One-Message Zero Knowledge and Non-Malleable Commitments

Nir Bitansky* Huijia Lin†

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.

*Tel Aviv University. Email: nirbitan@tau.ac.il. 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.

†University of Santa Barbra. Email: rachel.lin@cs.ucsb.edu. 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.

Contents

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]. 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.[1]

[1] If there were such a simulator, then due to weak soundness, the simulator should fail to find accepting proofs for no-instances

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 $\mathsf{H} : \{0,1\}^\lambda \to \{0,1\}^{\lambda/2}$ 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.[2] Here, K is again a fixed polynomial (e.g., quadratic) and relatively-efficient means mildly superpolynomial-time (e.g. quasipolynomial or subexponential).

Theorem 1.1 (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* $w \in \mathcal{W}_\lambda$. On the other hand, $\mathcal{W}_\lambda$ is relatively *dense* in $\{0,1\}^\lambda$, in the sense that a random $w \leftarrow \{0,1\}^\lambda$ is in $\mathcal{W}_\lambda$ with relatively high probability $\Delta = 2^{-o(\lambda)}$.[3] For concreteness, the reader may think of $T = 2^{\lambda^{.01}} \ll 2^{\lambda^{.99}} = \Delta^{-1}$.

Theorem 1.2 (Informal)**.** *Assuming (T, K, Δ)-incompressible problems, where $K \ll T \ll \Delta^{-1} \ll 2^{\lambda^{.99}}$, 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 $\mathrm{poly}(\Delta^{-1})$-time simulator.*[4]

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

$\bar{x} \notin \mathcal{L}$ 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 $\mathcal{L}$.

[2] To 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.

[3] To 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 form linear compression.

[4] Here (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 $\mathrm{poly}(T)$, and T may be superpolynomial and the blowup function is K.

[KS17, LPS17] based on subexponential Decision-Diffie-Hellman or Quadratic Residousity 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*, or only against *uniform* adversaries [LPS17]. 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 1.3 (Informal)**.** *Under the same assumptions as in Theorem 1.2 (or 1.1), as well as subexponential time-lock puzzles, there exist fully-concurrent one-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-fledge 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 then 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.

1.2.1 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 $T_{\mathrm{td}} \ll 2^{o(|x|)}$. 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 $T_{\mathrm{com}} \ll T_{\mathrm{td}}$, 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 (in particular, with description smaller than the function's input), even in time $\mathrm{poly}(T_{\mathrm{com}})$. 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* $\mathsf{2Ext} : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^\ell$, 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 $\mathsf{2Ext}(w, w') = x$ 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 $\mathsf{2Ext}(U, U') = x$. These properties are both satisfied by the classical Hadamard extractor [CG88, Vaz85]. See further details in Section 4.

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

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

Starting from a $(T_{\mathcal{W}}, K, \Delta)$-incompressible problem, we choose a mildly-hard commitment so that it is extractable in time $T_{\mathrm{com}} \ll T_{\mathcal{W}}$. 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 $\mathsf{S}_{\mathsf{pre}}$, subexponential in $\varphi(x)$, that produces $\mathsf{st}_{\varphi(x)}$, and a short postprocessing step $\mathsf{S}_{\mathsf{pos}}$ that takes polynomial time given the trapdoor state $\mathsf{st}_{\varphi(x)}$.

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 Section 4.

1.2.2 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 commitments, 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 $O(\log^* n)$ 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 $\mathsf{tg}' = (\mathsf{tg}_1, \cdots, \mathsf{tg}_{\gamma/2})$ of iNM. To commit to a value v, oNM computes $\gamma/2$ commitments to v using iNM with respect to tags $\mathsf{tg}_1, \cdots, \mathsf{tg}_{\gamma/2}$, 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 $\mathsf{tg}' = (\mathsf{tg}_1, \cdots, \mathsf{tg}_{\gamma/2})$, the committer C generates $\{\mathrm{nm}_j \leftarrow \mathsf{iNM}(\mathsf{tg}_j, v)\}_{j \in [\gamma/2]}$ 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 $\mathsf{tg}' = (\mathsf{tg}_1, \cdots, \mathsf{tg}_{\gamma/2})$ and giving a right commitment using $\tilde{\mathsf{tg}}' = (\tilde{\mathsf{tg}}_1, \cdots, \tilde{\mathsf{tg}}_{\gamma/2})$. If $\mathsf{tg}' \neq \tilde{\mathsf{tg}}'$, there must exist $i^\star$, such that, $\tilde{\mathsf{tg}}_{i^\star} \neq \mathsf{tg}_i$ for all i — the $i^\star$'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 $\{\widetilde{\mathrm{nm}}_j\}$ 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 $i^\star$'th one $\widetilde{\mathrm{nm}}_{i^\star}$ — 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 $\tilde{v}_{i^\star}$ committed in $\widetilde{\mathrm{nm}}_{i^\star}$ does not change throughout hybrids.

In the first hybrid, the left SPS argument (π_1, π_2) is simulated. To ensure that $\tilde{v}_{i^\star}$ 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 $\tilde{v}_{i^\star}$ 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 $\widetilde{\mathrm{nm}}_{i^\star}$ uses a tag $\tilde{\mathsf{tg}}_{i^\star}$ different from all left tags, its committed value $\tilde{v}_{i^\star}$ does not change through these hybrids. Note that this requires the non-malleability of iNM to hold against T_{iNM}-time attackers for $T_{\mathsf{iNM}} \gg T_{\mathrm{SPS}}$. 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 $\{\widetilde{\mathrm{nm}}_j\}$ 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 underling non-malleable commitment, i.e. $T_{\mathsf{iNM}} \gg T_{\mathrm{SPS}}$.

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

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 $\mathsf{tg}' = (\mathsf{tg}_1, \cdots, \mathsf{tg}_{\gamma/2})$, the committer C generates $c \leftarrow \mathsf{Com}(v)$, $\{\mathrm{nm}_j \leftarrow \mathsf{iNM}(\mathsf{tg}_j, v)\}_{j \in [\gamma/2]}$, and a 1ZK argument π showing that c and all iNM commitments commit to the same value. The 1ZK statement is given by $x = (c, \mathrm{nm}_1, \cdots, \mathrm{nm}_{\gamma/2})$ 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 $\tilde{v}_{i^\star}$ committed in $\widetilde{\mathrm{nm}}_{i^\star}$.

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 $\mathsf{st} \leftarrow \mathsf{S}_{\mathsf{pre}}(c)$, and ii) a polynomial $\mathrm{poly}(|x|, \bar{n})$-time postprocessing phase that generates the simulated proof $\widehat{\pi} \leftarrow \mathsf{S}_{\mathsf{pos}}(x, \mathsf{st})$. 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 $\tilde{v}_{i^\star}$ 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 $\tilde{v}_{i^\star}$:

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

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 $f(r), \mathrm{hc}(r) \oplus b$, 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 $T_{\mathsf{iNM}} \gg T_{\mathrm{SPS}}$ by leveraging the efficient postprocessing of 1ZK simulation. Recall that given a trapdoor state $\mathsf{st} \leftarrow \mathsf{S}_{\mathsf{pre}}(c)$ that depends only on the projection c, simulating the proof $\widehat{\pi} \leftarrow \mathsf{S}_{\mathsf{pos}}(x, \mathsf{st})$ 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 $\tilde{v}_{i^\star}$ on the right changes, there must exist a commitment c (committing to 0) such that conditioned on c occurring in the hybrids the value $\tilde{v}_{i^\star}$ 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 $\mathsf{st} \leftarrow \mathsf{S}_{\mathsf{pre}}(c)$ 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 $\mathsf{tg}_0^\star, \mathsf{tg}_1^\star$ of iNM for special use, and commitments are computed using to other tags $\mathsf{tg} \neq \mathsf{tg}_0^\star, \mathsf{tg}_1^\star$.

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:

$$\mathrm{nm}_0 \leftarrow \mathsf{iNM}(\mathsf{tg}_0^\star, v) \qquad \mathrm{nm}_1 \leftarrow \mathsf{iNM}(\mathsf{tg}_1^\star, v) \ ,$$

  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 $\mathsf{tg}_0^\star$ and $\mathsf{tg}_1^\star$ 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 $\widetilde{\mathrm{nm}}_{0,j}$ and $\widetilde{\mathrm{nm}}_{1,j}$ 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 $\{\widetilde{\mathrm{nm}}_{1,j}\}$ do not change. Similarly, as the left commitment nm_1 is simulated, one can switch to showing that values in $\{\widetilde{\mathrm{nm}}_{0,j}\}$ 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 $\mathsf{tg}_0^\star, \mathsf{tg}_1^\star$. That is, i) simulating the 1ZK argument on the left does not change the values that the attacker commits to in iNM commitments $\{\widetilde{\mathrm{nm}}_{0,j}, \widetilde{\mathrm{nm}}_{1,j}\}$ 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 $\mathsf{tg}_0^\star, \mathsf{tg}_1^\star$ 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 5.2 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 $(\mathrm{poly}(\frac{T'}{T}) + (1-\delta)^k))$-hard for $(T' \ll T)$-time attackers. As the number k of copies increases, the hardness approaches the limit of $\mathrm{poly}(\frac{T'}{T})$.

The work of [DJMW12a] 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 $f(x_1), \ldots, f(x_\ell)$, so that the combined bit is 2^{ℓ^ε}-unpredicatable; that is, the level of unpredicatbility increases at least subexponentially as more hardcore bits are combined and beyond the limit $\mathrm{poly}(\frac{T'}{T})$.

We show that amplifiable one-way functions are useful for constructing non-malleable commitments. They essentially allow us to construct commitment schemes $(\mathsf{Com}, \mathsf{Com}')$, 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 $f(x_1), \ldots, f(x_\ell)$ 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 form 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 one, 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 $\mathrm{poly}(\varepsilon^{-1})$ 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.

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 $\mathsf{M} = \{\mathsf{M}_\lambda\}_{\lambda \in \mathbb{N}}$, 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 $\lambda \in \mathbb{N}$, $\mathsf{M}_\lambda = \mathsf{M}$ 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., $|A_\lambda| \geq \lambda$ (via padding).
- We say that a function $f : \mathbb{N} \to \mathbb{R}$ 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 $f : \mathbb{N} \to \mathbb{R}$ 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 $\lambda \to \infty$.

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

Definition 2.1 ((T,μ)-Indistinguishability)**.** *Let $\mathcal{X}^{(b)} = \{X^{(b)}_\lambda\}_{\lambda\in\mathbb{N}}$ 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 $\mathrm{poly}(T)$-time distinguishers $\mathcal{D}$, and all large enough λ,*

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

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

- $(c, d) \leftarrow \mathsf{Com}(v, 1^\lambda)$*: Given 1^λ and $v \in \{0,1\}^*$, Com samples a commitment c and a decommitment string d.*
- $b = \mathsf{Open}(c, v, d)$*: 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 $\mathsf{Open}(c, v, d) = 1$.*

We make the following requirements:

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

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

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

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

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

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

(Over-)Extractability We say a commitment scheme is extractable by an extractor $\mathsf{Com.E}$, if $\mathsf{Com.E}$, on input a commitment c, i) extracts the unique committed value whenever c is valid (i.e., has a valid decommitment), and ii) outputs $\perp$ otherwise. In addition, we say the commitment scheme is over-extractable by $\mathsf{Com.E}$ if only the first condition is fulfilled. We first define the unique committed value of a commitment via the value function, and then define (over-)extractability.

Definition 2.3 (The Value Function)**.** *The value function $\mathsf{val} : \{0,1\}^* \to \{0,1\}^*$ of a commitment scheme $(\mathsf{Com}, \mathsf{Open})$ is defined as follows: For every string $c \in \{0,1\}^*$,*

$$\mathsf{val}(c) = \begin{cases} v & \exists d \in \{0,1\}^*,\ \mathsf{Open}(c, v, d) = 1 \\ \perp & \textit{otherwise} \end{cases} .$$

Definition 2.4 ((Over-)extractable commitment scheme)**.** *A commitment scheme $(\mathsf{Com}, \mathsf{Open})$ is $T_{\mathsf{Com.E}}$-extractable, if there exists a uniform $\mathrm{poly}(T_{\mathsf{Com.E}})$-time extractor $\mathsf{Com.E}$ and a negligible function μ, such that, for every c,*

$$\Pr\left[\mathsf{Com.E}(c) \neq \mathsf{val}(c)\right] \leq \mu(|c|) .$$

We say that $(\mathsf{Com}, \mathsf{Open})$ is over-extractable if the above condition only holds for valid c, whose $\mathsf{val}(c) \neq \perp$.

Tag-based Commitments We consider "tag-based" commitment schemes.

Definition 2.5 (Tag-based commitment scheme)**.** *A commitment scheme $(\mathsf{Com}, \mathsf{Open})$ is a tag-based scheme with t-bit tags if, in addition to 1^λ, Com also receive a "tag" (a.k.a. identity) $\mathsf{tg} \in \{0,1\}^{t(\lambda)}$ as input, $c \leftarrow \mathsf{Com}(\mathsf{tg}, v, 1^\lambda)$. We assume w.l.o.g that commitments generated by Com contains the tag used for generating them. For any sequence of fixed tags $\mathsf{tg} = \{\mathsf{tg}_\lambda\}_\lambda$, the corresponding $(\mathsf{Com}_{\mathsf{tg}}, \mathsf{Open}_{\mathsf{tg}}) = \left\{(\mathsf{Com}_{\mathsf{tg}_\lambda}, \mathsf{Open}_{\mathsf{tg}_\lambda})\right\}_\lambda$ 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 $\mathsf{NM} = (\mathsf{Com}, \mathsf{Open})$ 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 $m = m(\lambda)$, $\mathcal{A}_\lambda$ on input 1^λ, receives m commitments from an honest committer C to values $v_1, \ldots, v_m \in \{0,1\}^\lambda$, and sends m commitments to R to values $\tilde{v}_1, \ldots, \tilde{v}_m \in \{0,1\}^\lambda$. 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 $\mathsf{tg}_1, \mathsf{tg}_2, \ldots, \mathsf{tg}_m$ and $\tilde{\mathsf{tg}}_1, \tilde{\mathsf{tg}}_2, \ldots, \tilde{\mathsf{tg}}_m$ chosen adaptively by $\mathcal{A}_\lambda$ for each commitment. The values $\tilde{v}_j$ in the j'th right commitment $\tilde{c}_j$ is defined as

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

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 $\mathsf{MIM}^{\mathcal{A}}_{\mathsf{NM}}(v_1, \ldots, v_m, 1^\lambda)$ the above described man-in-the-middle experiment. Next, we define two flavours of non-malleability — first, the standard non-malleability with respect to commitments [DDN03, LPV08b], and then the notion of non-malleability with respect to extraction [LPS17].

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

Definition 2.6 (Non-Malleability)**.** *A commitment scheme* NM *for* t*-bit tags is concurrent* T*-non-malleable if for any non-uniform* $\mathrm{poly}(T)$*-time adversary* $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda\in\mathbb{N}}$ *and for every polynomial* $m = m(\lambda)$*, it holds that:*

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

Non-malleability with respect to Extraction Consider the man-in-the-middle execution $\mathsf{MIM}^{\mathcal{A}}_{\mathsf{NM}}(v_1, \ldots, v_m)$ with a commitment scheme NM that is over-extractable by an extractor $\mathsf{NM.E}$. Non-malleable with respect to extraction [LPS17] requires that the joint distribution of the view of the adversary $\mathcal{A}_\lambda$ and the value extracted from the right commitments using $\mathsf{NM.E}$ to be indistinguishable when $\mathcal{A}_\lambda$ receives commitments to different values on the left. More precisely, for the j'th right commitment $\tilde{c}_j$, we define the extracted value ev_j as

$$ev_j = \begin{cases} \perp & \text{if } \exists i,\ \mathsf{tg}_i = \widetilde{\mathsf{tg}}_j \\ \mathsf{NM.E}(\tilde{c}_j) & \text{otherwise} \end{cases}$$

Define $\mathsf{emim}^{\mathcal{A}}_{\mathsf{NM}}(v_1, \ldots, v_m)$ to be the random variable describing the view of $\mathcal{A}_\lambda$ together with the values $(ev_1, \cdots ev_m)$ extracted from the right commitments.

Definition 2.7 (Non-malleability with respect to extraction)**.** *A commitment scheme* NM *for* t*-bit tags is concurrent* T*-non-malleable with respect to extraction by* $\mathsf{NM.E}$ *if the following hold:*

1. NM *is over-extractable by* $\mathsf{NM.E}$*, and*
2. *for every non-uniform* $\mathrm{poly}(T)$*-time adversary* $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda\in\mathbb{N}}$ *and every polynomial* $m = m(\lambda)$*, it holds:*

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

The key difference between non-malleability with respect to commitment and non-malleability with respect to extraction is that the former considers the unique committed values whereas the latter considers the extracted values. By binding and over-extractability, they are the same when a right commitment is valid. Otherwise, the committed value is $\perp$, while the extracted value may not be.

One-one and one-many Non-malleability We also consider relaxed notions called one-one and one-many non-malleability with respect to commitment and extraction. In the one-one (a.k.a. standalone) setting,

the man-in-the-middle adversary $\mathcal{A}$ can only participate in one left and one right interaction, while in the one-many setting, $\mathcal{A}$ participates in one left and many right interactions. It is known that one-many non-malleability is equivalent to concurrent non-malleability with respect to commitment (or extraction).

Lemma 2.1 ([LPV08b, LPS17])**.** *If a commitment scheme is one-many T-non-malleable with respect to commitment (or extraction), it is also concurrently T-non-malleable with respect to commitment (or extraction respectively).*

Same-Tag Non-Malleability In this work, we also consider an intermediate notion of non-malleability stronger than one-one non-malleability, and weaker than concurrent non-malleability, called same-tag non-malleability. Roughly speaking, it requires one-many non-malleability to hold with respect to attackers restricted to using the same tag for all right commitments they send. More formally, same-tag non-malleability considers man-in-the-middle attackers $\mathcal{A} = \{\mathcal{A}_\lambda\}$ that on security parameter 1^λ, receives *one* left commitment nm to $v \in \{0,1\}^\lambda$ using tg, and gives many right commitments $\widetilde{\mathrm{nm}}_1, \cdots \widetilde{\mathrm{nm}}_m$ using the *same* tag $\widetilde{\mathsf{tg}}$. Both the left and right tags tg and $\widetilde{\mathsf{tg}}$ are chosen by $\mathcal{A}$.

Definition 2.8 (Same-tag Non-malleability)**.** *A commitment scheme NM for t-bit tags is T-same-tag-non-malleable with respect to commitment if for every $\mathrm{poly}(T)$-time non-uniform attacker $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda\in\mathbb{N}}$, restricted to using the same tag for all commitments it sends,*

$$\left\{\mathsf{mim}^{\mathcal{A}}_{\mathsf{NM}}(v, 1^\lambda)\right\}_{\lambda\in\mathbb{N}, v, u\in\{0,1\}^\lambda} \approx_c \left\{\mathsf{mim}^{\mathcal{A}}_{\mathsf{NM}}(u, 1^\lambda)\right\}_{\lambda\in\mathbb{N}, v, u\in\{0,1\}^\lambda} .$$

NM is T-same-tag-non-malleable with respect to extraction by $\mathsf{NM.E}$, if instead the following indistinguishability holds.

$$\left\{\mathsf{emim}^{\mathcal{A}}_{\mathsf{NM}}(v, 1^\lambda)\right\}_{\lambda\in\mathbb{N}, v, u\in\{0,1\}^\lambda} \approx_c \left\{\mathsf{emim}^{\mathcal{A}}_{\mathsf{NM}}(u, 1^\lambda)\right\}_{\lambda\in\mathbb{N}, v, u\in\{0,1\}^\lambda} .$$

2.3 Non-Interactive Witness-Indistinguishable Proofs

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

Definition 2.9 (NIWI)**.** *A non-interactive witness-indistinguishable proof system (P, V) for an **NP** relation $\mathcal{R}(x, w)$ consists of two polynomial-time algorithms:*

- $\pi \leftarrow \mathsf{P}(x, w, 1^\lambda)$*: Given an instance x, witness w, and security parameter 1^λ, P produces a proof π.*
- $b = \mathsf{V}(x, \pi)$*: 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 $\lambda \in \mathbb{N}, (x, w) \in \mathcal{R}$,

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

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

$$\mathsf{V}(x,\pi) \neq 1 \ .$$

T-Witness-Indistinguishability: For any sequence

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

It holds that

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

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 2.10 (Two-Source Extractor)**.** *A polynomial-time computable function $\mathsf{2Ext} : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^m$ 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*

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

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 $\mathsf{2Ext}$ we can efficiently sample uniformly random and independent sources X_1 and X_2 conditioned on $\mathsf{2Ext}(X_1, X_2) = y$.

Definition 2.11 (Efficient Reverse Sampling)**.** *A function $\mathsf{2Ext} : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^m$ is efficiently reverse-samplable if there exists a PPT that given $y \in \mathsf{Image}(\mathsf{2Ext})$ outputs a uniformly random pair x_1, x_2 such that $\mathsf{2Ext}(x_1, x_2) = y$.*

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

Lemma 2.2 ([CG88, Vaz85])**.** *The Hadamard extractor is a (k_1, k_2, ε)-two-source with error $\varepsilon \leq 2^{(-k_1-k_2+n+m)/2}$ and efficient reverse sampling.*

We will also rely on the following basic lemma, which roughly says that when the error is tiny (relative to the output size), reverse sampling conditioned on a specific extractor output does not bias any sufficiently-frequent event by too much.

Lemma 2.3 (Density Preservation Under Reverse Sampling). *Let* $\mathsf{2Ext} : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^m$ *be a* (k_1, k_2, ε)*-two-source-extractor, where* $\varepsilon < 2^{-m-1}$*. Let* $A \subseteq \{0,1\}^n$ *be of cardinality* $|A| > 2^{\max\{k_1,k_2\}}$ *and let* $z \in \{0,1\}^m$*. Then*

$$\Pr_{x,y \leftarrow \{0,1\}^n}[(x,y) \in A \times A \mid \mathsf{2Ext}(x,y) = z] \geq (1 - \varepsilon 2^{m+1}) \cdot \Pr_{x,y \leftarrow \{0,1\}^n}[(x,y) \in A \times A] \quad .$$

Proof. The lemma follows by a standard calculation and the definition of two-source extractors:

$$\Pr_{x,y \leftarrow \{0,1\}^n}[(x,y) \in A \times A \mid \mathsf{2Ext}(x,y) = z] =$$
$$\frac{\Pr_{x,y \leftarrow \{0,1\}^n}[\mathsf{2Ext}(x,y) = z \mid (x,y) \in A \times A]}{\Pr_{x,y \leftarrow \{0,1\}^n}[\mathsf{2Ext}(x,y) = z]} \Pr_{x,y \leftarrow \{0,1\}^n}[(x,y) \in A \times A] \geq$$
$$\frac{2^{-m} - \varepsilon}{2^{-m} + \varepsilon} \Pr_{x,y \leftarrow \{0,1\}^n}[(x,y) \in A \times A] \geq (1 - \varepsilon 2^{m+1}) \cdot \Pr_{x,y \leftarrow \{0,1\}^n}[(x,y) \in A \times A] \quad .$$

□

2.5 Derandomization

We define Nisan-Wigderson pseudorandom generators against nondeterministic circuits [NW94].

Definition 2.12 (Nondeterministic Circuits). *A nondeterministic boolean circuit* $C(x,w)$ *takes* x *as a primary input and* w *as a witness. We define* $C(x) := 1$ *if and only if there exists* w *such that* $C(x,w) = 1$*.*

Definition 2.13 (NW-Type PRGs against Nondeterministic Circuits). *An algorithm PRG* $: \{0,1\}^{d(n)} \to \{0,1\}^n$ *is an NW-generator against non-deterministic circuits of size* $t(n)$ *if it is computable in time* $2^{O(d(n))}$ *and any non-deterministic circuit* C *of size at most* $t(n)$ *distinguishes* $U \leftarrow \{0,1\}^n$ *from PRG*(s)*, where* $s \leftarrow \{0,1\}^{d(n)}$*, with advantage at most* $1/t(n)$*.*

We shall rely on the following theorem by Shaltiel and Umans [SU01] regarding the existence NW-type PRGs as above assuming worst-case hardness for non-deterministic circuits.

Theorem 2.1 ([SU01]). *Assume there exists a function* $f : \{0,1\}^n \to \{0,1\}$ *in* $\mathbf{E} = \boldsymbol{Dtime}(2^{O(n)})$ *with nondeterministic circuit complexity* $2^{\Omega(n)}$*. Then, for any polynomial* $t(\cdot)$*, there exists an NW-generator against non-deterministic circuits of size* $t(n)$ *PRG* $: \{0,1\}^{d(n)} \to \{0,1\}^n$*, where* $d(n) = O(\log n)$*.*

We remark that the above is a worst-case assumption in the sense that the function f needs to be hard in the worst-case (and not necessarily in the average-case). The assumption can be seen as a natural generalization of the assumption that **EXP** $\not\subseteq$ **NP**. We also note that there is a universal candidate for the corresponding PRG, by instantiating the hard function with any **E**-complete language under linear reductions. See further discussion in [BOV07].

3 Incompressible Problems

Following [BKP18], we consider a notion of incompressible problems. Here every security parameter λ, defines a search problem $\mathcal{W}_\lambda$ 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 3.1 (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 $\mathrm{poly}(T)$-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 $K = K(|\mathcal{A}_\lambda|)$,

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

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 3.1 (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 $T \geq \mathrm{poly}(K\Delta^{-1})$ the requirement becomes impossible.

Remark 3.2 (Probabilistic Adversaries). In common cryptographic settings, non-uniform adversaries can be assumed to be deterministic w.l.o.g (by fixing their randomness to that which maximizes their success probability). In the above definition, however, this is not the case. Indeed, crucially decouple description size and running time — we require that the number of solutions an adversary can find depends on its description size, but not on its running time and the amount of coins it tosses, which could be significantly larger. In particular, fixing a given number of coins increases the description of the adversary.

An Alternative Formulation of Incompressibility The above formulation closely follows the treatment in [BKP18]. Throughout most of this paper, it will be more convenient to consider an alternative formulation that says that an adversary of a give size S has a corresponding set $\mathcal{Z}$ of $K(S)$ solutions, so that it cannot sample solutions outside of this set except with negligible probability. We show below that this formulation is essentially equivalent.

Lemma 3.1 (Adversary's Set of Solutions)**.** *Let $\mathcal{W} = \{\mathcal{W}_\lambda\}_\lambda$ be a (T, K)-incompressible problem. Then for any non-uniform $\mathrm{poly}(T)$-time, polynomial-size, probabilistic adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}_\lambda$ there exists a negligible μ and a collection of sets $\mathcal{Z} = \{\mathcal{Z}_\lambda\}$, where $|\mathcal{Z}_\lambda| \leq O(K(|\mathcal{A}_\lambda|))$, such that for any $\lambda \in \mathbb{N}$,*

$$\Pr_{\mathcal{A}_\lambda}[w \in \mathcal{W}_\lambda \setminus \mathcal{Z}_\lambda \mid w \leftarrow \mathcal{A}_\lambda] \leq \mu(\lambda) \ .$$

Proof. Let $\widetilde{K} = K(|\mathcal{A}_\lambda| + \lambda) = O(K(|\mathcal{A}_\lambda|))$, define $\mathcal{Z}_\lambda$ to be $\widetilde{K}$ elements in $\mathcal{W}_\lambda$ that $\mathcal{A}_\lambda$ outputs with maximal probability. Assume toward contradiction that there exists a noticeable function $\varepsilon = \varepsilon(\lambda)$ such that for infinitely many $\lambda \in \mathbb{N}$,

$$\Pr_{\mathcal{A}_\lambda}[w \in \mathcal{W}_\lambda \setminus \mathcal{Z}_\lambda \mid w \leftarrow \mathcal{A}_\lambda] \geq \varepsilon \ . \tag{1}$$

We'll construct $\mathrm{poly}(T)$-time non-uniform polynomial-size adversary $\mathcal{B} = \{\mathcal{B}_\lambda\}_\lambda$ that breaks (T, K)-incompressibility. $\mathcal{B}_\lambda$ takes $2\widetilde{K}\left(4\widetilde{K}/\varepsilon\right)^2$ samples $w \leftarrow \mathcal{A}_\lambda$ and outputs all $w \in \mathcal{W}_\lambda$ (without repetitions). Note that $\mathcal{B}$ has running time $\mathrm{poly}(T)$ and description size $|B_\lambda| \leq |\mathcal{A}_\lambda| + \lambda$. We show that $\mathcal{B}$ outputs $K(|\mathcal{B}_\lambda|)$ distinct $w \in \mathcal{W}_\lambda$ with high probability.

We consider two cases. First, consider the case that

$$\min_{w \in \mathcal{Z}_\lambda} \Pr_{\mathcal{A}_\lambda}[w \leftarrow \mathcal{A}_\lambda] \geq \left(\varepsilon/4\widetilde{K}\right)^2 \ .$$

Here the expected number of samples required to collect all elements in $\mathcal{Z}_\lambda$ is at most $\widetilde{K}\left(4\widetilde{K}/\varepsilon\right)^2$. Thus, by Markov's inequality, after $2\widetilde{K}\left(4\widetilde{K}/\varepsilon\right)^2$ samples, $\mathcal{B}_\lambda$ has collected all $\widetilde{K}$ elements in $\mathcal{Z}_\lambda$, with probability at least $1/2$. Now, consider the case that

$$\min_{w \in \mathcal{Z}_\lambda} \Pr_{\mathcal{A}_\lambda}[w \leftarrow \mathcal{A}_\lambda] < \left(\varepsilon/4\widetilde{K}\right)^2 \ .$$

Since $\mathcal{Z}_\lambda$ includes elements in $\mathcal{W}_\lambda$ with maximal density, the above implies that for any $w \in \mathcal{W}_\lambda \setminus \mathcal{Z}_\lambda$,

$$\Pr_{\mathcal{A}_\lambda}[w \leftarrow \mathcal{A}_\lambda] < \left(\varepsilon/4\widetilde{K}\right)^2 \ . \tag{2}$$

By Equation 1, we know that the expected number of samples to collect $\widetilde{K}$ elements $w \in \mathcal{W}_\lambda \setminus \mathcal{Z}_\lambda$, with repetitions, is at most $\widetilde{K}/\varepsilon$. By Markov's inequality, after $2\widetilde{K}/\varepsilon$ samples, $\widetilde{K}$ such elements are collected with probability at least $1/2$. By Equation 2, the probability that these elements are not distinct is at most

$$\left(2\widetilde{K}/\varepsilon\right)^2 \cdot \left(\varepsilon/4\widetilde{K}\right)^2 \leq 1/4 \ .$$

Thus, in this case, $\mathcal{B}_\lambda$ outputs $\widetilde{K}$ distinct elements in $\mathcal{W}_\lambda$ with probability at least $1/4$. □

3.1 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.

We now recall the problems considered in [BKP18], and then discuss the possibility of additional candidates.

Keyless Hash Functions The standard notion of collision resistance for shrinking hash functions requires that a key for the hash function is sampled at random. Indeed, if one considers keyless hash functions, a non-uniform adversary can always have hardwired collisions. Bitansky, Kalai, and Paneth [BKP18] suggested a security notion for keyless hash functions, which postulates that hardwiring collisions is essentially the best that the adversary can do — an adversary with non-uniform description of size S and arbitrary polynomial running time, should not be able to find more than $K(S)$ collisions (for some blowup function K).

In a bit more details, they define two such notions of incompressibility. The first, called *multi-collision resistance*, says that it should be hard to sample a K-collision — K preimages $X_1, \ldots, X_K$ that all map to the same image Y. The second, called *strong multi-collision resistance*, says that its hard to find K plain collisions — pairs $(X_1, X_1'), \ldots, (X_K, X_K')$ such that each pair (X_i, X_i') maps to the same image Y_i.

We recall the formal definitions of such hash functions. In what follows, $\mathsf{H} = \{\mathsf{H}_\lambda\}_\lambda$ is an efficiently computable hash function mapping $\ell(\lambda) > \lambda$ bits to λ bits.

Definition 3.2 ((T,K)-Collision Resistance)**.** *A keyless hash function* $\mathsf{H} = \{\mathsf{H}_\lambda\}_\lambda$ *is* (T,K)*-collision-resistant if for any non-uniform* T*-time* polynomial-size *adversary* $\mathcal{A} = \{\mathcal{A}_\lambda\}_\lambda$*, there is a negligible function* μ*, such that for any* $\lambda \in \mathbb{N}$*, letting* $K = K(|\mathcal{A}_\lambda|)$*,*

$$\Pr\left[\begin{array}{c} \mathsf{H}_\lambda(X_1) = \cdots = \mathsf{H}_\lambda(X_K) \\ \forall i \neq j : X_i \neq X_j \end{array} \;\middle|\; (X_1,\ldots,X_K) \leftarrow \mathcal{A}_\lambda \right] \leq \mu(\lambda) \ .$$

Definition 3.3 (Strong (T,K)-Collision Resistance)**.** *A keyless hash function* $\mathsf{H} = \{\mathsf{H}_\lambda\}_\lambda$ *is strongly* (T,K)*-collision-resistant if for any non-uniform* T*-time* polynomial-size *adversary* $\mathcal{A} = \{\mathcal{A}_\lambda\}_\lambda$*, there is a negligible function* μ*, such that for any* $\lambda \in \mathbb{N}$*, letting* $K = K(|\mathcal{A}_\lambda|)$*,*

$$\Pr\left[\begin{array}{c} \forall i : \mathsf{H}_\lambda(X_i) = \mathsf{H}_\lambda(X'_i), X_i \neq X'_i \\ \forall i \neq j : \{X_i, X'_i\} \neq \left\{X_j, X'_j\right\} \end{array} \;\middle|\; (X_1, X'_1, \ldots, X_K, X'_K) \leftarrow \mathcal{A}_\lambda \right] \leq \mu(\lambda) \ .$$

As current candidates for such hash functions they suggest existing keyless hash functions such as SHA or AES-based hashing. They also show that random oracles satisfy multi-collision resistance even in the auxiliary-input model of Unruh [Unr07] where the adversary can store arbitrary polynomial auxiliary input about the random oracle.

From Multi-Collision Resistance to Incompressible Problems We now show how multi-collision resistant hash functions imply incompressible problems with subexponential density.

We start by providing an incompressible problem that follows directly from strong multi-collision resistance.

Claim 3.1. *Assume the existence of a keyless strongly* (T,K)*-collision resistant hash function* $\mathsf{H} = \{\mathsf{H}_\lambda\}_\lambda$ *mapping inputs of* $\ell(\lambda) = \lambda^{1/\varepsilon}$ *bits to* λ *bits for some* $0 < \varepsilon < 1$*. Then there exists a* (T,K)*-incompressible problem* $\mathcal{W} = \left\{\mathcal{W}_\lambda \subseteq \{0,1\}^{2\ell(\lambda)}\right\}_\lambda$ *with density* $\Delta(\ell) = \Omega(2^{-\ell^\varepsilon})$*.*

Proof Sketch. We define

$$\mathcal{W}_\lambda = \left\{\left\{X, X'\right\} \;\middle|\; X \neq X', \mathsf{H}_\lambda(X) = \mathsf{H}_\lambda(X')\right\} \ .$$

$\mathcal{W}$ is efficiently recognizable by a verifier $\mathcal{V}(X,X')$ that simply checks if $\mathsf{H}_\lambda(X) = \mathsf{H}_\lambda(X')$. The (T,K)-incompressibility of $\mathcal{W}$ follows directly from the strong (T,K)-resistance of H. Furthermore, the probability that two random X, X' collide and are distinct is at least $2^{-\ell^\varepsilon} - 2^{-\ell} = \Omega(2^{-\ell^\varepsilon})$. □

We now move to discuss an instantiation based on (weak) multi-collision resistance. Here the natural direction is to define the problem to be the set of preimages for some image Y whose preimage set is dense. Indeed, we certainly know that there exists at least one image whose preimage set has density at least $2^{-\ell^\varepsilon}$. The problem with this approach is that we might not know how to uniformly compute such an image Y, which is needed since we want the schemes constructed in this paper to be explicit.

We show how to circumvent this problem using a standard derandomization assumption. The basic idea is that for a random value X, the preimage set of $Y = \mathsf{H}_\lambda(X)$ has density $\Omega(2^{-\ell^\varepsilon})$ with high probability. We show that we can derandomize the choice of such X using appropriate NW-type pseudorandom generators as defined in Section 2.5.

Claim 3.2. *Assume the existence of a keyless* (T,K)*-collision resistant hash function* $\mathsf{H} = \{\mathsf{H}_\lambda\}_\lambda$ *mapping inputs of* $\ell(\lambda) = \lambda^{1/\varepsilon}$ *bits to* λ *bits for some* $0 < \varepsilon < 1$ *and an NW-type pseudorandom generator against non-deterministic circuits. Then there exists a* $(T, K\lambda^{O(1)})$*-incompressible problem* $\mathcal{W} = \left\{\mathcal{W}_\lambda \subseteq \{0,1\}^{2\ell(\lambda)}\right\}_\lambda$ *with density* $\Delta(\ell) = \Omega(2^{-\ell^\varepsilon})$*.*

Proof Sketch. Let $\mathrm{PRG} : \{0,1\}^{d(\lambda)} \to \{0,1\}^{\ell(\lambda)}$ be an NW-type PRG against nondeterministic circuits of size $t(\lambda)$, where t will be specified below. We define:

$$\mathcal{W}_\lambda = \left\{ X \;\middle|\; \exists s \in \{0,1\}^d : \mathsf{H}_\lambda(X) = \mathsf{H}_\lambda(\mathrm{PRG}(s)) \right\} \ .$$

First, since $d = O(\log \lambda)$, $\mathcal{W}$ is efficiently recognizable by a verifier $\mathcal{V}(X)$ that enumerates over all $\{X' = \mathrm{PRG}(s)\}_{s \in \{0,1\}^d}$ and checks if $\mathsf{H}_\lambda(X) = \mathsf{H}_\lambda(X')$. $(T, K2^d)$-incompressibility of $\mathcal{W}$ follows from the (T, K)-collision-resistance of H. Indeed, any adversary who outputs $K2^d$ distinct $X \in \mathcal{W}_\lambda$ with probability δ directly translates to an adversary who can output K preimages X for some specific $Y = \mathsf{H}_\lambda(\mathrm{PRG}(s))$ with probability $\delta/2^d = \delta/\lambda^{O(1)}$.

We move to proving that $\mathcal{W}$ has density $\Delta = \Omega(2^{-\ell^\varepsilon})$. First, note that for a random $X \leftarrow \{0,1\}^\ell$, the density δ_X of its preimage set $S_X := \{X' \mid \mathsf{H}_\lambda(X) = \mathsf{H}_\lambda(X')\}$ is at least $2^{-\ell^\varepsilon - 1}$ with probability at least $1/2$. We would like to establish a similar claim for the case that X is sampled pseudorandomly. We will show that the density in this case is at least $2^{-\ell^\varepsilon - 10}$ with probability at least $1/16$.

We will show that otherwise we can construct a nondeterministic circuit distinguisher D for the underlying NW-PRG. We will describe a randomized non-deterministic circuit, which is sufficient by later fixing the randomness to maximize the distinguishing advantage. The distinguisher D given a sample $X \in \{0,1\}^\ell$, which is either random or pseudorandom, samples a pairwise independent hash h from a family mapping $\{0,1\}^\ell$ to $T := \{0,1\}^{\ell - \ell^\varepsilon - 5}$, and nondeterministically checks whether there exists $X' \in S_X$ such that $h(X') = 0$.

We use the following standard facts regarding pairwise independent hashing:

$$\Pr_h \left[|h^{-1}(0) \cap S_X| \le \frac{|S_X|}{2|T|} \right] \le \frac{4|T|}{|S_X|} \ ,$$
$$\Pr_h \left[|h^{-1}(0) \cap S_X| > 0 \right] \le \frac{|S_X|}{|T|} \ .$$

In particular, if $\delta_X \ge 2^{-\ell^\varepsilon - 1}$, $|S_X| \ge 2^{\ell - \ell^\varepsilon - 1} = 16|T|$, and

$$\Pr_h \left[|h^{-1}(0) \cap S_X| < 1 \right] \le 1/4 \ .$$

In contrast, if $\delta_X \le 2^{-\ell^\varepsilon - 10}$, $|S_X| \le 2^{\ell - \ell^\varepsilon - 10} = |T|/32$, and

$$\Pr_h \left[|h^{-1}(0) \cap S_X| > 0 \right] \le 1/32 \ .$$

Fixing $t(\lambda) \ge \max\{|D|, 32\}$, we conclude that

$$\Pr_{\substack{s \leftarrow \{0,1\}^d \\ X \leftarrow \mathrm{PRG}(s)}} \left[\delta_X \ge 2^{-\ell^\varepsilon - 10} \right] \ge \frac{1}{2} - \frac{1}{4} - \frac{1}{32} - \frac{1}{t} \ge \frac{1}{16} \ .$$

This in particular implies that $\mathcal{W}$ has density Δ at least $2^{-\ell^\varepsilon - 10}$. □

Remark 3.3 (Shrinkage). Above we assumed (strong) multi-collision-resistant hash functions with polynomial shrinkage. In [BKP18], it is shown how strong (T, K)-collision-resistant functions with linear shrinkage can be transformed into (weak) $(T, O(K))$-collision-resistant functions with arbitrary polynomial shrinkage. They also show how (weak) (T, K)-collision-resistant functions with linear shrinkage can be transformed into (weak) $(T, \mathrm{quasipoly}(K))$-collision-resistant functions with arbitrary polynomial shrinkage (when $T \gg \mathrm{quasipoly}(K)$).

Beyond Keyless Hash Functions? We note that for our applications general incompressible problems suffice. This may give rise to other candidates in the future. In addition, incompressibility in general may be qualitatively weaker than multi-collision-resistance. In fact, unlike multi-collision resistant hash functions, we do not even know that incompressible functions imply one-way functions.

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 4.1 (1ZK)**.** *A one-message zero-knowledge argument system* (P, V) *for an* ***NP*** *relation* $\mathcal{R}(x, w)$ *consists of two polynomial-time algorithms:*

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

The system is parameterized by functions $T_{\mathsf{D}}(\cdot), T_{\mathsf{S}}(\cdot), T_{\mathsf{P}}(\cdot), K(\cdot)$*.*

We make the following requirements:

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

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

$(T_{\mathsf{D}}, T_{\mathsf{S}})$-Zero-Knowledge: There exists a uniform $\mathrm{poly}(T_{\mathsf{S}})$*-time simulator* S*, such that,*

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

(T_{P}, K)-Weak-Soundness: For any non-uniform $\mathrm{poly}(T_{\mathsf{P}})$*-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* $\lambda \in \mathbb{N}$*,*

$$\Pr_{\mathcal{A}_\lambda}\left[\begin{array}{c} x \notin \mathcal{L}(\mathcal{R}) \cup \mathcal{Z}_\lambda \\ \mathsf{V}(x, \pi, 1^\lambda) = 1 \end{array} \;\middle|\; (x, \pi) \leftarrow \mathcal{A}_\lambda\right] \leq \mu(\lambda) \ .$$

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 $\varphi : \{0,1\}^{>\ell} \to \{0,1\}^{\ell}$ 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 4.2 (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* $\varphi = \left\{\varphi_\lambda : \{0,1\}^{\geq \ell(\lambda)} \to \{0,1\}^{\ell(\lambda)}\right\}_\lambda$ *if it satisfies:*

Simulation Speedup: The system is $(T_{\mathsf{D}}, T_{\mathsf{S}})$*-zero-knowledge with a uniform simulator* $\mathsf{S} = (\mathsf{S}_{\mathsf{pre}}, \mathsf{S}_{\mathsf{pos}})$ *such that* $\mathsf{S}(x, 1^\lambda)$ *consists of two phases:*

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

Overall, $T_{\mathsf{S}}(|x|, \lambda) = \mathrm{poly}(T_{\mathsf{S}_{\mathsf{pre}}}(\ell(\lambda)), |x|)$ *depends only polynomially on* $|x|$ *(and superpolynomially on* $|\varphi_\lambda(x)|$*).*

$(T_{\mathsf{P}}, K, \varphi)$-Weak-Soundness: For any non-uniform $\mathrm{poly}(T_{\mathsf{P}})$*-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* $\lambda \in \mathbb{N}$*,*

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

We will also use a slightly more general soundness requirement that extends the basic soundness requirement to the case that the adversary outputs multiple statements and proofs.

Lemma 4.1 (Multi-Proof Weak Soundness)**.** *For any polynomial* $t = t(\lambda)$*,* $(T_{\mathsf{P}}, K, \varphi)$*-Weak-Soundness implies*

$(T_{\mathsf{P}}, K, \varphi, t)$-Weak-Soundness: For any non-uniform $\mathrm{poly}(T_{\mathsf{P}})$*-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* $|\mathcal{Z}_\lambda| \leq K(|\mathcal{A}_\lambda| + O(1))$*, such that for any* $\lambda \in \mathbb{N}$*,*

$$\Pr_{\mathcal{A}_\lambda}\left[\exists i \in [t] : \begin{array}{l} x_i \notin \mathcal{L}(\mathcal{R}), \varphi_\lambda(x_i) \notin \mathcal{Z}_\lambda \\ \mathsf{V}(x_i, \pi_i, 1^\lambda) = 1 \end{array} \;\middle|\; (x_1,\pi_1), \ldots, (x_t,\pi_t) \leftarrow \mathcal{A}_\lambda\right] \leq \mu(\lambda) \ .$$

The proof follows readily from the definitions.

Proof Sketch. For any $\mathcal{A}$ as in the conditions of the lemma, consider a new adversary $\mathcal{A}'$ that runs $\mathcal{A}$, and outputs $(x,_i \pi_i)$ for a random $i \in [t]$. We can then fix $\mathcal{Z}$ to be the corresponding system of sets given for $\mathcal{A}'$ by the (single-proof) weak soundness requirement. This set satisfies the required conditions since if $\mathcal{A}$ outputs as part of its list an accepting proof π for a statement x such that $\varphi_\lambda(x) \notin \mathcal{Z}_\lambda$ with probability ε, $\mathcal{A}'$ does so with probability ε/t. □

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 incompressability of the incompressible problem.

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

- An incompressible problem $\mathcal{W} = \left\{\mathcal{W}_\lambda \subseteq \{0,1\}^{4\ell(\lambda)}\right\}_\lambda$ with associated verifier $\mathcal{V}$, density Δ, and $(T_{\mathcal{W}}, K_{\mathcal{W}})$ incompressability, where $K_{\mathcal{W}} \ll T_{\mathcal{W}} \ll \Delta^{-1}$.
- A commitment scheme $(\mathsf{Com}, \mathsf{Open})$ that is T_{R}-hiding and $T_{\mathsf{Com.E}}$-extractable where $T_{\mathsf{R}} \ll T_{\mathsf{Com.E}} \ll T_{\mathcal{W}}$.
- A $T_{\mathsf{D}}^{\mathsf{niwi}}$-indistinguishable NIWI system for an **NP** language, specified in the construction below.
- A two-source extractor $\mathsf{2Ext} = \left\{\mathsf{2Ext} : \{0,1\}^{4\ell(\lambda)} \times \{0,1\}^{4\ell(\lambda)} \to \{0,1\}^{\ell(\lambda)}\right\}_\lambda$ 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** $\mathsf{P}(x, w, 1^\lambda)$**:**
  - Computes a commitment $c \leftarrow \mathsf{Com}(0^{8\ell})$.
  - Computes a NIWI proof π for the statement

    $$\psi_{x,c} := \begin{array}{l} \text{“}\textit{Either } x \in \mathcal{L}(\mathcal{R}) \textit{ or} \\ c \textit{ is a commitment to } (\mathrm{td}_1, \mathrm{td}_2) \in \mathcal{W}_\lambda \times \mathcal{W}_\lambda \textit{ such that } \mathsf{2Ext}(\mathrm{td}_1, \mathrm{td}_2) = \varphi_\lambda(x).\text{”} \end{array}$$

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

Theorem 4.1. *The above is a φ-tuned 1ZK for $\mathcal{R}$ that is $(T_{\mathsf{S}}, T_{\mathsf{D}})$-zero-knowledge and $(T_{\mathsf{P}}, K, \varphi)$-weakly sound for*

$$T_{\mathsf{S}} = \Delta^{-1}, T_{\mathsf{D}} = \min\left\{T_{\mathsf{R}}, T_{\mathsf{D}}^{\mathsf{niwi}}\right\}, \qquad T_{\mathsf{P}} = T_{\mathcal{W}}, K = O(K_{\mathcal{W}}) \ .$$

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 $T_{\mathsf{R}}, T_{\mathsf{Com.E}}, T_{\mathcal{W}}, T_{\mathsf{D}}^{\mathsf{niwi}}$ to be super-polynomial functions satisfying:

$$T_{\mathsf{R}} \ll T_{\mathsf{Com.E}} \ll T_{\mathcal{W}} \ll \Delta^{-1} = 2^{\ell(\lambda)^\delta} \ .$$

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 $T_{\mathsf{S}} \ll 2^{\ell(\lambda)}$. [5] On the other

[5] Note 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.

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.)

Proof of Theorem 4.1. The completeness of the system follows readily from the construction, we focus on showing weak soundness and φ-tuned zero-knowledge.

Weak Soundness Let $\mathsf{P}^* = \{\mathsf{P}^*_\lambda\}_\lambda$ be any non-uniform prover of polynomial size and $\mathrm{poly}(T_{\mathcal{W}})$ running time. We would like to establish the existence of $O(K(|\mathsf{P}^*_\lambda|))$-size sets $\mathcal{Z} = \{\mathcal{Z}_\lambda\}_\lambda$, such that except with negligible probability, P^* cannot sample an accepting pair $(x,(c,\pi))$ for no-instances $x \notin \mathcal{L}(\mathcal{R})$, unless $\varphi_\lambda(x) \in \mathcal{Z}_\lambda$.

Consider $\mathcal{W} \times \mathcal{W} = \{\mathcal{W}_\lambda \times \mathcal{W}_\lambda\}_\lambda$, and note that it is also a $(T_{\mathcal{W}}, K_{\mathcal{W}})$-incompressible problem. Consider an adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}_\lambda$ that tries to sample solutions in $\mathcal{W} \times \mathcal{W}$ as follows. It runs P^*_λ to sample $(x,(c,\pi))$ and then applies the extractor $\mathsf{Com.E}(c)$ to obtain $(\mathrm{td}_1, \mathrm{td}_2)$. Note that the description size of this adversary is bounded by $|\mathcal{A}_\lambda| \leq |\mathsf{P}^*_\lambda| + O(1)$, and that its running time is at most $T_{\mathsf{P}^*} + T_{\mathsf{Com.E}} \leq \mathrm{poly}(T_{\mathcal{W}})$. Thus, by Lemma 3.1, there exist sets $\mathcal{Z}^{\mathcal{A}} = \left\{\mathcal{Z}^{\mathcal{A}}_\lambda\right\}_\lambda$ of size $O(K_{\mathcal{W}}(|\mathcal{A}_\lambda|)) = O(K_{\mathcal{W}}(|\mathsf{P}^*_\lambda|))$ such that except with negligible probability, $\mathcal{A}_\lambda$ does not output $(\mathrm{td}_1, \mathrm{td}_2) \in \mathcal{W}_\lambda \times \mathcal{W}_\lambda$, unless $(\mathrm{td}_1, \mathrm{td}_2) \in \mathcal{Z}^{\mathcal{A}}_\lambda$.

We now define

$$\mathcal{Z}_\lambda := \left\{ y \;\middle|\; \begin{array}{l} y = \mathsf{2Ext}(\mathrm{td}_1, \mathrm{td}_2) \\ (\mathrm{td}_1, \mathrm{td}_2) \in \mathcal{Z}^{\mathcal{A}}_\lambda \end{array} \right\} .$$

It is left to note that by the soundness of the NIWI system, whenever P^* outputs an accepting pair $(x,(c,\pi))$, if $x \notin \mathcal{L}(\mathcal{R})$, c is a commitmnet to $(\mathrm{td}_1, \mathrm{td}_2) \in \mathcal{W}_\lambda \times \mathcal{W}_\lambda$. It follows that $\mathcal{Z} = \{\mathcal{Z}_\lambda\}_\lambda$ as defined above fulfills the $(T_{\mathcal{W}}, O(K_{\mathcal{W}}))$-weak-soundness property.

Zero Knowledge We now wish to establish a $\mathrm{poly}(\Delta^{-1})$-time simulator for the system. We will further see that the simulator can be decomposed into a relatively-efficient preprocessing simulator $\mathsf{S}_{\mathsf{pre}}$ and a poly-time postprocessing simulator $\mathsf{S}_{\mathsf{pos}}$ as required by the definition of φ-tuned zero-knowledge.

The Simulator $\mathsf{S}(x, 1^\lambda)$:

1. Computes $y = \varphi_\lambda(x)$.
2. Reverse samples $(\mathrm{td}_1, \mathrm{td}_2)$ such that $\mathsf{2Ext}(\mathrm{td}_1, \mathrm{td}_2) = y$ and $(\mathrm{td}_1, \mathrm{td}_2) \in \mathcal{W}_\lambda \times \mathcal{W}_\lambda$.
3. Computes a commitment $c = \mathsf{Com}(\mathrm{td}_1, \mathrm{td}_2)$.
4. Computes a NIWI proof π for the statement $\psi_{x,c}$, using $(\mathrm{td}_1, \mathrm{td}_2)$ and the decommitment string d for the commitments c, as the witness for the statement.
5. The simulated proof $\widehat{\pi}$ is set to be (c, π).

Decomposition S can be decomposed into $(\mathsf{S}_{\mathsf{pre}}, \mathsf{S}_{\mathsf{pos}})$, where $\mathsf{S}_{\mathsf{pre}}(y = \varphi_\lambda(x), 1^\lambda)$ performs Step 2 of the simulation and outputs $(\mathrm{td}_1, \mathrm{td}_2)$ and $\mathsf{S}_{\mathsf{pos}}(x, (\mathrm{td}_1, \mathrm{td}_2))$ performs steps 3-5. Clearly, $\mathsf{S}_{\mathsf{pos}}$ runs in time $\mathrm{poly}(\lambda, |x|)$. We will show that $\mathsf{S}_{\mathsf{pre}}$ runs in expected time $\mathrm{poly}(\Delta^{-1})$.[6]

Let $y = \varphi_\lambda(x)$. It suffices to show that

$$\Pr_{(\mathrm{td}_1, \mathrm{td}_2) \leftarrow \{0,1\}^{4\ell \times 2}} [(\mathrm{td}_1, \mathrm{td}_2) \in \mathcal{W}_\lambda \times \mathcal{W}_\lambda \mid \mathsf{2Ext}(\mathrm{td}_1, \mathrm{td}_2) = y] \geq \Delta^{O(1)} .$$

[6] To get strict time simulation, we can cutoff the simulation after $\lambda \mathrm{poly}(\Delta^{-1})$ steps, inducing a $2^{-\Omega(\lambda)}$ simulation error.

Indeed, by the definition of density,

$$|\mathcal{W}_\lambda| = \Delta 2^{4\ell} = 2^k, \text{ for } k = 4\ell - \log \Delta^{-1}.$$

Recalling that the extractor $\mathsf{2Ext}$ has error $2^{-\ell-2}$, we have by Lemma 2.3

$$\Pr_{(\mathrm{td}_1,\mathrm{td}_2)\leftarrow\{0,1\}^{4\ell\times 2}}[(\mathrm{td}_1,\mathrm{td}_2)\in\mathcal{W}_\lambda\times\mathcal{W}_\lambda \mid \mathsf{2Ext}(\mathrm{td}_1,\mathrm{td}_2)=y] \geq$$
$$\left(1-2^{-\ell-2}2^{\ell}\right)\Pr_{(\mathrm{td}_1,\mathrm{td}_2)\leftarrow\{0,1\}^{4\ell\times 2}}[(\mathrm{td}_1,\mathrm{td}_2)\in\mathcal{W}_\lambda\times\mathcal{W}_\lambda] = \frac{1}{2}\Delta^2 \ .$$

This establishes the required simulation time.

Indistinguishability We now show the validity of the simulation. We consider the following hybrids:

$\mathcal{H}_{\mathsf{P}}$**:** Here the proof (c,π) is generated as it is generated by the honest prover P — c is a commitment to $0^{8\ell}$ and the NIWI proof for $\psi_{x,c}$ is generated using the witness w for x.

$\mathcal{H}$**:** Here c is generated as a commitment $(\mathrm{td}_1,\mathrm{td}_2)\in\mathcal{W}_\lambda\times\mathcal{W}_\lambda$ such that $\varphi_\lambda(x)=\mathsf{2Ext}(\mathrm{td}_1,\mathrm{td}_2)$. (We've already established the existence of such $\mathrm{td}_1,\mathrm{td}_2$ above.) Then, by the T_{R}-hiding of the commitment

$$\mathcal{H}_{\mathsf{P}} \approx_{T_{\mathsf{R}}} \mathcal{H} \ .$$

$\mathcal{H}_{\mathsf{S}}$**:** Here the NIWI proof π is generated using $(\mathrm{td}_1,\mathrm{td}_2)$ as the witness. Namely, the distribution of proofs is as generated by the simulator S. Then, by the $T_{\mathsf{D}}^{\mathsf{niwi}}$-indistinguishability of the NIWI

$$\mathcal{H} \approx_{T_{\mathsf{D}}^{\mathsf{niwi}}} \mathcal{H}_{\mathsf{S}} \ .$$

Overall, we have established $(T_{\mathsf{S}},T_{\mathsf{D}})$ zero knowledge for $T_{\mathsf{S}}=\Delta^{-1}$ and $T_{\mathsf{D}}=\min\left\{T_{\mathsf{R}},T_{\mathsf{D}}^{\mathsf{niwi}}\right\}$. □

5 One-Message Non-Malleable Commitments

In this section, we construct one-message non-malleable commitments based on our notion of one-message zero-knowledge (and other more standard primitives).

At a high-level, our construction follows the paradigm in previous works [LP11, LPS17, KS17], staring from a basic scheme that is non-malleable w.r.t. commitment or w.r.t. extraction, but only for 4 tags, and increasing the number of tags by iteratively applying a tag-amplification transformation until it reaches an exponential number. More specifically, our *tag-amplification* transformation transforms a commitment scheme for γ tags that is non-malleabile w.r.t. commitment *or* extraction into a commitment scheme for $2^{\tilde{\Omega}(\gamma)}$ tags that is non-malleability w.r.t. both commitment *and* extraction. We plug-in existing four-tag non-malleable commitments w.r.t. extraction from subexponentially secure injective one-way functions and time-lock puzzles [LPS17] to instantiate the basis for the scheme. By performing sufficiently many iterations of the above transformation to the basic four-tag scheme, we obtain our full-fledged non-malleable commitments for an exponential number of tags.

The basic transformation gives rise to one-one (a.k.a standalone) non-malleable commitments. To obtain concurrent non-malleability, we present another *same-tag-concurrency to full-concurrency* transformation, that turn commitments satisfying a weak notion of concurrent non-malleability that we call *same-tag non-malleability* into a fully concurrently non-malleable ones. It turns out that the basic four-tag commitments

by [LPS17] is also same-tag non-malleable w.r.t. extraction, and the previous tag amplification transformation preserves the same-tag non-malleability. Therefore, by applying the same-tag to full-concurrency transformation at the end, we obtain concurrent non-malleability.

Finally, we provide new candidate 4-tag non-malleable commitments (w.r.t. commitment), based on hypothesis related to hardness amplification of one-way functions—that there exists one-way functions with hardcore bits whose unpredictability can be *ideally* amplified by combining (e.g., xoring) many ℓ hardcore bits, that is, attackers under certain time bound cannot predict the combined hardcore bit with probability beyond $1/2 + 2^{-\ell^\varepsilon}$. This type of ideal hardness amplification, however, is outside what current proof techniques can validate. On the other hand, there is no evidence suggesting that natural one-way functions, such as, discrete logarithm, RSA, Rabin do not have this ideal hardness amplification behavior. We thus put forward hypothesis on them, under which we can construct non-malleable commitments for constant number of tags. The new basic candidates give new instantiation of stand-alone non-malleable commitments without time-lock puzzles. However, they do not satisfy the same-tag notion and cannot be used to obtain concurrent non-malleability.

Organization of the section The tag amplification transformation is presented in Section 5.1. The same-tag to full-concurrency transformation is presented in Section 5.2. In Section 5.3, we explain how to start with basic non-malleable commitments for 4 tags and apply different transformations to obtain full-fledged (concurrent) non-malleable commitments, and analyze the complexity of the resulting schemes. In Section 5.4, we recall existing four-tag schemes, and in Section 5.5, we present our new candidates based on the hypothesis of ideal hardness amplification of natural one-way functions.

5.1 Tag Amplification

We show a transformation that turns an input *non-interactive* commitment scheme for γ tags, denoted by iNM, into an output *non-interactive* commitment scheme for $\gamma' = \binom{\gamma}{\gamma/2}$ tags, denoted by oNM; when $\gamma \geq 4$, the number of tags increases. If iNM is non-malleable with respect to *extraction*, the output scheme is non-malleable with respect to *commitment* and *extraction*.

Below we first focus on one-one non-malleability, and then extend to *same-tag non-malleability*—a weaker notion of one-many non-malleability where the attackers are restricted to using the same tag for all right commitments. The latter setting will be important later in Section 5.2 for obtaining concurrent non-malleability.

Our transformation is based on a similar idea to the one used by Khurana and Sahai [KS17]. The main difference is that their transformation makes use of a 2-message Super-Polynomial-time Simulation (SPS) ZK protocol, whereas our input and output protocols are both non-interactive. To achieve this, we use our one-message ZK argument constructed in the previous section to replace the SPS ZK protocol. We show that the weak soundness provided by our one-message ZK argument suffices for the transformation.

Building Blocks Our transformation will make use of several building blocks. Their (subexponential) hardness will be parameterized by constants $\rho, \delta, \varepsilon$ (the relation between these parameters is addressed below). Let λ denote the *global* security parameter, and also an upper bound on the length of the committed strings. We use different building blocks with different security parameters n and $\bar{n}$, both bounded by $\mathrm{poly}(\lambda)$. The use of different security parameter allows to carefully tune the levels of security of different components (see Remark 5.1).

- An input non-interactive commitment scheme $\mathsf{iNM} = (\mathsf{iNM.Com}, \mathsf{iNM.Open})$ for sets $\{\Lambda_n\}_n$ of $\gamma = \gamma(n)$ tags. It is $T_{\mathsf{iNM.E}}$-*over-extractable* by $\mathsf{iNM.E}$ and T_{iNM}-one-one non-malleable with respect to *extraction* by $\mathsf{iNM.E}$, where $T_{\mathsf{iNM.E}}(n) = 2^{2n}$, and $T_{\mathsf{iNM}}(n) \ll T_{\mathsf{iNM.E}}(n)$ can be arbitrary.

- A non-interactive *bit* commitment scheme $(\mathsf{Com}, \mathsf{Open})$ that is T_{R}-hiding and $T_{\mathsf{Com.E}}$-extractable by extractor $\mathsf{Com.E}$ (*without over-extraction*), for $T_{\mathsf{R}}(\bar{n}) = 2^{\bar{n}^\rho}$ and $T_{\mathsf{Com.E}}(\bar{n}) = 2^{\bar{n}}$. We require that the size of commitments $\ell = \ell(\bar{n})$ is linear in the security parameter $\bar{n}$; that is, $\ell(\bar{n}) = O(\bar{n})$. Such a scheme can be constructed from any $2^{\bar{n}^\rho}$-secure family of injective one-way functions, whose output length is linear in the input length, with Goldreich-Levin hardcore bit. See remark 5.2 on why we use a commitment scheme for bits instead of strings.

- A one-message φ-tuned zero-knowledge argument (P, V) for $\varphi = \{\varphi_{\bar{n}} : \{0,1\}^* \to \{0,1\}^{\ell(\bar{n})}\}$, where $\varphi_{\bar{n}}(x) = x[1, \cdots, \ell(\bar{n})]$. It is $(T_{\mathsf{D}}, T_{\mathsf{S}})$-zero-knowledge, and $(T_{\mathsf{P}}, K, \varphi)$-weakly sound, where $T_{\mathsf{D}}(\bar{n}) = T_{\mathsf{P}}(\bar{n}) = 2^{\bar{n}^\varepsilon}$ and $T_{\mathsf{S}}(\bar{n}, |x|) = T_{\mathsf{S}_{\mathsf{pre}}}(\ell(\bar{n})) + \mathrm{poly}(|x|, \bar{n}) = 2^{\ell(\bar{n})^\delta} + \mathrm{poly}(|x|, \bar{n})$. Such a 1ZK argument is constructed in Section 4.

Relation between different components. We require that the building blocks satisfy the relations appearing below on the right. We achieve this by setting the parameters $\varepsilon, \delta, \rho$ and $n, \bar{n}$ as appears below on the left.

$$\begin{array}{c} \bar{n}^{\min(\varepsilon,\rho)/2} = n = \omega(\log \lambda) \\ \rho > \delta > \varepsilon \end{array} \implies \begin{array}{c} T_{\mathsf{P}}(\bar{n}) = T_{\mathsf{D}}(\bar{n}) = 2^{\bar{n}^\varepsilon} \gg 2^{2n} = T_{\mathsf{iNM.E}}(n)\ ; \\ T_{\mathsf{R}}(\bar{n}) = 2^{\bar{n}^\rho} \gg 2^{2n} = T_{\mathsf{iNM.E}}(n)\ ; \\ T_{\mathsf{R}}(\bar{n}) = 2^{\bar{n}^\rho} \gg \Theta(2^{\ell(\bar{n})^\delta}) = T_{\mathsf{S}}(\bar{n}, |x| = \mathrm{poly}(\lambda, \bar{n})) \gg 2^{\bar{n}^\varepsilon} = T_{\mathsf{P}}(\bar{n})\ . \end{array} \tag{3}$$

Remark 5.1 (Use of different security parameters). The reason that we use separate security parameters n and $\bar{n}$ is because later we will apply the provided transformation iteratively many times, where the output scheme NM_i of the i'th iteration is the input scheme for the $(i+1)$'th iteration. The security parameter $\bar{n}_i$ used in NM_i in the i'th iteration is the security parameter $n_{i+1} = \bar{n}_i$ of the input scheme in the $(i+1)$'th iteration. We will set these parameters in such a way that satisfies the above relations. The length λ of the committed string or the global security parameter stays the same throughout all iterations.

Remark 5.2 (On Using Bit Commitments). We remark that the level of sub-exponential hardness of Com, $2^{\bar{n}^\rho}$, is larger than the simulation time of the 1ZK argument, $2^{\bar{n}^\rho} > \Theta(2^{\ell^\delta})$, where ℓ is the length of Com commitments. When Com is a bit commitment scheme, ℓ depends only on the security parameter $\ell = \ell(\bar{n})$. If Com further has linear-size commitments $\ell = O(\bar{n})$ and $\rho > \delta$, the relation in Equation (3) is satisfied for all sufficiently large $\bar{n}$, including the case that $\bar{n} \ll \lambda$. This will be significant later on when in order to amplify the number of tags from 4 to exponential, we iteratively apply the transformation for a super-constant number $L = \omega(1)$ of times, with increasingly large security parameters, $\bar{n}_1 < \bar{n}_2 < \cdots < \bar{n}_L$, satisfying that $\bar{n}_{i+1} = \bar{n}_i^{1/\min(\varepsilon,\rho)/2}$. To ensure that all security parameters are polynomially-bounded in λ, the smallest security parameter $\bar{n}$ has to be $\lambda^{o(1)}$.

We note that using directly string commitments would not have satisfied the above. The length of the commitments would depend, not only on the security parameter $\bar{n}$, but also on the length λ of committed strings, in particular $\ell \geq \lambda$. In this case, if Com is subexponentially secure and $\rho > \delta$, to ensure the relation in Equation (3), we need to set $\bar{n} = \Omega(\lambda)$.

The Transformation Given the above building blocks and functions $n(\lambda), \bar{n}(\lambda)$ satisfying the above conditions, we construct an output scheme $\mathsf{oNM} = (\mathsf{oNM.Com}, \mathsf{oNM.Open})$ as follows: Fix a security parameter $\lambda \in \mathbb{N}$.

The set of tags Λ': Each $\mathsf{tg}' \in \Lambda'$ is a subset of $\gamma/2$ tags of the input scheme, $\mathsf{tg}' = \{\mathsf{tg}_1, \cdots, \mathsf{tg}_{\gamma/2}\} \subset \Lambda$. There are in total $\gamma' = \binom{\gamma}{\gamma/2} \approx \frac{2^\gamma}{\sqrt{\gamma}}$ tags.

Commitment $\mathsf{oNM.Com}(\mathsf{tg}', v, 1^\lambda)$: On input $\mathsf{tg}' = \{\mathsf{tg}_1, \cdots, \mathsf{tg}_{\gamma/2}\}$, and string $v \in \{0,1\}^\lambda$, do:

- For every $i \in [\gamma/2]$, generate an iNM commitment to v using security parameter n and tg_i, $\mathrm{nm}_i \leftarrow \mathsf{iNM.Com}(\mathsf{tg}_i, v, 1^n)$. Let ρ_i be the random coins used.
- For every bit $j \in |v|$, generate a bit commitment to $v[j]$ using security parameter $\bar{n}$, $c_j \leftarrow \mathsf{Com}(v[j], 1^{\bar{n}})$. Let ρ_j^c be the random coins used.
- For every bit $j \in |v|$, generate a 1ZK proof showing that the bit committed in c_j equals to the j'th bit of strings committed in nm_i for all i. More precisely, the statement x_j, witness w_j, and NP relation $\mathcal{R}$ are specified below:

$$
\begin{aligned}
x_j &= (c_j, j, \{\mathrm{nm}_i\}_{i\in[\gamma/2]}) \,, \\
w_j &= ((b, \rho_j^c), \{(s_i, \rho_i)\}_{i\in[\gamma/2]}) \,, \\
\mathcal{R}(x_j, w_j) = 1 \quad &\text{iff} \quad \begin{array}{l} (b, \rho_j^c) \text{ is a valid decommitment to } c_j, \text{ and} \\ \forall i \in [\gamma/2],\ (s_i, \rho_i) \text{ is a valid decommitment to } \mathrm{nm}_i \text{ and } s_i[j] = b \end{array} .
\end{aligned}
$$

  The j'th proof is $\pi_j \leftarrow \mathsf{P}(x, w, 1^{\bar{n}})$. Recall that proof is generated with respect to projection $\varphi_{\bar{n}}(x) = x[1, \cdots, \ell(\bar{n})]$. In particular, $\varphi_{\bar{n}}(x_j)$ outputs the Com commitment c_j in x_j.[7]

Output the commitment nm' and decommitment string d:

$$
\mathrm{nm}' = \left(\mathsf{tg}', \{c_j\}_{j\in[\lambda]}\,,\ \{\mathrm{nm}_i\}_{i\in[\gamma/2]}\,,\ \{\pi_j\}_{j\in[\lambda]}\right) \,, \quad d = \left\{\rho_j^c\right\}_{j\in[\lambda]} \,.
$$

Decommitment $\mathsf{oNM.Open}(\mathrm{nm}', v, d)$**:** It accepts iff the following two conditions are satisfied:

- For every $j \in [\lambda]$, verify whether π_j is accepting $\mathsf{V}(x_j, \pi_j, 1^{\bar{n}}) = 1$.
- For every $j \in [\lambda]$, verify whether $(v[j], \rho_j^c)$ is a valid decommitment to c_j, where ρ_j^c is the j'th string in d, and c_j is the j'th Com commitment in nm'.

Extractor $\mathsf{oNM.E}(\mathrm{nm}')$**:** The extractor does:

- For every $j \in [\lambda]$, verify whether π_j is accepting $\mathsf{V}(x_j, \pi_j, 1^{\bar{n}}) = 1$. Abort and output $\perp$ if any proof is not accepting.
- For every $j \in [\lambda]$, extract the bit $v'[j]$ committed in c_j. Output v'.

Next we show that the above scheme is indeed non-malleable with respect to commitments and with respect to extraction. Later, we will extend the proof below to reason about non-malleability in the (same-tag) concurrent setting.

Theorem 5.1. *If $\varepsilon, \delta, \rho, n, \bar{n}$ satisfy the conditions in Equation (3), then oNM above is $T_{\mathsf{oNM}}(\bar{n})$-non-malleable with respect to commitment and $T_{\mathsf{oNM.E}}(\bar{n})$-extractable by $\mathsf{oNM.E}$ (without over-extraction) for*

$$
\begin{aligned}
&T_{\mathsf{oNM.E}}(\bar{n}) = O(\lambda T_{\mathsf{Com.E}}(\bar{n})) = O(\lambda 2^{\bar{n}}) < 2^{2\bar{n}} \,, \\
&T_{\mathsf{oNM}}(\bar{n}) = T_{\mathsf{iNM}}(n) \,.
\end{aligned}
$$

[7] *Jumping ahead, this means a malicious prover can only cheat on false statements x_j corresponding to a polynomial number of different c_j's.*

Since oNM.E is extractable without over-extraction, non-malleability with respect to commitment directly implies non-malleability with respect to extraction by oNM.E.

Corollary 5.1. *If $\varepsilon, \delta, \rho, n, \bar{n}$ satisfy the conditions in Equation (3), then* oNM *above is $T_{\mathsf{oNM.E}}(\bar{n})$-extractable by* oNM.E *and $T_{\mathsf{oNM}}(n)$-non-malleable with respect to extraction by* oNM.E*, for $T_{\mathsf{oNM.E}}$ and T_{oNM} specified in Theorem 5.1.*

Proof of Theorem 5.1. The run-time of the extractor oNM.E is of order $\lambda T_{\mathsf{Com.E}}$ which is dominated by $2^{2\bar{n}}$ as λ is dominated by $2^{\bar{n}}$.

To show that oNM is $(T_{\mathsf{oNM}}(\bar{n}) = T_{\mathsf{iNM}}(n))$-non-malleable with respect to commitment, fix an arbitrary $\mathrm{poly}(T_{\mathsf{iNM}}(n))$-time (polynomial-sized) non-uniform attacker $\mathcal{A} = \{\mathcal{A}_\lambda\}_\lambda$ that receives one left commitment and gives one right commitment. Also fix two arbitrary ensembles of messages $\{w_\lambda\}_\lambda$, $\{u_\lambda\}_\lambda$ of length λ. We want to show that

$$\left\{\mathsf{mim}^{\mathcal{A}}_{\mathsf{oNM}}(w_\lambda)\right\}_{\lambda\in\mathbb{N}} \approx_c \left\{\mathsf{mim}^{\mathcal{A}}_{\mathsf{oNM}}(u_\lambda)\right\}_{\lambda\in\mathbb{N}} .$$

Towards this, fix a security parameter λ, and let $w = w_\lambda$ and $u = u_\lambda$. For $v \in \{u, w\}$, we construct a sequence of hybrids $\mathcal{H}_0(v), \cdots, \mathcal{H}_3(v)$, where $\mathcal{H}_0(v)$ is an honest experiment with left committed value v and $\mathcal{H}_3(v)$ is independent of v. In each hybrid $\mathcal{H}_i(v)$, we use the same notations as in the construction above to denote components in the left commitment, and use tilde for components in the right commitment. Moreover,

- if the left and right tags are different $\mathsf{tg}' \neq \tilde{\mathsf{tg}}'$, we denote by $i^\star$ the smallest index, such that $\tilde{\mathsf{tg}}_{i^\star} \neq \mathsf{tg}_i$ for all $i \in [\gamma/2]$. Note that such an $i^\star$ exists whenever the left and right tags are different $\mathsf{tg}' \neq \tilde{\mathsf{tg}}'$; otherwise, we let $i^\star = \perp$.
- we denote by $\widetilde{ev}_{i^\star}$ the value extracted from the $i^\star$'th iNM commitment $\widetilde{\mathrm{nm}}_{i^\star}$, that is, $\widetilde{ev}_{i^\star} = \mathsf{iNM.E}(\widetilde{\mathrm{nm}}_{i^\star})$; $\widetilde{ev}_{i^\star} = \perp$ if $i^\star = \perp$ (or equivalently if $\mathsf{tg}' = \tilde{\mathsf{tg}}'$).

We proceed with the description of the hybrids and the corresponding proof of indistinguishability. Below, in hybrid $\mathcal{H}_i(v)$, let $\mathsf{mim}_i(v) = (\mathsf{View}_{\mathcal{A}}, \tilde{v})$ and $\mathsf{diff}_i(v) := (\mathsf{View}_{\mathcal{A}}, \tilde{v}_{i^\star})$.

$\mathcal{H}_0(v)$**:** This hybrid proceeds identically to an honest man-in-the-middle execution with $\mathcal{A}$, where the left committed value is v.

We first prove that to show that the view of $\mathcal{A}$ and the value $\tilde{v}$ it commits to on the right are indistinguishable between the hybrids, it suffices to maintain that the view $\mathsf{View}_{\mathcal{A}}$ of $\mathcal{A}$ and extracted value $\tilde{v}_{i^\star}$ are indistinguishable. This will be done based on the weak soundness of the 1ZK system.

Recall that the 1ZK scheme (P, V) is $(T_{\mathsf{P}}(\bar{n}), K, \varphi)$-weakly sound and the setting of parameters guarantees that $T_{\mathsf{P}}(\bar{n}) \gg T_{\mathsf{iNM.E}}(n)$ (see Equation (3)). We consider a wrapper adversary $\mathcal{A}'_\lambda$ that runs $\mathcal{A}_\lambda$ internally by generating the left commitment to v honestly for $\mathcal{A}_\lambda$ and outputs all 1ZK proofs $\{\tilde{\pi}_j\}$ in the right commitment. $\mathcal{A}'_\lambda$ runs in time $\mathrm{poly}(\lambda, n, \bar{n}) + T_{\mathsf{iNM}}(n) \ll T_{\mathsf{iNM.E}}(n)$ (as $T_{\mathsf{iNM.E}}(n) \gg \lambda + n + \bar{n}$; see Equation (3)). Therefore, by (multi-proof) weak-soundness (Lemma 4.1), there exists sets $\mathcal{Z}(v)$ of size $|\mathcal{Z}(v)| \leq K(|\mathcal{A}'_\lambda| + O(1)) \leq K(\mathrm{poly}(\lambda) + |\mathcal{A}_\lambda|)$, such that the following event occurs with negligible probability $\mu(\bar{n})$:

- **Event** $\mathsf{Cheat}[\mathcal{Z}(v)]$**:** There exists a $j \in [\lambda]$, such that, $\mathcal{A}$ generates an accepting proof $\tilde{\pi}_j$ for a false statement $\tilde{x}_j$, and the projection $\varphi_{\bar{n}}(\tilde{x}_j) = \tilde{c}_j$ is not in the set $\mathcal{Z}(v)$.

Claim 5.1. *There is a negligible function μ and an ensemble of sets $\{\mathcal{Z}(v)\}_\lambda$ of size $\mathrm{poly}(\lambda)$, such that, for every $\lambda \in \mathbb{N}$, the probability that event $\mathsf{Cheat}[\mathcal{Z}(v)]$ occurs in $\mathcal{H}_0(v)$ is at most $\mu(\lambda)$.*

Based on the above claim, we show an efficient non-uniform transformation that takes the view $\mathsf{View}_{\mathcal{A}}$ of $\mathcal{A}$ and the value $\widetilde{ev}_{i^\star}$ extracted from $\widetilde{\mathrm{nm}}_{i^\star}$ and reconstructs the right committed value $\tilde{v}$, using as auxiliary information $\mathsf{Aux}(v)$ that contains the bit committed to in every $\tilde{c}_j \in \mathcal{Z}(v)$ for every j, that is,

$$\mathsf{Aux}(\mathcal{Z}(v)) = \left\{ (c, b) \ : \ \begin{array}{l} b \text{ is the value committed in } c \text{ via } \mathsf{Com} \\ c \in \mathcal{Z}(v) \end{array} \right\} \tag{4}$$

The reconstruction function $\mathsf{Reconst}(\mathsf{View}_{\mathcal{A}}, \widetilde{ev}_{i^\star}, \mathsf{Aux})$, given a set $\mathsf{Aux} \supseteq \mathsf{Aux}(\mathcal{Z}(v))$, outputs $\tilde{v}$ such that:

$$\tilde{v}[j] = \begin{cases} \bot & \text{if } \mathsf{tg}' = \tilde{\mathsf{tg}}' \\ \tilde{b}_j & \text{if } (\tilde{c}_j, \tilde{b}_j) \in \mathsf{Aux} \\ \widetilde{ev}_{i^\star}[j] & \text{otherwise} \end{cases} . \tag{5}$$

Claim 5.2. *For every $\lambda \in \mathbb{N}$, every set $\mathcal{Z}$, and every string $v \in \{0,1\}^\lambda$, conditioned on $\mathsf{Cheat}[\mathcal{Z}]$ not occurring in $\mathcal{H}_0(v)$, it holds that for every $\mathsf{Aux} \supseteq \mathsf{Aux}(\mathcal{Z})$,*

$$\tilde{v} = \mathsf{Reconst}(\mathsf{View}_{\mathcal{A}}, \widetilde{ev}_{i^\star}, \mathsf{Aux}) .$$

Proof. Condition on $\mathsf{Cheat}[\mathcal{Z}]$ not occurring. Observe that $\mathsf{Reconst}$ sets the bit $\tilde{v}[j]$ to $\bot$ correctly when the left and right s are identical, and sets $\tilde{v}[j] = \tilde{b}_j$ correctly to the bit committed in $\tilde{c}_j$ and when the latter appears in Aux.

Second, when neither of the two cases occur, we have that $\tilde{c}_j \notin \mathcal{Z}$. Conditioned on $\mathsf{Cheat}[\mathcal{Z}]$ not occurring, the statement $\tilde{x}_j$ must be true, that is, all $\{\widetilde{\mathrm{nm}}_i\}_i$ are valid, and the j'th bits of their committed strings $\{\tilde{v}_i\}_i$ are equal to the bit $\tilde{b}_j$ committed in $\tilde{c}_j$. In this case, by the over-extractability of iNM, the extracted value is the actual committed value $\tilde{v}_{i^\star} = \mathsf{iNM.E}(\widetilde{\mathrm{nm}}_{i^\star})$. Thus $\widetilde{ev}_{i^\star}$ (defined above) equals $\tilde{v}_{i^\star}$ whenever $\mathsf{tg}' \neq \tilde{\mathsf{tg}}'$. Therefore, $\mathsf{Reconst}$ also outputs $\tilde{v}[j]$ correctly in the third case. □

This means in order to show that the view of $\mathcal{A}$ and its right committed value are indistinguishable in $\mathcal{H}_0(w)$ and $\mathcal{H}_0(u)$, that is, $\{\mathsf{mim}_0(w)\} \approx \{\mathsf{mim}_0(u)\}$, it suffices to show that

$$\begin{aligned} &\{\mathsf{View}_{\mathcal{A}}, \mathsf{Reconst}(\mathsf{View}_{\mathcal{A}}, \widetilde{ev}_{i^\star}, \mathsf{Aux}(\mathcal{Z}(w)) \cup \mathsf{Aux}(\mathcal{Z}(u))) \text{ in } \mathcal{H}_0(w)\} \\ \approx &\{\mathsf{View}_{\mathcal{A}}, \mathsf{Reconst}(\mathsf{View}_{\mathcal{A}}, \widetilde{ev}_{i^\star}, \mathsf{Aux}(\mathcal{Z}(w)) \cup \mathsf{Aux}(\mathcal{Z}(u))) \text{ in } \mathcal{H}_0(u)\} , \end{aligned}$$

Since $|\mathsf{Aux}(\mathcal{Z}(w)) \cup \mathsf{Aux}(\mathcal{Z}(u))| \leq 2K(\mathrm{poly}(\lambda) + |\mathcal{A}_\lambda|) = \mathrm{poly}(\lambda)$, the $\mathsf{Reconst}$ function with inputs as above is computable $\mathrm{poly}(\lambda)$-size circuits. Therefore, it suffices to show that the view of $\mathcal{A}$ and the value $\widetilde{ev}_{i^\star}$ extracted from $\widetilde{\mathrm{nm}}_{i^\star}$ are indistinguishable in $\mathcal{H}_0(w)$ and $\mathcal{H}_0(u)$, that is, $\{\mathsf{diff}_0(w)\} \approx \{\mathsf{diff}_0(u)\}$.

Claim 5.3. $\{\mathsf{diff}_0(w)\}_\lambda \approx \{\mathsf{diff}_0(u)\}_\lambda$ *implies* $\{\mathsf{mim}_0(w)\}_\lambda \approx \{\mathsf{mim}_0(u)\}_\lambda$.

Hence, in the following hybrids, we maintain that $\mathsf{diff}_i(v) \approx \mathsf{diff}_{i+1}(v)$.

$\mathcal{H}_1(v)$**:** This hybrid proceeds identically to $\mathcal{H}_0(v)$ except that all the zero-knowledge proofs $\{\widehat{\pi}_j\}$ in the left commitment are simulated, $\widehat{\pi}_j \leftarrow \mathsf{S}(x_j, 1^{\bar{n}})$. (P, V) is $(T_{\mathsf{D}}, T_{\mathsf{S}})$-zero-knowledge. The setting of parameters guarantees $T_{\mathsf{D}}(\bar{n}) \gg T_{\mathsf{iNM.E}}(n) \gg T_{\mathsf{iNM}}(n)$ (see equation 3). Thus, the indistinguishability of the simulation holds against $\mathcal{A}$ and the extractor iNM.E of iNM. As a result, the view of $\mathcal{A}$ and the values extracted from all iNM commitments on the right are indistinguishable in $\mathcal{H}_0(v)$ and $\mathcal{H}_1(v)$. In particular, this implies $\{\mathsf{diff}_0(v)\} \approx \{\mathsf{diff}_1(v)\}$.

$\mathcal{H}_2(v)$**:** This hybrid proceeds identically to $\mathcal{H}_0(v)$ except that all Com commitments $\{c_j\}$ on the left commits to 0, $c_j \leftarrow \mathsf{Com}(0, 1^{\bar{n}})$. Com is $T_{\mathsf{R}}(\bar{n})$-hiding. The setting of parameters guarantees that $T_{\mathsf{R}}(\bar{n}) \gg T_{\mathsf{S}}(\bar{n})$ and $T_{\mathsf{R}}(\bar{n}) \gg T_{\mathsf{iNM.E}}(n) \gg T_{\mathsf{iNM}}(n)$ (see equation 3). Thus the hiding of Com holds against $\mathcal{A}$ and the extractor iNM.E of iNM, even when all left 1ZK proofs are simulated. Therefore, the view of $\mathcal{A}$ and the values extracted from all iNM commitments on the right are indistinguishable in $\mathcal{H}_1(v)$ and $\mathcal{H}_2(v)$. This implies $\{\mathsf{diff}_1(v)\} \approx \{\mathsf{diff}_2(v)\}$.

$\mathcal{H}_{3,I}(v),\ 0 \le I \le \gamma/2$**:** This hybrid proceeds identically to $\mathcal{H}_2(v)$ except that the first I left iNM commitments $\{\mathrm{nm}_i\}_{i \in I}$ commit to 0, that is, $\mathrm{nm}_i \leftarrow \mathsf{iNM}(\mathsf{tg}_i, 0^\lambda, 1^n)$; the rest $\gamma/2 - I$ left iNM commitments still commit to v. By definition, $\mathcal{H}_{3,0} = \mathcal{H}_2$.

Recall that $i^\star$ is the smallest index such that $\tilde{\mathsf{tg}}_{i^\star} \neq \mathsf{tg}_i$ for all $i \in [\gamma/2]$, and $\perp$ if $\tilde{\mathsf{tg}}' = \mathsf{tg}'$. We claim that by the $T_{\mathsf{iNM}}(n)$-non-malleability of iNM, when changing the value committed to in the I'th left iNM commitment from v (in $\mathcal{H}_{3,I-1}$) to 0^λ (in $\mathcal{H}_{3,I}$), the view of $\mathcal{A}$ and the value $\widetilde{ev}_{i^\star}$ extracted from $\widetilde{\mathrm{nm}}_{i^\star}$ is indistinguishable.

Claim 5.4. *For all sequence* $\{I\}_\lambda$*, where* $1 < I \le \gamma/2$*,* $\{\mathsf{diff}_{3,I-1}(v)\}_\lambda \approx \{\mathsf{diff}_{3,I}(v)\}_\lambda$

Proof. Note that we cannot directly apply $T_{\mathsf{iNM}}(n)$-non-malleability, since in hybrids $\{\mathcal{H}_{3,I}\}$ the left 1ZK arguments are simulated in time polynomial in $T_{\mathsf{S}}(\bar{n}) \gg T_{\mathsf{P}}(\bar{n}) > T_{\mathsf{iNM}}(n)$ (see equation 3). To circumvent this, we rely on the fact that (P, V) is φ-tuned, in particular, simulation $\mathsf{S}(x_j, 1^{\bar{n}})$ consists of a preprocessing step $\mathsf{st}_j \leftarrow \mathsf{S}_{\mathsf{pre}}(c_j, 1^{\bar{n}})$ that depends only on the projection $\varphi(x_j) = c_j$ of the statement (and takes superpolynomial time), and a postprocessing step $\widehat{\pi}_j \leftarrow \mathsf{S}_{\mathsf{pos}}(x_j, \mathsf{st}_j)$ of $\mathrm{poly}(|x_j|, \bar{n})$ time.

Suppose toward contradiction that $\mathsf{diff}_{3,I-1}(v)$ and $\mathsf{diff}_{3,I}(v)$ are distinguishable. Since in both hybrids $\mathcal{H}_{3,I-1}(v)$ and $\mathcal{H}_{3,I-1}(v)$, all Com commitments $\{c_j\}$ on the left commit to zero, there must exist a fixed set of commitments $\{c_j^\star\}$ and a corresponding set of preprocessing states $\{\mathsf{st}_j\}$, such that conditioned on $\{c_j^\star\}$ occurring and $\{\mathsf{st}_j\}$ used in simulation, $\mathsf{diff}_{3,I-1}(v)$ and $\mathsf{diff}_{3,I}(v)$ are distinguishable. Moreover, conditioned on $\{c_j^\star, \mathsf{st}_j\}$, the only difference between $\mathcal{H}_{3,I-1}(v)$ and $\mathcal{H}_{3,I-1}(v)$ is the value committed to in nm_I on the left, and post-precessing of simulation takes only polynomial time using $\{\mathsf{st}_j\}$. Therefore, we can non-uniformly fix $\{c_j^\star, \mathsf{st}_j\}$, and get a $T_{\mathsf{iNM}}(n)$-adversary that breaks the non-malleability of nm_I, which is a contradiction. It follows that $\mathsf{diff}_{3,I-1}(v)$ and $\mathsf{diff}_{3,I}(v)$ are indistinguishable as required. □

Finally, note that hybrid $\mathcal{H}_{3,\gamma/2}(v)$ is independent of v, therefore $\{\mathsf{diff}_{3,\gamma/2}(w)\} \equiv \{\mathsf{diff}_{3,\gamma/2}(u)\}$.

It follows from a hybrid argument that $\{\mathsf{diff}_0(w)\} \approx \{\mathsf{diff}_0(u)\}$. By Claim 5.3, we have $\{\mathsf{mim}_0(w)\} \approx \{\mathsf{mim}_0(u)\}$. □

Extending to Same-Tag Non-Malleability The above proof focuses on showing that our transformation preserves non-malleability in the stand-alone setting (or even slightly strengthens it to exact extractability), starting from non-malleability with respect to extraction and ending with non-malleability with respect to both extraction and commitment. Essentially the same proof applies to show that our transformation also preserves non-malleability in a restricted one-many setting, *where the man-in-the-middle uses the same tags for all of its right commitments.* We call this *same-tag non-malleability*. (In fact, the transformation also preserves non-malleability in fully concurrent setting, since we do not need this property in this work, we focus on the same-tag non-malleability.)

Theorem 5.2. *If $\varepsilon, \delta, \rho, n, \bar{n}$ satisfy the conditions in Equation (3), and* iNM *is $T_{\mathsf{iNM.E}}$-over-extractable by* iNM.E *and T_{iNM}-same-tag-non-malleable with respect to extraction by* iNM.E*, then* oNM *above is $T_{\mathsf{oNM}}(\bar{n})$-same-tag non-malleable with respect to commitment and $T_{\mathsf{oNM.E}}(\bar{n})$-extractable by* oNM.E *(without over-extraction) for*

$$T_{\mathsf{oNM.E}}(\bar{n}) = O(\lambda T_{\mathsf{Com.E}}(\bar{n})) = O(\lambda 2^{\bar{n}}) < 2^{2\bar{n}} ,$$
$$T_{\mathsf{oNM}}(\bar{n}) = T_{\mathsf{iNM}}(n) .$$

Proof Sketch. The complexity of the extractor stays the same as analyzed above in proof of Theorem 5.1.

To show T_{oNM}-same-tag-non-malleability with respect to commitment, we follow the same steps. Consider the same hybrids $\mathcal{H}_0(v), \cdots \mathcal{H}_3(v)$ as defined above, but now each hybrid contains a one-many man-in-the-middle execution with $\mathcal{A}_\lambda$, where the right tags are the same $\tilde{\mathsf{tg}}'$. Note that whenever $\mathsf{tg}' \neq \tilde{\mathsf{tg}}'$, there is an index $i^\star$, such that, $\tilde{\mathsf{tg}}' = (\tilde{\mathsf{tg}}_1, \cdots \tilde{\mathsf{tg}}_{\gamma/2})$, $\mathsf{tg}' = (\mathsf{tg}_1, \cdots \mathsf{tg}_{\gamma/2})$, and $\tilde{\mathsf{tg}}_{i^\star} \neq \mathsf{tg}_i$ for all i. Therefore, in every right commitment k, the $i^\star$'th iNM commitment, denoted as $\widetilde{\mathsf{nm}}^k_{i^\star}$, uses tag $\mathsf{tg}_{i^\star} \neq \mathsf{tg}_i$ for all i. Denote by $\widetilde{ev}^k_{i^\star}$ the value extracted from this commitment, which is set to $\perp$ when $\mathsf{tg}' = \tilde{\mathsf{tg}}'$.

The rest of proof follows the same blueprint as before, except that, we reason about the set of extracted values $\{\widetilde{ev}^k_{i^\star}\}_k$. First, by relying on the $(T_{\mathsf{P}}, K, \varphi)$-weak-soundness of (P, V), we can show that the indistinguishability of the view of $\mathcal{A}$ and the values it commits to on the right, $\mathsf{mim}_0(w) \approx \mathsf{mim}_0(u)$, reduces to the indistinguishability of the view of $\mathcal{A}$ and the set of extracted values $\{\widetilde{ev}^k_{i^\star}\}_k$, that is, $\mathsf{diff}_0(w) \approx \mathsf{diff}_0(u)$. In hybrids $\mathcal{H}_1(v)$ and $\mathcal{H}_2(v)$, it follows again from the $(T_{\mathsf{D}}(\bar{n}), T_{\mathsf{S}}(\bar{n}))$-zero-knowledge property and the $T_{\mathsf{R}}(\bar{n})$-hiding of Com that $\mathsf{diff}_0(v) \approx \mathsf{diff}_1(v)$, and $\mathsf{diff}_1(v) \approx \mathsf{diff}_2(v)$. The proof is identical to before as it essentially uses complexity leveraging and it does not matter if there are multiple values to extract or the right or just one. In hybrids $\{\mathcal{H}_{3,I}(v)\}$, we now rely on the $T_{\mathsf{iNM}}(n)$-same-tag-non-malleability of iNM to show that $\mathsf{diff}_{3,I}(v) \approx \mathsf{diff}_{3,I+1}(v)$. We conclude the proof by observing that $\mathcal{H}_{3,\gamma/2}(w) = \mathcal{H}_{3,\gamma/2}(u)$, and thus $\mathsf{diff}_3(w) = \mathsf{diff}_3(u)$ and $\mathsf{diff}_0(w) \approx \mathsf{diff}_0(u)$. □

5.2 Same-Tag Concurrency to Full Concurrency

In this section, we present another transformation from an input scheme iNM for γ tags to an output scheme oNM for $\gamma' = \gamma - 2$ tags. If iNM satisfies *same-tag* non-malleability with respect to commitment, then oNM is concurrently non-malleable with respect to commitment.

In the literature, a similar non-malleability strengthening transformation was presented in [LPS17]. Their transformation uses sub-exponentially secure NIWI, time-lock puzzles, non-interactive commitments *but is only secure against uniform attackers*. Our transformation handles *non-uniform* attackers and is more general. In particular, we start with the notion of same-tag non-malleability that can be instantiated from different assumptions other than time lock puzzles. Differently from the previous transformation, here we will not be able to use our 1ZK arguments as a black box, but would rather rely on specific properties of our 1ZK construction.

Building Blocks Our transformation will use the following building blocks, all sub-exponentially secure with respect to parameters $\rho, \delta, \varepsilon$. Let λ denote the *global* security parameter, and also an upper bound on the length of the committed strings. As in the previous section, we will use different building blocks with different security parameters $n, \bar{n}$, both $\mathrm{poly}(\lambda)$-bounded.

- An input non-interactive commitment scheme $\mathsf{iNM} = (\mathsf{iNM.Com}, \mathsf{iNM.Open})$ for a set $\{\Lambda_n\}_n$ of $\gamma = \gamma(n)$ tags. It is $(T_{\mathsf{iNM.E}}(n) = 2^{2n})$-extractable by $\mathsf{iNM.E}$ (without over-extraction) and is $(T_{\mathsf{iNM}}(n) = 2^{n^\varepsilon})$-same-tag-non-malleable with respect to commitment.

- A non-interactive *bit* commitment scheme $(\mathsf{Com}, \mathsf{Open})$ that is $(T_{\mathsf{R}}(\bar{n}) = 2^{\bar{n}^\rho})$-hiding and $T_{\mathsf{Com.E}}(\bar{n}) = 2^{\bar{n}}$-extractable by extractor $\mathsf{Com.E}$ (without over-extraction). As in the previous section, we require that the size of commitments is linear in the security parameter $\ell(\bar{n}) = O(\bar{n})$.

- In Section 4, we presented a construction of φ-tuned 1ZK argument from an incompressible problem $\mathcal{W}$, a NIWI, and a non-interactive commitment scheme. As mentioned, we will not be able to use 1ZK arguments as a black box below. Roughly speaking, first, we need the 1ZK arguments to be *simulation sound*. We achieve this by instantiating the commitments inside 1ZK with the input non-malleable commitments, such that, when receiving simulated arguments using one tag, the man-in-the-middle cannot cheat in arguments using a different tag. Second, we need the 1ZK arguments to be "non-malleable" w.r.t. two commitment schemes. This can be achieved by letting the two commitment schemes be iNM with two fixed tags $\mathsf{tg}_1^\star$ $\mathsf{tg}_2^\star$, and using only other tags $\mathsf{tg} \neq \mathsf{tg}_1^\star, \mathsf{tg}_2^\star$ in the 1ZK arguments. This guarantees that when switching an 1ZK argument on the left from being honest to simulated, the values the attacker commits to on the right using iNM with $\mathsf{tg}_1^\star$ or $\mathsf{tg}_2^\star$ does not change. Similarly, when switching the value committed in a commitment of iNM with $\mathsf{tg}_1^\star$ or $\mathsf{tg}_2^\star$, the probability that the attacker cheats in 1ZK arguments on the right remain almost the same.

  Specifically, we instantiate the construction using the input commitment scheme iNM (the other components stay the same), where NIWI uses security parameter $\bar{n}$, and iNM uses security parameter n. Recalling that iNM uses tags in Λ, we accordingly extend the interface of 1ZK argument (P, V) and its simulator S to additionally accept a tag $\mathsf{tg} \in \Lambda_n$, that is,

$$\pi \leftarrow \mathsf{P}(\mathsf{tg}, x, w, 1^{\bar{n}}),\ b = \mathsf{V}(\mathsf{tg}, x, \pi, 1^{\bar{n}}),\ \widehat{\pi} \leftarrow \mathsf{S}(\mathsf{tg}, x, 1^{\bar{n}})\ .$$

  Fix a projection $\varphi = \{\varphi_{\bar{n}}(x) = x[1, \cdots, \ell(\bar{n})]\}$. We require that the underlying components have the following levels of security,

  - iNM is $T_{\mathsf{iNM}}(n)$-hiding and $T_{\mathsf{iNM.E}}(n)$-extractable as specified above,
  - NIWI is $(T_{\text{NIWI}}(\bar{n}) = 2^{\bar{n}^\varepsilon})$-indistinguishable, and
  - $\mathcal{W}$ is $(T_{\mathsf{iNM.E}}(n), K_{\mathcal{W}})$-incompressible with density $\Delta(d) = 2^{-d^\alpha}$[8].

  By the proof in Section 4, for every fixed sequence of tags $\mathsf{tg} = \{\mathsf{tg}_n\}$, the corresponding $\{(\mathsf{P}(\mathsf{tg}_n, \star), \mathsf{V}(\mathsf{tg}_n, \star))\}_n$ is $(T_{\mathsf{P}}(n), K, \varphi)$-weakly sound for $K = O(K_{\mathcal{W}})$ and the corresponding simulator $\{\mathsf{S}(\mathsf{tg}_n, \star)\}$ runs in time polynomial in $T_{\mathsf{S}}(\bar{n})$, where for appropriate δ,

$$T_{\mathsf{P}}(n) = T_{\mathsf{iNM.E}}(n) > T_{\mathsf{iNM}}(n)\ ,\ T_{\mathsf{S}}(\ell(\bar{n}))) = T_{\mathsf{S}_{\mathsf{pre}}}(\ell(\bar{n})) + \mathrm{poly}(|x|, \bar{n}) = 2^{\ell(\bar{n})^\delta} + \mathrm{poly}(|x|, \bar{n})\ . \tag{6}$$

[8] The length of solutions of $\mathcal{W}$ used is determined by the length of outputs of the projection, namely, $d = O(\ell(\bar{n}))$; see Section 4.

Relation between different components. We require the building blocks to satisfy the following relations, which is achieved if $\varepsilon, \delta, \rho$ and $n, \bar{n}$ satisfy the following conditions:

$$\begin{array}{c} \bar{n}^{\min(\varepsilon,\rho)/2} = n = \omega(\log \lambda) \\ \rho > \delta \end{array} \quad \Longrightarrow \quad \begin{array}{c} T_{\text{NIWI}} = 2^{\bar{n}^{\varepsilon}} \gg 2^{2n} = T_{\mathsf{iNM.E}} \; ; \\ T_{\mathsf{R}}(\bar{n}) = 2^{\bar{n}^{\rho}} \gg 2^{n} = T_{\mathsf{iNM.E}}(n) \; ; \\ T_{\mathsf{R}}(\bar{n}) = 2^{\bar{n}^{\rho}} \gg \Theta(2^{\ell(\bar{n})^{\delta}}) = T_{\mathsf{S}}(\bar{n}, |x| = \mathrm{poly}(\lambda, \bar{n})) \; . \end{array} \tag{7}$$

The Transformation Given the above building blocks and polynomially-bounded functions $n, \bar{n}$ satisfying the above conditions, we construct an output scheme $\mathsf{oNM} = (\mathsf{oNM.Com}, \mathsf{oNM.Open})$ as follows: Fix a security parameter $\lambda \in \mathbb{N}$.

The Set of Tags Λ': Let $\mathsf{tg}_0^\star$ and $\mathsf{tg}_1^\star$ be two arbitrary tags in Λ. The set of tags of oNM is $\Lambda' = \Lambda - \{\mathsf{tg}_0^\star, \mathsf{tg}_1^\star\}$. There are in total $\gamma' = \gamma - 2$ tags.

Commitment $\mathsf{oNM.Com}(\mathsf{tg}, v, 1^\lambda)$: On input a $\mathsf{tg} \in \Lambda'$, and string $v \in \{0,1\}^\lambda$[9], do:

- Generate two iNM commitments $\mathrm{nm}_0, \mathrm{nm}_1$ to v using tags $\mathsf{tg}_0^\star, \mathsf{tg}_1^\star$ and security parameter n, that is, for $i \in \{0,1\}$ $\mathrm{nm}_i \leftarrow \mathsf{iNM.Com}(\mathsf{tg}_i^\star, v, 1^n)$. Let ρ_i be the random coins used.
- For every bit $j \in |v|$, generate a bit commitment to $v[j]$ using security parameter $\bar{n}$, $c_j \leftarrow \mathsf{Com}(v[j], 1^{\bar{n}})$. Let ρ_j^c be the random coins used.
- For every bit $j \in |v|$, generate a 1ZK proof showing that the bit committed in c_j equals to the j'th bit of strings committed in $\mathrm{nm}_0, \mathrm{nm}_1$. More precisely, the statement x_j, witness w_j, and NP relation $\mathcal{R}$ are:

$$\begin{array}{rcl} x_j & = & (c_j, j, \mathrm{nm}_0, \mathrm{nm}_1) \; , \\ w_j & = & ((b, \rho_j^c), \{(s_i, \rho_i)\}_{i \in \{0,1\}}) \; , \\ \mathcal{R}(x_j, w_j) = 1 & \text{iff} & \begin{array}{l} (b, \rho_j^c) \text{ is a valid decommitment to } c_j, \text{ and} \\ \forall i \in \{0,1\}, (s_i, \rho_i) \text{ is a valid decommitment to } \mathrm{nm}_i \text{ and } s_i[j] = b \end{array} \; . \end{array}$$

  Recall that proofs are generated with respect to $\varphi_{\bar{n}}(x) = x[1, \cdots, \ell(\bar{n})]$. In particular, $\varphi_{\bar{n}}(x_j)$ outputs the Com commitment c_j in x_j.

  Furthermore, every proof π_j is generated using iNM with tg as the underlying commitment, that is, $\pi_j \leftarrow \mathsf{P}(\mathsf{tg}, x, w, 1^{\bar{n}})$. More precisely, π_j consists the following components:
  - A iNM commitment $\mathrm{nm}_{3,j}$ to 0^L of appropriate length L, using tg, $\mathrm{nm}_{3,j} \leftarrow \mathsf{iNM.Com}(\mathsf{tg}, 0^L, 1^n)$;
  - A NIWI proof wi_j that
    - either, x_j is true—we refer to this as *the honest statement,*
    - or, $\mathrm{nm}_{3,j}$ commits to a "trapdoor" which is two solutions $(\mathrm{td}_1, \mathrm{td}_2)$ of appropriate length to the incompressible problem $\mathcal{W}$, and $\mathsf{2Ext}(\mathrm{td}_1, \mathrm{td}_2) = \varphi_{\bar{n}}(x_j) = c_j$—we refer to this as *the cheating statement.*

Output the commitment nm' and decommitment string d:

$$\mathrm{nm}' = \left(\mathsf{tg}, \{c_j\}_{j \in [\lambda]}, \; \mathrm{nm}_0, \mathrm{nm}_1, \; \{\pi_j = (\mathrm{nm}_{3,j}, \mathrm{wi}_j)\}_{j \in [\lambda]}\right) \; , \quad d = \left\{\rho_j^c\right\}_{j \in [\lambda]} \; .$$

[9]We describe the scheme with respect to input string v of length exactly λ. It is easy to see how the scheme works with shorter strings.

Decommitment $\mathsf{oNM.Open}(\mathrm{nm}', v, d)$**:** It accepts iff the following two conditions are satisfied:

- For every $j \in [\lambda]$, verify whether π_j is accepting $\mathsf{V}(\mathsf{tg}, x_j, \pi_j, 1^{\bar{n}}) = 1$.
- For every $j \in [\lambda]$, verify whether $(v[j], \rho_j^c)$ is a valid decommitment to c_j, where ρ_j^c is the j'th string in d, and c_j is the j'th Com commitment in nm'.

Extractor $\mathsf{oNM.E}(\mathrm{nm}')$**:** The extractor does:

- For every $j \in [\lambda]$, verify whether π_j is accepting $\mathsf{V}(\mathsf{tg}, x_j, \pi_j, 1^{\bar{n}}) = 1$. Abort and output $\perp$ if any proof is not accepting.
- For every $j \in [\lambda]$, extract the bit $v'[j]$ committed in c_j. Output v'.

We show that the scheme is concurrently non-malleable w.r.t commitments.

Theorem 5.3. *If $\varepsilon, \delta, \rho, n, \bar{n}$ satisfy the conditions in Equation (7), then* oNM *above is $T_{\mathsf{oNM}}(n)$-non-malleable with respect to commitment and $T_{\mathsf{oNM.E}}(\bar{n})$-extractable by* $\mathsf{oNM.E}$ *for*

$$T_{\mathsf{oNM.E}}(\bar{n}) = O(\lambda T_{\mathsf{Com.E}}(\bar{n})) = O(\lambda 2^{\bar{n}}) \ll 2^{2\bar{n}} ,$$
$$T_{\mathsf{oNM}}(\bar{n}) = T_{\mathsf{iNM}}(n) .$$

Proof. The run-time of the extractor $\mathsf{oNM.E}$ is of order $\lambda T_{\mathsf{Com.E}}$ which is dominated by $2^{2\bar{n}}$ as $\lambda \ll 2^{\bar{n}}$.

By Lemma 2.1, to show that oNM is $T_{\mathsf{oNM}}(\bar{n}) = T_{\mathsf{iNM}}(n)$-concurrent non-malleable with respect to commitment, it suffices to show that it is $T_{\mathsf{oNM}}(\bar{n}) = T_{\mathsf{iNM}}(n)$-one-many non-malleable with respect to commitment.

Fix any $\mathrm{poly}(T_{\mathsf{iNM}}(n))$-time (polynomial-sized) non-uniform attacker $\mathcal{A} = \{\mathcal{A}_\lambda\}$ that receives one left commitment and gives many, $m = m(\lambda)$, right commitment (under many different tags). Also fix two arbitrary ensembles of messages $\{w_\lambda\}_\lambda$, $\{u_\lambda\}_\lambda$ of length λ. We want to show that

$$\left\{\mathsf{mim}^{\mathcal{A}}_{\mathsf{oNM}}(w_\lambda)\right\}_{\lambda\in\mathbb{N}} \approx_c \left\{\mathsf{mim}^{\mathcal{A}}_{\mathsf{oNM}}(u_\lambda)\right\}_{\lambda\in\mathbb{N}} .$$

Towards this, fix a security parameter λ, and $w = w_\lambda$, $u = u_\lambda$. We construct hybrids $\mathcal{H}_0(v), \cdots, \mathcal{H}_4(v)$ for $v \in \{w, u\}$ Each hybrid contains a man-in-the-middle execution with $\mathcal{A}_\lambda$, where it receives a single left commitment, and gives many m right commitments. We denote components in the left commitment using notations as in the construction, and use tilde and superscript k for components in the right commitment $k \in [m]$. In particular,

- We denote by $\tilde{v}_0^k, \tilde{v}_1^k$ the values committed in the iNM commitments $\widetilde{\mathrm{nm}}_0^k, \widetilde{\mathrm{nm}}_1^k$ using tags $\mathsf{tg}_0^\star, \mathsf{tg}_1^\star$ in the kth right commitment. Note that $\tilde{v}_b^k = \mathsf{iNM.E}(\widetilde{\mathrm{nm}}_b^k)$, as $\mathsf{iNM.E}$ does not over-extract.
- We denote by $\tilde{v}_{3,j}^k$ the value committed in the iNM commitment $\widetilde{\mathrm{nm}}_{3,j}^k$ contained in the j'th proof $\tilde{\pi}_j^k$ in the kth right commitment, and again $\tilde{v}_{3,j}^k = \mathsf{iNM.E}(\widetilde{\mathrm{nm}}_{3,j}^k)$; it is set to $\perp$ if $\mathsf{tg} = \mathsf{tg}^k$.

Throughout all hybrids, we will maintain that certain so called cheating events never occur, except with negligible probability. Conditioned on them not occurring, the right committed values $\{\tilde{v}^k\}$ can be reconstructed from either $\{\tilde{v}_0^k\}$ or $\{\tilde{v}_1^k\}$. This is the key for showing that the view and the right committed values are indistinguishable in neighboring hybrids.

Invariants in All Hybrids: We maintain that in every hybrid the following cheating condition occurs with negligible probability, with respect to some polynomial-sized set $\mathcal{Z}$ that depends on the hybrid. We say that a right commitment is successful if it is accepting and has a tag $\widetilde{\mathsf{tg}}^k$ different from that of the left commitment.

- **Event** $\mathsf{Cheat}'[\mathcal{Z}]$**:** In some successful right commitment $k \in [m]$, there is a proof $\tilde{\pi}_j^k = (\widetilde{\mathrm{nm}}_{3,j}^k, \widetilde{\mathrm{wi}}_j^k)$ for $j \in [\lambda]$, such that, the value committed in the iNM commitment $\tilde{v}_{3,j}^k = \mathsf{iNM.E}(\widetilde{\mathrm{nm}}_{3,j}^k)$ is a "trapdoor"—that is, $\tilde{v}_{3,j}^k = (\mathrm{td}_1, \mathrm{td}_2)$ are two solutions of the incompressible problem $\mathcal{W}$ s.t. $\mathsf{2Ext}(\mathrm{td}_1, \mathrm{td}_2) = \tilde{c}_j^k$—but $\tilde{c}_j^k \notin \mathcal{Z}$.

We observe that if $\mathsf{Cheat}'[\mathcal{Z}]$ occurs with negligible probability, then the following event $\mathsf{Cheat}[\mathcal{Z}]$ also occurs with negligible probability.

- **Event** $\mathsf{Cheat}[\mathcal{Z}]$**:** There exists a successful right commitment $k \in [m]$ and proof $\tilde{\pi}_j^k$ for $j \in [\lambda]$, such that, $\tilde{\pi}_j^k$ is an accepting proof for a false statement $\tilde{x}_j^k$, and $\tilde{c}_j^k \notin \mathcal{Z}$.

Claim 5.5. *For every $\lambda \in \mathbb{N}$, in any man-in-the-middle execution with $\mathcal{A}_\lambda$, and for every set $\mathcal{Z}$, conditioned on $\mathsf{Cheat}'[\mathcal{Z}]$ not occurring with respect to the right commitments that $\mathcal{A}_\lambda$ sends, then $\mathsf{Cheat}[\mathcal{Z}]$ does not occur with respect to these right commitments either.*

Proof. To show that $\mathsf{Cheat}[\mathcal{Z}]$ does not occur, we show that for every k, j, whenever $\tilde{\pi}_j^k = (\widetilde{\mathrm{nm}}_{3,j}^k, \widetilde{\mathrm{wi}}_j^k)$ is an accepting proof for a false statement $\tilde{x}_j^k$, it holds that $\tilde{c}_j^k \in \mathcal{Z}(v)$. By the perfect soundness of the NIWI, when $\widetilde{\mathrm{wi}}_j^k$ is accepting, its statement is true. That is, Either 1) $\tilde{x}_j^k$ is true, or 2) $\widetilde{\mathrm{nm}}_{3,j}^k$ is a valid commitment to a trapdoor. Given that the statement $\tilde{x}_j^k$ is false, $\widetilde{\mathrm{nm}}_{3,j}^k$ must be a valid commitment to a trapdoor. Since $\mathsf{Cheat}'[\mathcal{Z}]$ does not occur, this implies that $\tilde{c}_j^k \in \mathcal{Z}(\lambda)$. □

Furthermore, conditioned on event $\mathsf{Cheat}[\mathcal{Z}]$ not occurring, we claim that the committed values $\tilde{v}^k$ in any right commitment can be reconstructed from the committed value $\tilde{v}_b^k$ of either of iNM commitments $\widetilde{\mathrm{nm}}_0^k, \widetilde{\mathrm{nm}}_1^k$, using auxiliary information that contains the bits committed in all commitments in $\mathcal{Z}(v)$ (identical to that in Equation 4), that is,

$$\mathsf{Aux}(\mathcal{Z}(v)) = \left\{ (c, b) \ : \ \begin{array}{l} b \text{ is the bit committed to in } c \text{ via } \mathsf{Com} \\ c \in \mathcal{Z}(v) \end{array} \right\}$$

The reconstruction procedure $\mathsf{Reconst}(\mathsf{View}_{\mathcal{A}}, \tilde{v}_b^k, \mathsf{Aux})$ using auxiliary information $\mathsf{Aux} \supseteq \mathsf{Aux}(\mathcal{Z}(v))$ outputs $\tilde{v}^k$ satisfying the following (identical to that in Equation 5).

$$\tilde{v}^k[j] = \begin{cases} \perp & \text{if } \mathsf{tg} = \widetilde{\mathsf{tg}} \\ \tilde{b}_j & \text{if } (\tilde{c}_j, \tilde{b}_j) \in \mathsf{Aux} \\ \tilde{v}_b^k[j] & \text{otherwise} \end{cases} .$$

Claim 5.6. *For every $\lambda \in \mathbb{N}$, in any man-in-the-middle execution with $\mathcal{A}_\lambda$, conditioned on $\mathsf{Cheat}[\mathcal{Z}]$ not occurring with respect to the right commitments that $\mathcal{A}_\lambda$ sends, it holds that for every $\mathsf{Aux} \supseteq \mathsf{Aux}(\mathcal{Z})$, and every right commitment $k \in [m]$,*

$$\tilde{v}^k = \mathsf{Reconst}(\mathsf{View}_{\mathcal{A}}, \tilde{v}_0^k, \mathsf{Aux}) = \mathsf{Reconst}(\mathsf{View}_{\mathcal{A}}, \tilde{v}_1^k, \mathsf{Aux}) .$$

Proof. Fix an arbitrary $b \in \{0, 1\}$. We show that conditioned on $\mathsf{Cheat}[\mathcal{Z}]$ not occurring on the right, $\tilde{v}^k = \mathsf{Reconst}(\mathsf{View}_{\mathcal{A}}, \tilde{v}_b^k, \mathsf{Aux})$ for every k. First, observe that $\mathsf{Reconst}$ sets the bit $\tilde{v}^k[j]$ to $\perp$ correctly when the left and right tags are identical, and sets $\tilde{v}^k[j] = \tilde{b}_j^k$ correctly to the bit committed in $\tilde{c}_j^k$ when the latter appears in Aux.

Second, when neither of the two cases occur, we have that $\tilde{c}_j^k \notin \mathcal{Z}$. Conditioned on $\mathsf{Cheat}[\mathcal{Z}]$ not occurring, the statement $\tilde{x}_j^k$ must be true, that is, $\widetilde{\mathrm{nm}}_0^k$ and $\widetilde{\mathrm{nm}}_1^k$ are valid commitments to strings $\tilde{v}_0^k$ and $\tilde{v}_1^k$ and the bit committed in $\tilde{c}_j^k$ is equal to both $\tilde{v}_0^k[j]$ and $\tilde{v}_1^k[j]$. Thus $\mathsf{Reconst}$ also outputs $\tilde{v}^k[j]$ correctly in the third case. □

Let us briefly describe how these invariants will be used in the hybrids below. Consider hybrid experiments $\mathsf{Exp}_0, \mathsf{Exp}_1$. Assume that in each experiment Exp_β, event $\mathsf{Cheat}'[\mathcal{Z}_\beta]$ almost never occurs with respect to some polynomial-size set $\mathcal{Z}_\beta$. Then, to show that the view of $\mathcal{A}_\lambda$ and the values it commits to on the right are indistinguishable in these two experiments, it suffices to show that the view of $\mathcal{A}_\lambda$ and values it commits to in the iNM commitments using $\mathsf{tg}_b^\star$ for either $b = 0$ or $b = 1$ are indistinguishable.

$$\begin{aligned}&\{\mathsf{View}_{\mathcal{A}}, \left\{\mathsf{Reconst}(\mathsf{View}_{\mathcal{A}}, \tilde{v}_b^k, \mathsf{Aux}(\mathcal{Z}_0) \cup \mathsf{Aux}(\mathcal{Z}_1))\right\}_k \text{ in } \mathsf{Exp}_0\}\\ \approx&\{\mathsf{View}_{\mathcal{A}}, \left\{\mathsf{Reconst}(\mathsf{View}_{\mathcal{A}}, \tilde{v}_b^k, \mathsf{Aux}(\mathcal{Z}_0) \cup \mathsf{Aux}(\mathcal{Z}_1))\right\}_k \text{ in } \mathsf{Exp}_1\} \ ,\end{aligned}$$

Since $|\mathcal{Z}_0 \cup \mathcal{Z}_1|$ is of polynomial size, the $\mathsf{Reconst}$ function with inputs as above is (non-uniformly) computable by a polynomial-size circuit. Therefore, it suffices to show that

$$\{\mathsf{View}_{\mathcal{A}}, \{\tilde{v}_b^k\}_k \text{ in } \mathsf{Exp}_0\} \approx \{\mathsf{View}_{\mathcal{A}}, \{\tilde{v}_b^k\}_k \text{ in } \mathsf{Exp}_1\} \ .$$

We use this approach to show the indistinguishability of neighboring hybrids — through showing the indistinguishability of view and values in the iNM commitments using $\mathsf{tg}_b^\star$, for some $b \in \{0, 1\}$. In hybrid $\mathcal{H}_i(v)$, we denote by $\mathsf{nmz}_i(v)$ the tuple $(\mathsf{View}_{\mathcal{A}}, \{\tilde{v}_0^k\}_k)$, $\mathsf{nmo}_i(v)$ the tuple $(\mathsf{View}_{\mathcal{A}}, \{\tilde{v}_1^k\}_k)$, and $\mathsf{mim}_i(v)$ the tuple $(\mathsf{View}_{\mathcal{A}}, \{\tilde{v}^k\}_k)$.

We proceed with the description of hybrids.

$\mathcal{H}_0(v)$**:** This hybrid proceeds identically to an honest main-in-the-middle execution with $\mathcal{A}$, where the left committed value is v.

We show that there exists a polynomial-sized set $\mathcal{Z}(v)$, such that, event $\mathsf{Cheat}'[\mathcal{Z}(v)]$ occurs with negligible probability.

Claim 5.7. *There is a negligible function μ and an ensemble of sets $\{\mathcal{Z}(v)\}_\lambda$ of size $\mathrm{poly}(\lambda)$, such that, for every $\lambda \in \mathbb{N}$, the probability that event $\mathsf{Cheat}'[\mathcal{Z}(v)]$ occurs in $\mathcal{H}_0(v)$ is at most $\mu(\lambda)$.*

Proof. Consider a wrapper adversary $\mathcal{A}'_\lambda$ that runs $\mathcal{A}_\lambda$ internally by generating the left commitment to v honestly for $\mathcal{A}_\lambda$ and outputs all 1ZK proofs $\left\{\tilde{\pi}_j^k\right\}_{k\in[m],j\in[\lambda]}$. $\mathcal{A}'_\lambda$ runs in time $\mathrm{poly}(\lambda, n, \bar{n}) + T_{\mathsf{iNM}}(n) \ll T_{\mathsf{iNM.E}}(n)$ (as $T_{\mathsf{iNM.E}}(n) \gg \lambda + n + \bar{n}$; see Equation (7)), and $|\mathcal{A}'_\lambda| = |\mathcal{A}_\lambda| + \mathrm{poly}(\lambda)$. The $(T_{\mathsf{iNM.E}}(n), K_{\mathcal{W}})$-incompressibility of $\mathcal{W}$ implies that there is a set $\mathcal{Z}(v)$, such that, with overwhelming probability, for every k, j, either the value extracted $\tilde{v}_{3,j}^k = \mathsf{iNM.E}(\widetilde{\mathrm{nm}}_{3,j}^k)$ is not a valid trapdoor, or it is a valid trapdoor and $\tilde{c}_j^k \in \mathcal{Z}(v)$. In addition, the size of the set is bounded by $K(|\mathcal{A}'_\lambda|)$, which is polynomial. (The details of this argument closely follow the proof of weak-soundness (P, V) in Theorem 4.1 and its multi-proof variant in Lemma 4.1.) □

$\mathcal{H}_1(v)$**:** This hybrid proceeds identically to $\mathcal{H}_0(v)$ except that all the zero-knowledge proofs $\{\widehat{\pi}_j^k\}$ in the left commitment are simulated, $\widehat{\pi}_j^k \leftarrow \mathsf{S}(\mathsf{tg}, x_j, 1^{\bar{n}})$. We follow the simulation procedure and consider a sequence of sub-hybrids, $\{\mathcal{H}_{1,J}(v)\}_{J\in[2\lambda]}$:

- $\mathcal{H}_{1,J}(v)$ for $0 \leq J \leq \lambda$. In this hybrid, for every $j \in [J]$, find a trapdoor by running the preprocessing stage of simulation $(\mathrm{td}_{1,j}, \mathrm{td}_{2,j}) \leftarrow \mathsf{S}_{\mathsf{pre}}(\varphi(x_j) = c_j, 1^{\bar{n}})$. Then, commit to the trapdoor in the corresponding iNM commitment, $\mathrm{nm}_{3,j} \leftarrow \mathsf{iNM}(\mathsf{tg}, (\mathrm{td}_{1,j}, \mathrm{td}_{2,j}), 1^n)$. The rest of the iNM commitments $\{\mathrm{nm}_{3,j}\}_{J<j\leq\lambda}$ commit to zero as in $\mathcal{H}_0(v)$. By definition $\mathcal{H}_{1,0} = \mathcal{H}_0$.

  We show the following two claims with respect to $\mathcal{H}_{1,\lambda}(v)$.

  Claim 5.8. *Let $\{\mathcal{Z}(v)\}_\lambda$ be the ensemble of polynomial-sized sets established by Claim 5.7. There is a negligible function μ, such that, for every $\lambda \in \mathbb{N}$, the probability that event $\mathsf{Cheat}'[\mathcal{Z}(v)]$ occurs in $\mathcal{H}_{1,\lambda}(v)$ is at most $\mu(\lambda)$.*

  Claim 5.9. *The view of $\mathcal{A}$ and values it commits to on the right are indistinguishable in $\mathcal{H}_0(v)$ and $\mathcal{H}_{1,\lambda}(v)$,*

$$\{\mathsf{mim}_0(v)\}_\lambda \approx \{\mathsf{mim}_{1,\lambda}(v)\}_\lambda \,.$$

  Proof of Claim 5.8. Suppose for contradiction that $\mathsf{Cheat}'[\mathcal{Z}(v)]$ occurs with some inverse polynomial probability $1/p(\lambda)$ in $\mathcal{H}_{1,\lambda}(v)$. Then, following Claim 5.7, there must exist $0 \leq J < \lambda$, such that, the probabilities that $\mathsf{Cheat}'[\mathcal{Z}(v)]$ occurs in $\mathcal{H}_{1,J}$ and $\mathcal{H}_{1,J+1}$ differ by at least $1/(p(\lambda)\lambda)$. By definition of $\mathsf{Cheat}'[\mathcal{Z}(v)]$ there must exist $k \in [m]$ and $j \in [\lambda]$, such that, the probabilities that $\mathsf{Cheat}'[\mathcal{Z}(v)]$ occurs with respect to this particular k, j differ by at least $1/(p(\lambda)\lambda^2 m)$. Recall that $\mathsf{Cheat}'[\mathcal{Z}(v)]$ occurs with respect to k, j when 1) the right commitment k is accepting and has a different tag $\mathsf{tg} \neq \widetilde{\mathsf{tg}}^k$ from the left commitment, 2) the committed value $\tilde{v}^k_{3,j}$ is a trapdoor, 3) but $\tilde{c}^k_j \notin \mathcal{Z}(v)$.

  On the other hand, the only difference between $\mathcal{H}_{1,J}$ and $\mathcal{H}_{1,J+1}$ lies in the value committed in the iNM commitment $\mathrm{nm}_{3,J}$ on the left. Recall that $\mathrm{nm}_{3,J}$ uses tg of the left commitment, while $\widetilde{\mathrm{nm}}^k_{3,j}$ uses $\widetilde{\mathsf{tg}}^k$ of the right commitment k. Since the probabilities that $\mathsf{Cheat}'[\mathcal{Z}(v)]$ occurs with respect to k, j are inverse polynomially far, we get a contradiction to the T_{iNM}-one-one non-malleability with respect to commitment of iNM according to which the view of $\mathcal{A}$ and the value committed in $\widetilde{\mathrm{nm}}^k_{3,j}$ should remain indistinguishable when changing $\mathrm{nm}_{3,J}$, as long as $\mathsf{tg} \neq \widetilde{\mathsf{tg}}^k$. □

  Proof of Claim 5.9. To prove the claim, we show that the view of $\mathcal{A}$ and values $\{\tilde{v}^k_0\}_k$ committed in $\{\widetilde{\mathrm{nm}}^k_0\}$ on the right are indistinguishable in $\mathcal{H}_0(v)$ and $\mathcal{H}_{1,\lambda}(v)$, that is,

$$\{\mathsf{nmz}_0(v)\}_\lambda \approx \{\mathsf{nmz}_{1,\lambda}(v)\}_\lambda \,.$$

  Suppose the above holds. By Claim 5.7 and 5.8, the probabilities that $\mathsf{Cheat}'[\mathcal{Z}(v)]$ occur in $\mathcal{H}_0(v)$ and $\mathcal{H}_{1,\lambda}(v)$ are negligible. By Claim 5.5 and 5.6, conditioned on $\mathsf{Cheat}'[\mathcal{Z}(v)]$ not occurring, we have that for every right commitment k,

$$\tilde{v}^k = \mathsf{Reconst}(\mathsf{View}_{\mathcal{A}}, \tilde{v}^k_0, \mathcal{Z}(v)) \,.$$

  Therefore, if $\mathsf{nmz}_0(v)$ and $\mathsf{nmz}_{1,\lambda}(v)$ are indistinguishable, so are $\mathsf{mim}_0(v)$ and $\mathsf{mim}_{1,\lambda}(v)$.

  Suppose for contradiction that there exists a distinguisher D, that distinguishes $\mathsf{nmz}_0(v)$ and $\mathsf{nmz}_{1,\lambda}(v)$ with inverse polynomial advantage $1/p(\lambda)$. There must exist $0 \leq J < \lambda$ such that, D distinguishes $\mathsf{nmz}_{1,J}(v)$ and $\mathsf{nmz}_{1,J+1}(v)$ with inverse polynomial advantage $1/p(\lambda)\lambda$. Note that the only difference between $\mathcal{H}_{1,J}$ and $\mathcal{H}_{1,J+1}$ lies in the value committed in the iNM commitment $\mathrm{nm}_{3,J}$ on the left. Recall that $\mathrm{nm}_{3,J}$ uses tg of the left commitment, while commitments $\{\widetilde{\mathrm{nm}}^k_0\}$

on the right use $\mathsf{tg}_0^\star \neq \mathsf{tg}$. Therefore, it follows from the T_{iNM}-same-tag-non-malleability with respect to commitment of iNM that the view of $\mathcal{A}$ and the values $\{\tilde{v}_0^k\}$ committed in $\{\widetilde{\mathrm{nm}}_0^k\}$ are indistinguishable, which concludes the proof. □

- $\mathcal{H}_{1,J'+\lambda}(v)$ for $J' \in [\lambda]$. This hybrid proceeds identically to $\mathcal{H}_{1,\lambda}$, except that, the first J' NIWI proofs $\{\mathrm{wi}_j\}_{j\in[J']}$ are generated by proving that the corresponding iNM commitments $\{\mathrm{nm}_{3,j}\}_{j\in[J']}$ commit to trapdoors. The rest of the NIWI proofs $\{\mathrm{wi}_j\}_{J'<j\leq\lambda}$ are generated by proving that the honest statements $\{x_j\}_{J'<j\leq\lambda}$ are true. We show the following

  Claim 5.10. *Let $\{\mathcal{Z}(v)\}_\lambda$ be the ensemble of polynomial-sized sets established by Claim 5.7. There is a negligible function μ, such that, for every $\lambda \in \mathbb{N}$, the probability that event $\mathsf{Cheat}'[\mathcal{Z}(v)]$ occurs in $\mathcal{H}_{1,2\lambda}(v)$ is at most $\mu(\lambda)$.*

  Proof. Suppose for contradiction that $\mathsf{Cheat}'[\mathcal{Z}(v)]$ occurs with some inverse polynomial probability $1/p(\lambda)$ in $\mathcal{H}_{1,2\lambda}(v)$. As argued in proof of Claim 5.8, there must exist $0 \leq J' < \lambda$, $k \in [m]$ and $j \in \lambda$, such that, the probabilities that $\mathsf{Cheat}'[\mathcal{Z}(v)]$ occurs with respect to this particular k, j in $\mathcal{H}_{1,\lambda+J'}(v)$ and $\mathcal{H}_{1,\lambda+J'+1}(v)$ differ by at least $1/(p(\lambda)\lambda^2 m)$. Recall that $\mathsf{Cheat}'[\mathcal{Z}(v)]$ occurs with respect to k, j when 1) the right commitment k is accepting and has a different tag $\mathsf{tg} \neq \widetilde{\mathsf{tg}}^k$ from the left commitment, 2) the committed value $\tilde{v}_{3,j}^k = \mathsf{iNM.E}(\widetilde{\mathrm{nm}}_{3,j}^k)$ is a trapdoor, 3) but $\tilde{c}_j^k \notin \mathcal{Z}(v)$.
  On the other hand, the only difference between $\mathcal{H}_{1,\lambda+J'}$ and $\mathcal{H}_{1,\lambda+J'+1}$ lies which witness is used for generating the J''th NIWI proof $\mathrm{wi}_{J'}$ on the left. The indistinguishability of NIWI holds against $T_{\text{NIWI}}(\bar{n})$-time distinguishers. By the setting of parameters, we have that $T_{\text{NIWI}}(\bar{n}) \gg T_{\mathsf{iNM.E}}(n) \gg T_{\mathsf{iNM}}(n)$ (see Equation 7). Therefore, the view of $\mathcal{A}$ and the value $\tilde{v}_{3,j}^k = \mathsf{iNM.E}(\widetilde{\mathrm{nm}}_{3,j}^k)$ committed in $\widetilde{\mathrm{nm}}_{3,j}^k$ are indistinguishable in $\mathcal{H}_{1,\lambda+J'}$ and $\mathcal{H}_{1,\lambda+J'+1}$. This implies that the probabilities that $\mathsf{Cheat}'[\mathcal{Z}(v)]$ occurs with respect to k, j are negligibly close, which gives a contradiction. □

  Claim 5.11. *The view of $\mathcal{A}$ and values it commits to on the right are indistinguishable in $\mathcal{H}_{1,\lambda}(v)$ and $\mathcal{H}_{1,2\lambda}(v)$,*

$$\{\mathsf{mim}_{1,\lambda}(v)\}_\lambda \approx \{\mathsf{mim}_{1,2\lambda}(v)\}_\lambda \,.$$

  Proof. By Claim 5.8 and 5.10, in both hybrids $\mathcal{H}_{1,\lambda}(v)$ and $\mathcal{H}_{1,2\lambda}(v)$, the probabilities that $\mathsf{Cheat}'[\mathcal{Z}(v)]$ occurs are negligible. Thus, following the same argument as in the proof of Claim 5.9, to show the claim, it suffices to show that the view of $\mathcal{A}$ and values $\{\tilde{v}_0^k\}_k$ committed in $\{\widetilde{\mathrm{nm}}_0^k\}$ on the right are indistinguishable in $\mathcal{H}_{1,\lambda}(v)$ and $\mathcal{H}_{1,2\lambda}(v)$, that is,

$$\{\mathsf{nmz}_{1,\lambda}(v)\}_\lambda \approx \{\mathsf{nmz}_{1,2\lambda}(v)\}_\lambda \,.$$

  Suppose for contradiction that there exists a distinguisher D, that distinguishes them with inverse polynomial advantage $1/p(\lambda)$. There must exist $0 \leq J' < \lambda$ such that, D distinguishes $\mathsf{nmz}_{1,\lambda+J'}(v)$ and $\mathsf{nmz}_{1,\lambda+J'+1}(v)$ with inverse polynomial advantage $1/p(\lambda)\lambda$.
  Note that the only difference between $\mathcal{H}_{1,\lambda+J'}$ and $\mathcal{H}_{1,\lambda+J'+1}$ lies in the witness used for generating the J''th NIWI proof $\mathrm{wi}_{J'}$ on the left. By the $T_{\text{NIWI}}(\bar{n})$-indistinguishability of NIWI and the fact that $T_{\text{NIWI}}(\bar{n}) \gg T_{\mathsf{iNM.E}}(n) \gg T_{\mathsf{iNM}}(n)$ (see Equation 7), we have that the view of $\mathcal{A}$ and the values $\{\tilde{v}_0^k\}_k$ committed in $\{\widetilde{\mathrm{nm}}_0^k\}_k$ are indistinguishable, which gives a contradiction. □

$\mathcal{H}_2(v)$**:** This hybrid proceeds identically to $\mathcal{H}_0(v)$ except that all Com commitments $\{c_j\}$ on the left commits to 0, $c_j \leftarrow \mathsf{Com}(0, 1^{\bar{n}})$. Com is hiding against $\mathrm{poly}(T_{\mathsf{R}}(\bar{n}))$ time adversaries. The setting of parameters guarantees that $T_{\mathsf{R}}(\bar{n}) \gg T_{\mathsf{S}}(\bar{n})$ and $T_{\mathsf{R}}(\bar{n}) \gg T_{\mathsf{iNM.E}}(n) \gg T_{\mathsf{iNM}}(n)$ (see equation 7). Thus the hiding of Com holds against $\mathcal{A}$ and the extractor $\mathsf{iNM.E}$ of iNM, even when all the left 1ZK proofs are simulated. Therefore, the view of $\mathcal{A}$ and the values $\{\tilde{v}_0^k, \tilde{v}_1^k, \tilde{v}_{3,j}^k\}_{k,j}$ it commits to in all iNM commitments $\{\widetilde{\mathrm{nm}}_0^k, \widetilde{\mathrm{nm}}_1^k, \widetilde{\mathrm{nm}}_{3,j}^k\}_{k,j}$ on the right are indistinguishable in $\mathcal{H}_{1,2\lambda}(v)$ and $\mathcal{H}_2(v)$. This directly implies that the probabilities that $\mathsf{Cheat}'[\mathcal{Z}(v)]$ occurs in these two hybrids differ at most by a negligible amount. Thus,

Claim 5.12. *Let $\{\mathcal{Z}(v)\}_\lambda$ be the ensemble of polynomial-sized sets established by Claim 5.7. There is a negligible function μ, such that, for every $\lambda \in \mathbb{N}$, the probability that event $\mathsf{Cheat}'[\mathcal{Z}(v)]$ occurs in $\mathcal{H}_2(v)$ is at most $\mu(\lambda)$.*

Given that $\mathsf{Cheat}'[\mathcal{Z}(v)]$ almost never occurs in either $\mathcal{H}_{1,2\lambda}(v)$ or $\mathcal{H}_2(v)$, by the same argument as in Claim 5.9, to show the indistinguishability of $\mathsf{mim}_{1,2\lambda}(v)$ and $\mathsf{mim}_2(v)$, it suffices to show that of $\mathsf{nmz}_{1,2\lambda}(v)$ and $\mathsf{nmz}_2(v)$. The latter follows directly from the fact that the view of $\mathcal{A}$ and the values committed in all iNM commitments on the right are indistinguishable.

Claim 5.13. *The view of $\mathcal{A}$ and values it commits to on the right are indistinguishable in $\mathcal{H}_{1,2\lambda}(v)$ and $\mathcal{H}_2(v)$,*

$$\{\mathsf{mim}_{1,2\lambda}(v)\}_\lambda \approx \{\mathsf{mim}_2(v)\}_\lambda \,.$$

$\mathcal{H}_3(v)$**:** This hybrid proceeds identically to $\mathcal{H}_2(v)$ except that the left iNM commitment nm_1 using $\mathsf{tg}_1^\star$ commits to 0^λ, $\mathrm{nm}_1 \leftarrow \mathsf{iNM}(\mathsf{tg}_1^\star, 0^\lambda, 1^n)$. Note that $\mathsf{tg}_1^\star$ used for nm_1 on the left is different from $\mathsf{tg}_0^\star$ used for $\widetilde{\mathrm{nm}}_0^k$, and $\{\mathsf{tg}^k\}_k$ for $\{\widetilde{\mathrm{nm}}_{3,j}^k\}_{k,j}$ on the right. We show that by the $T_{\mathsf{iNM}}(n)$-one-one non-malleability with respect to commitment of iNM, the probability that event $\mathsf{Cheat}'[\mathcal{Z}]$ occurs is negligible in $\mathcal{H}_3(v)$, and by the $T_{\mathsf{iNM}}(n)$-same-tag non-malleability with respect to commitment of iNM, the view of $\mathcal{A}$ and right committed values are indistinguishable in $\mathcal{H}_2(v)$ and $\mathcal{H}_3(v)$.

Claim 5.14. *Let $\{\mathcal{Z}(v)\}_\lambda$ be the ensemble of polynomial-sized sets established by Claim 5.7. There is a negligible function μ, such that, for every $\lambda \in \mathbb{N}$, the probability that event $\mathsf{Cheat}'[\mathcal{Z}(v)]$ occurs in $\mathcal{H}_3(v)$ is at most $\mu(\lambda)$.*

Claim 5.15. *The view of $\mathcal{A}$ and values it commits to on the right are indistinguishable in $\mathcal{H}_2(v)$ and $\mathcal{H}_3(v)$,*

$$\{\mathsf{mim}_2(v)\}_\lambda \approx \{\mathsf{mim}_3(v)\}_\lambda \,.$$

Proof of Claim 5.14. We cannot directly apply $T_{\mathsf{iNM}}(n)$-one-one non-malleability with respect to commitment of iNM, since in hybrids $\mathcal{H}_2(v)$ and $\mathcal{H}_3(v)$ the left 1ZK arguments are simulated in time polynomial in $T_{\mathsf{S}}(\bar{n}) \gg T_{\mathsf{iNM}}(n)$ (see equation 6). To circumvent this, we rely on the fact that (P, V) is φ-tuned. In particular, simulation $\mathsf{S}(x_j)$ consists of a preprocessing step $\mathsf{st}_j \leftarrow \mathsf{S}_{\mathsf{pre}}(c_j, 1^{\bar{n}})$ that depends only on the projection $\varphi(x_j) = c_j$ of the statement (and takes superpolynomial time), and a postprocessing step $\widehat{\pi}_j \leftarrow \mathsf{S}_{\mathsf{pos}}(x_j, \mathsf{st}_j)$ that is $\mathrm{poly}(|x_j|, \bar{n})$-time.

Suppose for contradiction that $\mathsf{Cheat}'[\mathcal{Z}(v)]$ occurs with some inverse polynomial probability $1/p(\lambda)$ in $\mathcal{H}_3(v)$. As shown in Claim 5.12, this event occurs with only negligible probability in $\mathcal{H}_2(v)$. Thus, there must be $k \in [m]$ and $j \in \lambda$, such that, the probabilities that $\mathsf{Cheat}'[\mathcal{Z}(v)]$ occurs with respect to

this particular k, j in $\mathcal{H}_2(v)$ and $\mathcal{H}_3(v)$ differ by at least $1/(p(\lambda)\lambda m)$. Recall that $\mathsf{Cheat}'[\mathcal{Z}(v)]$ occurs with respect to k, j when 1) the right commitment k is accepting and has a different tag $\mathsf{tg} \neq \widetilde{\mathsf{tg}}^k$ from the left commitment, 2) the committed value $\tilde{v}^k_{3,j}$ is a trapdoor, 3) but $\tilde{c}^k_j \notin \mathcal{Z}(v)$.

Since the only difference between $\mathcal{H}_2(v)$ and $\mathcal{H}_3(v)$ is the value committed to in nm_1 using $\mathsf{tg}^\star_1$, there must exist a fixed set of commitments $\{c^\star_j\}$ and corresponding preprocessed states $\{\mathsf{st}_j\}$, such that conditioned on $\{c^\star_j\}$ occurring and $\{\mathsf{st}_j\}$ used in simulation in $\mathcal{H}_2(v)$ and $\mathcal{H}_3(v)$, the probabilities that $\mathsf{Cheat}'[\mathcal{Z}(v)]$ occurs with respect to k, j still differ by at least $1/(p(\lambda)\lambda m)$. Conditioned on $\{c^\star_j, \mathsf{st}_j\}$, postprocessing of simulation can be performed in polynomial time. Therefore, since the probabilities of $\mathsf{Cheat}'[\mathcal{Z}(v)]$ occurring inverse-polynomially far, we can get a non-uniform adversary that breaks the $T_{\mathsf{iNM}}(n)$-one-one non-malleability with respect to commitment of iNM by distinguishing the view of $\mathcal{A}$ and the value $\tilde{v}^k_{3,j}$ committed in $\widetilde{\mathrm{nm}}^k_{3,j}$ using $\mathsf{tg}^k \neq \mathsf{tg}^\star_1$, according to the value of nm_1. □

Proof of Claim 5.15. By Claim 5.12 and 5.14, in both hybrids $\mathcal{H}_{2,\lambda}(v)$ and $\mathcal{H}_3(v)$, the probabilities that $\mathsf{Cheat}'[\mathcal{Z}(v)]$ occurs are negligible. Therefore, as argued before, it suffices to show that the view of $\mathcal{A}$ and values $\{\tilde{v}^k_0\}_k$ committed in $\{\widetilde{\mathrm{nm}}^k_0\}$ on the right are indistinguishable, that is,

$$\{\mathsf{nmz}_2(v)\}_\lambda \approx \{\mathsf{nmz}_3(v)\}_\lambda .$$

Suppose for contradiction that, $\mathsf{nmz}_2(v)$ and $\mathsf{nmz}_3(v)$ are distinguishable. By the same argument as in proof of Claim 5.14, there must exist a fixed set of commitments $\{c^\star_j\}$ and preprocessed states $\{\mathsf{st}_j\}$ related to them, such that, conditioned on $\{c^\star_j\}$ occurring and $\{\mathsf{st}_j\}$ used in simulation in $\mathcal{H}_2(v)$ and $\mathcal{H}_3(v)$, $\mathsf{nmz}_2(v)$ and $\mathsf{nmz}_3(v)$ are still distinguishable. Conditioned on $\{c^\star_j, \mathsf{st}_j\}$, postprocessing of simulation of 1ZK arguments can be done efficiently using $\{\mathsf{st}_j\}$, and the only difference between $\mathcal{H}_2(v)$ and $\mathcal{H}_3(v)$ is the value committed in nm_1 using $\mathsf{tg}^\star_1$ on the left. Therefore, it follows from the $T_{\mathsf{iNM}}(n)$-same-tag non-malleability with respect to commitment of iNM that the view of $\mathcal{A}$ and values it commits to in $\{\widetilde{\mathrm{nm}}^k_0\}$ using $\mathsf{tg}^\star_0 \neq \mathsf{tg}^\star_1$ do not change, which gives a contradiction. □

$\mathcal{H}_4(v) = \mathcal{H}_3(0^\lambda)$**:** This hybrid proceeds identically to $\mathcal{H}_3(v)$ except that the left iNM commitment nm_0 using $\mathsf{tg}^\star_0$ commits to 0^λ, $\mathrm{nm}_0 \leftarrow \mathsf{iNM.Com}(\mathsf{tg}^\star_0, 0^\lambda, 1^n)$. By Claim 5.14, we thus immediately have that,

Claim 5.16. *Let $\{\mathcal{Z}(0^\lambda)\}_\lambda$ be the ensemble of polynomial-sized sets established by Claim 5.7 with respect to string 0^λ. There is a negligible function μ, such that, for every $\lambda \in \mathbb{N}$, the probability that event $\mathsf{Cheat}'[\mathcal{Z}(0^\lambda)]$ occurs in $\mathcal{H}_3(0^\lambda)$ is at most $\mu(\lambda)$.*

Note that $\mathsf{tg}^\star_0$ used for nm_0 on the left is different from $\mathsf{tg}^\star_1$ used for $\{\widetilde{\mathrm{nm}}^k_1\}_k$ on the right. We show that by $T_{\mathsf{iNM}}(n)$-same-tag non-malleability with respect to commitment of iNM, the view of $\mathcal{A}$ and right committed values are indistinguishable in $\mathcal{H}_3(v)$ and $\mathcal{H}_3(0^\lambda)$.

Claim 5.17. *The view of $\mathcal{A}$ and values it commits to on the right are indistinguishable in $\mathcal{H}_3(v)$ and $\mathcal{H}_3(0^\lambda)$,*

$$\{\mathsf{mim}_3(v)\}_\lambda \approx \{\mathsf{mim}_3(0^\lambda)\}_\lambda .$$

Proof. By Claim 5.14 the probability of $\mathsf{Cheat}'[\mathcal{Z}(v)]$ occurring in hybrid $\mathcal{H}_3(v)$ is negligible, and by 5.16, the probability of $\mathsf{Cheat}'[\mathcal{Z}(0^\lambda)]$ occurring in hybrid $\mathcal{H}_3(0^\lambda)$ is negligible. By Claim 5.5

and 5.6, conditioned on $\mathsf{Cheat}'[\mathcal{Z}(v)]$ not occurring in $\mathcal{H}_3(v)$ and $\mathsf{Cheat}'[\mathcal{Z}(0^\lambda)]$ not occurring in $\mathcal{H}_3(0^\lambda)$, we have that in both hybrids for every right commitment k,

$$\tilde{v}^k = \mathsf{Reconst}(\mathsf{View}_{\mathcal{A}}, \tilde{v}_1^k, \mathcal{Z}(v) \cup \mathcal{Z}(0^\lambda)) \ .$$

Therefore, it suffices to show that the view of $\mathcal{A}$ and values $\{\tilde{v}_1^k\}_k$ committed in $\{\widetilde{\mathrm{nm}}_1^k\}$ on the right are indistinguishable, that is,

$$\{\mathsf{nmo}_3(v)\}_\lambda \approx \{\mathsf{nmo}_3(0^\lambda)\}_\lambda \ .$$

The only difference between $\mathcal{H}_3(v)$ and $\mathcal{H}_3(0^\lambda)$ is the value committed in nm_0 using $\mathsf{tg}_1^\star$ on the left. By the same argument as in proof of Claim 5.15, the $T_{\mathsf{iNM}}(n)$-same-tag non-malleability with respect to commitment of iNM implies that the view of $\mathcal{A}$ and values it commits to in $\{\widetilde{\mathrm{nm}}_1^k\}$ using $\mathsf{tg}_1^\star \neq \mathsf{tg}_0^\star$ do not change, which concludes the claim. □

Finally, it follows from a hybrid argument that the view of $\mathcal{A}$ and the right committed values are indistinguishable in $\mathcal{H}_0(v)$ and $\mathcal{H}_4(v) = \mathcal{H}_3(0^\lambda)$. Since this holds for any $v \in \{u, w\}$, we have that $\{\mathsf{mim}_0(w)\} \approx \{\mathsf{mim}_0(u)\}$.

□

5.3 From Non-Malleability for 4 Tags to Full-Fledged Non-Malleability

Starting from the building blocks in Section 5.1, where the *initial* input commitment scheme iNM has 4 tags $|\Lambda_\lambda| = 4$ and is subexponentially non-malleable with respect to extraction, we show how to iteratively apply tag-amplification scheme for $L = O(\log^* \lambda)$ times to obtain a sequence of output commitment schemes $\mathsf{NM}_0, \mathsf{NM}_1, \cdots, \mathsf{NM}_L$, where the *final* output scheme $\mathsf{oNM} = \mathsf{NM}_L$ has at least 2^λ tags and is non-malleable with respect to commitment.

Non-Malleability with respect to Extraction for 4 Tags to Non-Malleability with respect to Commitments Fix a global security parameter λ, which also bounds the length of committed strings. Proceed as follows:

Base Case: When $i = 1$, NM_0 is set to the initial input scheme iNM. It has $\gamma_0 = 4$ tags, and is $T_{\mathsf{NM.E}_0}$-*over-extractable* and T_{NM_0}-one-one non-malleable with respect to *extraction* by $\mathsf{NM.E}_0$, where

$$T_{\mathsf{NM.E}_0}(\bar{n}_0) = 2^{2\bar{n}_0} \ , \qquad T_{\mathsf{NM}_0}(\bar{n}_0) = 2^{\bar{n}_0^\varepsilon} \text{ for some } \varepsilon > 0 \ .$$
$$\text{Set } \ \bar{n}_0 = (\log \lambda)^{2/\varepsilon} \ , \qquad \text{then } T_{\mathsf{NM}_0}(\bar{n}_0) = 2^{\log^2 \lambda} \ .$$

Invariant: For every iteration $i-1$, the output scheme NM_{i-1} of iteration $i-1$ (or the base case $i = 1$) has γ_{i-1} tags, and is $T_{\mathsf{NM.E}_{i-1}}$-*over-extractable* and $T_{\mathsf{NM}_{i-1}}$-one-one non-malleable w.r.t *extraction* by $\mathsf{NM.E}_{i-1}$, where

$$T_{\mathsf{NM.E}_{i-1}}(\bar{n}_{i-1}) = 2^{2\bar{n}_{i-1}} \ , \qquad T_{\mathsf{NM}_{i-1}}(\bar{n}_{i-1}) = T_{\mathsf{NM}_0}(\bar{n}_0) = 2^{\log^2 \lambda} \ .$$

Iteration i: Apply the tag-transformation to NM_{i-1}, using security parameter $\bar{n}_{i-1}$ for NM_{i-1} and security parameter $\bar{n}_i$ for other building blocks Com and the 1ZK arguments (P, V), where

$$\bar{n}_i = \bar{n}_{i-1}^{1/\alpha} = \bar{n}_0^{1/\alpha^i} = (\log \lambda)^{2/\varepsilon\alpha^i} \ , \text{ for } \alpha = \min(\varepsilon, \rho)/2$$

NM_i has $\gamma_i = \binom{\gamma_{i-1}}{\gamma_{i-1}/2}$ tags. By Theorem 5.1, the output scheme NM_i is T_{NM_i}-non-malleable with respect to commitment and $T_{\mathsf{NM.E}_i}$-extractable by $\mathsf{NM.E}_i$ (without over-extraction) for

$$T_{\mathsf{NM.E}_i}(\bar{n}_i) < 2^{2\bar{n}_i} ,$$
$$T_{\mathsf{NM}_i}(\bar{n}_i) = T_{\mathsf{iNM}_{i-1}}(\bar{n}_{i-1}) = T_{\mathsf{NM}_0}(\bar{n}_0) = 2^{\log^2 \lambda} .$$

Furthermore, by Corollary 5.1, NM_i is also $T_{\mathsf{NM}_i}(\bar{n}_i)$-non-malleable with respect to extraction by $\mathsf{NM.E}_i$. Therefore NM_i satisfies the above invaraint.

Final Output Scheme: Whenever NM_i has at least $\gamma_i \geq 2^\lambda$ tags, terminate the iteration and outputs $\mathsf{oNM} = \mathsf{NM}_i$. Since the number of tags grows exponentially, the total number L of iterations is bounded by $O(\log^* \lambda)$. Therefore, the security parameter $\bar{n}_L$ of the final scheme is bounded by a polynomial in λ.

On Growth of the Complexity of Schemes. If naively apply the tag amplification for a super-constant number of times, the complexity of the output schemes would grow by a polynomial factor each time, and by a super-polynomial factor overall. To see this, recall that a scheme, say NM_i, output by the transformation consists of λ commitments of Com, $\gamma_{i-1}/2$ commitments of NM_{i-1}, and λ 1ZK arguments for the statements $\{x_j\}_j$ that the j'th bits of strings committed using NM_i are consistent with the bit committed in the j'th Com commitment. Thus, the complexity of each 1ZK argument is polynomial in the complexity of NM_{i-1} and Com. Since there are λ 1ZK arguments, the complexity of NM_i is higher than that of iNM_{i-1} by at least a multiplicative factor of λ. After $L = \omega(1)$ iterations, the complexity becomes super-polynomial.

As in other tag amplification techniques in the literature [LP11, KS17, LPS17], this blow-up can be avoided with a simple modification. For any iteration $i \geq 2$, the tag amplification is applied to an input scheme NM_{i-1} produced by the transformation itself. To verify that a commitment of NM_{i-1} commits to string v, it requires i) verifying all 1ZK arguments in it and ii) that the Com commitments commit to v bit by bit. Thus, when transforming NM_{i-1} to NM_i, verifying the statement x_j can be decomposed into a public part—that verifies that all 1ZK arguments in all commitments of NM_{i-1} are accepting—and a private part—the j'th Com commitment in NM_i is consistent with the j'th Com commitment inside every commitment of NM_{i-1}. The key observation is that the public part can be verified publicly, and the 1ZK arguments only need to prove about the private part, which takes a fixed polynomial time $\mathrm{poly}(\lambda)$. With this modification, the complexity of NM_i is larger than that of NM_{i-1} only by an additive polynomial factor $\mathrm{poly}(\lambda)$ and a multiplicative factor of $\gamma_{i-1}/2$. Since $\prod_{i\in[L]}(\gamma_{i-1}/2)$ is bounded by a polynomial, the complexity of the final scheme is polynomial.

Remark 5.3 (On growth of security loss and distinguishing advantage). We also discuss the growth of security loss and distinguishing advantage when applying the tag amplification for a super-constant number of times. Recall that in the proof of Theorem 5.1, we proved the non-malleability of an output scheme iNM_i via a sequence of hybrids, where the indistinguishability of neighboring hybrids either reduces to the non-malleability of the input scheme iNM_{i-1}, or to the security of Com and 1ZK arguments. The security reduction to the former participates in a man-in-the-middle execution of iNM_{i-1}, while internally emulating a man-in-the-middle execution of iNM_i for the attacker. This reduction incurrs only an additive polynomial security loss. Therefore, for any fixed adversary $\mathcal{A}_\lambda$, across $L = O(\log^* \lambda)$ iterations, the non-malleability of iNM_L against $\mathcal{A}$ relies on the non-malleability of iNM_0 against adversary of time $T_{\mathcal{A}} + \mathrm{poly}(\lambda)$, which is easily accommodated. On the other hand, the security reduction to Com and 1ZK has large security losses, which are accommodated by appropriately scaling the security parameters used with Com and 1ZK arguments in different iterations. Finally, the distinguishing gap of iNM_i is the sum of the distinguishing

gaps of neighboring hybrids. Since the total number of hybrids across all iterations is bounded by $\mathrm{poly}(\lambda)$. The overall distinguishing gap is still negliglble.

From Same-Tag Non-Malleability with respect to Extraction for 4 Tags to Concurrent Non-Malleability with respect to Commitment To obtain concurrent non-malleability, we start from same-tag non-malleability with respect to extraction for 4 tags and apply the following two steps.

Input Scheme: The initial input scheme iNM has $\gamma_0 = 4$ tags, and is $T_{\mathsf{iNM.E}_0}$-over-extractable and T_{iNM}-*same-tag* non-malleable with respect to extraction by iNM, where

$$T_{\mathsf{iNM.E}}(n) = 2^{2n}\ , \qquad T_{\mathsf{iNM}}(n) = 2^{n^{\varepsilon}} \text{ for some } \varepsilon > 0\ .$$

Step 1—Iteratively amplify number of tags: Apply the tag amplification transformation to iNM for $L = O(\log^* \lambda)$ times as above to obtain an output scheme NM for at least 2^λ tags. NM is T_{NM}-non-malleable with respect to commitment and $T_{\mathsf{NM.E}}$-extractable by $\mathsf{NM.E}$ (without over-extraction) for

$$\bar{n} = (\log \lambda)^{2/\varepsilon\alpha^L}\ , \text{ for } \alpha = \min(\varepsilon, \rho)/2, \qquad T_{\mathsf{NM.E}}(\bar{n}) < 2^{2\bar{n}}\ , \qquad T_{\mathsf{NM}}(\bar{n}) = 2^{\log^2 \lambda}\ .$$

Step 2—Strengthen Non-Malleability: Apply the non-malleablity strengthening transformation to NM, using security parameter $\bar{n}$ with NM, and security parameter $\bar{n}' = \bar{n}^{1/\alpha}$ for α described above for other components, namely, Com and NIWI in 1ZK arguments. By Theorem 5.3, the output scheme oNM is T_{oNM}-concurrently non-malleable with respect to commitment and $T_{\mathsf{oNM.E}}$-extractable by $\mathsf{oNM.E}$ (without over-extraction) for

$$T_{\mathsf{oNM.E}}(\bar{n}') < 2^{2\bar{n}'}\ , \qquad T_{\mathsf{oNM}}(\bar{n}') = T_{\mathsf{NM}}(\bar{n}) = 2^{\log^2 \lambda}\ .$$

5.4 Same-Tag Non-Malleability for 4-Tags from Time-Lock Puzzles

The work of Lin, Pass, and Soni (LPS) [LPS17] presented a non-interactive commitment scheme for γ tags, for any constant γ, that is non-malleable with respect to extraction. Their scheme relies on sub-exponentially secure injective one-way functions and sub-exponentially secure time-lock puzzles. Roughly speaking, the latter are puzzles that can be solved by "brute-force" in time 2^t, but cannot be solved significantly faster in *parallel time/depth* 2^{t^ϵ}. The most popular instantiation of time-lock puzzles is the repeated squaring assumption introduced by Rivest, Shamir, and Wagner that 2^t repeated squarings mod $N = pq$ cannot be solved significantly faster, in 2^{t^ϵ} parallel time. See [LPS17] for formal definitions of subexponentially secure time-lock puzzles and the repeated squarings assumption.

Let us briefly review the LPS construction. From sub-exponentially secure injective one-way functions, one can obtain a sub-exponentially secure commitment scheme Com^s. By complexity leveraging, one can instantiate Com^s with different security parameters, to obtain a family of γ schemes $\{\mathsf{Com}^s_i\}_{i\in[\gamma]}$ satisfying that for $i > j$, Com^s_i is "harder" than Com^s_j *in the axis of time*, that is, Com^s_i remains hiding in *time* sufficient for extracting from Com^s_j. Moreover, the extraction procedure is highly parallelizable and has a fixed polynomial parallel-time/depth. Similarly, when starting from subexponentially secure time lock puzzles, by complexity leveraging, can obtain a family of γ commitment schemes $\{\mathsf{Com}^d_i\}_{i\in[\gamma]}$ s.t. for any $i > j$, Com^d_i is "harder" than Com^d_j in the axis of parallel-time/depth, that is, Com^d_i remains hiding in *parallel-time/depth* sufficient for extraction from Com^d_j. In addition, one can make sure that every Com^s commitment remains hiding in *time* for extracting from any Com^d commitment.

To construct a non-malleable commitment scheme NM, their key idea is combining a Com^s and Com^d scheme with opposite strength:

$$\mathsf{NM}(\mathsf{tg}, v, 1^\lambda) \ : \ \mathsf{Com}^s_{\mathsf{tg}}(s_1, 1^\lambda), \mathsf{Com}^d_{\gamma-\mathsf{tg}}(s_2, 1^\lambda) \ , \ \text{for } s_1 \leftarrow \{0,1\}^{|v|}, \ s_2 = v \oplus s_1 \ .$$

To see why this works, consider two cases. First, if the left tag i is smaller than the right tag j, the $\mathsf{Com}^d_{\gamma-i}$ commitment on the left remains hiding in *parallel-time/depth* for extracting from both $\mathsf{Com}^d_{\gamma-j}$ and Com^s_j (recall that the latter can be extracted in poly parallel-time). Therefore the left committed value remains, while the right is extracted. Otherwise, if the left tag i is larger than the right tag j, the Com^s_i commitment on the left remains hiding in *time* for extracting from both Com^s_j and $\mathsf{Com}^d_{\gamma-j}$. Thus again, the left committed value remains hidden, while the right is extracted.

Here, we make the observation that the LPS non-malleable commitment scheme NM is in fact same-tag non-malleable with respect to extraction. This essentially follows from the same proof as described above, with the modification that we now need to reason about extraction from multiple right commitments using the same tg. The above argument states that the left committed value remains hiding while the right committed value is extracted. The complexity of extraction is determined by the right tag tg. When there are multiple, $m = \mathrm{poly}(\lambda)$, right commitments using the same tg, the complexity simply increases by $m = \mathrm{poly}(\lambda)$ folds, and the left committed value remains hidden. Hence, NM is same-tag non-malleable with respect to extraction. We remark that the restriction on right commitments having the same tag is necessary here. Otherwise, consider the scenario where the left commitment uses tag i, which is smaller than the tag $j_1 > i$ for one right commitment, and larger than the other $j_2 < i$. In this case, extracting both right committed values, would break hiding of both Com^s_i and $\mathsf{Com}^d_{\gamma-i}$.

Theorem 5.4 ([LPS17])**.** *Assume the existence of a subexponentially secure time-lock puzzle and subexponentially secure injective one-way functions. For any constant γ, there is a commitment scheme NM for γ tags that is 2^{n^ε}-same-tag non-malleable with respect to extraction for some $\varepsilon > 0$.*

5.5 Non-Malleability for 4 Tags from Amplifiable One-Way Functions

In this section, we present a new non-malleable commitment scheme for a constant number of tags from *amplifiable one-way functions*. Hardness amplification of one-way functions have been extensively studied. Celebrated results showed that direct product (i.e., parallel repetition) or XOR (i.e., XOR of parallel repetition) can strengthen a weakly-hard function to a strongly-hard one. For instance, Yao's XOR-lemma states that if a boolean function f is δ-hard to compute for T-time algorithms (meaning that every T-time algorithm $\mathcal{M}$ computes f wrong for at least a δ fraction of the inputs), then the k-fold XOR'ed function $f^{\oplus k}(x_1, \cdots, x_k) = f(x_1) \oplus \cdots \oplus f(x_k)$ is $(\frac{1}{2} - (\mathrm{poly}(\frac{T'}{T}) + (1-\delta)^k))$-hard for T'-time algorithms. The hardness of the function $f^{\oplus k}$ strengthens optimally with k, but hits the limit at $\frac{1}{2} - \mathrm{poly}(\frac{T'}{T})$. A fascinating question is whether this limit is inherent. The work of [DJMW12b] constructed artificial one-way functions for which direct product does not amplify the hardness beyond negligible. However, no similar evidence is known for natural one-way functions such as discrete logarithm, RSA, Rabin, and others. In this work, we put forward the notion of *amplifiable one-way functions*, and use them to construct non-malleable commitments with a constant number of tags.

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 $f(x_1), \ldots, f(x_\ell)$, so that the combined bit is 2^{ℓ^ε}-unpredicatable; that is, the level of unpredicatbility increases at least subexponentially as more hardcore bits are combined, beyond the limit $\mathrm{poly}(\frac{T'}{T})$.

Such one-way functions are useful for constructing non-malleable commitments because they essentially allow us to construct a set of commitment schemes, such that comparing any two of them $(\mathsf{Com}, \mathsf{Com}')$, Com is "harder" than Com' *in the axis of time*, that is, Com remains hiding in time needed for extracting from Com', whereas Com' is "harder" than Com *in the axis of distinguishing advantage*, that is, the maximum distinguishing advantage of Com' is smaller than the probability that one can guess a decommitment of Com. As shown in [LPS17], such commitments that are harder than each other in different measures are non-malleable with respect to one another. Note that the security of Com' is weaker than that of Com in terms of the maximum attacker run-time it tolerates, but has a smaller distinguishing advantage.

Below, we first recall the definitions of one-way function families and hard-core bits, and then introduce the notion of amplifiable one-way functions.

Definition 5.1 (Family of One-Way Functions)**.** *A family of one-way functions consists of an ensemble of sets $\mathcal{F} = \{\mathcal{F}_\lambda\}$, where $\mathcal{F}_\lambda$ contains efficiently computable functions f mapping from domain $\mathcal{X}_f$ to range $\mathcal{Y}_f$. Moreover, the following property holds:*

T-one-wayness: For any non-uniform $\mathrm{poly}(T)$-time, polynomial-size, probabilistic adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}_\lambda$ there exists a negligible μ, such that for every $\lambda \in \mathbb{N}$,

$$\Pr[f(\mathcal{A}_\lambda(f, h, y)) = y \ : \ (f, h) \leftarrow \mathcal{F}_\lambda \ , \ x \leftarrow \mathcal{X}_f \ , \ y = f(x)] \leq \mu(\lambda) \ .$$

We say that $\mathcal{F}$ is injective*, if for every λ, every f in $\mathcal{F}_\lambda$, f is injective. We say that $\mathcal{F}$ is* uniformly samplable *if there is a efficient uniform sampling algorithm $\mathsf{samp}_\mathcal{F}$ that on input 1^λ samples uniformly from the set $\mathcal{F}_\lambda$.*

We remark that we restrict our attention to one-way function families that are uniformly samplable.

Definition 5.2 (Hard-Core Bit)**.** *Let $\mathcal{F}$ be a family of one-way functions. We say that a predicate $h : \{0,1\}^* \to \{0,1\}$ is a hardcore bit of $\mathcal{F}$, if it satisfies the following:*

T-unpredictability: For any non-uniform $\mathrm{poly}(T)$-time, polynomial-size, probabilistic adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}_\lambda$ there exists a negligible μ, s.t. for every $\lambda \in \mathbb{N}$,

$$\Pr[\mathcal{A}_\lambda(f, h, y) = h(x) \ : \ (f, h) \leftarrow \mathcal{F}_\lambda \ , \ x \leftarrow \mathcal{X}_f \ , \ y = f(x)] \leq \frac{1}{2} + \mu(\lambda) \ .$$

Definition 5.3 ((T, δ)-hardness-amplifiable One-Way Functions)**.** *We say that a family $\mathcal{F}$ of one-way functions is* (T, δ)-hardness-amplifiable*, if there are two efficiently computable functions $h : \{0,1\}^* \to \{0,1\}$ and $C : \{0,1\}^* \to \{0,1\}$,*

- *for any non-uniform $\mathrm{poly}(T)$-time, polynomial-size, probabilistic adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}_\lambda$, every sufficiently large polynomial ℓ, and sufficiently large $\lambda \in \mathbb{N}$*

$$\Pr\left[\mathcal{A}_\lambda(f, h, y_1, \cdots, y_\ell) = C(h(x_1), \cdots, h(x_\ell)) \ : \ \begin{array}{l}(f, h) \leftarrow \mathcal{F}_\lambda \ , \\ \forall i \in [\ell], \ x_i \leftarrow \mathcal{X}_f \ , \ y_i = f(x_i)\end{array}\right] \leq \frac{1}{2} + \delta(\lambda, \ell(\lambda)) \ .$$

We say that $\mathcal{F}$ is subexponentially-hardness-amplifiable if it is (T, δ)-hardness-amplifiable for $T(\lambda) = 2^{\lambda^\varepsilon}$ and $\delta(\lambda, \ell) = 2^{-\ell^\varepsilon}$ for some constant $\varepsilon > 0$.

Below, we refer to $C(h(x_1) \cdots h(x_\ell))$ the ℓ-way combined hardcore bit.

We formalize the hardness amplification assumption with respect to natural one-way functions, such as, discrete logarithm and RSA. Similar assumptions can be made w.r.t. other natural one-way functions as well.

Discrete Logarithm Amplification Assumption A discrete logarithm function $f_{G,g,m}$ is described by a cyclic group G with generator g and order $m = \mathrm{poly}(\lambda)$:

$$f_{G,g,m} \,:\, \mathbb{Z}_m \to G, \qquad f(x) = g^x \,.$$

The descrete logarithm function family $\mathcal{DL} = \{\mathcal{DL}_\lambda\}$ contains an ensemble of such functions $\mathcal{DL}_\lambda = \{f_{G_\lambda,g_\lambda,m_\lambda}\}$ defined by an ensemble $\{(G_\lambda, g_\lambda, m_\lambda)\}_{\lambda\in\mathbb{N}}$ of cyclic groups that can be uniformly and efficiently generated $(G_\lambda, g_\lambda, m_\lambda) = \mathsf{samp}_{\mathcal{DL}}(1^\lambda)$.

Assumption 5.1. *The family of discrete logarithm functions is sub-exponentially amplifiable.*

Note also that $\mathcal{DL}$ is injective and uniformly samplable, as the function family contains a single function for each λ.

RSA Amplification Assumption A RSA function $f_{N,e}$ is described by a product N of two equal-length primes $N = pq$ and an exponent $e \in \mathbb{Z}^*_{\Phi(N)}$,

$$f_{N,e} \,:\, \mathbb{Z}^*_N \to \mathbb{Z}^*_N, \qquad f(x) = x^e \mod N \,.$$

The RSA function family $\mathcal{RSA} = \{\mathcal{RSA}_\lambda\}$ contains an ensemble of sets of such functions $\mathcal{RSA}_\lambda = \{f_{N,e} \,:\, N = pq \text{ for } p, q \text{ primes of length } \lambda,\ e \in \mathbb{Z}^*_{\Phi(N)}\}$.

Assumption 5.2. *The family of RSA functions is sub-exponentially hardness amplifiable.*

Note that $\mathcal{RSA}$ is injective, and the function family can be efficiently and uniformly sampled, as one can sample products of primes of certain length and elements in $\mathbb{Z}^*_m$ for any m efficiently and uniformly.

Theorem 5.5. *Assume the existence of a family $\mathcal{F}$ of injective one-way functions functions that is uniformly samplable and subexponentially amplifiable. For any constant γ, there is a commitment scheme NM for γ tags that is 2^{λ^ε}-non-malleable w.r.t. commitment.*

Proof. Let h and C be the hardcore bit and combiner w.r.t. which $\mathcal{F}$ is $(2^{\lambda^\varepsilon}, 2^{\ell^\varepsilon})$-hardness amplifiable. Let samp be the sampler that samples functions in $\mathcal{F}$ uniformly and randomly. Let λ^c be an upper bound on the run-time of $\mathsf{samp}(1^\lambda)$ and time for computing any f in $\mathcal{F}_\lambda$.

We now present our non-malleable commitment scheme $\mathsf{NM} = (\mathsf{Com}, \mathsf{Open})$ for γ tags. Fix a global security parameter λ, which is also an upper bound on the length of the input strings. The construction uses the following security parameters depending on λ.

$$n_0 = \lambda \,, \qquad \forall \mathsf{tg} \in [\gamma],\ n_{\mathsf{tg}} = n_{\mathsf{tg}-1}^{2/\varepsilon} = \lambda^{(2c/\varepsilon)^{\mathsf{tg}}} \tag{8}$$

$$\ell_0 \text{ is a sufficient large polynomial s.t. } \ell_0(\lambda) > n_\gamma \lambda + 1, \qquad \forall \mathsf{tg} \in [\gamma],\ \ell_{\mathsf{tg}}(\lambda) = \ell_{\mathsf{tg}-1}(\lambda)^{3/\varepsilon} \tag{9}$$

Commitment $\mathsf{Com}(\mathsf{tg}, v, 1^\lambda)$**:** Given tg, the commitment samples a one-way function f from $\mathcal{F}$ with security parameter $n_{\gamma-\mathsf{tg}}$, and uses the ℓ_{tg}-way combined hardcore bits of f to hide v bit by bit. More specifically, sample $f = \mathsf{samp}(1^{n_{\gamma-\mathsf{tg}}}; \rho)$ using randomness ρ, and for every bit $v[j]$, do

- For every $i \in [\ell_{\mathsf{tg}}]$, sample $x_{j,i} \leftarrow \mathcal{X}_f$, and compute $y_{j,i} = f(x_{j,i})$.
- Hide $v[j]$ using the combined hardcore bit $c_j = C(h(x_{j,1}), \cdots, h(x_{j,\ell_{\mathsf{tg}}})) \oplus v[j]$.

The final commitment c and decommitment d is

$$c = \left(\mathsf{tg}, f, \{y_{j,i}, c_j\}_{j\in[\lambda], i\in[\ell_{\mathsf{tg}}]}\right) \,, \qquad d = \left(\rho, \{x_{j,i}\}_{j\in[\lambda], i\in[\ell_{\mathsf{tg}}]}\right) \,.$$

Decommitment $\mathsf{Open}(c, v, d, 1^\lambda)$**:** Verify the following:

- Verify that f is indeed sampled from $\mathcal{F}$ using randomness ρ, that is $f = \mathsf{samp}(1^{n_{\gamma-\mathsf{tg}}}; \rho)$. (This ensures that f is injective.)
- For every $j \in [\lambda]$ and $i \in [\ell_{\mathsf{tg}}]$, verify that the x-strings contained in d are preimages of the y-strings in c, that is, $f(x_{j,i}) = y_{j,i}$ for all j, i.
- For every $j \in [\lambda]$, compute the combined hardcore bit $b_j = C(h(x_{j,1}), \cdots, h(x_{j,\ell_{\mathsf{tg}}}))$. Verify that $v[j] = c[j] \oplus b_j$.

Output 0 if any of the above verification fails and 1 otherwise.

The perfect binding property of NM follows from the fact that when f is indeed sampled from the family $\mathcal{F}$, it is injective, and hence the preimages are the y-strings in c is unique, which uniquely determines the committed value v. We show in the following two claims useful properties of NM that will be instrumental for proving its non-malleability.

Claim 5.18. *For every λ, tag $\mathsf{tg} \in [\gamma]$, and every string c with prefix tg, it takes at most $T_{\mathsf{tg}}(\lambda)$ time to find the value $v = \mathsf{val}(c)$ committed in c, where $T_{\mathsf{tg}}(\lambda) = \mathrm{poly}(\lambda) 2^{n^c_{\gamma-\mathsf{tg}}}$.*

Proof. To find a decommitment (v, d) of $c = \left(\mathsf{tg}, f, \{y_{j,i}, c_j\}_{j\in[\lambda], i\in[\ell_{\mathsf{tg}}]}\right)$, it suffices to find $d = \left(\rho, \{x_{j,i}\}_{j\in[\lambda], i\in[\ell_{\mathsf{tg}}]}\right)$ from which v can be computed.

Finding ρ s.t. $f = \mathsf{samp}(1^{n_{\gamma-\mathsf{tg}}}; \rho)$ takes at most $2^{n^c_{\gamma-\mathsf{tg}}}$ time, since samp takes at most $n^c_{\gamma-\mathsf{tg}}$ time. If no such ρ exists, the commitment is invalid and $v = \perp$. If such a ρ is found, f is injective and for every y in the image of f, there exists a unique preimages. For each $y_{j,i}$, the unique preimage $x_{j,i}$ can be found againt in $2^{n^c_{\gamma-\mathsf{tg}}}$ time, as f can be computed in $n^c_{\gamma-\mathsf{tg}}$ time. If for any j, i, no such preimage is found, c is invalid and $v = \perp$. Otherwise, for every j, recover the j'th bit of v as $v[j] = c_j \oplus C(h(x_{j,1}), \cdots, h(x_{j,\mathsf{tg}}))$.

The above procedure takes at most $\mathrm{poly}(\lambda) 2^{n^c_{\gamma-\mathsf{tg}}}$ time. □

Claim 5.19. *For every tag $\mathsf{tg} \in [\gamma]$, every λ, and every valid commitment c in the support of $\mathsf{Com}(\mathsf{tg}, \star, 1^\lambda)$, the probability that a randomly sampled string is a decommitment of c is universal $P^{\mathrm{decom}}_{\mathsf{tg}}(\lambda)$ independent of c, and $P^{\mathrm{decom}}_{\mathsf{tg}}(\lambda) \geq 2^{-\ell_{\mathsf{tg}}(\lambda)^2}$.*

Proof. Let m denote the length of decommitment to c. When $c = (\mathsf{tg}, f, \{y_{j,i}, c_j\}_{j,i})$ is valid, to guess (v, d), it suffices to guess correctly $d = (\rho, \{x_{j,i}\}_{j\in[\lambda], i\in[\ell_{\mathsf{tg}}]})$, from which the committed value v can be computed. When c is valid, f is injective and the preimages of y-strings are unique, thus the probabilty of guessing correctly $x_{j,i}$ is exactly $2^{-|x_{j,i}|} \geq 2^{-n^c_{\gamma-\mathsf{tg}}}$. Since samp samples functions in $\mathcal{F}_{n_{\gamma-\mathsf{tg}}}$ uniformly and randomly, the probabilty of guessing the randomness ρ that leads to f is exactly the inverse of the size of the function set $1/|\mathcal{F}_{n_{\gamma-\mathsf{tg}}}| \geq 2^{-n^c_{\gamma-\mathsf{tg}}}$. Therefore, as shown below, the probability of guessing a decommitment is independent of the specific commitment c.

$$\begin{aligned} P^{\mathrm{decom}}_{\mathsf{tg}}(\lambda) =& \Pr[\text{guess } \rho \text{ s.t. } f = \mathsf{samp}(1^{n_{\gamma-\mathsf{tg}}};\ \rho)] \ \times\ \Pr[\forall j, i,\ \text{guess } x_{j,i} \text{ s.t. } f(x_{j,i}) = y_{j,i}] \\ =& \frac{1}{|\mathcal{F}_{\gamma-\mathsf{tg}}|} \times \prod_{i,j} 2^{-|x_{j,i}|} \ . \end{aligned}$$

Furthermore, this universal probability is lower bounded by

$$P^{\mathrm{decom}}_{\mathsf{tg}}(\lambda) \ \geq 2^{-n^c_{\gamma-\mathsf{tg}}} \times \left(2^{-n^c_{\gamma-\mathsf{tg}}}\right)^{\ell_{\mathsf{tg}}\times\lambda} \ \geq \ 2^{-\ell^2_{\mathsf{tg}}} \ .$$

The inequalities are guaranteed by the setting of parameters; see Equation (8). □

Next, we show that NM is 2^{λ^ε}-non-malleable with respect to commitment. Fix an arbitrary $\mathrm{poly}(2^{\lambda^\varepsilon})$-time (polynomial-sized) non-uniform attacker $\mathcal{A} = \{\mathcal{A}_\lambda\}$ that receives one left commitment and sends one right commitment. Also fix two arbitrary ensembles of messages $\{v_\lambda\}_\lambda$, $\{u_\lambda\}_\lambda$ of length λ. We want to show that

$$\left\{\mathsf{mim}^{\mathcal{A}}_{\mathsf{NM}}(v_\lambda)\right\}_{\lambda\in\mathbb{N}} \approx_c \left\{\mathsf{mim}^{\mathcal{A}}_{\mathsf{NM}}(u_\lambda)\right\}_{\lambda\in\mathbb{N}} .$$

Towards this, consider the following two events in the man-in-the-middle execution:

- Event left-tag-smaller: The left tg is smaller than the right $\widetilde{\mathsf{tg}}$. In this case, the left one-way function f is sampled with a larger security parameter than the right one-way function $\tilde{f}$, $n_{\gamma-\mathsf{tg}} > n_{\gamma-\widetilde{\mathsf{tg}}}$. We will show that the right committed value can be extracted by brute force in time that does not hurt the hiding of the left commitment. Thus,

  Lemma 5.1. *For any polynomial-time and polynomial-sized distinguisher $\{\mathcal{D}_\lambda\}$,*

  $$|\Pr[\mathcal{D}(\mathsf{mim}^{\mathcal{A}}_{\mathsf{NM}}(v)) = 1 \ \wedge \ \textit{left-tag-smaller}] - \Pr[\mathcal{D}(\mathsf{mim}^{\mathcal{A}}_{\mathsf{NM}}(u)) = 1 \ \wedge \ \textit{left-tag-smaller}]| \leq \mathrm{negl}(\lambda) .$$

- Event left-tag-larger: The left tg is larger than the right $\widetilde{\mathsf{tg}}$. In this case, the left one-way function f is sampled with smaller security parameter than the right one-way function $\tilde{f}$, $n_{\gamma-\mathsf{tg}} < n_{\gamma-\widetilde{\mathsf{tg}}}$. But, and the left committed value v is hidden using combination of more hardcore bits than the right committed value $\ell_{\mathsf{tg}} > \ell_{\widetilde{\mathsf{tg}}}$. We will show that it follows from the amplification property of $\mathcal{F}$ that

  Lemma 5.2. *For any polynomial-time and polynomial-sized distinguisher $\{\mathcal{D}_\lambda\}$,*

  $$|\Pr[\mathcal{D}(\mathsf{mim}^{\mathcal{A}}_{\mathsf{NM}}(v)) = 1 \ \wedge \ \textit{left-tag-larger}] - \Pr[\mathcal{D}(\mathsf{mim}^{\mathcal{A}}_{\mathsf{NM}}(u)) = 1 \ \wedge \ \textit{left-tag-larger}]| \leq \mathrm{negl}(\lambda) .$$

It follows directly from the above two lemmas that $\mathsf{mim}^{\mathcal{A}}_{\mathsf{NM}}(v)$ and $\mathsf{mim}^{\mathcal{A}}_{\mathsf{NM}}(u)$ are indistinguishable.

Proof of Lemma 5.1. In this event left-tag-smaller, the left one-way function f is sampled with a larger security parameter $n_{\gamma-\mathsf{tg}}$ than the right one $\tilde{f}$, sampled with $n_{\gamma-\widetilde{\mathsf{tg}}}$. By claim 5.18, one can find the right committed value $\tilde{v}$ in time $\mathrm{poly}(2^{n^c_{\gamma-\widetilde{\mathsf{tg}}}})$. On the other hand, left commitment is hiding against adversary of time $2^{n^\varepsilon_{\gamma-\mathsf{tg}}} \gg 2^{n^c_{\gamma-\widetilde{\mathsf{tg}}}}$, as the combined hard-core bits used for hiding the left committed values are $(2^{n^\varepsilon_{\gamma-\mathsf{tg}}}, 2^{\ell^\varepsilon_{\mathsf{tg}}})$-unpredictable. Therefore, the left commitment remains hiding even when the right committed value is found by brute force. Thus, the view of $\mathcal{A}$ and the right committed value is indistinguishable when the left tag is smaller. □

Proof of Lemma 5.2. Suppose for contradiction that there exists a distinguisher $\mathcal{D}$ that violates the statement of the lemma,

$$|\Pr[\mathcal{D}(\mathsf{mim}^{\mathcal{A}}_{\mathsf{NM}}(v)) = 1 \ \wedge \ \text{left-tag-larger}] - \Pr[\mathcal{D}(\mathsf{mim}^{\mathcal{A}}_{\mathsf{NM}}(u)) = 1 \ \wedge \ \text{left-tag-larger}]| \geq 1/\mathrm{poly}(\lambda) .$$

Consider hybrids $\mathcal{H}_0 \cdots \mathcal{H}_\lambda$, where in $\mathcal{H}_j$, the value committed on the left is $u[1\cdots j]v[j+1\cdots\lambda]$. Denote by mim_j the view of $\mathcal{A}$ and the value it commits to on the right in hybrid $\mathcal{H}_j$. Then, there must exist $J \in [\lambda]$, s.t.

$$\left|\Pr\left[\mathcal{D}\left(\mathsf{mim}_{J-1}\right) = 1 \ \wedge \ \text{left-tag-larger}\right] - \Pr\left[\mathcal{D}\left(\mathsf{mim}_{J}\right) = 1 \ \wedge \ \text{left-tag-larger}\right]\right| \geq 1/\mathrm{poly}(\lambda) .$$

Consider two experiments $\mathsf{Exp}_0^{\mathcal{M}}$ and $\mathsf{Exp}_1^{\mathcal{M}}$ defined with J, u, v and an arbitrary *polynomial-time* polynomial-sized distinguisher $\mathcal{M}$.

Experiment $\mathsf{Exp}_d^{\mathcal{M}}$ *for* $d \in \{0,1\}$ $\mathcal{M}$ receives as input $(f, \{y_{J,i}\}_{i\in[\ell_{\mathsf{tg}}]}, c_J)$, where f is a one-way function sampled randomly from $\mathcal{F}_{n_{\gamma-\mathsf{tg}}}$, $\{y_{J,i}\}_i$ are images of randomly sampled inputs $\{x_{J,i}\}_i$ of f, and $c_J = b^\star \oplus w[J]$ hides $w[J]$ using the combined hard-core bit $b^\star = C(h(x_{J,1}), \cdots, h(x_{J,\ell_{\mathsf{tg}}}))$ with $w[J] = v[J]$ if $d = 0$ and $w[J] = u[J]$ if $d = 1$. It follows directly from the amplification property of $\mathcal{F}$ that $b^\star$ is $2^{-\ell_{\mathsf{tg}}^{\varepsilon}}$-unpredicatable for all $\mathrm{poly}(2^{n_{\gamma-\mathsf{tg}}^{\varepsilon}})$-time adversaries. Thus,

Claim 5.20. *Let $\mathcal{M}$ be any $\mathrm{poly}(2^{n_{\gamma-\mathsf{tg}}^{\varepsilon}})$-time and polynomial-sized distinguisher. The advantage of $\mathcal{M}$ in distinguishing Exp_0 and Exp_1 is at most $2 \times 2^{-\ell_{\mathsf{tg}}^{\varepsilon}}$.*

Using the above experiments, we first show the following claim:

Claim 5.21. *The probabilities that $\mathcal{A}_\lambda$ sends a* valid *commitment on the right in $\mathcal{H}_{J-1}$ and $\mathcal{H}_J$ differ at most by a negligible amount.*

$$|P_{J-1} - P_J| \leq \mathrm{negl}(\lambda) \ , \ \textit{where } P_j = \Pr[\textit{right commitment in } \mathcal{H}_j \textit{ is valid}] \ .$$

Proof. Suppse for contradiction that $|P_{J-1} - P_J| \geq 1/\mathrm{poly}(\lambda)$. We derive a contradiction by constructing a $\mathrm{poly}(2^{\lambda^{\varepsilon}})$-time attacker $\mathcal{B}$ that distinguishes Exp_0 and Exp_1 with advantage higher than $2 \times 2^{-\ell_{\mathsf{tg}}^{\varepsilon}}$, which contradicts Claim 5.20.

In $\mathsf{Exp}_d^{\mathcal{B}}$, $\mathcal{B}$ upon receiving $(f, \{y_{J,i}\}_i, c_J)$ proceeds as follows:

- Step 1: $\mathcal{B}$ internally runs $\mathcal{A}$ and emulates the left commitment for $\mathcal{A}$ as follows. i) It forwards f to $\mathcal{A}$; ii) it forwards $\{y_{J,i}\}_i$ and c_J as the y-strings and commitment to the J'th bit, iii) for every other bit $j \neq J$, it honestly samples images $\{y_{j,i}\}_i$ of random inputs $\{x_{j,i}\}_i$, use the combined hardcore bit $b_j = C(h(x_{j,0}) \cdots h(x_{j,\ell_{\mathsf{tg}}}))$ to hide $w[j]$, $c_j = w[j] \oplus b_j$, where $w[j] = v[j]$ if $j < J$ and $w[j] = u[j]$ if $j > J$. $\mathcal{A}$ upon receiving the emulated left commitment sends a right commitment $\tilde{c} = (\tilde{f}, \{\tilde{y}_{j,i}, \tilde{c}_j\}_{j,i})$.

  Note that in Exp_0, $\mathcal{B}$ receives c_J that hides $v[j]$ and emulates perfectly hybrid $\mathcal{H}_{J-1}$ for $\mathcal{A}$. On the other hand, in Exp_1, $\mathcal{B}$ receives c_J that hides $u[j]$ and emulates perfectly hybrid $\mathcal{H}_J$ for $\mathcal{A}$.

- Step 2: $\mathcal{B}$ guesses the decommitment of $\tilde{c}$ at random $(\tilde{v}, \tilde{d}) \leftarrow \{0,1\}^m$, where m is the length of decommitment. If $(\tilde{v}, \tilde{d})$ is not a valid decommitment, that is, $\mathsf{Open}(\tilde{c}, \tilde{v}, \tilde{d}, 1^\lambda) = 0$, $\mathcal{B}$ outputs 0.

- Step 3: If $(\tilde{v}, \tilde{d})$ is a valid decommitment, that is, $\mathsf{Open}(\tilde{c}, \tilde{v}, \tilde{d}, 1^\lambda) = 1$, $\mathcal{B}$ outputs 1.

Let us analyze the advantage of $\mathcal{B}$ in distinguishing Exp_0 and Exp_1. We note that if the right commitment is invalid, the probability of guessing correctly is zero, whereas if the right commitment is valid, as shown by Claim 5.19, the probability of guessing the decommitment is $P_{\widetilde{\mathsf{tg}}}^{\mathrm{decom}}$. Therefore, the advantage of $\mathcal{B}$ is

$$\begin{aligned}
&|\Pr[\mathcal{B} \text{ output 1 in } \mathsf{Exp}_0] - \Pr[\mathcal{B} \text{ output 1 in } \mathsf{Exp}_1]| \\
\geq &|P_{J-1} \times P_{\widetilde{\mathsf{tg}}}^{\mathrm{decom}} - P_J \times P_{\widetilde{\mathsf{tg}}}^{\mathrm{decom}}| \geq 2^{-\ell_{\widetilde{\mathsf{tg}}}^2} |P_{J-1} - P_J| > 2 \times 2^{-\ell_{\mathsf{tg}}^{\varepsilon}} \ .
\end{aligned}$$

The last inequality follows from the setting of parameters that $\ell_{\mathsf{tg}}^{\varepsilon} > \ell_{\widetilde{\mathsf{tg}}}^2$ for any $\mathsf{tg} > \widetilde{\mathsf{tg}}$; see Equation (8). □

Using the fact that the probabilities that $\mathcal{A}$ gives a valid commitment on the right is almost the same in $\mathcal{H}_{J-1}$ and $\mathcal{H}_J$, we show the following two claims:

Claim 5.22. *There exists a negligible function μ, such that,*

$$\Pr\left[\mathcal{D}\left(\mathsf{mim}_{J-1}\right)=1 \ \wedge \ \textit{right commitment } \underline{\textit{valid}} \ \wedge \ \textit{left-tag-larger in } \mathcal{H}_{J-1}\right] - \Pr\left[\mathcal{D}\left(\mathsf{mim}_{J}\right)=1 \ \wedge \ \textit{right commitment } \underline{\textit{valid}} \ \wedge \ \textit{left-tag-larger in } \mathcal{H}_{J}\right] \leq \mu(\lambda) \ .$$

Claim 5.23. *There exists a negligible function μ, such that,*

$$\Pr\left[\mathcal{D}\left(\mathsf{mim}_{J-1}\right)=1 \ \wedge \ \textit{right commitment } \underline{\textit{invalid}} \ \wedge \ \textit{left-tag-larger in } \mathcal{H}_{J-1}\right] - \Pr\left[\mathcal{D}\left(\mathsf{mim}_{J}\right)=1 \ \wedge \ \textit{right commitment } \underline{\textit{invalid}} \ \wedge \ \textit{left-tag-larger in } \mathcal{H}_{J}\right] \leq \mu(\lambda) \ .$$

These two claims would directly imply $\mathcal{D}$ cannot distinguish mim_{J-1} from mim_J in event left-tag-larger, which contradicts with the hypothesis and concludes the lemma.

Proof of Claim 5.22. Assume for contradiction that

$$\Pr\left[\mathcal{D}\left(\mathsf{mim}_{J-1}\right)=1 \ \wedge \ \text{right commitment valid} \ \wedge \ \text{left-tag-larger in } \mathcal{H}_{J-1}\right] - \Pr\left[\mathcal{D}\left(\mathsf{mim}_{J}\right)=1 \ \wedge \ \text{right commitment valid} \ \wedge \ \text{left-tag-larger in } \mathcal{H}_{J}\right] \geq 1/\mathrm{poly}(\lambda) \ .$$

Then, we derive a contradiction by constructing a $\mathrm{poly}(2^{\lambda^\varepsilon})$-time attacker $\mathcal{C}$ that distinguishes Exp_0 and Exp_1 with advantage higher than $2\times 2^{-\ell_{\mathsf{tg}}^{\varepsilon}}$.

In $\mathsf{Exp}_d^{\mathcal{C}}$, $\mathcal{C}$ upon receiving $(f,\{y_{J,i}\}_i,c_J)$ proceeds identically to $\mathcal{B}$ in proof of Claim 5.21, except for the last step. More specifically, in step 1, $\mathcal{C}$ runs $\mathcal{A}_\lambda$ internally and emulates the left commitment for it as $\mathcal{B}$ does; in step 2, upon $\mathcal{A}_\lambda$ sending the right commitment $\tilde{c}$, $\mathcal{C}$ tries to guess a decommitment $(\tilde{v},\tilde{d})$, and outputs a random bit if the guess is incorrect; and

- Step 3: if $(\tilde{v},\tilde{d})$ is a valid decommitment of $\tilde{c}$, that is, $\mathsf{Open}(\tilde{c},\tilde{v},\tilde{d},1^{n_{\gamma-\tilde{\mathsf{tg}}}})=1$. $\mathcal{C}$ feeds the view $\mathsf{View}_{\mathcal{A}}$ of $\mathcal{A}$ and the value $\tilde{v}$ to $\mathcal{D}$ and outputs what $\mathcal{D}$ returns.

Let us analyze the probability that $\mathcal{C}$ outputs 1 in Exp_d.

$$\begin{aligned}
&\Pr[\mathcal{C}\text{ outputs 1 in }\mathsf{Exp}_d]\\
=&\frac{1}{2}\left(\Pr[\tilde{c}\text{ invalid in }\mathsf{Exp}_d]+\Pr[\tilde{c}\text{ valid}\wedge(\tilde{v},\tilde{d})\text{ not decommitment of }\tilde{c}\text{ in }\mathsf{Exp}_d]\right)\\
&+\Pr[\mathcal{D}(\mathsf{View}_{\mathcal{A}},\tilde{v})=1 \ \wedge \ \tilde{c}\text{ valid} \ \wedge \ (\tilde{v},\tilde{d})\text{ decommitment of }\tilde{c}\text{ in }\mathsf{Exp}_d]\\
=&\frac{1}{2}(1-P_{J-1+d}P_{\tilde{\mathsf{tg}}}^{\mathrm{decom}})+\Pr\left[\mathcal{D}\left(\mathsf{mim}_{J-1+d}\right)=1 \ \wedge \ \text{right commitment valid} \ \wedge \ \text{left-tag-larger}\right]\times P_{\tilde{\mathsf{tg}}}^{\mathrm{decom}}
\end{aligned}$$

The last equality follows from the fact that whenever the right commitment is valid, the probability of guessing the decommitment is $P_{\tilde{\mathsf{tg}}}^{\mathrm{decom}}$ as shown in Claim 5.19. Therefore, the advantage of $\mathcal{C}$ is

$$\begin{aligned}
&\left|\Pr[\mathcal{C}\text{ outputs 1 in }\mathsf{Exp}_0]-\Pr[\mathcal{C}\text{ outputs 1 in }\mathsf{Exp}_1]\right|\\
=&\Bigg|\frac{P_{\tilde{\mathsf{tg}}}^{\mathrm{decom}}}{2}(P_J-P_{J-1})+\\
&P_{\tilde{\mathsf{tg}}}^{\mathrm{decom}}\times\Pr\left[\mathcal{D}\left(\mathsf{mim}_{J-1}\right)=1 \ \wedge \ \text{right commitment valid} \ \wedge \ \text{left-tag-larger}\right]-\\
&P_{\tilde{\mathsf{tg}}}^{\mathrm{decom}}\times\Pr\left[\mathcal{D}\left(\mathsf{mim}_{J}\right)=1 \ \wedge \ \text{right commitment valid} \ \wedge \ \text{left-tag-larger}\right]\Bigg|\\
\geq&\ \frac{P_{\tilde{\mathsf{tg}}}^{\mathrm{decom}}}{\mathrm{poly}(\lambda)}\geq 2^{-\ell_{\tilde{\mathsf{tg}}}^{2}}/\mathrm{poly}(\lambda)>2\times 2^{-\ell_{\tilde{\mathsf{tg}}}^{\varepsilon}} \ .
\end{aligned}$$

The last inequality follows from Claim 5.21 that P_J and P_{J-1} differ at most by a negligible amount, Claim 5.19 that $P^{\mathsf{decom}}_{\widetilde{\mathsf{tg}}} > 2^{-\ell^2_{\widetilde{\mathsf{tg}}}}$, and $\ell^2_{\widetilde{\mathsf{tg}}} < \ell^{\varepsilon}_{\mathsf{tg}}$ for $\widetilde{\mathsf{tg}} < \mathsf{tg}$ by the setting of parameter as in Equation (8). This gives a contradiction. □

Proof of Claim 5.23. When the right commitment is invalid, the random variable mim_{J-1} (or mim_J) contains the view $\mathsf{View}_{\mathcal{A},J-1}$ (or $\mathsf{View}_{\mathcal{A},J}$) of $\mathcal{A}$ in $\mathcal{H}_{J-1}$ (or $\mathcal{H}_J$) and the committed value $\tilde{v} = \bot$. Thus, we need to show:

$$\Big|\Pr\left[\mathcal{D}\left(\mathsf{View}_{\mathcal{A},J-1}, \bot\right) = 1 \ \wedge \ \text{right commitment invalid} \ \wedge \ \text{left-tag-larger in } \mathcal{H}_{J-1}\right] - \Pr\left[\mathcal{D}\left(\mathsf{View}_{\mathcal{A},J}, \bot\right) = 1 \ \wedge \ \text{right commitment invalid} \ \wedge \ \text{left-tag-larger in } \mathcal{H}_{J}\right]\Big| \leq \mathrm{negl}(\lambda)$$

It suffice to show that, for any polynomial-time and polynomial-sized distinguisher $\mathcal{D}'$,

$$\Big|\Pr\left[\mathcal{D}'\left(\mathsf{View}_{\mathcal{A},J-1}\right) = 1 \ \wedge \ \text{right commitment invalid} \ \wedge \ \text{left-tag-larger in } \mathcal{H}_{J-1}\right] - \Pr\left[\mathcal{D}'\left(\mathsf{View}_{\mathcal{A},J}\right) = 1 \ \wedge \ \text{right commitment invalid} \ \wedge \ \text{left-tag-larger in } \mathcal{H}_{J}\right]\Big| \leq \mathrm{negl}(\lambda)\,. \quad (10)$$

It follows directly from the hiding of the left commitment that,

$$\left|\Pr\left[\mathcal{D}'\left(\mathsf{View}_{\mathcal{A},J-1}\right) = 1 \ \wedge \ \text{left-tag-larger in } \mathcal{H}_{J-1}\right] - \Pr\left[\mathcal{D}'\left(\mathsf{View}_{\mathcal{A},J}\right) = 1 \ \wedge \ \text{left-tag-larger in } \mathcal{H}_{J}\right]\right| \leq \mathrm{negl}(\lambda)\,.$$

In addition, it follows from Claim 5.22 that the following probability difference is also negligible,

$$\Big|\Pr\left[\mathcal{D}'\left(\mathsf{View}_{\mathcal{A},J-1}\right) = 1 \ \wedge \ \text{right commitment } \underline{\text{valid}} \ \wedge \ \text{left-tag-larger in } \mathcal{H}_{J-1}\right] - \Pr\left[\mathcal{D}'\left(\mathsf{View}_{\mathcal{A},J}\right) = 1 \ \wedge \ \text{right commitment } \underline{\text{valid}} \ \wedge \ \text{left-tag-larger in } \mathcal{H}_{J}\right]\Big| \leq \mathrm{negl}(\lambda)\,.$$

Note that the probability difference in Equation 10 is upper bounded by the sum of the above probability difference, which is negligible. □

□

□

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.

[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.

[DJMW12a] 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.

[DJMW12b] Yevgeniy Dodis, Abhishek Jain, Tal Moran, and Daniel Wichs. Counterexamples to hardness amplification beyond negligible. 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.

[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.

[NW94] Noam Nisan and Avi Wigderson. Hardness vs randomness. *J. Comput. Syst. Sci.*, 49(2):149–167, 1994.

[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 *Progressin Cryptology - VIETCRYPT 2006, First International Conferenceon 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.

[SU01] Ronen Shaltiel and Christopher Umans. Simple extractors for all min-entropies and a new pseudo-random generator. In *42nd Annual Symposium on Foundations of Computer Science, FOCS 2001, 14-17 October 2001, Las Vegas, Nevada, USA*, pages 648–657, 2001.

[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.