tags that is both non-malleable w.r.t. extraction and w.r.t. commitment.

Previously, similar tag-amplification techniques were presented by Lin and Pass [6] and Wee [9]. Our transformation follows the same blueprint, but differ at two important aspects. First, our transformation starts with and preserves non-malleable w.r.t. extractability, which is not considered in their work. Second, their amplification techniques incur a constant additive overhead in the round complexity of the protocol, whereas our transformation keeps the number of rounds invariant at 2. To do so, our amplification step combines ideas from previous works with the new idea of using our depth-and-size robust commitments to create different 2-round sub-protocols that are mutually "non-malleable" when executed in parallel, in the sense that the security of one sub-protocol remains intact even when the security of another is violated by force.

C. Overview of Our 2-Round Tag-Amplification

Similar to [6], [9], the transformation proceeds in two steps:

- First, amplify the security of a scheme from (*one-one*) non-malleability w.r.t. extraction to *one-many* non-malleability w.r.t. extraction and commitment, which, following a proof in [5], implies *concurrent* (or many-many) non-malleability w.r.t. extraction and commitment. (This is why our final protocol can be made concurrently non-malleable.) Here, one-many and concurrent non-malleability w.r.t. extraction or commitment naturally generalize standard non-malleability to the setting where the man-in-the-middle concurrently receives one or many commitments on the left and gives many commitments on the right, and ensures that the joint distribution of the values extracted from or committed in right commitments is independent of the value(s) committed in the left.
- Next, apply the "log-n trick" by Dolev, Dwork and Naor [19] to amplify the number of tags supported from l to 2^{l-1} at the price of losing concurrent security, yielding a protocol that is (*one-one*) non-malleable w.r.t. extraction and commitment.

The main technical challenges lie in the first step. We briefly review the LP approach. At a high-level, they construct one-many non-malleable commitment following the Fiat-Shamir paradigm: The receiver starts by setting up a *hidden* "trapdoor" t. The sender commits to a value v using an arbitrary (potentially malleable) 2-message commitment scheme, followed by committing to 0^n using a (one-one) non-malleable commitment and proving using *many* witness-indistinguishable proofs of knowledge (WIPOK) that either it knows a decommitment to v *or* it knows a decommitment of the non-malleable commitment to the trapdoor t; the former, called the honest witness, is used by the honest committer, while the latter, called the fake witness, is used for simulation.

The LP protocol arranges all components — the trapdoor-setup, commitment to v, non-malleable commitment (for trapdoor), and every WIPOK — *sequentially*. To compress the protocol into 2 rounds, we run all components in *parallel*, and replace multiple WIPOK proofs with a single 2-round ZAP proof.

Unfortunately, arranging all components in parallel renders the proof of one-many non-malleability in LP invalid. They designed a sequence of hybrids in which different components in the (single) left interaction are gradually switched from being honestly generated to simulated, while maintaining two invariants regarding the (many) right interactions. First, the *soundness* condition states that the man-in-the-middle never commits to a trapdoor in any right interaction. Second, in every right interaction, there is always a WIPOK that can be rewound to extract the value committed to in this interaction, without rewinding the left component being changed; the value extracted must be a valid decommitment since the fake witness does not exist

by the soundness invariant — this establishes *strong extractability*. The second invariant is true because the LP protocol contains sufficiently many sequential WIPOKs so that there is always a proof that does not interleave with the left-component being changed. The first invariant, on the other hand, relies not only on the non-malleability of the input commitment scheme, but also on its "robustness" to other components that have a small fixed k number of interactions (such as 2-message commitment and WIPOK). The robustness captures "non-malleability" w.r.t. other protocols, and is achieved by embedding more than k rewinding slots in the input commitment scheme.

In our 2-round protocol, we cannot afford to have many rewinding slots for extraction, nor for establishing non-malleability between different components. Naturally, we resort to our size-and-depth robust commitments, which can be made mutually non-malleable w.r.t. extraction by setting the appropriate profiles of size-and-depth robustness. We embed a family of 4 such schemes in different components of the protocol, and mimic the LP proof in the following (overly-simplified) manner: In every hybrid, in the left interaction, either a size-and-depth robust commitment or the non-malleable commitment is changed, while on the right, values are extracted from a *different* size-and-depth robust commitment and from the non-malleable commitment. This idea works, although we need to overcome several major challenges; we describe the challenges and show how to overcome them in the next section.

Finally, in the above transformation, the hardness of size-and-depth robust commitments must be set appropriately according to that of the non-malleable commitment scheme.

This concludes the overview of our construction of 2-message concurrent non-malleable commitment. Due to the lack of space, we refer the reader to the full version [34] for formal description of the protocols and security proofs. Furthermore, the non-malleable commitment scheme produced by the above transformation has weaker security than the input scheme. As a result, to iteratively apply the tag-amplification technique for $O(\log^* n)$ times, we need $O(\log^* n)$ levels of depth- and size-robustness. This can be easily instantiated using subexponentially secure non-interactive commitment schemes and time-lock puzzles as stated in Theorem 1.

III. Strengthening Non-malleability

We now give more details on how to transform a commitment scheme $\langle C, R\rangle$ that is only (one-one) non-malleable w.r.t. extraction, into a commitment scheme $\langle \widehat{C}, \widehat{R}\rangle$ that is (one-many) non-malleable w.r.t. commitment and w.r.t. extraction. As discussed above, we start with the following bare-bone protocol inspired by [6].

A Bare-Bone Protocol $\langle \widehat{C}, \widehat{R}\rangle$: The receiver sends a *puzzle* puzz (whose solution s is the required trapdoor), together with the first message a_{NM} of $\langle C, R\rangle$ and the first message a_{ZAP} of ZAP. The committer commits to v using a non-interactive commitment scheme Com, sends the second message b_{NM} of $\langle C, R\rangle$ committing to a random string $r1$, and the second message b_{ZAP} of ZAP proving that either i) $c1$ commits to v or ii) $(a_{\mathsf{NM}}, b_{\mathsf{NM}})$ commits to a solution s of the puzzle puzz (which is efficiently verifiable).

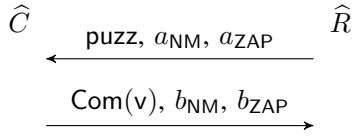

To show the security, *ideally*, we would like different components — puzz, $\langle C, R\rangle$, Com, and ZAP — to be *mutually non-malleable*. Informally speaking, we say that a primitive P is more secure than a primitive Q, denoted as $\mathsf{P} \succ \mathsf{Q}$, if the security of P holds even when security of Q is broken by force; P and Q are mutually non-malleable if $\mathsf{P} \prec\succ \mathsf{Q}$. The ideal configuration is illustrated in Figure 2 (i). Towards realizing as many constraints in the ideal configuration as possible, the first idea is using three size-and-depth robust commitment schemes $\mathsf{ECom}_1, \mathsf{ECom}_4, \mathsf{ECom}_3$ to implement Com and puzz, and augment ZAP so that they become mutually non-malleable. But, we run into problems with respect to the input non-malleable commitment $\langle C, R\rangle$.

Challenge 1: $\langle C, R\rangle$ is only secure against adversaries which have both bounded depth *AND* bounded size. This is the case for the basic schemes described in 1, as well as the schemes produced by the transformation in this section.) This type of *AND* security means either a primitive P is more secure than $\langle C, R\rangle$ or less, but cannot be mutually non-malleable. Though through a more careful analysis, we can remove some constraints w.r.t. the non-malleable commitment, it still requires $\langle C, R\rangle \prec\succ \mathsf{puzz}$, in order to show the security of the bare-bone protocol.

Challenge 2: In addition, constructing a puzzle from size-and-depth robust commitment ECom_4 is not straightforward. If we naively use $\mathsf{puzz} = \mathsf{ECom}_4(s)$ as a puzzle, a malicious man-in-the-middle can send an invalid commitment, which has no solution; this would make the security proof stuck. To prevent this, one straightforward

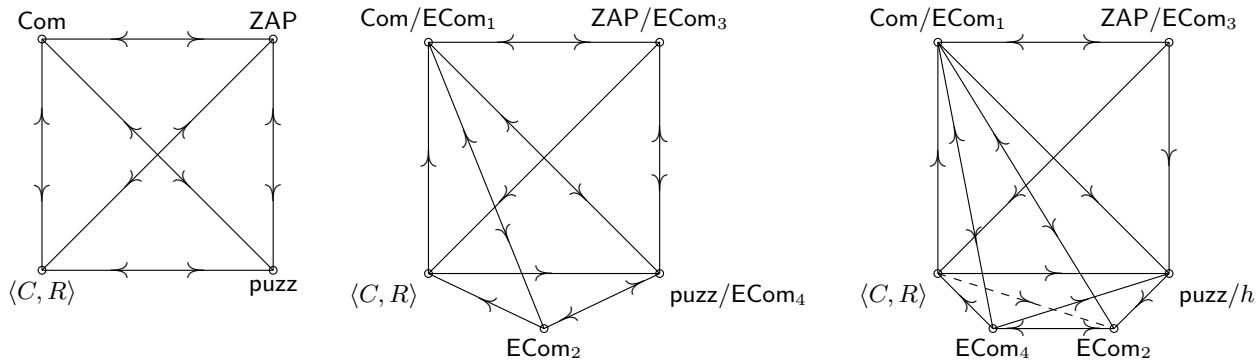

(i) Ideal Configuration (ii) Assume NIWI (iii) Assume CRH

Figure 2: The relation between different primitives. **(i)**: The ideal configuration where all primitives are mutually non-malleable to each other; however, it cannot be instantiated. **(ii)** A sufficient configuration; it can be instantiated assuming NIWI. **(iii)**: A sufficient configuration, which can be instantiated assuming collision resistant hash functions or one-way permutations. (The dashed line is by transitivity.)

approach is asking the receiver to send two puzzles and prove using NIWI that at least one of them is well-formed. However, this requires relying the exsitence of NIWI.

To resolve Challenge 1, we modify the bare bone protocol using an additional size-and-robust commitment ECom_2. The key idea is creating a "buffer" between $\langle C, R\rangle$ and puzz, by setting the following relation: $\mathsf{ECom}_2 \succ \langle C, R\rangle$, $\langle C, R\rangle \succ$ puzz, and $\mathsf{ECom}_2 \prec\!\succ$ puzz, as illustrated in Figure 2 (ii). Note that now the non-malleable commitment does not need to satisfy mutual non-malleability with either ECom_2 or puzz. On the other hand, the mutual non-malleability of ECom_2 and puzz helps the security proof to go through.

However, to fulfill the relation $\mathsf{ECom}_2 \prec\!\succ$ puzz, it seems necessary to instantiate puzz using a size-and-depth robust commitment scheme, which however would involve using NIWI. To avoid this, we would like to set puzz to be, for example, a randomly chosen collision resistant hash (CRH) function h, or a randomly chosen image $y = f(s)$ of a one-way permutation (OWP), whose corresponding solutions are respectively a collision of h and a preimage of y. These puzzles have the advantage that their validity are efficiently verifiable and hence NIWI can be disposed. But, a problem with using, say, h as the puzzle is that, it cannot be mutually non-malleable with ECom_2. To resolve this, we use a $h \succ \mathsf{ECom}_2$, and to compensate for the fact that $h \not\succ \mathsf{ECom}_2$, we use non-uniformity in the proof as follows: When reducing to the security of ECom_2, the reduction instead of finding a collisioin of h by force, receives a collision as a non-uniform advice. This can be done since the puzzle h is sent in the first message completely before the ECom_2 commitment.

Unfortunately, instantiating the puzzles using CRH or OWP creates another problem: Given that $\langle C, R\rangle \succ$ puzz $= h$ and $h \succ \mathsf{ECom}_2$, it actually implies that $\langle C, R\rangle \succ \mathsf{ECom}_2$. This transitivity holds because the h is only secure against attackers with bounded size. (If h were replaced with another size-and-depth robust commitment ECom', then transitivity does not hold in general.) But this means $\langle C, R\rangle$ needs to be mutually non-malleable with ECom_2 again. To solve this problem, we again use the idea of creating "buffers". More specifically, we set the following relation: $\mathsf{ECom}_4 \succ \langle C, R\rangle$, $\langle C, R\rangle \succ$ puzz, puzz $\succ \mathsf{ECom}_2$, and $\mathsf{ECom}_2 \prec\!\succ \mathsf{ECom}_4$, as illustrated in Figure 2 (iii). Now transitivity implies that $\langle C, R\rangle \succ \mathsf{ECom}_2$, but $\langle C, R\rangle$ no longer need to be simultanously weaker than ECom_2, and only needs to be weaker than the new "buffer" ECom_4. Moreover, the mutual non-malleability between ECom_2 and ECom_4 helps the proof to go through.

Commitment Scheme $\langle \widehat{C}, \widehat{R}\rangle$: With building blocks described above, namely $\mathsf{ECom}_1, \cdots, \mathsf{ECom}_4$, h, $\langle C, R\rangle$ satisyfing relations in Figure 2 (iii), we now describe our construction of $\langle \widehat{C}, \widehat{R}\rangle$ that is both concurrently non-malleable w.r.t. extraction, and concurrently non-malleable w.r.t. commitment.

The committer $\widehat{C}$ and the receiver $\widehat{R}$ receive the security parameter 1^n and identity $\mathsf{id} \in \{0,1\}^{t(n)}$ as common input. Furthermore, $\widehat{C}$ gets a private input $v \in \{0,1\}^n$ which is the value to be committed.

- Commit stage - First round:

  1) $\widehat{R}$ samples a hash function h from $\mathcal{H}$.

2) $\widehat{R}$ samples the first message a_{ZAP} of ZAP.
3) $\widehat{R}$ generates the first message a_{NM} of $\langle C, R\rangle$ using the honest receiver R with identity id.
4) $\widehat{R}$ sends $(h, a_{\mathsf{ZAP}}, a_{\mathsf{NM}})$ as the first round message to $\widehat{C}$.

- Commit stage - Second round:

1) a) $\widehat{C}$ computes a commitment $c1$ to the value v using ECom_1. Let $d1$ be the corresponding decommitment string.
   b) $\widehat{C}$ computes a commitment $c3$ to the decommitment $(v, d1)$ of $c1$ using ECom_3.
2) a) $\widehat{C}$ computes a commitment $c2$ to a random string $r1$ using ECom_2.
   b) Given a_{NM}, $\widehat{C}$ computes the second message b_{NM} of $\langle \hat{C}, \hat{R}\rangle$ using the honest committer C with identity id to commit to a random string $r2$.
   c) $\widehat{C}$ computes a commitment $c4$ to a random string $r3$ using ECom_4.
3) Given a_{ZAP}, $\widehat{C}$ computes the second message b_{ZAP} of ZAP to prove the following OR-statement:
   a) *either* there exists a string $\bar{v}$ such that $c1$ is a commitment to $\bar{v}$ and $c3$ commits to a decommitment of $c1$.
   b) *or* there exists a string $\bar{s} = (x_1, x_2)$ such that $c2$ is a commitment to $\bar{s}$ and $c4$ commits to a decommitment of $c2$ and $(a_{\mathsf{NM}}, b_{\mathsf{NM}})$ commit to a decommitment of $c4$ and $h(x_1) = h(x_2)$.

   $\widehat{C}$ proves the statement (a) by using a decommitment of $c3$ to $(v, d1)$ — decommitment of $c1$ to v — as the witness.
4) $\widehat{C}$ sends $(c1, c2, c3, c4, b_{\mathsf{NM}}, b_{\mathsf{ZAP}})$ as the second message to $\widehat{R}$ and keeps the decommitment $(v, d1)$ private.

- Reveal stage:
  On receiving $(v, d1)$ from $\widehat{C}$, $\widehat{R}$ accepts the decommitment if the ZAP proof is accepting and if $\mathsf{EOpen}_1(c1, v, d1) = 1$. Otherwise, it rejects.

We refer to the entire transcript of the interaction as the commitment c, and we say a commitment c is *accepting* if the ZAP proof contained in c is accepting. According to the reveal stage, the value $\mathsf{val}(c)$ of a commitment c is the value committed under $c1$ (contained in c) if c is accepting. Otherwise, $\mathsf{val}(c)$ is $\perp$.

Since the depth-and-size robust commiment ECom_1 is over-extractable, so is the scheme $\langle \hat{C}, \hat{R}\rangle$. Let $o\mathcal{E}_1$ be the extractor for ECom_1, then the following machine $\widehat{o\mathcal{E}}_{\mathsf{NM}}$ is an extractor for $\langle \hat{C}, \hat{R}\rangle$.

- Extraction - Extractor $\widehat{o\mathcal{E}}_{\mathsf{NM}}$:
  On receiving a commitment c and identity id, $\widehat{o\mathcal{E}}_{\mathsf{NM}}$ first verifies the ZAP proof and outputs $\perp$ if the proof is not accepting. Otherwise, it runs the extractor $o\mathcal{E}_1$ on $c1$ and outputs the extracted value v'.

Theorem 3. *$\langle \widehat{C}, \widehat{R}\rangle$ is a 2-round perfectly binding, commitment scheme for identities of length $t(n)$, satisfying concurrent non-malleability w.r.t. extraction by $\widehat{o\mathcal{E}}_{\mathsf{NM}}$, and concurrent non-malleability (w.r.t. commitment).*

Due to the lack of space, We refer the reader to the full version [34] for the proof of the theorem.

Acknowledge

We thank Stefano Tessaro for many helpful discussions on memory-hard functions and time-lock puzzles.

REFERENCES

[1] D. Dolev, C. Dwork, and M. Naor, "Nonmalleable cryptography," *SIAM Journal on Computing*, vol. 30, no. 2, pp. 391–437, 2000.

[2] R. Pass and A. Rosen, "Concurrent non-malleable commitments," in *46th FOCS*. IEEE Computer Society Press, Oct. 2005, pp. 563–572.

[3] B. Barak, "Constant-round coin-tossing with a man in the middle or realizing the shared random string model," in *43rd FOCS*. IEEE Computer Society Press, Nov. 2002, pp. 345–355.

[4] R. Pass and A. Rosen, "New and improved constructions of non-malleable cryptographic protocols," in *37th ACM STOC*, H. N. Gabow and R. Fagin, Eds. ACM Press, May 2005, pp. 533–542.

[5] H. Lin, R. Pass, and M. Venkitasubramaniam, "Concurrent non-malleable commitments from any one-way function," in *TCC 2008*, ser. LNCS, R. Canetti, Ed., vol. 4948. Springer, Heidelberg, Mar. 2008, pp. 571–588.

[6] H. Lin and R. Pass, "Non-malleability amplification," in *41st ACM STOC*, M. Mitzenmacher, Ed. ACM Press, May / Jun. 2009, pp. 189–198.

[7] O. Pandey, R. Pass, and V. Vaikuntanathan, "Adaptive one-way functions and applications," in *CRYPTO 2008*, ser. LNCS, D. Wagner, Ed., vol. 5157. Springer, Heidelberg, Aug. 2008, pp. 57–74.

[8] R. Pass and H. Wee, "Constant-round non-malleable commitments from sub-exponential one-way functions," in *EUROCRYPT 2010*, ser. LNCS, H. Gilbert, Ed., vol. 6110. Springer, Heidelberg, May 2010, pp. 638–655.

[9] H. Wee, "Black-box, round-efficient secure computation via non-malleability amplification," in *51st FOCS*. IEEE Computer Society Press, Oct. 2010, pp. 531–540.

[10] V. Goyal, "Constant round non-malleable protocols using one way functions," in *43rd ACM STOC*, L. Fortnow and S. P. Vadhan, Eds. ACM Press, Jun. 2011, pp. 695–704.

[11] H. Lin and R. Pass, "Constant-round non-malleable commitments from any one-way function," in *43rd ACM STOC*, L. Fortnow and S. P. Vadhan, Eds. ACM Press, Jun. 2011, pp. 705–714.

[12] V. Goyal, C.-K. Lee, R. Ostrovsky, and I. Visconti, "Constructing non-malleable commitments: A black-box approach," in *53rd FOCS*. IEEE Computer Society Press, Oct. 2012, pp. 51–60.

[13] M. Ciampi, R. Ostrovsky, L. Siniscalchi, and I. Visconti, "4-round concurrent non-malleable commitments from one-way functions," Cryptology ePrint Archive, Report 2016/621, 2016, http://eprint.iacr.org/2016/621.

[14] V. Goyal, S. Richelson, A. Rosen, and M. Vald, "An algebraic approach to non-malleability," in *55th FOCS*. IEEE Computer Society Press, Oct. 2014, pp. 41–50.

[15] V. Goyal, O. Pandey, and S. Richelson, "Textbook non-malleable commitments," in *48th ACM STOC*, D. Wichs and Y. Mansour, Eds. ACM Press, Jun. 2016, pp. 1128–1141.

[16] M. Ciampi, R. Ostrovsky, L. Siniscalchi, and I. Visconti, "Concurrent non-malleable commitments (and more) in 3 rounds," in *CRYPTO 2016, Part III*, ser. LNCS, M. Robshaw and J. Katz, Eds., vol. 9816. Springer, Heidelberg, Aug. 2016, pp. 270–299.

[17] R. Pass, "Unprovable security of perfect NIZK and non-interactive non-malleable commitments," in *TCC 2013*, ser. LNCS, A. Sahai, Ed., vol. 7785. Springer, Heidelberg, Mar. 2013, pp. 334–354.

[18] V. Goyal, D. Khurana, and A. Sahai, "Breaking the three round barrier for non-malleable commitments," in *57th FOCS*, I. Dinur, Ed. IEEE Computer Society Press, Oct. 2016, pp. 21–30.

[19] C. Dwork and M. Naor, "Zaps and their applications," in *41st FOCS*. IEEE Computer Society Press, Nov. 2000, pp. 283–293.

[20] R. L. Rivest, A. Shamir, and D. A. Wagner, "Time-lock puzzles and timed-release crypto," Massachusetts Institute of Technology, Cambridge, MA, USA, Tech. Rep., 1996.

[21] T. May, "Timed-release crypto," 1993.

[22] D. Boneh and M. Naor, "Timed commitments," in *CRYPTO 2000*, ser. LNCS, M. Bellare, Ed., vol. 1880. Springer, Heidelberg, Aug. 2000, pp. 236–254.

[23] N. Bitansky, S. Goldwasser, A. Jain, O. Paneth, V. Vaikuntanathan, and B. Waters, "Time-lock puzzles from randomized encodings," in *ITCS 2016*, M. Sudan, Ed. ACM, Jan. 2016, pp. 345–356.

[24] B. Barak, S. J. Ong, and S. Vadhan, "Derandomization in cryptography," Cryptology ePrint Archive, Report 2005/365, 2005, http://eprint.iacr.org/2005/365.

[25] J. Groth, R. Ostrovsky, and A. Sahai, "Non-interactive zaps and new techniques for NIZK," in *CRYPTO 2006*, ser. LNCS, C. Dwork, Ed., vol. 4117. Springer, Heidelberg, Aug. 2006, pp. 97–111.

[26] N. Bitansky and O. Paneth, "ZAPs and non-interactive witness indistinguishability from indistinguishability obfuscation," in *TCC 2015, Part II*, ser. LNCS, Y. Dodis and J. B. Nielsen, Eds., vol. 9015. Springer, Heidelberg, Mar. 2015, pp. 401–427.

[27] B. Barak and R. Pass, "On the possibility of one-message weak zero-knowledge," in *TCC 2004*, ser. LNCS, M. Naor, Ed., vol. 2951. Springer, Heidelberg, Feb. 2004, pp. 121–132.

[28] P. Rogaway, "Formalizing human ignorance," in *Progress in Cryptology - VIETCRYPT 06*, ser. LNCS, P. Q. Nguyen, Ed., vol. 4341. Springer, Heidelberg, Sep. 2006, pp. 211–228.

[29] C. Dwork and M. Naor, "Pricing via processing or combatting junk mail," in *CRYPTO'92*, ser. LNCS, E. F. Brickell, Ed., vol. 740. Springer, Heidelberg, Aug. 1993, pp. 139–147.

[30] M. Jakobsson and A. Juels, "Proofs of work and bread pudding protocols," in *Proceedings of the IFIP TC6/TC11 Joint Working Conference on Secure Information Networks: Communications and Multimedia Security*, ser. CMS '99. Deventer, The Netherlands, The Netherlands: Kluwer, B.V., 1999, pp. 258–272. [Online]. Available: http://dl.acm.org/citation.cfm?id=647800.757199

[31] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system, 2008," 2012. [Online]. Available: http://www.bitcoin.org/bitcoin.pdf

[32] D. Khurana and A. Sahai, "Birthday simulation and applications," 2017.

[33] ——, "Two-message non-malleable commitments from standard sub-exponential assumptions," Cryptology ePrint Archive, Report 2017/291, 2017, http://eprint.iacr.org/2017/291.

[34] 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.

[35] S. Kiyoshima, "Round-efficient black-box construction of composable multi-party computation," in *CRYPTO 2014, Part II*, ser. LNCS, J. A. Garay and R. Gennaro, Eds., vol. 8617. Springer, Heidelberg, Aug. 2014, pp. 351–368.