

One-Message Zero Knowledge and Non-Malleable Commitments *

Nir Bitansky^{1†} and Huijia Lin^{2‡}

¹ Tel Aviv University, email nirbitan@tau.ac.il

² University of Santa Barbra, email rachel.lin@cs.ucsb.edu

Abstract. We introduce a new notion of *one-message zero-knowledge (1ZK) arguments* that satisfy a *weak soundness* guarantee — the number of false statements that a polynomial-time non-uniform adversary can convince the verifier to accept is not much larger than the size of its non-uniform advice. The zero-knowledge guarantee is given by a simulator that runs in (mildly) super-polynomial time.

We construct such 1ZK arguments based on the notion of multi-collision-resistant *keyless* hash functions, recently introduced by Bitansky, Kalai, and Paneth (STOC 2018). Relying on the constructed 1ZK arguments, subexponentially-secure time-lock puzzles, and other standard assumptions, we construct *one-message fully-concurrent non-malleable commitments*. This is the first construction that is based on assumptions that do not already incorporate non-malleability, as well as the first based on (subexponentially) falsifiable assumptions.

1 Introduction

Zero-knowledge proofs [GMR89] are a cornerstone of modern cryptography. Their birth was enabled by introducing two new concepts to classical proofs — *interaction and randomness*. Indeed, both were shown [GO94] to be essential — for non-trivial languages, zero-knowledge proofs (or their computationally-sound counterparts known as *arguments*) require a randomized verifier that exchanges at least three messages with the prover. In particular, unlike classical proofs, zero-knowledge proofs cannot be transferred, published, nor stored.

One setting in which this barrier can be circumvented is when a trusted setup (such as a *common random string*) is available [BFM88]. In the absence of a trusted setup, a natural approach to the problem is to relax the requirements of zero-knowledge protocols. Along this vein, Dwork and Naor [DN07] showed that for witness-indistinguishable (WI) proofs, two messages suffice, and by now, we know how to achieve them with no interaction at all [BOV07, GOS12].

*The full version of this extended abstract can be found on Eprint [BL18].

†Member of the Check Point Institute of Information Security. Supported by the Alon Young Faculty Fellowship and by Len Blavatnik and the Blavatnik Family foundation.

‡Supported by NSF grants CNS-1528178, CNS-1514526, CNS-1652849 (CAREER), a Hellman Fellowship, the Defense Advanced Research Projects Agency (DARPA) and Army Research Office (ARO) under Contract No. W911NF-15-C-0236, and a subcontract No. 2017-002 through Galois. 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.

Pass [Pas03] considered a stronger notion — zero-knowledge with a super-polynomial simulator (SPS). Indeed, WI proofs stand at the extreme of this notion, as they admit an exponential-time simulator (that can find a witness for the underlying statement by brute force). In contrast, based on subexponential hardness assumptions, Pass constructed two-message arguments where the zero-knowledge simulator runs in subexponential, or even quasi-polynomial time (without violating the hardness of the underlying language). Such SPS zero-knowledge has proven instrumental for central applications such as concurrent computation [Pas03, PS04, BS05, MMY06, CLP16, GGJS12, GKP17, BGI⁺17, BGJ⁺17] and non-malleable commitments [KS17].

While Pass’ proofs break the three-message barrier, they still consist of two messages and do not enjoy the merits of completely non-interactive proofs. Following the introduction of non-interactive WI (NIWI) proofs, Barak and Pass [BP04] investigated the possibility that SPS zero-knowledge can also be made non-interactive (with no trusted setup). They observed that non-interactive proofs (or arguments) that satisfy the usual notion of soundness and have a T_{SPS} -time simulator are impossible to achieve against non-uniform adversaries, except for languages \mathcal{L} decidable in time T_{SPS} . Indeed, if the simulator cannot decide \mathcal{L} , there must *exist* proofs π for false statements $x \notin \mathcal{L}$, and a non-uniform prover can have such proofs hardwired in its code. Accordingly, Barak and Pass define a notion of SPS zero-knowledge protocols satisfying a *weak* notion of soundness that only holds against efficient *uniform* provers. They show how to construct such protocols based on keyless hash functions that are collision-resistant against subexponential uniform adversaries (or more general uniform sampling problems).

This Work: Weak Soundness against Non-Uniform Provers. We introduce a new notion of weak soundness for one-message zero-knowledge (1ZK) *that also captures non-uniform adversaries*.

The notion is inspired by the notion of multi-collision resistance for keyless hash functions, introduced recently in [BKP18]. Roughly speaking, it requires that an efficient non-uniform adversary *cannot do more than hardwire false statements with their accepting proofs in its code*. That is, any non-uniform adversary, with description of polynomial size S and arbitrary polynomial running time $T \gg S$, should not be able to find (i.e., output in one shot) more than $K(S)$ false statements $x \notin \mathcal{L}$ together with an accepting proof π , where K is some blowup function (for concreteness, the reader may think of $K(S) = S^2$ throughout this introduction). In other words, *false statements with their accepting proofs cannot be significantly compressed*.

The zero-knowledge requirement is the same SPS requirement as before — the simulator is allowed to be mildly super-polynomial (and in particular, cannot

decide the underlying language \mathcal{L}). We note that even with such weak soundness, the SPS relaxation is essential — languages \mathcal{L} that are hard on average cannot have an efficient simulator.³

1.1 Results and Discussion

We construct 1ZK arguments satisfying the new notion of weak soundness based on the notion of multi-collision resistance and generalizations thereof. Then, relying on such arguments, we construct one-message (concurrently) non-malleable commitments, which has been a long standing problem. We now elaborate on each of these results.

Constructing 1ZK Arguments We show how to construct 1ZK arguments from keyless hash functions that satisfy the notion of multi-collision resistance recently introduced in [BKP18]. Such a hash function $H : \{0, 1\}^\lambda \rightarrow \{0, 1\}^{\lambda/2}$ guarantees that no relatively-efficient adversary with non-uniform description of polynomial size S can find more than $K(S)$ collisions in the underlying function.⁴ Here, K is again a fixed polynomial (e.g., quadratic) and relatively-efficient means mildly superpolynomial-time (e.g. quasipolynomial or subexponential).

Theorem 1 (Informal). *Assuming multi-collision-resistant keyless hash functions, injective one-way functions, and non-interactive witness-indistinguishable proofs, all subexponentially-secure, there exist 1ZK arguments for NP with weak soundness and a subexponential-time simulator.*

As noted in [BKP18], while non-standard, multi-collision resistance is a falsifiable and relatively simple assumption. As candidates they suggest existing keyless hash functions such as SHA, or AES-based hashing, and point out directions for investigating additional candidates. We can, in fact, rely on a more general notion of *incompressible problems*, for which additional candidates may be found. At high-level, a (T, K, Δ) -incompressible problem is a collection $\mathcal{W} = \{\mathcal{W}_\lambda\}_\lambda$ of efficiently recognizable sets (one set for each security parameter λ) satisfying the following. On one hand, no T -time adversary with non-uniform description of polynomial size S can find more than $K(S)$ solutions $w \in \mathcal{W}_\lambda$.

³If there were such a simulator, then due to weak soundness, the simulator should fail to find accepting proofs for no-instances $\bar{x} \notin \mathcal{L}$ sampled from any efficiently samplable distribution. In contrast, for yes-instance $x \in \mathcal{L}$, it should succeed by the zero-knowledge guarantee. Thus, such a simulator would violate the average-case hardness of \mathcal{L} .

⁴To be exact, in [BKP18], they call this notion strong multi-collision resistance. They define (weak) multi-collision resistance as the problem of finding multiple inputs that all map to the same image. Throughout the introduction, we ignore this difference. In the body, we show that we can rely on either one, relying in addition on standard derandomization assumptions.

On the other hand, \mathcal{W}_λ is relatively *dense* in $\{0, 1\}^\lambda$, in the sense that a random $w \leftarrow \{0, 1\}^\lambda$ is in \mathcal{W}_λ with relatively high probability $\Delta = 2^{-o(\lambda)}$.⁵ For concreteness, the reader may think of $T = 2^{\lambda \cdot 0.1} \ll 2^{\lambda \cdot 0.99} = \Delta^{-1}$.

Theorem 2 (Informal). *Assuming (T, K, Δ) -incompressible problems, where $K \ll T \ll \Delta^{-1} \ll 2^{\lambda \cdot 0.99}$, and subexponentially-secure injective one-way functions and non-interactive witness-indistinguishable proofs, there exist 1ZK arguments for NP with (T, K) -weak soundness and a $\text{poly}(\Delta^{-1})$ -time simulator.*⁶

We also define and construct, under the same assumptions, a more general notion that we call φ -tuned 1ZK that admits a more flexible tradeoff between the level of soundness and simulation time, and will be useful when applying these arguments. We defer the details to the technical overview below.

One-Message Non-Malleable Commitments The question of the round complexity of non-malleable commitments [DDN03] has been long pursued. The past two decades have seen impressive progress [Bar02, PR05a, PR05b, LPV08a, LP09, PPV08, PW10, Wee10, Goy11, LP11, GLOV12, GRRV14, GPR16, COSV16, COSV17, Khu17], culminating in two recent constructions of *two-message* non-malleable commitments [KS17, LPS17] based on subexponential Decision-Diffie-Hellman or Quadratic Residuosity in the first, and subexponential *time-lock puzzles* [RSW00] in the second (which achieves also full concurrency).

Yet, one-message non-malleable commitments have remained somewhat elusive. So far, they have only been constructed starting from a non-falsifiable assumption that already incorporates non-malleability called *adaptive injective one-way functions*, against *uniform* adversaries [LPS17], or for a restricted class of algebraic mauling functions and entropic plaintexts [KY18]. Indeed, one-message non-malleable commitments would give rise to powerful features that cannot be achieved with interaction, such as the ability to publish them on public ledgers, transfer them from one hand to another, or store them for future use.

Relying on 1ZK arguments with weak soundness, we construct one-message fully-concurrent non-malleable commitments against *non-uniform* adversaries.

Theorem 3 (Informal). *Under the same assumptions as in Theorem 2 (or 1), as well as subexponential time-lock puzzles, there exist fully-concurrent one-*

⁵To get subexponential density, we need to multi-collision-resistant hash functions with polynomial, rather than linear, shrinkage. In [BKP18], it is shown how polynomial compression can be achieved from linear compression.

⁶Here (T, K) -weak soundness refers to the expected generalization of the weak soundness notion discussed above where the prover may run in time at most $\text{poly}(T)$, and T may be superpolynomial and the blowup function is K .

message non-malleable commitments against all efficient non-uniform adversaries.

We actually prove a more general theorem that transforms commitments satisfying a notion of *four-tag non-malleability* into full-fledged non-malleable commitments as stated in the above theorem. (More specifically, the former refers to non-malleability w.r.t. four tags, whereas full-fledged non-malleability can handle an exponential number of tags.) Such four-tag (or constant-tag) commitments are constructed in [LPS17] based on sub-exponentially secure time-lock puzzles and injective one-way functions. In addition, we present new candidate four-tag (or constant-tag) non-malleable commitments from a new assumption regarding *injective one-way functions that are amenable to hardness amplification*, which can replace time-lock puzzles in the above theorem. This yields new candidates from natural one-way functions such as discrete logarithms, RSA, or Rabin. See further details in the technical overview below.

On the Underlying Assumptions The assumptions that we rely on, most notably incompressible problems, are not standard. Nevertheless, we do find them simple and plausible. Bitansky, Kalai, and Paneth give evidence that multi-collision resistance may hold for existing cryptographic hash functions and in particular does not require any special algebraic structure — they show that this property is satisfied by random oracles, even in the auxiliary-input model [Unr07] (where the adversary may first store arbitrary polynomial information about the oracle).

We also note that all of our assumptions are *subexponentially-falsifiable* (i.e., falsifiable w.r.t. sub-exponential time adversaries). Here we note that Pass [Pas13] showed that non-malleable commitments in less than three messages cannot be shown secure based on black-box reductions to polynomially-falsifiable assumptions.

A more conservative view of our results would be that to rule out the existence of one-message non-malleable commitments, one must show that incompressible problems do not exist. That is, any efficiently recognizable, somewhat dense, set must have a non-trivial sampler (where by non-trivial we mean that it can output more samples than its non-uniform size). In particular, one would have to show that for any keyless hash function, it is possible to compress collisions. This would also constitute a strong (and non-contrived) separation between random oracles and any keyless hash function.

Using Weak Soundness Weak soundness is the best one could hope for when considering one-message zero-knowledge without trusted setup and non-uniform cheating provers, *but when is it useful?* Generally speaking, weak soundness could be leveraged in settings where a prover does not fully determine proven statements, namely, *statements have some non-trivial entropy*.

This gives some intuition on why weak soundness is useful in our application of non-malleable commitments. Roughly speaking, to maul a commitment c to a value v , the attacker is required to generate a new commitment c' to a related value v' , and prove that the new commitment is well-formed. As long as the attacker does not always produce a fixed commitment c' , or rather a commitment c' from some fixed polynomial-size set \mathcal{Z} , proven statements are sufficiently entropic and weak soundness kicks in. In contrast, mauling c into c' from such a set \mathcal{Z} would not constitute a meaningful attack — the distribution of the value v' in the commitment c' cannot depend on the committed value v in c , or a reduction that has the set \mathcal{Z} hardcoded could break the hiding of c . See more details in the technical overview below.

It is plausible that weak soundness will be found useful in other settings with entropic statements or in different man-in-the-middle attack models.

Robustness beyond Human Ignorance When considering the possibility of integrating non-interactive zero-knowledge in real-world systems, the need for a trusted common reference string may present a serious hurdle (certainly in decentralized applications whose essence is to avoid central trust). The system of Barak and Pass [BP04], when instantiated, say, with SHA256, already avoids the need for central trust and suggests a meaningful guarantee of *soundness in the face of human ignorance* (a term coined by Rogaway [Rog06]). Namely, as long as humanity fails to find collisions in SHA256, it will also fail to find accepting proofs for false statements. However, the moment even a single collision in SHA256 is found, the Barak and Pass system would completely lose soundness — it will be possible to easily prove *any false statement*.

Our system has a more robust guarantee — finding a few collisions only allows finding a few false statements with accepting proofs, and the mapping from collisions to false statements is deterministic and efficiently computable.

1.2 Technical Overview

We now give an overview of the main ideas and techniques behind our results.

Throughout this overview, it will be convenient to consider a slight variant of incompressible problems requiring that for any efficient adversary \mathcal{A} with a non-uniform description of polynomial size S , there exists a set \mathcal{Z} of size at most $K(S)$, such that \mathcal{A} cannot find solutions $w \in \mathcal{W} \setminus \mathcal{Z}$. In the body, we show that this variant is indeed equivalent to requiring that the adversary fails to find more than K solutions w . We consider a similar variant for the definition of weak soundness, where the adversary cannot output a false statement and accepting proof (x, π) , except for statements x from some size- K set.

One-Message Zero-Knowledge

The starting point for our construction is the Barak-Pass [BP04] construction against uniform provers. They follow the common [FLS99] paradigm in which the prover provides a WI proof that

“Either $x \in \mathcal{L}$ or the prover knows some trapdoor”.

The trapdoor should be such that it is too hard for an efficient prover to compute, but only mildly hard, so that a super-polynomial simulator can obtain it relatively fast in time $T_{\text{td}} \ll 2^{o(|x|)}$. The hardness of obtaining the trapdoor, and the soundness of the proof, guarantee the soundness of the argument, whereas as the WI property, along with the simulator’s ability to find the trapdoor, give rise to SPS simulation. To realize this idea, the prover sends a commitment c and proves that $x \in \mathcal{L}$ or c is a commitment to the trapdoor. The commitment is only mildly hard — the committed value could be extracted by brute force in time $T_{\text{com}} \ll T_{\text{td}}$, which does not suffice to find the trapdoor. Therefore, violating soundness requires violating the hardness of finding a trapdoor in T_{td} .

The question is *what could be the trapdoor*. Focusing on uniform provers, Barak and Pass rely on problems that are hard for uniform algorithms. For instance finding collisions of certain keyless hash functions is conjectured to be hard for uniform algorithms (or more generally, algorithms whose description is smaller than the function’s input), even in time $\text{poly}(T_{\text{com}})$. This of course miserably fails against non-uniform provers who could simply have such a trapdoor (e.g., a collision) hardwired in their code and use it to cheat.

Leveraging Incompressible Problems Recall that we are only interested in a weak notion of soundness — we wish to guarantee that there is only a small set of false statements for which the prover may give false proofs (where small is some polynomial $K(S)$ in the prover’s non-uniform description size S). A first natural idea is to simply replace the trapdoor problem with an incompressible problem \mathcal{W} (for instance, replace collision-resistance against uniform adversaries with multi-collision resistance against non-uniform ones).

This first attempt, however, fails. The problem is that any *single* solution in \mathcal{W} allows to efficiently generate accepting proofs for *all* statements x . Thus, a non-uniform attacker with one such hardwired solution, can convince the verifier of accepting any number of false statement, thereby violating the weak soundness requirement. The problem stems from the fact that in such a protocol, the concept of a useful trapdoor is completely detached from the proven statement x . We solve this by binding trapdoors and statements, so that, finding accepting proofs for different false statements requires finding different solutions in \mathcal{W} . Thus, an attacker who can only find a small set of solutions, can only generate proofs for a small number of corresponding false statements.

More specifically, we aim to achieve two goals. First, every trapdoor $w \in \mathcal{W}$ is associated with a specific statement $x = f(w)$ determined by some efficiently computable function f — this would ensure that the prover could only provide accepting proofs for false statements from a small set $\mathcal{X} = f(\mathcal{Z})$ determined by the small set \mathcal{Z} of trapdoors it may be able to find. Second, we would like to guarantee that for any $x \in \mathcal{L}$, the simulator would be able to reverse sample a trapdoor $w \in \mathcal{W}$ such that $x = f(w)$, and it should do so relatively fast.

We achieve the above combinatorial properties as follows. For instances x of size ℓ , we choose f to be a *two-source extractor* $2\text{Ext} : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^\ell$, where n is a parameter dictated by the quality of the extractor (in our actual construction $n = 4\ell$). We then choose our incompressible problem to be pairs of solutions $\mathcal{W} \times \mathcal{W} \subseteq \{0, 1\}^n \times \{0, 1\}^n$ for some underlying incompressible problem \mathcal{W} . It is easy to see that the product of incompressible problems is itself an incompressible problem, and so weak soundness is obtained according to the above reasoning. Furthermore, by choosing an appropriate extractor, we can guarantee that as long as \mathcal{W} has density $\Delta \geq 2^{-o(\ell)}$, for any $x \in \{0, 1\}^\ell$, it is possible to sample $(w, w') \in \mathcal{W}$ such that $2\text{Ext}(w, w') = x$ in time $O(\Delta^{-2})$, as required.

The above is satisfied by any extractor with the following two properties. First, it has an exponentially small error — for independent sources with min-entropy $n - o(\ell)$, the output is $2^{-\ell - \Omega(1)}$ -close to uniform. Second, it admits efficient reverse sampling — for any x , it is possible to efficiently sample from the uniform distribution on $U_n \times U'_n$ conditioned on $2\text{Ext}(U, U') = x$. These properties are both satisfied by the classical Hadamard extractor [CG88, Vaz85]. See further details in the full version of this paper.

To recap, the final proof (c, π) consists of a commitment c to a string of length $2n$, and a NIWI that

“Either $x \in \mathcal{L}$ or c is a commitment to $(w, w') \in \mathcal{W} \times \mathcal{W}$ such that $2\text{Ext}(w, w') = x$ ”.

Starting from a $(T_{\mathcal{W}}, K, \Delta)$ -incompressible problem, we choose a mildly-hard commitment so that it is extractable in time $T_{\text{com}} \ll T_{\mathcal{W}}$. The resulting system is then $(T_{\mathcal{W}}, K)$ -weakly-sound and has a Δ^{-2} -time simulator. In particular, for the discussed setting of parameters $K \ll T \ll \Delta^{-1} \ll 2^{\ell^{99}}$, we get a subexponential-time simulator.

φ -Tuned 1ZK We also consider a generalization of the 1ZK definition that admits a more flexible soundness vs. simulation-time tradeoff. Specifically, we parameterize our system by a projection function $\varphi(x)$ and obtain the following augmented guarantees:

- *Weaker Soundness*: we are only guaranteed that the prover produces accepting proofs for false statements x whose projection $\varphi(x)$ is taken from a small set \mathcal{Z} (but x itself is not restricted to any small set).
- *Faster Simulation*: simulation time is only subexponential in $|\varphi(x)|$ and not in $\ell = |x|$. Furthermore, fixing any projection y , there is a corresponding trapdoor state st_y that allows simulating any $x \in \varphi^{-1}(y)$ in *polynomial time*. A bit more formally, simulation for x can be split into a long preprocessing step S_{pre} , subexponential in $|\varphi(x)|$, that produces $\text{st}_{\varphi(x)}$, and a short postprocessing step S_{pos} that takes polynomial time given the trapdoor state $\text{st}_{\varphi(x)}$.

Note that the above is indeed a generalization of the previous notion when considering the identity as the projection φ . As we shall see later on, the flexibility of choosing φ differently, with the above tradeoff, will be useful in our application to non-malleable commitments. The construction of such φ -tuned 1ZK is identical to the construction described above only that we require that the trapdoor (w, w') fixes $\varphi(x)$ rather than x . See further details in the full version.

One-Message Non-Malleable Commitments

We now give an overview of how to use our 1ZK arguments to construct one-message non-malleable commitments. We adopt a standard formulation of non-malleable commitments where players have identities, and the commitment protocol depends on the identity of the committer, which is referred to as the *tag* of the interaction. Non-malleability [DDN03] ensures that 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*, as long as the tags of the left and right commitments are different. More formally, for any two values u and w , the values the man-in-the-middle commits to after receiving left commitments to u or w , along with the commitments it sees on the left, are indistinguishable. The notion of *concurrent non-malleability* [DDN03, PR05a] further requires that no attacker can “maul” a set of left commitments into a set of right commitments so that the joint distribution of right committed values depends on the left committed values.

The number γ of tags a scheme supports can be viewed as a quantitative measure of how non-malleable it is: A γ -tag non-malleable commitment gives a family of γ commitment schemes — each with a hardwired tag — that are “mutually non-malleable” to each other. Therefore, the fewer tags, the easier it is to construct a corresponding non-malleable commitment. Indeed, as shown by [LPS17], non-interactive non-malleable commitments for a constant number of tags can be constructed from subexponentially-secure injective one-way functions and time-lock puzzles [RSW00]. Full-fledged non-malleable commit-

ments, in contrast, have an exponential number of tags $\gamma = 2^\lambda$. Thus, the main challenge lies in increasing the number of tags from a constant to exponential.

Techniques for amplifying the number of tags have been explored in the literature [DDN03, LP11, KS17, LPS17]. They show that a non-malleable commitment scheme for γ tags can be transformed into one for $2^{\tilde{\Omega}(\gamma)}$ tags. Thus, starting from constant-tag non-malleable commitments, applying the transformation iteratively for $O(\log^* \lambda)$ times yields non-malleable commitments for exponentially many tags. However, all existing tag-amplification techniques crucially rely on interaction — even if the initial constant-tag non-malleable commitments are non-interactive, the transformation increases the message-complexity to at least two. For instance, the tag-amplification technique of Khurana and Sahai makes use of 2-message SPS zero-knowledge arguments. In this work, we show how to replace the 2-message SPS ZK arguments with our 1ZK arguments, which gives a non-interactive tag-amplification technique, and hence non-interactive non-malleable commitments.

Two-Message Tag-Amplification We start with reviewing the Khurana and Sahai (KS) 2-message tag-amplification technique, which transforms a non-interactive input scheme iNM for γ tags into a 2-message output scheme oNM for $\binom{\gamma}{\gamma/2} = 2^{\Omega(\gamma)}$ tags. Each tg' of oNM consists of a subset of $\gamma/2$ tags $\text{tg}' = (\text{tg}_1, \dots, \text{tg}_{\gamma/2})$ of iNM. To commit to a value v , oNM computes $\gamma/2$ commitments to v using iNM with respect to tags $\text{tg}_1, \dots, \text{tg}_{\gamma/2}$, followed by a 2-message SPS argument that all commitments are consistent. More precisely,

KS 2-message tag-amplification—oNM:

- The receiver R sends the first message π_1 of a 2-message SPS argument.
- To commit to v using $\text{tg}' = (\text{tg}_1, \dots, \text{tg}_{\gamma/2})$, the committer C generates $\{\text{nm}_j \leftarrow \text{iNM}(\text{tg}_j, v)\}_{j \in [\gamma/2]}$ and the second message π_2 of a 2-message SPS argument that all iNM commitments commit to the same value.

The committed value is defined to be the value committed in nm_1 .

To see that oNM is non-malleable, consider a man-in-the-middle receiving a left commitment using $\text{tg}' = (\text{tg}_1, \dots, \text{tg}_{\gamma/2})$ and giving a right commitment using $\tilde{\text{tg}}' = (\tilde{\text{tg}}_1, \dots, \tilde{\text{tg}}_{\gamma/2})$. If $\text{tg}' \neq \tilde{\text{tg}}'$, there must exist i^* , such that, $\tilde{\text{tg}}_{i^*} \neq \text{tg}_i$ for all i — the i^* th right iNM commitment uses a tag different from all left tags.

Then, they reduce the non-malleability of oNM to the non-malleability of iNM. To do so, they rely on the soundness of the 2-message SPS argument to argue that in *left-honest* man-in-the-middle executions, the attacker must send consistent iNM commitments $\{\widetilde{\text{nm}}_j\}$ on the right, or else it would fail in the SPS argument. (Here by *left-honest*, we mean the proofs on the left are honestly generated and not simulated.) Thus, to show that the right committed values do not change in two left-honest executions with different left committed values u or

w , it suffices to show that the value committed in any right iNM commitment — in particular, the i^* ’th one $\widetilde{\text{nm}}_{i^*}$ — does not change (in a distinguishable manner). To show this, they gradually simulate components in the left commitment in a sequence of hybrids, while maintaining that \tilde{v}_{i^*} committed in $\widetilde{\text{nm}}_{i^*}$ does not change throughout hybrids.

In the first hybrid, the left SPS argument (π_1, π_2) is simulated. To ensure that \tilde{v}_{i^*} does not change, they rely on *complexity leveraging* to make simulated proofs “harder to distinguish” than extracting from the commitment iNM; that is, the indistinguishability of SPS simulation holds even when \tilde{v}_{i^*} is extracted by brute force. Once the left SPS argument is simulated, the left iNM commitments are switched to committing to 0 in following hybrids. By the non-malleability of iNM and the fact that $\widetilde{\text{nm}}_{i^*}$ uses a tag $\widetilde{\text{tg}}_{i^*}$ different from all left tags, its committed value \tilde{v}_{i^*} does not change through these hybrids. Note that this requires the non-malleability of iNM to hold against T_{iNM} -time attackers for $T_{\text{iNM}} \gg T_{\text{SPS}}$. Using SPS ZK where simulation-time only depends on the underlying security parameter (and not the size of the instance), the above can be satisfied by appropriately choosing the relation between the iNM security parameter n and the SPS security parameter \bar{n} .

Non-Interactive Tag-Amplification To obtain non-interactive tag-amplification, a natural idea is replacing the 2-message SPS in the KS transformation with our 1ZK argument. However, two challenges arise:

- Challenge 1: Our 1ZK is only weakly sound. Thus, the man-in-the-middle attacker is able to generate an accepting 1ZK argument $\tilde{\pi}$ even when the right iNM commitments $\{\widetilde{\text{nm}}_j\}$ are inconsistent (i.e., committing to different values).
- Challenge 2: In our basic 1ZK, the simulation time is subexponential in the length of the statement $|x|$ (and the security parameter). This makes it difficult to guarantee that the simulator cannot break the underlying non-malleable commitment, i.e. $T_{\text{iNM}} \gg T_{\text{SPS}}$.

Specifically, the statement x concerns the consistency of $\gamma/2$ iNM commitments, and thus the simulation time is at least $T_{\text{SPS}} = 2^{(\gamma \times \ell_{\text{nm}}/2)^\epsilon}$, where $\ell_{\text{nm}} = \ell_{\text{nm}}(n)$ is the length of iNM commitments and could scale polynomially with the security parameter n of iNM. It could well be that $T_{\text{iNM}} \ll T_{\text{SPS}}$.

In a nutshell, to solve the first problem, we rely on the weak soundness of 1ZK to argue that whenever the right iNM commitments are not consistent (that is, the statement is false), the right commitments are taken from a small “apriori known” set, and their underlying values can be non-uniformly hardcoded into the

reduction. To solve the second problem, we make the security of iNM independent of the simulation time, by introducing an extra commitment under another scheme Com and using the φ -tuned version of 1ZK to reduce the simulation time to only depend on the length of commitments in Com, instead of commitments in iNM.

The Actual Tag-Amplification and Resulting Scheme oNM:

To commit to v using $\text{tg}' = (\text{tg}_1, \dots, \text{tg}_{\gamma/2})$, the committer C generates $c \leftarrow \text{Com}(v)$, $\{\text{nm}_j \leftarrow \text{iNM}(\text{tg}_j, v)\}_{j \in [\gamma/2]}$, and a 1ZK argument π showing that c and all iNM commitments commit to the same value. The 1ZK statement is given by $x = (c, \text{nm}_1, \dots, \text{nm}_{\gamma/2})$ and we consider its projection $\varphi(x) = c$ that only fixes the Com commitment c .

The committed value is defined to be the value committed in c .

Let us see how the above two problems are resolved.

Resolving Challenge 1: The weak soundness of φ -tuned 1ZK guarantees that for any attacker \mathcal{A} of polynomial size S , there is a set \mathcal{Z} consisting of a polynomial number $K(S)$ of Com commitments c (the so called projections) such that \mathcal{A} cannot prove a false statement x where the corresponding commitment c is not in \mathcal{Z} . This means that in left-honest man-in-the-middle executions, one of the following two cases occurs: Either the right Com commitment \tilde{c} and the iNM commitments are all consistent, or the commitment \tilde{c} belongs to \mathcal{Z} . In the latter case, the right committed value must belong to the polynomial-sized set $\{\tilde{v} : \tilde{v} \text{ is the value in } \tilde{c} \in \mathcal{Z}\}$, which can be hardwired non-uniformly into the reduction. In the first case, showing the indistinguishability of the right committed values again reduces to showing that of \tilde{v}_{i^*} committed in $\tilde{\text{nm}}_{i^*}$.

Resolving Challenge 2: Recall that φ -tuned 1ZK enjoys a simulation speedup. Specifically, simulation consists of i) a $2^{|c|^\delta}$ -time preprocessing phase that depends only on the projection c and computes a trapdoor state $\text{st} \leftarrow S_{\text{pre}}(c)$, and ii) a polynomial $\text{poly}(|x|, \bar{n})$ -time postprocessing phase that generates the simulated proof $\hat{\pi} \leftarrow S_{\text{pos}}(x, \text{st})$. With this speed-up, let us examine again the sequence of hybrids where the left Com and iNM commitments are gradually switched to committing to 0, while the 1ZK argument on the left is simulated. We need to ensure that \tilde{v}_{i^*} does not change.

To change the Com commitment, we require that its hiding holds even in the presence of 1ZK simulation and (brute-force) extraction from \tilde{v}_{i^*} :

$$T_{\text{Com}} \gg T_{\text{SPS}} = 2^{|c|^\delta} + \text{poly}(|x|, \bar{n}) \quad \text{and} \quad T_{\text{Com}} \gg T_{\text{iNM.E}}$$

The latter can be satisfied by setting the security parameter \bar{n} of Com to be sufficiently larger than the security parameter n of iNM. The former is more

subtle as it requires Com to be at least $2^{|c|^\delta}$ -secure, where $|c|$ is the length of Com commitments. Such a commitment scheme for strings of length ℓ , can be instantiated by the classical Blum-Micali bit commitment scheme [BM84] (recall that a commitment to b is $f(r), \text{hc}(r) \oplus b$, where hc is a hardcore bit of an injective one-way function f), instantiated with any 2^{k^ρ} -hard injective one-way function, and sufficiently large security parameter $k > \Omega(\ell^{\delta/\rho-\delta})$.

Next, when changing the left iNM commitments, we can circumvent the requirement that $T_{\text{iNM}} \gg T_{\text{SPS}}$ by leveraging the efficient postprocessing of 1ZK simulation. Recall that given a trapdoor state $\text{st} \leftarrow S_{\text{pre}}(c)$ that depends only on the projection c , simulating the proof $\hat{\pi} \leftarrow S_{\text{pos}}(x, \text{st})$ takes only polynomial time. When changing the values committed in left iNM commitments, the left Com commitment c is independent — it is by now a commitment to 0. If in two neighboring hybrids, the value \tilde{v}_{i^*} on the right changes, there must exist a commitment c (committing to 0) such that conditioned on c occurring in the hybrids the value \tilde{v}_{i^*} still changes. With respect to this specific c , 1ZK simulation can now be done in polynomial time, given as non-uniform advice the preprocessed state $\text{st} \leftarrow S_{\text{pre}}(c)$ depending on c . This suffices for the security reduction, as now, the non-malleability of iNM is detached from the 1ZK simulation time.

A Subtle Issue The above description captures the main idea, but misses a subtle issue. Roughly speaking, in order to apply our tag-amplification iteratively, across different iterations, we need to increase the level of security of the Com schemes used in each iteration. In particular, the security parameter k for the one-way functions underlying Com needs to grow polynomially in each iteration. If we start with $k > \ell^{\delta/(\rho-\delta)} = \ell^{\Omega(1)}$, after a super-constant number of iterations (out of the $\log^* n$ iterations needed), k would grow to be super-polynomial in ℓ .

To avoid this, we modify the scheme oNM to have a separate 1ZK argument for each bit commitment c_j (committing to a bit v_j of the committed value), proving that all iNM commitments are consistent with it, in the sense that, the j 'th bit of their committed strings equals to the bit committed in c_j . By doing so, c_j only needs to be $2^{|c_j|^\delta}$ -secure, independent of the length ℓ of committed values. Thus, we no longer need to set k to be $k = \ell^{\Omega(1)}$, but instead to $k = \ell^{o(1)}$. Though k still increases through $O(\log^* n)$ iterations, it is always kept polynomial in ℓ . See Section for a formal description of the final transformation.

Achieving Concurrency Applying our non-interactive tag amplification to the 4-tag non-malleable commitments of [LPS17] gives a full-fledged non-interactive non-malleable commitment, which however, is only stand-alone (i.e., one-one) but not concurrently non-malleable. This is because the basic commitments of [LPS17] are not concurrently non-malleable.

To obtain concurrent non-malleability, we give another transformation from non-malleable commitments in a restricted concurrent setting, called *same-tag concurrency* into *fully concurrent* ones. Roughly speaking, in the same-tag concurrent setting, we require non-malleability to hold with respect to attackers who always use the *same tag* in all commitments on the right. We observe that the 4-tag commitments of [LPS17] actually are same-tag non-malleable, and our tag amplification preserves this property. Therefore, by applying the same-tag to full-concurrency transformation after tag amplification, we obtain concurrent non-malleability.

Our transformation is inspired by the *2-round* non-malleability strengthening transformation in [LPS17], but works in one message and is simpler and more modular; in particular, the transformation of [LPS17] relies directly on time-lock puzzles, whereas we work with any non-malleable commitment satisfying the intermediate notion of same-tag non-malleability.

At a high level, starting from a same-tag non-malleable input scheme iNM, our transformation follows the Naor-Yung paradigm for constructing CCA encryption, producing an output scheme oNM as follows. oNM fixes two arbitrary tags tg_0^* , tg_1^* of iNM for special use, and commitments are computed using to other tags $\text{tg} \neq \text{tg}_0^*, \text{tg}_1^*$.

The Same-Tag to Fully-Concurrent Transformation and Resulting Scheme oNM (Simplified):

- On input v and tag tg , the committer C commits to v using iNM with the two special tags:

$$\text{nm}_0 \leftarrow \text{iNM}(\text{tg}_0^*, v) \quad \text{nm}_1 \leftarrow \text{iNM}(\text{tg}_1^*, v) ,$$

and proves that both iNM commitments commit to the same value v . The proof is computed using a *simulation-sound* variant of our 1ZK argument relative to the tag tg .

To argue the concurrent non-malleability of oNM, it suffices to argue one-many non-malleability [LPV08b] (that is, the man-in-the-middle receives a single commitment on the left and gives many commitments on the right.)

The two commitments of iNM using special tags tg_0^* and tg_1^* are the counterparts of the as two public-key encryptions in the Naor-Yung paradigm, and the proof of non-malleability follows similarly to the proof of CCA security. The simulation soundness of 1ZK ensures that the man-in-the-middle attacker can only send consistent $\widetilde{\text{nm}}_{0,j}$ and $\widetilde{\text{nm}}_{1,j}$ in every right commitment j , *even when the left 1ZK argument is simulated*. Therefore, as the left commitment nm_0 is simulated (by committing to 0), one can argue that the right committed values

do not change by showing that values in $\{\widetilde{\text{nm}}_{1,j}\}$ do not change. Similarly, as the left commitment nm_1 is simulated, one can switch to showing that values in $\{\widetilde{\text{nm}}_{0,j}\}$ do not change. Here same-tag non-malleability is essential for arguing that the joint distribution of all right committed values does not change (in a distinguishable manner).

To achieve simulation-soundness, we open the construction of our 1ZK arguments. Recall that these arguments rely on a basic commitment scheme, a NIWI, and an incompressible language. We show that by replacing the basic commitment scheme with a non-malleable one (such as the input scheme iNM), our 1ZK arguments become simulation-sound. For this approach to work, we additionally need “mutual non-malleability” between the commitment in our simulation-sound 1ZK arguments and the iNM commitments using $\text{tg}_0^*, \text{tg}_1^*$. That is, i) simulating the 1ZK argument on the left does not change the values that the attacker commits to in iNM commitments $\{\widetilde{\text{nm}}_{0,j}, \widetilde{\text{nm}}_{1,j}\}$ on the right, and ii) changing the values committed in the iNM commitments on the left does not allow the attacker to break (weak) soundness on the right. Such “mutual non-malleability” is achieved again relying on the same-tag non-malleability of iNM and the fact that the iNM commitments use two special tags $\text{tg}_0^*, \text{tg}_1^*$ different from the tags we use for iNM commitments in 1ZK arguments.

The above discussion is overly-simplified. Indeed, this transformation also has to deal with the challenges presented before in the tag-amplification transformation. They are dealt with using similar techniques. See Section ?? for details.

New Candidate Constant-Tag Non-Malleable Commitments As explained above, our transformations start from non-malleable commitments for a constant number of tags, which were previously known based on time-lock puzzles [LPS17]. We also provide new candidate constant-tag non-malleable commitments, based on a new assumption on hardness amplification of (injective) one-way functions.

Known results on hardness amplification have shown ways of strengthening weak one-way functions to strong ones, via direct product lemmas or XOR lemmas. However, these results have a common weakness — hardness does not amplify beyond negligible. Concretely, starting from a function f that is δ -hard against T -time attackers, the k -fold combined function f' is $(\text{poly}(\frac{T'}{T}) + (1 - \delta)^k)$ -hard for $(T' \ll T)$ -time attackers. As the number k of copies increases, the hardness approaches the limit of $\text{poly}(\frac{T'}{T})$.

The work of [DJMW12] showed that this limit is inherent for certain contrived one-way functions, but there is no evidence that this limit should bound natural one-way functions, such as, discrete logarithm, RSA, or Rabin. We put forward the notion of *amplifiable one-way functions and hardcore bits*: Roughly

speaking, we say that a one-way function f is amplifiable, if there is a way to combine (e.g. XOR), say ℓ , hardcore bits, corresponding to ℓ independent images $f(x_1), \dots, f(x_\ell)$, so that the combined bit is 2^{ℓ^ε} -unpredictable; that is, the level of unpredictability increases at least subexponentially as more hardcore bits are combined and beyond the limit $\text{poly}(\frac{T'}{T})$.

We show that amplifiable one-way functions are useful for constructing non-malleable commitments. They essentially allow us to construct commitment schemes $(\text{Com}, \text{Com}')$, such that, Com is “harder” than Com' *in the time axis* — Com remains hiding in time needed for extracting from Com' , whereas Com' is “harder” than Com *in the distinguishing axis* — the maximum distinguishing advantage of Com' is smaller than the probability that one can guess a decommitment of Com . As shown in [LPS17], commitments that are harder than each other under different measures are essentially non-malleable. This yields new candidate constant-tag non-malleable commitments with one-way functions that are believed to have amenable hardness amplification behavior, such as, discrete logarithm, RSA, or Rabin.

1.3 Concurrent Work

In concurrent and independent work, Holmgren and Lombardi [HL18] study *one-way product functions*, which are related to our notion of amplifiable one-way functions. Their notion requires that ℓ independent images $f(x_1), \dots, f(x_\ell)$ cannot be inverted simultaneously by efficient algorithms, except with exponentially small probability in the input size. They show how to use such functions in different parameter regimes to obtain several applications ranging from collision-resistant hashing to correlation intractability (when combined with indistinguishability obfuscation). (The exact inversion probability and choice of ℓ depends on the specific application. Most of their applications are in the regime where ℓ is small, e.g. constant, and the inversion probability is at most $2^{-n-\omega(\log n)}$.)

While their one-way product functions and our amplifiable one-way functions are very related, there are some notable differences. For once, we make a stronger requirement than the hardness of inversion, namely, the hardness of predicting a combined hardcore bit. (Note that this gap cannot be bridged by the classic Goldreich-Levin theorem, where the adversary’s distinguishing advantage ε translates to a reduction running in time at least $\text{poly}(\varepsilon^{-1})$ to invert the underlying function.) On the other hand, since we allow ℓ to grow polynomially, our notion could potentially hold for one-way functions where a single copy is only mildly hard to invert, whereas for many of their applications (like collision-resistant hashing), ℓ is required to be small, and accordingly the one-way function has to be hard to invert except with exponentially small probability.

Organization. The rest of this extended abstract is organized as follows. In Section 2, we give some of the basic definitions used in the paper, including the definition of non-malleable commitments that we achieve. In Section 3, we define the notion of incompressible problems. In Section 4, we define and construct our new notion of one-message zero knowledge. Our constructions of non-malleable commitments, as well as all proofs, can be found in the full version of the paper.

2 Preliminaries

We rely on the following standard computational concepts:

- We model algorithms as (possibly probabilistic and possibly interactive) Turing machines. A *non-uniform* algorithm M is given by a family of algorithms $M = \{M_\lambda\}_{\lambda \in \mathbb{N}}$, where λ is a security parameter, and each M_λ corresponds to an input size $n(\lambda)$ and has description-size related to λ .
 - M is T -time, if for every $\lambda \in \mathbb{N}$, M_λ performs at most $T(\lambda)$ steps.
 - M is S -size if for every $\lambda \in \mathbb{N}$, M_λ has description size at most $S(\lambda)$.
 Throughout, we assume w.l.o.g. that the description-size of a non-uniform algorithm is bounded by its running time $S(\lambda) \leq T(\lambda)$ for all λ . A *uniform* algorithm M is a special-case of a non-uniform algorithm where for all $\lambda \in \mathbb{N}$, $M_\lambda = M$ is a single, constant-size, algorithm. A PPT is a probabilistic polynomial-time uniform algorithm. By default, algorithms in cryptographic schemes are PPTs.
- We model T -time adversaries as arbitrary non-uniform T -time algorithms $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$. Efficient adversaries have polynomial time. Throughout this work, we consider polynomial-size adversaries, and assume w.l.o.g. that their sizes are at least λ , i.e., $|\mathcal{A}_\lambda| \geq \lambda$ (via padding).
- We say that a function $f : \mathbb{N} \rightarrow \mathbb{R}$ is negligible if for all constants $c > 0$, there exists $N \in \mathbb{N}$ such that for all $n > N$, $f(n) < n^{-c}$. We sometimes denote negligible functions by negl .
- We say that a function $f : \mathbb{N} \rightarrow \mathbb{R}$ is noticeable if there exists a constant $c > 0$ and $N \in \mathbb{N}$ such that for all $n > N$, $f(n) \geq n^{-c}$.
- For two functions $T(\lambda), T'(\lambda)$, we write that $T' \ll T$ if $T' = T^{o(1)}$, when $\lambda \rightarrow \infty$.

In this paper, we will sometimes consider security of primitives against general $\text{poly}(T)$ -time adversaries, as illustrated in the definition of T -indistinguishability below.

Definition 1 ((T, μ) -Indistinguishability). Let $\mathcal{X}^{(b)} = \{X_\lambda^{(b)}\}_{\lambda \in \mathbb{N}}$ for $b \in \{0, 1\}$ be two ensembles of random variables indexed by $\lambda \in \mathbb{N}$. We say that $\mathcal{X}^{(0)}$

and $\mathcal{X}^{(1)}$ are (T, μ) -indistinguishable for functions T, μ , if for all $\text{poly}(T)$ -time distinguishers \mathcal{D} , and all large enough λ ,

$$\left| \Pr[\mathcal{D}(X_\lambda^{(0)}) = 1] - \Pr[\mathcal{D}(X_\lambda^{(1)}) = 1] \right| \leq \mu(\lambda)^{\Omega(1)}.$$

We say that $\mathcal{X}^{(0)}$ and $\mathcal{X}^{(1)}$ are T -indistinguishable if it is (T, μ) -indistinguishable for some negligible function μ . We say that they are computational indistinguishable if they are T -indistinguishable for every polynomial T .

We denote the above notions of indistinguishability by $\mathcal{X}^{(0)} \approx_{T, \mu} \mathcal{X}^{(1)}$, $\mathcal{X}^{(0)} \approx_T \mathcal{X}^{(1)}$, and $\mathcal{X}^{(0)} \approx \mathcal{X}^{(1)}$, respectively.

2.1 Commitments

We define non-interactive commitments.

Definition 2 (Commitment Scheme). A non-interactive commitment scheme consists of two polynomial-time algorithms $(\text{Com}, \text{Open})$, with the following syntax:

- $(c, d) \leftarrow \text{Com}(v, 1^\lambda)$: Given 1^λ and $v \in \{0, 1\}^*$, Com samples a commitment c and a decommitment string d .
- $b = \text{Open}(c, v, d)$: Given a commitment c , value v , and decommitment string d , Open outputs a bit b , where $b = 1$ indicates acceptance. We say that a commitment c is valid, if there exists a decommitment (v, d) , such that $\text{Open}(c, v, d) = 1$.

We make the following requirements:

Correctness: For any $\lambda \in \mathbb{N}$, $v \in \{0, 1\}^*$,

$$\Pr[\text{Open}(c, v, d) = 1 : (c, d) \leftarrow \text{Com}(v, 1^\lambda)] = 1 .$$

Binding: For any string c , values v, v' , and decommitment strings d, d' ,

$$\text{if } \text{Open}(c, v, d) = \text{Open}(c, v', d') = 1 \text{ then } v = v' .$$

T -hiding: For any polynomial $n = n(\lambda)$,

$$\left\{ \text{Com}(v, 1^\lambda) \right\}_{\lambda \in \mathbb{N}, v, v' \in \{0, 1\}^{n \times 2}} \approx_T \left\{ \text{Com}(v', 1^\lambda) \right\}_{\lambda \in \mathbb{N}, v, v' \in \{0, 1\}^{n \times 2}} .$$

Tag-based Commitments We consider “tag-based” commitment schemes.

Definition 3 (Tag-based commitment scheme). A commitment scheme $(\text{Com}, \text{Open})$ is a tag-based scheme with t -bit tags if, in addition to 1^λ , Com also receive a “tag” (a.k.a. identity) $\text{tg} \in \{0, 1\}^{t(\lambda)}$ as input, $c \leftarrow \text{Com}(\text{tg}, v, 1^\lambda)$. We assume w.l.o.g that commitments generated by Com contains the tag used for generating them. For any sequence of fixed tags $\text{tg} = \{\text{tg}_\lambda\}_\lambda$, the corresponding $(\text{Com}_{\text{tg}}, \text{Open}_{\text{tg}}) = \left\{ (\text{Com}_{\text{tg}_\lambda}, \text{Open}_{\text{tg}_\lambda}) \right\}_\lambda$ satisfy correctness, binding, and hiding as defined for plain commitment schemes. By default, a tag-based commitment scheme has t -bit tags for some polynomial t .

2.2 Non-Malleable Commitments

The Man-in-the-Middle (MIM) Execution: Let $\text{NM} = (\text{Com}, \text{Open})$ be a commitment scheme for t -bit tags, and $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$ an arbitrary non-uniform adversary. For a security parameter λ , and $m = m(\lambda)$, \mathcal{A}_λ on input 1^λ , receives m commitments from an honest committer C to values $v_1, \dots, v_m \in \{0, 1\}^\lambda$, and sends m commitments to R to values $\tilde{v}_1, \dots, \tilde{v}_m \in \{0, 1\}^\lambda$. The commitments received by the adversary are called *the left commitments* and those sent are called *the right commitments*. The left and right commitments use $t = t(\lambda)$ -bit tags $\text{tg}_1, \text{tg}_2, \dots, \text{tg}_m$ and $\tilde{\text{tg}}_1, \tilde{\text{tg}}_2, \dots, \tilde{\text{tg}}_m$ chosen adaptively by \mathcal{A}_λ for each commitment. The values \tilde{v}_j in the j 'th right commitment \tilde{c}_j is defined as

$$\tilde{v}_j = \begin{cases} \perp & \text{if } \exists i, \text{tg}_i = \tilde{\text{tg}}_j \\ \text{val}(\tilde{c}_j) & \text{otherwise} \end{cases} .$$

That is, \tilde{v}_j is either the unique committed value if the commitment \tilde{c}_j is valid and uses a tag different from all left tags, or \perp otherwise. (Recall that by binding, \tilde{v}_j is uniquely defined whenever \tilde{c}_j is valid.)

We denote by $\text{MIM}_{\text{NM}}^{\mathcal{A}}(v_1, \dots, v_m, 1^\lambda)$ the above described man-in-the-middle experiment.

Non-Malleability with respect to Commitment Let $\text{mim}_{\text{NM}}^{\mathcal{A}}(v_1, \dots, v_m, 1^\lambda)$ denote the random variable that describes the view of \mathcal{A}_λ (consisting of all left commitments) and the values $\tilde{v}_1, \dots, \tilde{v}_m$ it commits to on the right in the above man-in-the-middle experiment.

Definition 4 (Non-Malleability). A commitment scheme NM for t -bit tags is concurrent T -non-malleable if for any non-uniform $\text{poly}(T)$ -time adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$ and for every polynomial $m = m(\lambda)$, it holds that:

$$\left\{ \text{mim}_{\text{NM}}^{\mathcal{A}}(v_1, \dots, v_m, 1^\lambda) \right\}_{\lambda \in \mathbb{N}, v_1, \dots, v_m, v'_1, \dots, v'_m \in \{0, 1\}^\lambda} \\ \approx_c \left\{ \text{mim}_{\text{NM}}^{\mathcal{A}}(v'_1, \dots, v'_m, 1^\lambda) \right\}_{\lambda \in \mathbb{N}, v_1, \dots, v_m, v'_1, \dots, v'_m \in \{0, 1\}^\lambda} .$$

2.3 Non-Interactive Witness-Indistinguishable Proofs

We define non-interactive witness-indistinguishable proofs (NIWIs).

Definition 5 (NIWI). A non-interactive witness-indistinguishable proof system (P, V) for an NP relation $\mathcal{R}(x, w)$ consists of two polynomial-time algorithms:

- $\pi \leftarrow P(x, w, 1^\lambda)$: Given an instance x , witness w , and security parameter 1^λ , P produces a proof π .
- $b = V(x, \pi)$: Given a proof π for instance x , V outputs a bit b , where $b = 1$ indicates acceptance.

We make the following requirements:

Completeness: For every $\lambda \in \mathbb{N}$, $(x, w) \in \mathcal{R}$,

$$\Pr_{\mathcal{P}}[V(x, \pi) = 1 : \pi \leftarrow P(x, w, 1^\lambda)] = 1 .$$

Soundness: For every $x \notin \mathcal{L}(\mathcal{R})$ and $\pi \in \{0, 1\}^*$:

$$V(x, \pi) \neq 1 .$$

T-Witness-Indistinguishability: For any sequence

$$\mathcal{I} = \left\{ (\lambda, x, w_0, w_1) : \begin{array}{l} \lambda \in \mathbb{N}, x, w_0, w_1 \in \{0, 1\}^{\text{poly}(\lambda)}, \\ (x, w_0), (x, w_1) \in \mathcal{R} \end{array} \right\}$$

It holds that

$$\left\{ \pi_0 \leftarrow P(x, w_0, 1^\lambda) \right\}_{(\lambda, x, w_0, w_1) \in \mathcal{I}} \approx_T \left\{ \pi_1 \leftarrow P(x, w_1, 1^\lambda) \right\}_{(\lambda, x, w_0, w_1) \in \mathcal{I}} .$$

Barak, Ong, and Vadhan [BOV07] constructed NIWIs based on NIZK and the worst-case assumption that there exists a problem solvable in deterministic time $2^{O(n)}$ with non-deterministic circuit complexity $2^{\Omega(n)}$ (or more generally the existence of hitting set generators that fool non-deterministic distinguishers). Groth, Ostrovsky, and Sahai [GOS12] then constructed NIWIs based on standard assumptions on bilinear maps such as the Decision Linear Assumption, the Symmetric External Diffie Hellman assumption, or the Subgroup Decision Assumption. Bitansky and Paneth [BP15] constructed NIWIs from indistinguishability obfuscation and one-way permutations.

2.4 Two-Source Extractors

We rely on the standard notion of two-source extractors.

Definition 6 (Two-Source Extractor). *A polynomial-time computable function $2\text{Ext} : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^m$ is a (k_1, k_2, ε) -two-source extractor, if for any two independent sources X_1, X_2 with min-entropies at least k_1 and k_2 , respectively, it holds that*

$$\|2\text{Ext}(X_1, X_2) - U_m\|_1 \leq \varepsilon ,$$

where U_m is the uniform distribution over $\{0, 1\}^m$.

We also require *efficient reverse sampling*, which says that given any y in the image of the extractor 2Ext we can efficiently sample uniformly random and independent sources X_1 and X_2 conditioned on $2\text{Ext}(X_1, X_2) = y$.

Definition 7 (Efficient Reverse Sampling). *A function $2\text{Ext} : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^m$ is efficiently reverse-samplable if there exists a PPT that given $y \in \text{Image}(2\text{Ext})$ outputs a uniformly random pair x_1, x_2 such that $2\text{Ext}(x_1, x_2) = y$.*

Two source extractors with efficient reverse sampling and an exponentially small error are known based on the Hadamard code over an appropriate field.

3 Incompressible Problems

Following [BKP18], we consider a notion of incompressible problems. Here every security parameter λ , defines a search problem \mathcal{W}_λ with superpolynomially many solutions $w \in \mathcal{W}_\lambda$. Since the problem is fixed, a non-uniform adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}$ may always have hardwired solutions $w \in \mathcal{W}_\lambda$ in its code. We require, however, that it is *impossible to significantly compress solutions* — an adversary with description size at most S and bounded running time T , larger than S , should fail to produce more than S solutions (or $K(S)$ solutions for some polynomial blowup function $K(\cdot)$).

Definition 8 (Incompressible Problem). *An incompressible problem \mathcal{W} is associated with a polynomial-time verifier algorithm \mathcal{V} and a collection of sets $\{\mathcal{W}_\lambda\}_\lambda$, such that $\mathcal{W}_\lambda \subseteq \{0, 1\}^\ell$ for some polynomial $\ell = \ell(\lambda)$, and for any $w \in \{0, 1\}^\ell$, $\mathcal{V}(w) = 1$ if and only if $w \in \mathcal{W}_\lambda$. For any function $T = T(\lambda) \geq \lambda$ and polynomial K , we make the following incompressibility requirement.*

(T, K)-Incompressibility: for any non-uniform poly(T)-time, polynomial-size, probabilistic adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}$, there is a negligible function μ , such that for any $\lambda \in \mathbb{N}$, letting $K = K(|\mathcal{A}_\lambda|)$,

$$\Pr_{\mathcal{A}_\lambda} \left[\begin{array}{l} W \subseteq \mathcal{W}_\lambda \\ |W| \geq K \end{array} \middle| W \leftarrow \mathcal{A}_\lambda \right] \leq \mu(\lambda) .$$

We say that \mathcal{W} has density $\Delta = \Delta(\lambda)$, if for every sufficiently large $\lambda \in \mathbb{N}$, letting $\ell = \ell(\lambda)$, it holds that $|\mathcal{W}_\lambda| \geq \Delta 2^\ell$. We say that \mathcal{W} has subexponential density if it has density $\Delta = 2^{-\ell^\varepsilon}$ for some constant ε .

Remark 1 (Parameters). The parameters T, K, Δ that we consider will always be such that

$$K \leq T \ll K \Delta^{-1} .$$

Indeed, when $T < K$ the requirement trivializes and when $T \geq \text{poly}(K \Delta^{-1})$ the requirement becomes impossible.

Candidates Candidates for incompressible problems were introduced in [BKP18]. The problems addressed there come from *keyless* (shrinking) hash functions where collisions are incompressible in some sense. We can rely on more general incompressible problems, which may give rise to additional candidates. The problems considered in [BKP18] and a discussion of additional possible candidates can be found in the full version of the paper.

4 One-Message Zero Knowledge

In this section, we give a new definition of a one-message zero-knowledge (1ZK) system, and construct such a system based on incompressible problems. The definition relaxes both the zero knowledge requirement and soundness. Here the zero knowledge definition is the standard super-polynomial simulation (SPS) definition [Pas03]. The soundness definition is new and roughly says that a (relatively) efficient adversary of description size S shouldn't be able to sample more than S (or $K(S)$ for some polynomial blowup K) false statements x together with an accepting proof π . As discussed in the introduction, both of these relaxations are necessary.

We proceed to the formal definition.

Definition 9 (1ZK). A one-message zero-knowledge argument system (P, V) for an NP relation $\mathcal{R}(x, w)$ consists of two polynomial-time algorithms:

- $\pi \leftarrow P(x, w, 1^\lambda)$: Given an instance x , witness w , and security parameter 1^λ , P produces a proof π .

- $b = V(x, \pi, 1^\lambda)$: Given a proof π for instance x , V outputs a bit b , where $b = 1$ indicates acceptance.

The system is parameterized by functions $T_D(\cdot), T_S(\cdot), T_P(\cdot), K(\cdot)$.

We make the following requirements:

Completeness: For every $\lambda \in \mathbb{N}, (x, w) \in \mathcal{R}$,

$$\Pr_{\mathcal{P}}[V(x, \pi, 1^\lambda) = 1 : \pi \leftarrow P(x, w, 1^\lambda)] = 1 .$$

(T_D, T_S) -Zero-Knowledge: There exists a uniform $\text{poly}(T_S)$ -time simulator S , such that,

$$\left\{ \pi \leftarrow P(x, w, 1^\lambda) \right\}_{\substack{(x,w) \in \mathcal{R} \\ \lambda \in \mathbb{N}}} \approx_{T_D} \left\{ \hat{\pi} \leftarrow S(x, 1^\lambda) \right\}_{\substack{(x,w) \in \mathcal{R} \\ \lambda \in \mathbb{N}}} .$$

(T_P, K) -Weak-Soundness: For any non-uniform $\text{poly}(T_P)$ -time, polynomial-size, probabilistic adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}_\lambda$ there exists a negligible μ and a collection of sets $\mathcal{Z} = \{\mathcal{Z}_\lambda\}_\lambda$, where $|\mathcal{Z}_\lambda| \leq K(|\mathcal{A}_\lambda|)$, such that for any $\lambda \in \mathbb{N}$,

$$\Pr_{\mathcal{A}_\lambda} \left[V(x, \pi, 1^\lambda) = 1 \mid (x, \pi) \leftarrow \mathcal{A}_\lambda \right] \leq \mu(\lambda) .$$

φ -Tuning: Relaxed Soundness and Speeding-up Simulation We in fact consider a more general definition that allows to get faster simulators on the account of relaxing soundness. Here the argument system is associated with a non-expanding (typically, shrinking) projection function $\varphi(\cdot)$ defined over instances x . Soundness is relaxed and guarantees that the adversary could only output accepting pairs (x, π) for false statements whose projection $\varphi(x)$ falls in a set of size at most $K(S)$. Simulation is performed in two steps — a first preprocessing step that depends only on $\varphi(x)$, and a postprocessing step that depends on the instance x itself and the state produced in the preprocessing phase. The preprocessing phase takes superpolynomial time, but only depends on $\ell := |\varphi(x)|$ and not on $|x|$; the postprocessing phase takes polynomial time.

Note that the previous basic definition is indeed a special case of this definition by considering the identity projection (in this case the entire simulation is done in the preprocessing phase, and takes superpolynomial time in $|x|$). We gain from this definitions in scenarios where $\varphi : \{0, 1\}^{>\ell} \rightarrow \{0, 1\}^\ell$ is a shrinking projection — here when $\ell \ll |x|$, simulation can become significantly faster; furthermore, in settings where $\varphi(x)$, and its preprocessing are known ahead of time (but x isn't), we can get efficient simulation. On the other hand, we will only

get the above relaxed soundness guarantee. In our application to non-malleable commitments, relaxed soundness will be enough, and we'll indeed benefit from the above simulation speedup.

We proceed with the definition.

Definition 10 (φ -tuned 1ZK). A one-message zero-knowledge argument system (P, V) for an **NP** relation $\mathcal{R}(x, w)$ is φ -tuned for a polynomial-time projection function $\varphi = \{\varphi_\lambda : \{0, 1\}^{\geq \ell(\lambda)} \rightarrow \{0, 1\}^{\ell(\lambda)}\}_\lambda$ if it satisfies:

Simulation Speedup: The system is (T_D, T_S) -zero-knowledge with a uniform simulator $S = (S_{\text{pre}}, S_{\text{pos}})$ such that $S(x, 1^\lambda)$ consists of two phases:

- $\text{st} \leftarrow S_{\text{pre}}(\varphi_\lambda(x), 1^\lambda)$ is a preprocessing phase whose running time $T_{S_{\text{pre}}}(\ell(\lambda))$ depends on $\ell(\lambda) = |\varphi_\lambda(x)|$, but not on $|x|$.
- $\hat{\pi} \leftarrow S_{\text{pos}}(x, \text{st})$ is a postprocessing phase that takes time $\text{poly}(|x| + \lambda)$.

Overall, $T_S(|x|, \lambda) = \text{poly}(T_{S_{\text{pre}}}(\ell(\lambda)), |x|)$ depends only polynomially on $|x|$ (and superpolynomially on $|\varphi_\lambda(x)|$).

(T_P, K, φ) -Weak-Soundness: For any non-uniform $\text{poly}(T_P)$ -time, polynomial-size, probabilistic adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}_\lambda$ there exists a negligible μ and a collection of sets $\mathcal{Z} = \{\mathcal{Z}_\lambda\}_\lambda$, where $|\mathcal{Z}_\lambda| \leq K(|\mathcal{A}_\lambda|)$, such that for any $\lambda \in \mathbb{N}$,

$$\Pr_{\mathcal{A}_\lambda} \left[\begin{array}{l} x \notin \mathcal{L}(\mathcal{R}), \varphi_\lambda(x) \notin \mathcal{Z}_\lambda \\ V(x, \pi, 1^\lambda) = 1 \end{array} \mid (x, \pi) \leftarrow \mathcal{A}_\lambda \right] \leq \mu(\lambda) .$$

4.1 Construction

We now construct a φ -tuned 1ZK based on incompressible problems and other standard primitives. The parameters of the construction are derived from those of the underlying building blocks, and in particular on the density and incompressibility of the incompressible problem.

Building Blocks In what follows, let $\varphi = \{\varphi_\lambda : \{0, 1\}^{\geq \ell(\lambda)} \rightarrow \{0, 1\}^{\ell(\lambda)}\}_\lambda$ be a polynomial-time projection. Our transformation will make use the following building blocks:

- An incompressible problem $\mathcal{W} = \{\mathcal{W}_\lambda \subseteq \{0, 1\}^{4\ell(\lambda)}\}_\lambda$ with associated verifier \mathcal{V} , density Δ , and $(T_{\mathcal{W}}, K_{\mathcal{W}})$ incompressibility, where $K_{\mathcal{W}} \ll T_{\mathcal{W}} \ll \Delta^{-1}$.
- A commitment scheme $(\text{Com}, \text{Open})$ that is T_R -hiding and $T_{\text{Com.E}}$ -extractable where $T_R \ll T_{\text{Com.E}} \ll T_{\mathcal{W}}$.
- A T_D^{niwi} -indistinguishable NIWI system for an **NP** language, specified in the construction below.

- A two-source extractor $2\text{Ext} = \{2\text{Ext} : \{0, 1\}^{4\ell(\lambda)} \times \{0, 1\}^{4\ell(\lambda)} \rightarrow \{0, 1\}^{\ell(\lambda)}\}_\lambda$ with error $\varepsilon(\lambda) = 2^{-\ell(\lambda)-2}$ for sources of min-entropies $k_1 = k_2 > 4\ell(\lambda) - \log \Delta^{-1}$, and efficient reverse sampling.

The Proof System We now describe the system (P, V) for an **NP** relation \mathcal{R} .

- **The prover** $P(x, w, 1^\lambda)$:
 - Computes a commitment $c \leftarrow \text{Com}(0^{8\ell})$.
 - Computes a NIWI proof π for the statement

$$\psi_{x,c} :=$$

“Either $x \in \mathcal{L}(\mathcal{R})$ or

c is a commitment to $(\text{td}_1, \text{td}_2) \in \mathcal{W}_\lambda \times \mathcal{W}_\lambda$ such that $2\text{Ext}(\text{td}_1, \text{td}_2) = \varphi_\lambda(x)$.”

The prover uses the witness w to compute π .

- Overall the proof consists of (c, π) .
- **The verifier** $V(x, (c, \pi), 1^\lambda)$:
 - Applies the NIWI verifier to verify the statement $\psi_{x,c}$.

Theorem 4. *The above is a φ -tuned IZK for \mathcal{R} that is (T_S, T_D) -zero-knowledge and (T_P, K, φ) -weakly sound for*

$$T_S = \Delta^{-1}, T_D = \min \{T_R, T_D^{\text{niwi}}\}, \quad T_P = T_W, K = O(K_W) .$$

A Concrete Setting of Parameters. A natural setting of parameters that will be considered throughout this paper is subexponential $\Delta(\ell) = 2^{-\ell^\delta}$. We can accordingly set $T_R, T_{\text{Com.E}}, T_W, T_D^{\text{niwi}}$ to be super-polynomial functions satisfying:

$$T_R \ll T_{\text{Com.E}} \ll T_W \ll \Delta^{-1} = 2^{\ell(\lambda)^\delta} .$$

Indeed, the main tradeoff is between the simulation time T_S and the density Δ of the incompressible problem \mathcal{W} . On one hand, we aim for a short as possible simulation time $T_S \ll 2^{\ell(\lambda)}$.⁷ On the other hand, shorter simulation time requires higher density, which strengthens the corresponding incompressibility assumption. (In terms of existing candidates for incompressible problems based on fixed hash functions, subexponential density corresponds to polynomially-compressing hash functions.)

⁷Note that when φ is the identity, a witness for $x \in \{0, 1\}^{\ell(\lambda)}$ can already be found by brute force in time $2^{O(\ell(\lambda))}$, in which case the zero-knowledge requirement collapses to witness indistinguishability.

Acknowledgments

We thank Ilan Komargodski for pointing out [KY18].

References

- Bar02. Boaz Barak. Constant-round coin-tossing with a man in the middle or realizing the shared random string model. In *43rd Symposium on Foundations of Computer Science (FOCS 2002), 16-19 November 2002, Vancouver, BC, Canada, Proceedings*, pages 345–355, 2002.
- BFM88. Manuel Blum, Paul Feldman, and Silvio Micali. Non-interactive zero-knowledge and its applications (extended abstract). In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing, May 2-4, 1988, Chicago, Illinois, USA*, pages 103–112, 1988.
- BGI⁺17. Saikrishna Badrinarayanan, Sanjam Garg, Yuval Ishai, Amit Sahai, and Akshay Wadia. Two-message witness indistinguishability and secure computation in the plain model from new assumptions. In *Advances in Cryptology - ASIACRYPT 2017 - 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part III*, pages 275–303, 2017.
- BGJ⁺17. Saikrishna Badrinarayanan, Vipul Goyal, Abhishek Jain, Dakshita Khurana, and Amit Sahai. Round optimal concurrent MPC via strong simulation. In *Theory of Cryptography - 15th International Conference, TCC 2017, Baltimore, MD, USA, November 12-15, 2017, Proceedings, Part I*, pages 743–775, 2017.
- BKP18. Nir Bitansky, Yael Tauman Kalai, and Omer Paneth. In *Proceedings of the 50th Annual ACM Symposium on Theory of Computing, STOC 2018, Los-Angeles, CA, USA, June 25-29, 2018*, 2018.
- BL18. Nir Bitansky and Huijia Lin. One-message zero knowledge and non-malleable commitments. *IACR Cryptology ePrint Archive*, 2018:613, 2018.
- BM84. Manuel Blum and Silvio Micali. How to generate cryptographically strong sequences of pseudo-random bits. *SIAM J. Comput.*, 13(4):850–864, 1984.
- BOV07. Boaz Barak, Shien Jin Ong, and Salil P. Vadhan. Derandomization in cryptography. *SIAM J. Comput.*, 37(2):380–400, 2007.
- BP04. Boaz Barak and Rafael Pass. On the possibility of one-message weak zero-knowledge. In *Theory of Cryptography, First Theory of Cryptography Conference, TCC 2004, Cambridge, MA, USA, February 19-21, 2004, Proceedings*, pages 121–132, 2004.
- BP15. Nir Bitansky and Omer Paneth. Zaps and non-interactive witness indistinguishability from indistinguishability obfuscation. In *Theory of Cryptography - 12th Theory of Cryptography Conference, TCC 2015, Warsaw, Poland, March 23-25, 2015, Proceedings, Part II*, pages 401–427, 2015.
- BS05. Boaz Barak and Amit Sahai. How to play almost any mental game over the net - concurrent composition via super-polynomial simulation. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2005), 23-25 October 2005, Pittsburgh, PA, USA, Proceedings*, pages 543–552, 2005.
- CG88. Benny Chor and Oded Goldreich. Unbiased bits from sources of weak randomness and probabilistic communication complexity. *SIAM J. Comput.*, 17(2):230–261, 1988.
- CLP16. Ran Canetti, Huijia Lin, and Rafael Pass. Adaptive hardness and composable security in the plain model from standard assumptions. *SIAM J. Comput.*, 45(5):1793–1834, 2016.

- COSV16. Michele Ciampi, Rafail Ostrovsky, Luisa Siniscalchi, and Ivan 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*, pages 270–299, 2016.
- COSV17. Michele Ciampi, Rafail Ostrovsky, Luisa Siniscalchi, and Ivan Visconti. Four-round concurrent non-malleable commitments from one-way functions. In *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part II*, pages 127–157, 2017.
- DDN03. Danny Dolev, Cynthia Dwork, and Moni Naor. Nonmalleable cryptography. *SIAM Review*, 45(4):727–784, 2003.
- DJMW12. Yevgeniy Dodis, Abhishek Jain, Tal Moran, and Daniel Wichs. Counterexamples to hardness amplification beyond negligible. In *Theory of Cryptography - 9th Theory of Cryptography Conference, TCC 2012, Taormina, Sicily, Italy, March 19-21, 2012. Proceedings*, pages 476–493, 2012.
- DN07. Cynthia Dwork and Moni Naor. Zaps and their applications. *SIAM J. Comput.*, 36(6):1513–1543, 2007.
- FLS99. Uriel Feige, Dror Lapidot, and Adi Shamir. Multiple noninteractive zero knowledge proofs under general assumptions. *SIAM J. Comput.*, 29(1):1–28, 1999.
- GGJS12. Sanjam Garg, Vipul Goyal, Abhishek Jain, and Amit Sahai. Concurrently secure computation in constant rounds. In *Advances in Cryptology - EUROCRYPT 2012 - 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings*, pages 99–116, 2012.
- GKP17. Sanjam Garg, Susumu Kiyoshima, and Omkant Pandey. On the exact round complexity of self-composable two-party computation. In *Advances in Cryptology - EUROCRYPT 2017 - 36th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Paris, France, April 30 - May 4, 2017, Proceedings, Part II*, pages 194–224, 2017.
- GLOV12. Vipul Goyal, Chen-Kuei Lee, Rafail Ostrovsky, and Ivan Visconti. Constructing non-malleable commitments: A black-box approach. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012*, pages 51–60, 2012.
- GMR89. Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM J. Comput.*, 18(1):186–208, 1989.
- GO94. Oded Goldreich and Yair Oren. Definitions and properties of zero-knowledge proof systems. *Journal of Cryptology*, 7(1):1–32, December 1994.
- GOS12. Jens Groth, Rafail Ostrovsky, and Amit Sahai. New techniques for noninteractive zero-knowledge. *J. ACM*, 59(3):11, 2012.
- Goy11. Vipul Goyal. Constant round non-malleable protocols using one way functions. In *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pages 695–704, 2011.
- GPR16. Vipul Goyal, Omkant Pandey, and Silas Richelson. Textbook non-malleable commitments. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 1128–1141, 2016.
- GRRV14. Vipul Goyal, Silas Richelson, Alon Rosen, and Margarita Vald. An algebraic approach to non-malleability. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 41–50, 2014.

- HL18. Justin Holmgren and Alex Lombardi. Cryptographic hashing from strong one-way functions. *IACR Cryptology ePrint Archive*, 2018:385, 2018.
- Khu17. Dakshita Khurana. Round optimal concurrent non-malleability from polynomial hardness. In *Theory of Cryptography - 15th International Conference, TCC 2017, Baltimore, MD, USA, November 12-15, 2017, Proceedings, Part II*, pages 139–171, 2017.
- KS17. Dakshita Khurana and Amit Sahai. How to achieve non-malleability in one or two rounds. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 564–575, 2017.
- KY18. Ilan Komargodski and Eylon Yogev. Another step towards realizing random oracles: Non-malleable point obfuscation. In *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part I*, pages 259–279, 2018.
- LP09. Huijia Lin and Rafael Pass. Non-malleability amplification. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 189–198, 2009.
- LP11. Huijia Lin and Rafael Pass. Constant-round non-malleable commitments from any one-way function. In *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pages 705–714, 2011.
- LPS17. Huijia Lin, Rafael Pass, and Pratik Soni. Two-round and non-interactive concurrent non-malleable commitments from time-lock puzzles. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 576–587, 2017.
- LPV08a. Huijia Lin, Rafael Pass, and Muthuramakrishnan Venkitasubramaniam. Concurrent non-malleable commitments from any one-way function. In *Theory of Cryptography, Fifth Theory of Cryptography Conference, TCC 2008, New York, USA, March 19-21, 2008.*, pages 571–588, 2008.
- LPV08b. Huijia Lin, Rafael Pass, and Muthuramakrishnan Venkitasubramaniam. Concurrent non-malleable commitments from any one-way function. In *Theory of Cryptography, Fifth Theory of Cryptography Conference, TCC 2008, New York, USA, March 19-21, 2008.*, pages 571–588, 2008.
- MMY06. Tal Malkin, Ryan Moriarty, and Nikolai Yakovenko. Generalized environmental security from number theoretic assumptions. In *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006, Proceedings*, pages 343–359, 2006.
- Pas03. Rafael Pass. Simulation in quasi-polynomial time, and its application to protocol composition. In *EUROCRYPT*, pages 160–176, 2003.
- Pas13. Rafael Pass. Unprovable security of perfect NIZK and non-interactive non-malleable commitments. In *Theory of Cryptography - 10th Theory of Cryptography Conference, TCC 2013, Tokyo, Japan, March 3-6, 2013. Proceedings*, pages 334–354, 2013.
- PPV08. Omkant Pandey, Rafael Pass, and Vinod Vaikuntanathan. Adaptive one-way functions and applications. In *Advances in Cryptology - CRYPTO 2008, 28th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2008. Proceedings*, pages 57–74, 2008.
- PR05a. Rafael Pass and Alon Rosen. Concurrent non-malleable commitments. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2005), 23-25 October 2005, Pittsburgh, PA, USA, Proceedings*, pages 563–572, 2005.
- PR05b. Rafael Pass and Alon Rosen. New and improved constructions of non-malleable cryptographic protocols. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, pages 533–542, 2005.

- PS04. Manoj Prabhakaran and Amit Sahai. New notions of security: achieving universal composability without trusted setup. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing, Chicago, IL, USA, June 13-16, 2004*, pages 242–251, 2004.
- PW10. Rafael Pass and Hoeteck Wee. Constant-round non-malleable commitments from sub-exponential one-way functions. In *Advances in Cryptology - EUROCRYPT 2010, 29th Annual International Conference on the Theory and Applications of Cryptographic Techniques, French Riviera, May 30 - June 3, 2010. Proceedings*, pages 638–655, 2010.
- Rog06. Phillip Rogaway. Formalizing human ignorance. In *Progress in Cryptology - VIETCRYPT 2006, First International Conference on Cryptology in Vietnam, Hanoi, Vietnam, September 25-28, 2006, Revised Selected Papers*, pages 211–228, 2006.
- RSW00. Ronald L. Rivest, Adi Shamir, and David A. Wagner. Time-lock puzzles and timed-release crypto. Technical Report MIT/LCS/TR-684, MIT, February 2000.
- Unr07. Dominique Unruh. Random oracles and auxiliary input. In *Advances in Cryptology - CRYPTO 2007, 27th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2007, Proceedings*, pages 205–223, 2007.
- Vaz85. Umesh V. Vazirani. Towards a strong communication complexity theory or generating quasi-random sequences from two communicating slightly-random sources (extended abstract). In *Proceedings of the 17th Annual ACM Symposium on Theory of Computing, May 6-8, 1985, Providence, Rhode Island, USA*, pages 366–378, 1985.
- Wee10. Hoeteck Wee. Black-box, round-efficient secure computation via non-malleability amplification. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*, pages 531–540, 2010.