

Non-Malleable Codes Against Bounded Polynomial Time Tampering

Marshall Ball¹, Dana Dachman-Soled², Mukul Kulkarni²,
Huijia Lin³, and Tal Malkin¹

¹ Columbia University
{marshall,tal}@cs.columbia.edu

² University of Maryland
danadach@ece.umd.edu, mukul@umd.edu

³ University of Washington
rachel@cs.washington.edu

Abstract. We construct efficient non-malleable codes (NMC) that are (computationally) secure against tampering by functions computable in any *fixed* polynomial time. Our construction is in the plain (no-CRS) model and requires the assumptions that (1) **E** is hard for **NP** circuits of some exponential $2^{\beta n}$ ($\beta > 0$) size (widely used in the derandomization literature), (2) sub-exponential trapdoor permutations exist, and (3) **P**-certificates with sub-exponential soundness exist.

While it is impossible to construct NMC secure against *arbitrary* polynomial-time tampering (Dziembowski, Pietrzak, Wichs, ICS '10), the existence of NMC secure against $O(n^c)$ -time tampering functions (for any *fixed* c), was shown (Cheraghchi and Guruswami, ITCS '14) via a probabilistic construction. An explicit construction was given (Faust, Mukherjee, Venturi, Wichs, Eurocrypt '14) assuming an *untamperable* CRS with length longer than the runtime of the tampering function. In this work, we show that under computational assumptions, we can bypass these limitations. Specifically, under the assumptions listed above, we obtain non-malleable codes in the plain model against $O(n^c)$ -time tampering functions (for any fixed c), with codeword length independent of the tampering time bound.

Our new construction of NMC draws a connection with non-interactive non-malleable commitments. In fact, we show that in the NMC setting, it suffices to have a much weaker notion called *quasi non-malleable commitments* — these are non-interactive, non-malleable commitments in the plain model, in which the adversary runs in $O(n^c)$ -time, whereas the honest parties may run in longer (polynomial) time. We then construct a 4-tag quasi non-malleable commitment from any sub-exponential OWF and the assumption that **E** is hard for some exponential size **NP**-circuits, and use tag amplification techniques to support an exponential number of tags.

1 Introduction

Non-Malleable Codes (NMC) were introduced by Dziembowski, Pietrzak, and Wichs [25] as a modification of error correcting codes, with the goal of achieving security against adversarial tampering functions, that may change every part of a codeword. Informally, a NMC against a class \mathcal{F} guarantees that if a codeword is tampered via the application of a function $f \in \mathcal{F}$, then the decoding of the tampered codeword will either be exactly the original message, or completely unrelated to the original message. As noted in [25], it is impossible to construct NMC against arbitrary tampering functions, since non-malleability can always be broken by a tampering function which first decodes the codeword to learn the underlying message, then re-encodes a related message. In particular, there can be no efficient NMC against arbitrary polynomial time tampering. Thus, to achieve feasibility, we must restrict the class of tampering functions.

A natural way to restrict tampering adversaries is via well-studied computational complexity measures. Several recent works have followed this approach and have developed strong connections between NMC and techniques from computational complexity. For example, Ball et al. [5] constructed NMC against bounded depth circuits with constant fan-in (which includes NC^0), several works [13, 6, 3] constructed NMC against AC^0 relying on different complexity theoretic techniques, and some works considered (restricted variants of) NMC against space-bounded tampering [26, 6]. Specifically, the work of Faust et al. [26] considers space-bounded tampering adversaries in the random oracle model and achieves the security notion of *leaky* continuous non-malleability. The work of Ball et al. [6] is information-theoretic, considers streaming, space-bounded tampering adversaries and achieves standard non-malleability. The current work continues this line of research.

In this paper, we focus on the task of constructing NMC against *bounded* polynomial time tampering, namely tampering functions that are computable in an arbitrary *fixed* polynomial time. This is a very natural class to consider given the impossibility result for (unbounded) polynomial time, and indeed, some of the first works in this line of research have already considered this class. We discuss these next, along with the motivation and goals for our current work.

Cheraghchi and Guruswami [14] gave probabilistic constructions of efficient codes for circuits of size $O(n^c)$ (where an efficient randomized procedure outputs a “good” code with high probability). Faust et.al [27] gave an improved (in terms of the dependence on the error bound) construction against the same class, which is explicit, but relies on a model including an untamperable CRS (common reference string). The presence of CRS is undesirable, as not only must the CRS be generated by a trusted party, the CRS is also a non-tamperable component of the scheme. Moreover, both of these works can be viewed as using limited (t -wise) independence to partially derandomize probabilistic constructions. This approach inherently leads to a CRS whose length is at least as long as the bound on the size of the tampering circuits — meaning the tampering circuits cannot even read the entire CRS. We additionally

note that if we allowed other size parameters, in particular the codeword size, to be as large as the runtime of the tampering function, then achieving non-malleability would become trivial. Finally, we note that constructions of NMC against bounded polynomial-time adversaries are trivial in the ideal permutation model, where it is assumed that all parties have access to an ideal, invertible permutation. Since Feistel-based constructions in the random oracle model are indifferentiable from ideal permutations (and indeed ideal cipher) [18, 19, 20], the above results hold in the random oracle model as well and can be instantiated in practice based on e.g. SHA-3. However, whereas in the random oracle/ideal permutation/ideal cipher model, non-malleability comes for free, in this work we seek constructions that are based solely on *hardness* assumptions that do not have a non-malleability flavor.

This motivates the following question:

Can we construct efficient NMC against bounded polynomial time adversaries, in the plain model (i.e. without CRS or random oracles)? Ideally, with codeword length that is independent of the runtime of the adversarial tampering function?

As we elaborate next, we achieve this by moving to computational security and restricting our attention to uniform adversaries (while [14, 27] gave statistical guarantees against non-uniform adversaries). In addition, as explained shortly below, we allow uniform bounded polynomial time tampering adversaries to have an inverse polynomial advantage (as in [14]) as opposed to having only negligible advantage (as in [27]). We emphasize that to the best of our knowledge, there is no transformation that either (a) eliminates the CRS in the NMC of [27] to achieve security against uniform (or non-uniform) adversaries or (b) fully derandomizes the monte carlo construction of [14], *even* under derandomization assumptions. Our techniques highlight interesting new connections to complexity theory.

1.1 Our Results

Our construction requires a complexity theoretic assumption that some language in the complexity class **E** (the class of languages that can be decided by Turing machines running in time $2^{O(n)}$) is hard for **NP** circuits of some exponential $2^{\beta n}$ (for $\beta > 0$) size. As surveyed later, such assumptions are widely used in the derandomization literature, often referred to as *derandomization assumptions*, and have connection with cryptography. Our construction also relies on the following cryptographic assumptions: the existence of subexponential trapdoor permutations and **P**-certificates (**P**-cert) with sub-exponential soundness. **P**-certs (introduced by [15]) are “succinct” non-interactive arguments for languages $\mathcal{L} \in \mathbf{P}$, with proof length which is a fixed polynomial, independent of the time it takes to decide \mathcal{L} (see full version of this paper [4] for a formal definition). We provide more background on these assumptions in Section 1.2 below.

Theorem 1 (Informal). *Assuming*

- \mathbf{E} is hard for \mathbf{NP} circuits with some exponential size (namely $2^{\beta n}$ for some constant $\beta > 0$)
- Existence of sub-exponential trapdoor permutation
- Existence of \mathbf{P} -cert with sub-exponential soundness

for every constant c_A , there is an efficient construction of NMC in the plain (no-CRS) model against uniform, bounded polynomial n^{c_A} -time tampering adversaries, with inverse polynomial indistinguishability (for any polynomial time non-uniform distinguisher). Furthermore, the codeword size is a fixed polynomial independent of n^{c_A} .

A few remarks are in order. First, to formalize that a tampered codeword, if not copied from the original codeword, must decode to an independent value, the definition of non-malleability requires that the decoded values, u_1 and u_2 , obtained from tampering codewords of different values, v_1 and v_2 respectively, must be indistinguishable (u_b is replaced by **same** in the case of copying). Our NMC achieves inverse polynomial distinguishing advantage against polynomial-time non-uniform distinguishers.

Second, as mentioned before, it is important that the length of the codeword is smaller than the time-bound n^{c_A} of the tampering functions; otherwise, achieving non-malleability becomes trivial. Here, we achieve the ideal case, where the *length* of the codeword is bounded by a fixed polynomial, independent of n^{c_A} . As the adversarial time bound grows, the only parameter that grows is the *run time* of encode/decode. Moreover, this dependence is *necessary* as discussed earlier, since non-malleability is trivially impossible when the class of tampering functions includes the encode/decode functions.

Finally, we note that the assumption of the existence of sub-exponential trapdoor permutation in Theorem 1, can be replaced with the assumption of the existence of ZAPs (public coin, two message witness indistinguishable protocols) [24] with witness indistinguishability against sub-exponential adversaries and the existence of sub-exponential one-way functions (OWF). Note that ZAPs can be constructed from bilinear maps [39], which are not known to imply trapdoor permutations.

Connection between NMC and Non-Malleable Commitments. Our construction of NMC draws a connection with another important notion of non-malleability – non-malleable commitments [21, 51]. The only difference between NMC and non-interactive non-malleable commitments is that the former can be decoded efficiently, whereas decommitment of the latter cannot be done efficiently. A few prior works leverage this connection, showing that NMC can be used to obtain improved non-malleable commitments [37, 11], and using techniques from the non-malleable commitment literature to obtain NMC [12, 58]. However, the latter direction—tapping into the wealth of techniques from the non-malleable commitment literature to construct NMC—has been largely unexplored, perhaps due to the fact that NMCs are typically unconditionally secure.

In our NMC construction, we begin with the framework of Ball et al. [6], which provides a generic way to construct NMC against tampering classes \mathcal{F} for which sufficiently strong average-case hardness results are known, but requires a CRS. We show how to remove the CRS for particular tampering classes, including the class of bounded, poly-time adversaries. One modification is replacing the public key encryption scheme in the framework of [6] (whose public keys are contained in the CRS) with a non-interactive, non-malleable commitment scheme **NMCom** in the plain model.

At a very high (and overly simplified) level, our NMC, like [6], follows the Naor-Yung [56] paradigm that achieves CCA security of encryption, by composing two instances $\text{Encrypt}(\text{PK}, v), \text{Encrypt}(\text{PK}', v)$ of a public key encryption scheme, followed by a NIZK proof of the equality of encrypted values. In the context of NMC, we replace one instance of encryption with an encoding $E(v)$ that is decodable in some polynomial t time, but has certain complexity theoretic hardness (specified shortly) against the class of circuits of smaller $t' < t$ size. We further replace the other instance of encryption with a non-malleable commitment c to v . Following [56, 6], we provide a reduction that can turn any successful tampering adversary A against NMC, into an adversary B able to “maul” an encoding $E(v)$ of v into a non-malleable commitment \tilde{c} to a related value \tilde{v} . The challenge lies in ensuring that the reduction is “simple”, namely, can be implemented by a circuit of size t' . Then the complexity theoretic hardness that we rely on is that it is impossible for such a circuit to compress an encoding $E(v)$ into a much *shorter* string \tilde{c} correlated to v (despite that the correlation may take exponential time to verify). Such an encoding, E , can be based on the incompressible functions of Applebaum et al. [1], which can be constructed based on assumptions that are widely used in the derandomization literature. (For more details on the NMC construction see the technical overview in Section 1.3).

Connection between Complexity Theory and Non-Malleable Commitments. Another contribution of this work, is to develop new connections between complexity theory and non-malleable cryptography. We show that derandomization assumptions can be employed to build a new primitive we call Quasi Non-Malleable Commitments, which is weaker than standard non-malleable commitments, but nevertheless suffices for constructing NMC. This allows us to avoid adding the assumptions needed for standard non-interactive **NMCom** such as time-lock puzzles or hardness amplifiable injective one-way functions [50, 10].

Recall that in the non-malleable codes setting, encode/decode can be in a larger complexity class than the adversary, and so standard non-interactive **NMCom** is an overkill. This motivates our definition of *Quasi Non-Malleable Commitments* in which the adversary runs in $O(n^{c_{com}})$ -time, whereas the honest parties may run in longer (polynomial) time. To construct Quasi-**NMCom** from assumptions widely used in the derandomization literature, observe that these assumptions allow us to construct polynomial-time computable functions ψ for which non-deterministic advice does not help speed up the computation. This stands in stark contrast to the case of inversion of a one-way function ρ , which becomes easy with non-deterministic advice (as the advice can contain

a pre-image). Following the framework of [50], we construct two types of commitments that are *harder* than each other in different hardness “axes” — namely “BP-time” (corresponding to probabilistic Turing machines) and “non-deterministic (ND)-size” (corresponding to **NP**-circuits—see Sections 1.2 and 2.4). Specifically, one type of commitment com_1 are the standard schemes based on one-way functions ρ , and the other com_2 is based on the function ψ given by derandomization assumptions. The com_1 is much harder to break than com_2 in the axis of “BP-time”, as inverting one-way function ρ is much harder than computing ψ using probabilistic Turing machines. On other hand, com_2 is much harder to break than com_1 in the axis of “ND-size”, where both inverting ρ and computing ψ can be done in poly-size, but computing ψ is significantly harder.

From such mutually harder commitment schemes, we obtain a 4-tag Quasi-NMCom. Then, based on tag-amplification techniques in the literature [46, 10], we increase the number of tags supported to an exponential. It turns out that the *quasi*-setting makes amplification hard, which requires us to introduce a notion of “Double-Agent” adversaries. Informally, double-agent adversaries are probabilistic uniform Turing machines with “large” time complexity, that can also be represented as a distribution over circuits with “small” size complexity (see Section 2.1 for additional details). Post-amplification, our final Quasi-NMCom construction employs the same assumptions as Theorem 1. We believe these techniques may be useful for other applications in similar quasi-settings.

1.2 Background on Assumptions

In this section we provide some background on the assumptions that we use.

*On **P**-certificates.* **P**-certificates were introduced by [15] in pursuit of constant-round concurrent zero-knowledge. Loosely, a **P**-certificate is a non-interactive proof system that allows a prover to convince an efficient verifier of the validity of any statement in **P** via a short proof. In particular, both the proof length and the run-time of the verifier are bounded by some fixed polynomial, but the system should work any language in **P** (the prover’s efficiency should be comparable to the statement). CS-proofs [54] imply **P**-certificates, but unlike the former, the latter assumption is falsifiable.

*On “**E** requires circuits of exponential size.”* A fundamental family of questions in theoretical computer science is when and where randomness helps (vs strictly deterministic procedures). While it is widely believed that $\text{BPP} = \mathbf{P}$ (i.e., any efficient, randomized decision procedure can be efficiently simulated by a deterministic procedure), whether the equality indeed holds is still an open problem. This particular question ($\text{BPP} = \mathbf{P}$?) and others in the domain of *derandomization* have deep connections to cryptography.

In the 1980s, Yao [70] showed that one-way permutations suffice to create pseudorandom generators (PRG) for poly-time computation. PRGs expand a small sequence of uniform random bits to a long sequence of pseudorandom bits

that “fool” classes of procedures in the sense their behavior is essentially the same as if they were given truly random bits. In this sense, PRGs give a canonical technique for derandomizing decision procedures: running the procedure on multiple outputs of the PRG in parallel and taking majority of the obtained result. Later, it was shown that essentially minimal cryptographic assumptions (one-way functions) suffice for constructing PRGs [42].

However, while most cryptography implies non-trivial derandomization, there seem to be inherent barriers to statements of the converse form. In fact, the so-called “cryptographic” PRG’s yield, in two aspects, much more than what is required for derandomization since (a) the output of such PRGs fool *any* polynomial time procedure (including procedures that run in much more time than the PRG itself) and (b) such PRGs guarantee that the behavior of poly-time procedures is only negligibly different from their behavior on true randomness. On the other hand, one-way functions are not known to imply $\mathbf{P} = \mathbf{BPP}$ because known constructions only “stretch” random bits into polynomially many random bits (whereas exponential stretch is required for canonical simulation).

Capitalizing on these observations, Nisan and Wigderson [57] gave a generic means of constructing PRGs which “fool” a certain class of circuits \mathcal{C} , from any function that is hard-on-average for a slightly enlarged class of circuits. In particular, this in some sense reduces the problem of explicit derandomization to proving strong circuit lower bounds on explicit functions. To this end, later work showed that, in fact, simply assuming there is a language in \mathbf{E} that does not have circuits size $2^{\beta n}$ for some $\beta > 0$ (for almost all n), is sufficient to derandomize BPP [43, 67]. Moreover, because \mathbf{E} has complete problems, this yields explicit pseudorandom generators. However, for reasons alluded to above, this assumption is, to best of our knowledge, incomparable to standard cryptographic assumptions.

This latter (worst-case) conjecture and its generalization has appeared in a variety of contexts pertaining to derandomization and other questions in computational complexity [53, 57, 2, 43, 67, 41, 47, 65, 66, 29, 55, 68, 64, 35, 40, 22]. The conjectures we are concerned with in this work take the following form (following [1]):

Assumption 1 (\mathbf{E} is hard for exponential size X -circuit) *There exists a problem \mathbf{L} in $\mathbf{E} = \mathbf{DTIME}(2^{O(n)})$ and a constant $\beta > 0$, such that for every sufficiently large n , X -circuits of size $2^{\beta n}$ fail to compute the characteristic of \mathbf{L} on inputs of length n .*

where X -circuits can be circuits of type, $\{\text{non-deterministic, co-non-deterministic, NP, } \Sigma_i\}$. See Section 2.4, for definitions of these types of circuits. While these types of assumptions are independently interesting, in this work we will utilize some surprising implications outside of derandomization.

Recently, Applebaum et al. [1] presented (explicit) constructions of poly-time computable incompressible functions based on the assumption that \mathbf{E} is hard for exponential size non-deterministic circuits (based on the extractors for samplable distributions of Trevisan and Vadhan [68]). Loosely, a function,

ψ is incompressible for a class if no procedure in that class can “shrink” an input to the function, x , such that $\psi(x)$ can later be recovered. Note that, to our knowledge, it is not known how to construct incompressible functions from standard cryptographic assumptions (unlike the case of derandomization).

Barak et al. [8] observed that similar assumptions can be used to construct cryptographic primitives. In particular, they showed that if $\mathbf{E} = \mathbf{DTIME}(2^{O(n)})$ contains a function with *co-non-deterministic* circuit complexity $2^{\Omega(n)}$, then there exists (explicit) non-interactive witness indistinguishable proof systems for $\mathbf{L} \in \mathbf{NP}$ (additionally assuming the existence of trapdoor permutations). They also showed that the same assumption can be used to construct a non-interactive bit commitment scheme from a one-way function.

In this work, we use the above results and demonstrate new connections between these assumptions and non-malleable cryptography. In particular we show that if Assumption 1 holds for \mathbf{NP} -circuits and (sub-exponential) one-way functions exist, then we can construct quasi-non-malleable commitment schemes. We combine our construction of such commitment schemes along with NIZK proofs based on the NIWI of [8], as well as the incompressible functions of [1], to obtain our main result: a family of efficient non-malleable codes secure against tampering by uniform algorithms running in time $O(n^c)$.

1.3 Technical Overview

We begin by recalling (a simplified version of) the template for constructing non-malleable codes against complexity class \mathcal{F} (based on the Naor-Yung double encryption paradigm [56]) introduced in the work of Ball et al. [6]:

The CRS contains a public key PK for an encryption scheme $\mathcal{E} = (\text{Gen}, \text{Encrypt}, \text{Decrypt})$, and a CRS crs for a simulation-sound, non-interactive zero knowledge proof (NIZK). For $b \in \{0, 1\}$, let \mathcal{D}_b denote disjoint distributions over $x_1 \dots x_n \in \{0, 1\}^n$ such that $\psi(x_1 \dots x_n) = b$, where ψ is poly-time computable, yet every $f \in \mathcal{F}$ has low correlation with ψ .

To encode a bit b :

1. Randomly choose string $x_1 \dots x_n$ from \mathcal{D}_b
2. Compute $c \leftarrow \text{Encrypt}_{\text{PK}}(b)$.
3. Compute a NIZK proof T of “consistency”: $\exists b' \in \{0, 1\}$ s.t. $x_1 \dots x_n$ is in the support of $\mathcal{D}_{b'}$ and b' is the plaintext underlying c .
4. Output $(x_1 \dots x_n, c, T)$.

To decode $(x_1 \dots x_n, c, T)$:

1. Verify the NIZK proof T .
2. If it accepts, output $\psi(x_1 \dots x_n)$.

The proof of [6] proceeds (loosely) as follows: In the first hybrid they switch to simulated proof T' . Then they switch c , in the “challenge” encoding to an encryption of garbage c' , and next switch to an alternative decoding algorithm *in*

\mathcal{F} , which requires the trapdoor SK (corresponding to the public key PK which is contained in the CRS). If, in the final hybrid, decodings of tampered encodings depend on b , a circuit in \mathcal{F} can be constructed, whose output is correlated with the hard function ψ , reaching a contradiction. While [6] do in fact show that the CRS can be removed for constructions against certain classes \mathcal{F} of tampering, naively, their approach requires a CRS in two seemingly inherent ways: First, the CRS allows the use of the secret key trapdoor SK in the alternate decoding algorithm and second, it allows the use of a simulation-sound NIZK, which requires CRS.

In this work, we make two crucial observations that allow us to eliminate the CRS from the above construction. First, we consider a stronger notion of hardness for ψ , known as *incompressibility* (in fact, this hardness notion was already implicitly used in [6] for their multi-bit construction). Briefly, if a function ψ is incompressible by circuit class \mathcal{C} , it means that for $t \ll n$, for any (computationally unbounded) Boolean function $F : \{0, 1\}^t \rightarrow \{0, 1\}$ and any $C : \{0, 1\}^n \rightarrow \{0, 1\}^t \in \mathcal{C}$, the output of $F \circ C(x_1, \dots, x_n)$ is uncorrelated with $\psi(x_1, \dots, x_n)$ (over uniform choice of x_1, \dots, x_n). Now, since F is allowed to be computationally unbounded, we may consider an F that decrypts the ciphertext $c = \text{Encrypt}_{\text{PK}}(b)$ by brute force search. To elaborate, instead of using SK to efficiently decrypt c in complexity class \mathcal{C} , the alternate decoding algorithm D' is split into two parts $D' = D'_2 \circ D'_1$, where D'_1 can be implemented in \mathcal{F} , but has short output length, whereas D'_2 is computationally unbounded. Specifically, D'_1 checks the proof T and then outputs the entire ciphertext c (which is fine so long as the length of c is sufficiently smaller than n), and, due to the incompressibility property of ψ , we must still have that the output of $D' = D'_2 \circ D'_1$ is uncorrelated with $\psi(x_1, \dots, x_n)$. This eliminates the need of providing a trapdoor to the alternate decoding algorithm and so instead of using a public key encryption scheme, we may use a non-interactive statistically binding commitment scheme, which can be constructed from injective one-way function or from derandomization assumptions and any one-way function [8].⁴

Next, we note that it is possible to construct a NIZK proof system in the plain (no-CRS) model (i.e. “One-Message Zero Knowledge”), with soundness against uniform adversaries. To do so, one first constructs a non-interactive witness indistinguishable proof system (NIWI) in the plain model (based on standard cryptographic assumptions and derandomization assumptions [8]) and then converts from witness indistinguishability to full zero knowledge using the well-known FLS paradigm [28]. Specifically, the simulator will be given a trapdoor witness based on a problem that is computationally hard for *uniform* PPT adversaries such as finding a collision in a keyless collision resistant hash function. The problem with this approach is that in the proof

⁴ As we will see, in our setting of non-malleable codes against polynomially-bounded adversaries, our construction requires such derandomization assumptions in any case and so only standard one-way function is required in addition. However, for simplicity we will assume injective one-way function in the remainder of the exposition in this section.

sketch outlined above, we actually require *simulation-sound* NIZK, as opposed to regular NIZK. In simulation-sound NIZK, the soundness properties must hold, even after the adversary sees a simulated proof of a false statement. Whereas various constructions of (one-time) simulation-sound NIZK rely on embedding a trapdoor within the CRS (cf. [63, 52]), our approach to achieve the simulation-soundness property without CRS is to replace the commitment c (which replaced the encryption $\text{Encrypt}_{\text{pk}}(b)$ as described above) with a *non-interactive, non-malleable* commitment scheme. Unfortunately, currently known non-interactive, non-malleable commitment schemes require somewhat non-standard assumptions such as time-lock puzzles or hardness amplifiable injective one-way functions [50, 10], whereas our goal is to minimize assumptions. As we will see, the fact that our commitment scheme is only required to be non-malleable against adversaries in a restricted circuit class \mathcal{F} , allows us to obtain non-interactive, non-malleable commitments, while reducing assumptions.

Instantiating the Paradigm In this work we construct non-malleable codes against the class \mathcal{F} of uniform, polynomial-bounded tampering functions. Crucially, we will do so (1) *without* relying on CRS (2) with codeword length that is independent of the polynomial time bound (note that if the codeword length is longer than the polynomial time bound then the adversary does not even get to read the entire input, also it’s trivial to construct these) and (3) while reducing computational assumptions, to the extent possible.

Specifically, in addition to standard cryptographic assumptions, we will assume standard derandomization-type assumptions such as those discussed in the previous section. We also require the notion of \mathbf{P} -certificates, which seem to be necessary to implement the above high-level paradigm, as we discuss next. To see why this is so, note that the statement proved in NIZK proof T , involves proving that $\psi(x_1, \dots, x_n)$ is equal to some value. Note that ψ is a polynomial-time computable function, but that intrinsic in the approach is choosing ψ that is hard to compute in the specific polynomial time bound $T(n)$ corresponding to tampering class \mathcal{F} . Moreover, we require that the size of the proof T be independent of the polynomial time bound $T(n)$, and so in particular the size of the proof T must be independent of the time required to compute ψ . This is now exactly the notion of a \mathbf{P} -certificate.

We also note that given the above paradigm for encoding of a single bit, it is straightforward to obtain a scheme for the encoding of multiple bits (by individually encoding each bit and then using a single proof T to “wrap” together the individual encodings). The only restriction will be that the number of bits, m , that are encoded, multiplied by the length of a bit commitment, λ , should be sufficiently smaller than n , the input length of the function ψ . See Section 3 for additional details.

Instantiation of ψ Recall that for the above approach to work, we must instantiate ψ with a function that is incompressible against polynomially-bounded adversaries. Fortunately, such a construction was given by [1], based on a derandomization-type assumption. See Section 4 for additional details.

Instantiation of NMCom In fact, as discussed previously, we note that we do not need full-fledged NMCom, but only *Quasi* NMCom, i.e. NMCom with the following two relaxations: (1) The commitment scheme is only secure against bounded-poly (in fact “Double-Agent”) adversary and distinguisher (2) The complexity of the honest sender/receiver may be greater than the complexity of the adversary. To construct Quasi-NMCom, we adopt the approach of [50] to initially construct a commitment scheme with small number of tags, and use the fact that the derandomization assumptions that we employ in this work are believed to hold even against *non-deterministic* adversaries. In particular, we employ the well-studied assumption that **E** is hard for adversaries represented as exponential size **NP**-circuits—or circuits with access to a SAT oracle (See Sections 1.2 and 2.4 for further discussions on these assumptions). To construct our NMCom scheme, we start off with two different types of commitments, Type 0 and Type 1 such that if we get a Type 0 on left, we can extract from Type 1 on the right without breaking the security of Type 0 and vice versa. Each commitment consists of an input x to a Boolean function ψ' (with logarithmic input length) that is hard for **NP**-circuits of size $2^{\epsilon_3 \cdot \text{input length}}$ to compute as well as the output y of an injective OWF ρ , which is hard for ppt adversaries running in time $2^{\text{input length}^{\epsilon_3}}$.⁵ Each of these can be considered as a commitment to a bit (given x , the output of $\psi'(x)$ is the committed value and given y , a hardcore bit of ρ) and the final committed value is the xor of the two bits committed.

Type 0: input length $c_1 \log(n)$ to ψ' , input length $n^{\epsilon'_1}$ to ρ .
Type 1: input length $c_2 \log(n)$ to ψ' , input length $n^{\epsilon'_2}$ to ρ .

Set $c_2 > c_1 > \epsilon'_1 > \epsilon'_2$ so that (1) $n^{c_1} < n^{\epsilon_3 \cdot c_2}$ and (2) $2^{n^{\epsilon'_2}} < 2^{n^{\epsilon'_3 \cdot \epsilon'_1}}$. We now consider the two possible cases:

Type 0 on left, Type 1 on right. Extract by inverting the injective OWF ρ in deterministic time $2^{n^{\epsilon'_2}}$ and computing ψ' in deterministic time n^{c_2} . Note that this does not allow breaking injective OWF ρ with input length $n^{\epsilon'_1}$, which is secure against time $2^{n^{\epsilon'_3 \cdot \epsilon'_1}} > 2^{n^{\epsilon'_2}}$.

Type 1 on left, Type 0 on right. Extract by computing ψ' in deterministic time n^{c_1} and inverting the injective OWF ρ with an **NP**-circuit of size $n^{\epsilon'_1}$. Note that this does not allow breaking hardness of ψ' with input length $c_2 \log(n)$, which is secure against **NP**-circuits of size $n^{\epsilon_3 \cdot c_2} > n^{c_1}$.

See Figure 1 for a summary and [4] for additional details.

The above 2-tag commitment scheme can then be straightforwardly extended to work for 4 tags, at which point amplification techniques from [46] can be

⁵ For this exposition, we assume for simplicity that ψ' can be computed in deterministic time $2^{\text{input length}}$ and that the injective OWF has linear circuit size. Recall that we do not require injective OWF and that any statistically binding, non-interactive commitment scheme is sufficient, but that for simplicity we assuming injective OWF in this exposition.

	Input length to ψ'	Hardness of ψ'		Input length to ρ	Hardness of ρ	
		D	ND		D	ND
Type 0	$c_1 \cdot \log(n)$	$n^{\epsilon_3 \cdot c_1}$	$n^{\epsilon_3 \cdot c_1}$	$n^{\epsilon'_1}$	$2^{n^{\epsilon'_3 \cdot \epsilon'_1}}$	$n^{\epsilon'_1}$
Type 1	$c_2 \cdot \log(n)$	$n^{\epsilon_3 \cdot c_2}$	$n^{\epsilon_3 \cdot c_2}$	$n^{\epsilon'_2}$	$2^{n^{\epsilon'_3 \cdot \epsilon'_2}}$	$n^{\epsilon'_2}$

Fig. 1. ψ' and ρ are the functions described in the paragraph above. D stands for deterministic and ND stands for “non-deterministic” hardness. We set parameters so that $c_2 > c_1 > \epsilon'_1 > \epsilon'_2$.

applied to obtain NMCom with number of tags exponential in the security parameter. The analysis of the amplified scheme is somewhat different than in prior work, since our analysis must carefully take into account that some assumptions are inherently uniform (One-Message Zero Knowledge) and some assumptions (hardness of ψ') are inherently non-uniform (the adversary in the proof is so limited that it does not have time to *generate* new commitments on its own and therefore must receive them as non-uniform advice when reducing security to the hardness of computing ψ'). To solve this problem, we introduce the notion of “Double Agent” adversaries (as discussed in the introduction) and provide a proof of security of our amplified NMCom scheme against this class of adversaries. See [4] for additional details.

1.4 Related Work

Non-Malleable Codes. Non-malleable codes (NMC) were introduced by Dziembowski, Pietrzak and Wichs in their seminal work [25]. While there has been a long line of important results for various tampering classes, due to space limitations, we discuss here only the results most relevant to this work.

As discussed extensively in the introduction, Faust et.al [27] constructed efficient information-theoretically secure NMC in the CRS model, resilient against tampering function classes \mathcal{F} which can be represented as circuits of size $\text{poly}(n)$. Another important result by Cheraghchi and Guruswami [14] showed the existence of information theoretically secure NMC against tampering families \mathcal{F} of size $|\mathcal{F}| \leq 2^{2^{\alpha n}}$ with optimal rate $1 - \alpha$. They achieve error $\epsilon \in O(1/\text{poly}(n))$ as the run-time of the encoding and decoding algorithms is proportional to $\text{poly}(1/\epsilon)$ where ϵ is the error probability.

Ball et.al [5] constructed efficient information theoretic secure NMC against n^δ -local tampering functions, for any constant $\delta > 0$. This class includes tampering functions, which can be represented as constant depth circuits with bounded fan-in i.e NC^0 circuits. Chattopadhyay and Li [13] constructed NMC against AC^0 tampering functions from seedless non-malleable extractors, although the codeword length of this construction is super-polynomial in the message length n . Faust et.al [26] considered non-malleable codes against space bounded tampering adversaries in the random oracle model. The construction achieves a new notion of *leaky* continuous non-malleable codes (with self-destruct

property), where the adversary is allowed to learn some bounded $\log(|m|)$ bits of information about the underlying message m .

Recently, Ball et.al [6] presented a generic framework to construct NMC against tampering function classes for which average-case hardness bounds are known. They also instantiated their framework to construct the first efficient, computationally secure multi-bit NMC against tampering functions which can be represented as constant-depth circuits with unbounded fan-in (AC^0 tampering), as well as against tampering functions which can be represented as bounded depth decision tree. Additionally, they showed that the framework can be used to construct information-theoretic NMC against space-bounded streaming tampering. Information-theoretic secure, efficient NMC against AC^0 tampering were subsequently constructed by [3].

Derandomization and Cryptography The connection between derandomization techniques with cryptography was first explored by Barak et.al. [8], who constructed one-message witness indistinguishable proof systems (non-interactive commitment scheme) in the plain model based on trapdoor permutations (one-way functions) in addition to the derandomization assumptions. Recently, Applebaum et.al. [1] constructed incompressible functions against the class of bounded polynomial time functions from similar assumptions.

Non-Malleable Commitments Non-malleable commitments have been studied extensively since their introduction by [21] in their seminal paper. Great progress has been made in reducing the interaction between the sender and the receiver, while minimizing computational assumptions. We list just some of the results in this line of work [7, 59, 60, 48, 61, 36, 49, 38, 37, 16, 17, 45]. Recently, Lin, Pass, and Soni [50] gave a construction of a non-interactive, fully-concurrent, non-malleable commitment scheme secure against uniform adversaries based on sub-exponential non-interactive commitment schemes, non-interactive witness indistinguishable proof systems (NIWI), uniform collision resistant hash functions, and time-lock puzzles [62]. When replacing the uniform collision resistant hash functions with a family of collision resistant hash functions, their protocol becomes 2-round. Khurana and Sahai [46] constructed 2-round non-malleable commitments with bounded concurrency from standard sub-exponential assumptions. Bitansky, and Lin [10] gave a construction of a non-interactive, fully-concurrent, non-malleable commitment scheme from multi-collision-resistant keyless hash functions, sub-exponentially-secure time-lock puzzles, and other standard assumptions.

2 Definitions

2.1 Notation

When comparing distribution ensembles $\mathcal{D} = \{\mathcal{D}_n\}_{n \in \mathbb{N}}$, $\mathcal{D}' = \{\mathcal{D}'_n\}_{n \in \mathbb{N}}$, we use the notation $\mathcal{D} \stackrel{\mathcal{G}, \mathcal{S}}{\approx} \mathcal{D}'$, where \mathcal{G} , \mathcal{S} are sets, to indicate that for sufficiently large

n , every distinguisher $D \in \mathcal{G}$ distinguishes \mathcal{D}_n from \mathcal{D}'_n with probability at most $p(n)$, for some $p(\cdot) \in \mathcal{S}$. When comparing functions p, p' , we use the notation $p(n) \stackrel{\mathcal{S}}{\approx} p'(n)$, where \mathcal{S} is a set, to indicate that $|p(n) - p'(n)| \in \mathcal{S}$.

We consider “Double-Agent” adversaries A in computational classes denoted by $\mathbf{BPTIME}(T(n)) \cap \mathbf{SIZE}(t(n))$, for some functions $T(\cdot), t(\cdot)$. Intuitively, this computational class contains probabilistic uniform Turing machines A with “large” time complexity $T(n)$, that can also be represented as a distribution over circuits with “small” size complexity $t(n)$. Informally, this is possible since A can be split into subroutines in such a way that subroutines that require “large” time complexity can all be replaced with non-uniform advice. Formally, $A \in \mathbf{BPTIME}(T(n)) \cap \mathbf{SIZE}(t(n))$ if the following hold:

- $A = (A_1, A_2)$.
- $A_1 \in \mathbf{BPTIME}(T(n)), A_2 \in \mathbf{BPTIME}(t(n))$.
- A_1 receives only security parameter 1^n as input and produces output of length at most $t(n)$.
- A_2 receives the input of A as its input, along with the output of A_1 .

Note that, since A_1 takes only security parameter as input, the output of A_1 , can be viewed as non-uniform advice to A_2 . Thus, we can convert such a uniform adversary $A = (A_1, A_2)$ into a distribution over non-uniform circuits of size $t(n)$ with identical behavior to A .

2.2 Non-Malleable Codes

Definition 1 (Coding Scheme). Let $\Sigma, \hat{\Sigma}$ be sets of strings, and $\kappa, \hat{\kappa} \in \mathbb{N}$ be some parameters. A coding scheme consists of two algorithms (E, D) with the following syntax:

- The encoding algorithm (perhaps randomized) takes input a message in Σ and outputs a codeword in $\hat{\Sigma}$.
- The decoding algorithm takes input a codeword in $\hat{\Sigma}$ and outputs a message in Σ .

We require that for any message $\text{msg} \in \Sigma$, $\Pr[D(E(\text{msg})) = \text{msg}] = 1$.

Definition 2 ($O(1/p(n))$ -Non-malleability [25]). Let n be the security parameter, \mathcal{F} be some family of functions. For each function $f \in \mathcal{F}$, and $\text{msg} \in \Sigma$, define the tampering experiment:

$$\mathbf{Tamper}_{\text{msg}}^f \stackrel{\text{def}}{=} \left\{ \begin{array}{l} c \leftarrow E(\text{msg}), \tilde{c} := f(c), \tilde{\text{msg}} := D(\tilde{c}). \\ \text{Output} : \tilde{\text{msg}}. \end{array} \right\},$$

where the randomness of the experiment comes from the encoding algorithm. We say a coding scheme (E, D) is $O(1/p(n))$ -non-malleable with respect to \mathcal{F} if for each $f \in \mathcal{F}$, there exists a PPT simulator Sim such that for any message $\text{msg} \in \Sigma$, we have

$$\mathbf{Tamper}_{\text{msg}}^f \stackrel{PPT, O(1/p(n))}{\approx} \mathbf{Ideal}_{\text{Sim}, \text{msg}} \stackrel{\text{def}}{=} \left\{ \begin{array}{l} \tilde{\text{msg}} \cup \{\text{same}^*\} \leftarrow \text{Sim}^{f(\cdot)}. \\ \text{Output : msg if output of Sim is same}^*; \\ \text{otherwise } \tilde{\text{msg}}. \end{array} \right\}$$

Definition 3 ($O(1/p(n))$ -Medium Non-malleability). Let n be the security parameter, \mathcal{F} be some family of functions. Fix $\text{msg} \in \Sigma$. Let $c \leftarrow \mathbf{E}(\text{msg})$ and let $g(\cdot, \cdot)$ be a predicate such that for every $f \in \mathcal{F}$,

$$\Pr[g(c, f(c)) = 1] \wedge \mathbf{D}(f(c)) \neq \text{msg} \leq \text{negl}(n).$$

For g as above, each function $f \in \mathcal{F}$, and $\text{msg} \in \Sigma$, define the tampering experiment

$$\mathbf{MediumNM}_{\text{msg}, g}^f \stackrel{\text{def}}{=} \left\{ \begin{array}{l} c \leftarrow \mathbf{E}(\text{msg}), \tilde{c} := f(c), \tilde{\text{msg}} := \mathbf{D}(\tilde{c}) \\ \text{Output : same}^* \text{ if } g(c, \tilde{c}) = 1, \text{ and } \tilde{\text{msg}} \text{ otherwise.} \end{array} \right\}$$

The randomness of this experiment comes from the randomness of the encoding algorithm. We say that a coding scheme (\mathbf{E}, \mathbf{D}) is $O(1/p(n))$ -medium non-malleable with respect to \mathcal{F} if there exists a g as above and for any $\text{msg}, \text{msg}' \in \Sigma$ and for each $f \in \mathcal{F}$, we have:

$$\{\mathbf{MediumNM}_{\text{msg}, g}^f\}_{n \in \mathbb{N}} \stackrel{PPT, O(1/p(n))}{\approx} \{\mathbf{MediumNM}_{\text{msg}', g}^f\}_{n \in \mathbb{N}}$$

It is straightforward to check that medium non-malleability implies standard non-malleability.

2.3 Non-Interactive Commitment Scheme

Definition 4 (Commitment Scheme). A (non-interactive) commitment scheme for the message space $\{0, 1\}^m$, is a pair $(\text{Com}, \text{Open})$ such that:

- For all $\text{msg} \in \{0, 1\}^m$, $(c, d) \leftarrow \text{Com}(m)$ is the commitment/opening pair for the message msg .
- $\text{Open}(\text{msg}, c, d) \rightarrow \{0, 1\}$, where 1 indicates that d is a valid opening of c to msg and 0 is returned otherwise.

The commitment scheme must satisfy the standard correctness requirement,

$$\forall m \in \mathbb{N}, \forall \text{msg} \in \{0, 1\}^m, \Pr[\text{Open}(\text{msg}, \text{Com}(\text{msg})) = 1] = 1$$

where the probability is taken over the randomness of Com .

We will consider *statistically* binding commitment schemes. For the formal definitions of the Hiding and Binding properties, see [4].

Well-formed Commitments: Let $\text{val}(\cdot)$ be a function which takes an arbitrary commitment c as an input. val outputs msg if \exists unique msg such that $\text{Open}(\text{msg}, c, \cdot) = 1$, and outputs \perp otherwise.

Definition 5 (Tag-based Commitment Scheme [50]). A commitment scheme $(\text{Com}, \text{Open})$ is a tag-based commitment scheme with $\tau(m)$ number of tags if, in addition to the message msg , the sender (committer) and receiver also receive a “tag” of length $\text{poly}(\log(\tau(m)))$ as common input.

If $\tau(m)$ is exponential in security parameter m , we omit the prefix $\tau(m)$ and refer to the commitment scheme as simply a tag-based commitment scheme.

Man In The Middle Execution (MIM): Let $(\text{Com}, \text{Open})$ be a tag-based commitment scheme, and A an adversary. For security parameter m , consider the following interactions by $A(1^m)$:

- *Left interaction:* $A(1^m)$ interacts with the sender and receives commitment to a message msg of length m using identity tag as $c \leftarrow \text{Com}(\text{msg}, \text{tag})$.
- *Right interaction:* $A(1^m)$ interacts with the receiver and tries to commit to related message $\tilde{\text{msg}}$ using identity $\tilde{\text{tag}}$ of its choice. Specifically, for the commitment \tilde{c} sent to the receiver, let $\tilde{\text{msg}} = \text{val}(\tilde{c})$. Furthermore, if $\tilde{\text{tag}} = \text{tag}$, then we set $\tilde{\text{msg}} = \perp$.

Let $\text{mim}_C^A(\text{msg})$ denote the random variable that describes $\tilde{\text{msg}}$ that A commits to in the right interaction along with its output in the MIM execution $\text{MIM}_C^A(\text{msg})$ as described above.

Definition 6 ($O(1/p(m))$ -Non-Malleability against \mathcal{G} [50]).

A tag-based commitment scheme $\mathcal{C} = (\text{Com}, \text{Open})$ is said to be $O(1/p(m))$ -non-malleable against \mathcal{G} if $\forall A \in \mathcal{G}$, the following ensembles are indistinguishable,

$$\left\{ \text{mim}_C^A(\text{msg}_0) \right\}_{m \in \mathbb{N}, \text{msg}_0 \in \{0,1\}^m} \stackrel{\mathcal{G}, O(1/p(m))}{\approx} \left\{ \text{mim}_C^A(\text{msg}_1) \right\}_{m \in \mathbb{N}, \text{msg}_1 \in \{0,1\}^m}.$$

2.4 Incomputable and Incompressible Functions

Definition 7 (Incomputable Function [1]). A function $\psi : \{0,1\}^n \rightarrow \{0,1\}^m$ is incomputable by a function class \mathcal{C} if ψ is not contained in \mathcal{C} . We say that f is ϵ -incomputable by \mathcal{C} if for every function $C : \{0,1\}^n \rightarrow \{0,1\}^m$ in \mathcal{C} , $\Pr[C(x) = f(x)] \leq \frac{1}{2^m} + \epsilon$ for uniform random $x \leftarrow \{0,1\}^n$.

Definition 8 (Incompressible Function [23]). A function $f : \{0,1\}^n \rightarrow \{0,1\}^m$ is incompressible by a function class $\mathcal{C} : \{0,1\}^n \rightarrow \{0,1\}^\ell$ if for every function $D : \{0,1\}^\ell \rightarrow \{0,1\}^m$, there exists $x \in \{0,1\}^n$ such that $D(C(x)) \neq f(x)$. We say that f is ϵ -incompressible by \mathcal{C} if for every function $D : \{0,1\}^\ell \rightarrow \{0,1\}^m$, $\Pr[D(C(x)) = f(x)] \leq \frac{1}{2^m} + \epsilon$ for uniform random $x \leftarrow \{0,1\}^n$. We say that f is ℓ -incompressible (resp. (ℓ, ϵ) -incompressible) by a class \mathcal{C} if for every $C : \{0,1\}^n \rightarrow \{0,1\}^\ell$ in \mathcal{C} , f is incompressible (resp. ϵ -incompressible) by C .

Definition 9 (Non-deterministic Circuits and NP Circuits [1]). A non-deterministic circuit C has additional “non-deterministic input wires.” We say that the circuit C evaluates to 1 on x if and only if there exists an assignment to the non-deterministic wires that makes C output 1 on x . An oracle circuit $C^{(\cdot)}$ is a circuit which in addition to the standard gates uses an additional gate (potentially with large fan-in). When instantiated with specific boolean function A , C^A is the circuit in which the additional gate is A . Given boolean function $A(x)$, an A -circuit is a circuit that is allowed to use A gates in addition to the standard gates. An **NP**-circuit is a **SAT**-circuit (where **SAT** is the satisfiability function).

The size of all circuits is the total number of wires and gates.

We now present commonly used assumptions in the derandomization literature to *explicitly* construct pseudorandom generators. [2, 57, 67, 47, 65, 66, 29, 55, 68, 64, 35, 40, 8, 22]:

Assumption 2 (E is hard for exponential size X-circuits) *There exists a problem L in $E = \text{DTIME}(2^{O(n)})$ and a constant $\beta > 0$, such that for every sufficiently large n , X -circuits of size $2^{\beta n}$ fail to compute the characteristic function of L on inputs of length n , where $X \in \{\text{non-deterministic, co-non-deterministic, NP}\}$.*

Theorem 2 (Theorem 1.3, 1.10 [1]). *If E is hard for exponential size X -circuits, where $X \in \{\text{non-deterministic, co-non-deterministic, NP}\}$ (Assumption 2), then for every constant $c > 1$ there exists a constant $a > 1$ such that for every sufficiently large n , and every r such that $a \log n \leq r \leq n$ there is a function $\psi : \{0, 1\}^r \rightarrow \{0, 1\}$ that is n^{-c} -incomputable for size n^c X -circuits, Furthermore, ψ is computable in time $\text{poly}(n^c)$ (or $\text{poly}(n)$).*

We define NIZK without CRS against uniform adversaries and NIWI in [4]. In the remainder of this section, we focus on instantiations of the above primitives.

Theorem 3 ([8]). *Assume that E is hard for exponential size co-non-deterministic circuits and that (subexponentially secure) trapdoor permutations (resp. one-way functions) exist. Then every language in **NP** has a (subexponentially indistinguishable) NIWI proof system (resp. non-interactive commitment scheme).*

Moreover, by correctly setting the output length of the commitment scheme in terms of the security parameter n , we obtain a non-interactive perfectly binding and computationally hiding commitment scheme, such that given a commitment c , the committed message (i.e., $\text{val}(c)$) can be computed by a 2^{n^ϵ} -time algorithm, where n is the security parameter and ϵ is some constant.

To go from NIWI to NIZK, one can apply the well-known FLS technique [28]. The simulator is provided with a trapdoor via non-uniform advice, which is not known to the uniform adversary in the real world. Note that we choose the

trapdoor such that it *can* be obtained by a uniform adversary running in super-polynomial (sub-exponential) time. Formally, [9] show how to construct NIZK without CRS against uniform adversaries under the following assumptions:

Assumption A: There exists a NIWI proof system for every language $L \in \mathbf{NP}$ with WI against sub-exponential adversaries.

Assumption B: There exists a non-interactive perfectly binding and computationally hiding commitment scheme, such that given a commitment, the message can be computed by a 2^{n^ϵ} -time algorithm, where n is the security parameter and ϵ is some constant.

Assumption C: There exists a language $\Delta \in P$ and constants $\epsilon_1 < \epsilon_2 < 1$ such that:

Δ is hard to sample in time $2^{n^{\epsilon_1}}$: For every probabilistic $2^{n^{\epsilon_1}}$ -time algorithm A , the probability that $A(1^n) \in \Delta \cap \{0, 1\}^n$ is negligible.

Δ is easy to sample in time $2^{n^{\epsilon_2}}$: There exists a $2^{n^{\epsilon_2}}$ algorithm S_Δ such that for every $n \in N$, $\Pr[S_\Delta(1^n) \in \Delta \cap \{0, 1\}^n] = 1$.

Theorem 4 ([9]). *Under Assumptions A, B and C, there exists a NIZK argument system without CRS for \mathbf{NP} with soundness against sub-exponential uniform adversaries and zero-knowledge against sub-exponential adversaries.*

Lemma 1. *If \mathbf{E} is hard for exponential size non-deterministic circuits and \mathbf{P} -cert with soundness against sub-exponential adversaries exists, then Assumption C is true.*

The proof of the lemma can be found in [4].

Corollary 1. *Assuming that \mathbf{E} is hard for exponential size (co-)non-deterministic circuits, the existence of sub-exponential trapdoor permutations, and the existence of \mathbf{P} -cert with soundness against sub-exponential adversaries, there exists a NIZK argument system without CRS for \mathbf{NP} with soundness against sub-exponential uniform adversaries and zero knowledge against sub-exponential adversaries.*

3 Construction for Multi-Bit Messages

Let $\mathcal{C} = (\text{Com}, \text{Open})$ be a tag-based, non-interactive commitment scheme that is perfectly binding (see Definition 2.3). Let $\Pi^{\text{NI}} = (\text{P}^{\text{NI}}, \text{V}^{\text{NI}}, \text{Sim}^{\text{NI}})$ be a non-interactive simulatable proof system. Let $\mathcal{S} = (\text{Gen}, \text{Sign}, \text{Ver})$ be a one-time signature scheme. Let D_0, D_1 be disjoint distributions over $\{0, 1\}^n$. For $\mathbf{b} := b^1, \dots, b^m \in \{0, 1\}^m$, $D_{\mathbf{b}}$ denotes a draw from the product distribution $(D_{b^1}, \dots, D_{b^m})$. We define the following language:

Language \mathcal{L} : $s := ([\mathbf{x}^i]_{i \in [m]}, \mathbf{c}, \text{tag}) \in \mathcal{L}$ iff $\exists \mathbf{b} := b^1, \dots, b^m \in \{0, 1\}^m$ such that for $i \in [m]$, $\mathbf{x}^i = (x_1^i, \dots, x_n^i)$ is in the support of D_{b^i} and \mathbf{c} is a commitment to \mathbf{b} under tag .

$E(\mathbf{b} := b^1, \dots, b^m)$:

1. Choose $(\text{vk}, \text{SK}) \leftarrow \text{Gen}(1^{n'})$, where $n' \ll n$. We assume WLOG $|\text{vk}| = n'$.
2. Sample $\bar{\mathbf{x}} := \mathbf{x}^1, \dots, \mathbf{x}^m \leftarrow D_{\mathbf{b}}$, where for $i \in [m]$, $\mathbf{x}^i = x_1^i, \dots, x_n^i$.
3. Compute $(\mathbf{c}, \mathbf{d}) \leftarrow \text{Com}(\mathbf{b}, \text{tag} := \text{vk})$.
4. Compute a non-interactive, simulatable proof T proving $([\mathbf{x}^i]_{i \in [m]}, \mathbf{c}, \text{vk}) \in \mathcal{L}$.
5. Compute $\sigma \leftarrow \text{Sign}(\text{SK}, ([\mathbf{x}^i]_{i \in [m]}, \mathbf{c}, T))$.
6. Output $\text{CW} := (\text{vk}, [\mathbf{x}^i]_{i \in [m]}, \mathbf{c}, T, \sigma)$.

$D(\text{CW})$:

1. Parse $\text{CW} := (\text{vk}, [\mathbf{x}^i]_{i \in [m]}, \mathbf{c}, T, \sigma)$
2. Check that $\text{Ver}(\text{vk}, \sigma, ([\mathbf{x}^i]_{i \in [m]}, \mathbf{c}, T)) = 1$.
3. Check that \mathbf{V}^{NI} outputs 1 on proof T .
4. If yes, output $[b^i]_{i \in [m]}$ such that for all $i \in [m]$, x_1^i, \dots, x_n^i is in the support of D_{b^i} . If not, output $\mathbf{0}$.

Fig. 2. NON-MALLEABLE CODE (E, D), SECURE AGAINST \mathcal{F} TAMPERING.

$E_1(\text{td}, \mathbf{b} := b^1, \dots, b^m)$:

1. Choose $(\text{vk}, \text{SK}) \leftarrow \text{Gen}(1^{n'})$
2. Sample $\bar{\mathbf{x}} := \mathbf{x}^1, \dots, \mathbf{x}^m \leftarrow D_{\mathbf{b}}$, where for $i \in [m]$, $\mathbf{x}^i = x_1^i, \dots, x_n^i$.
3. Compute $(\mathbf{c}, \mathbf{d}) \leftarrow \text{Com}(\mathbf{b}, \text{tag} := \text{vk})$.
4. Simulate, using td , a non-interactive proof T' proving $s := ([\mathbf{x}^i]_{i \in [m]}, \mathbf{c}, \text{vk}) \in \mathcal{L}$.
5. Compute $\sigma \leftarrow \text{Sign}(\text{SK}, ([\mathbf{x}^i]_{i \in [m]}, \mathbf{c}, T'))$.
6. Output $\text{CW} := (\text{vk}, [\mathbf{x}^i]_{i \in [m]}, \mathbf{c}, T', \sigma)$.

Fig. 3. ENCODING ALGORITHM WITH SIMULATED PROOF.

The construction is presented in Figure 2:

Let $\Psi(p, x, y, z)$ be defined as a function that takes as input a predicate p , and variables x, y, z . If $p(x, y) = 1$, then Ψ outputs the m -bit string $\mathbf{0}$. Otherwise, Ψ outputs z .

Theorem 5. *Let (E, D) , E_1 , E_2 , D' and g be as defined in Figures 2, 3, 4, 5 and 6. Let \mathcal{F} be a computational class. If, for every pair of m -bit messages $\mathbf{b}_0, \mathbf{b}_1$ and for every tampering function $f \in \mathcal{F}$, all of the following hold:*

– *Simulation of proofs.*

1. $\Pr[g(\text{CW}_0, f(\text{CW}_0)) = 1] \stackrel{\text{negl}(n)}{\approx} \Pr[g(\text{CW}_1, f(\text{CW}_1)) = 1]$,

$E_2(\text{td}, \mathbf{b} := b^1, \dots, b^m)$:

1. Choose $(\text{vk}, \text{sk}) \leftarrow \text{Gen}(1^{n'})$
2. Sample $\bar{\mathbf{x}} := \mathbf{x}^1, \dots, \mathbf{x}^m \leftarrow D_{\mathbf{b}}$, where for $i \in [m]$, $\mathbf{x}^i = x_1^i, \dots, x_n^i$.
3. Compute $(\mathbf{c}', \mathbf{d}') \leftarrow \text{Com}(\mathbf{0}, \text{tag} := \text{vk})$.
4. Simulate, using td , a non-interactive proof T' proving $s := ([\mathbf{x}^i]_{i \in [m]}, \mathbf{c}', \text{vk}) \in \mathcal{L}$.
5. Compute $\sigma \leftarrow \text{Sign}(\text{sk}, ([\mathbf{x}^i]_{i \in [m]}, \mathbf{c}', T'))$.
6. Output $\text{CW} := (\text{vk}, [\mathbf{x}^i]_{i \in [m]}, \mathbf{c}', T', \sigma)$.

Fig. 4. ENCODING ALGORITHM WITH SIMULATED PROOF AND COMMITMENTS.

$D'(\text{CW}) := D'_2(D'_1(\text{CW}))$:

$D'_1(\text{CW})$:

1. Parse $\text{CW} := (\text{vk}, [\mathbf{x}^i]_{i \in [m]}, \mathbf{c}, T, \sigma)$
2. Check that $\text{Ver}(\text{vk}, \sigma, ([\mathbf{x}^i]_{i \in [m]}, \mathbf{c}, T)) = 1$.
3. Check that \mathbf{V}^{NI} outputs 1 on proof T
4. If not, output \perp , where \perp is a special symbol.
5. If yes, output $(\mathbf{c}, \text{tag} := \text{vk})$.

$D'_2(\mathbf{c}, \text{tag} := \text{vk})$:

1. If $\mathbf{c} = \perp$, output $[0]_{i \in [m]}$ and terminate.
2. Otherwise, check if there exists a string \mathbf{d} and a string $\tilde{\mathbf{b}}$ such that $\text{Open}(\mathbf{d}, \mathbf{c}, \text{vk}, \tilde{\mathbf{b}}) = 1$. If yes, output $\tilde{\mathbf{b}}$. Otherwise, output $[0]_{i \in [m]}$.

Fig. 5. ALTERNATE DECODING PROCEDURE D' .

$g(\text{CW}, \text{CW}^*)$:

1. Parse $\text{CW} = (\text{vk}, [\mathbf{x}^i]_{i \in [m]}, \mathbf{c}, T, \sigma)$, $\text{CW}^* = (\text{vk}^*, [\mathbf{x}^{*i}]_{i \in [m]}, \mathbf{c}^*, T^*, \sigma^*)$.
2. If $\text{vk} = \text{vk}^*$ and $\text{Ver}(\text{vk}^*, \sigma^*, ([\mathbf{x}^{*i}]_{i \in [m]}, \mathbf{c}^*, T^*)) = 1$ then output 1. Otherwise output 0.

Fig. 6. THE PREDICATE $g(\text{CW}, \text{CW}^*)$.

2. $\Psi(g, \text{CW}_0, f(\text{CW}_0), D(f(\text{CW}_0)))$ $PPT, \text{negl}(n)$
 $\Psi(g, \text{CW}_1, f(\text{CW}_1), D(f(\text{CW}_1)))$, \approx
 where $f \in \mathcal{F}$, $\text{CW}_0 \leftarrow E(\mathbf{b}_0)$ and $\text{CW}_1 \leftarrow E_1(\text{td}, \mathbf{b}_0)$.
 – **Simulation of Commitments.**

1. $\Pr[g(\text{CW}_1, f(\text{CW}_1)) = 1] \stackrel{\text{negl}(n)}{\approx} \Pr[g(\text{CW}_2, f(\text{CW}_2)) = 1],$
 2. $\Psi(g, \text{CW}_1, f(\text{CW}_1), \text{D}(f(\text{CW}_1))) \stackrel{PPT, \text{negl}(n)}{\approx} \Psi(g, \text{CW}_2, f(\text{CW}_2), \text{D}(f(\text{CW}_2))),$
- where $f \in \mathcal{F}$, $\text{CW}_1 \leftarrow \text{E}_1(\text{td}, \mathbf{b}_0)$ and $\text{CW}_2 \leftarrow \text{E}_2(\text{td}, \mathbf{b}_0)$.
- **Simulation Soundness.**

$$\Pr[\text{D}(f(\text{CW}_2)) \neq \text{D}'(f(\text{CW}_2)) \wedge g(\text{CW}_2, f(\text{CW}_2)) = 0] \in O(1/n^c),$$

where $f \in \mathcal{F}$, $\text{CW}_2 \leftarrow \text{E}_2(\text{td}, \mathbf{b}_0)$.

- **Hardness of $D_{\mathbf{b}}$ relative to Alternate Decoding.**

1. $\Pr[g(\text{CW}_2, f(\text{CW}_2)) = 1] \stackrel{PPT, O(1/n^c)}{\approx} \Pr[g(\text{CW}_3, f(\text{CW}_3)) = 1],$
2. For every Boolean function, represented by a circuit F over m variables,

$$F \circ \text{D}'(f(\text{CW}_2)) \stackrel{\text{stat}, O(1/n^c)}{\approx} F \circ \text{D}'(f(\text{CW}_3)),$$

where $\text{CW}_2 \leftarrow \text{E}_2(\text{td}, \mathbf{b}_0)$, and $\text{CW}_3 \leftarrow \text{E}_2(\text{td}, \mathbf{b}_1)$.

Then the construction presented in Figure 2 is a $O(1/n^c)$ -non-malleable code for class \mathcal{F} .

We present the proof of Theorem 5 in the full version [4].

4 Multi-Bit NMC Against Bounded Poly Adversaries

We describe the underlying components required to instantiate the generic construction. The tampering class \mathcal{F} corresponds to (uniform) tampering functions that run in time $O(n^{c_A})$, where n is security parameter. The length of the encoding is $L := O(n^{c_\ell})$, for some fixed constant c_ℓ . Therefore, the tampering function is allowed to run in time L^{c_A/c_ℓ} with respect to the input length L .

Let n be the input length for the hard distribution described in Section 4.1. We fix polynomials $t_\psi(n) = n^{c_\psi}$, $t_{\text{com}}(n) = n^{c_{\text{com}}}$ where c_ψ, c_{com} are constants (both greater than c_A) and superpolynomial time bounds $T_{\text{com}}(n)$, $T'_{\text{NIZK}}(n)$, $T_{\text{ZK}}(n)$. such that

- $c_\psi \ll c_{\text{com}}$,
- $T'_{\text{NIZK}}(n) \ll T_{\text{com}}(n)$,
- $T_{\text{ZK}}(n)$ is subexponential.

The distribution described in Section 4.1 is hard for $t_\psi(n)$ -time adversaries. $m \cdot \lambda \ll n$ is the length of the m -bit commitment using the commitment scheme described in Section 4.2, n is set such that $m \cdot \lambda + n' \leq (m+1) \cdot \lambda \in o(n)$ (so n is asymptotically larger than the length of the commitment $m \cdot \lambda$ plus the length of the tag n'). These commitments are hiding for polynomial-time adversaries and quasi-non-malleable for adversaries in $\mathbf{BPTIME}(T_{\text{com}}(n)) \cap \mathbf{SIZE}(t_{\text{com}}(n))$. The the non-interactive simulatable proof system in Section 4.3 has soundness against uniform, poly-time adversaries and zero knowledge against $T_{\text{ZK}}(n)$ time adversaries.

4.1 The Hard Distribution D_b (instance length n , hard against $t_\psi(n)$ -time adversaries)

Theorem 6 ([1]). *If E is hard for exponential size nondeterministic circuits, then for every constant $c_\psi > 1$, there exists a constant $d > 1$ such that for every sufficiently large n , there is a function $\psi : \{0, 1\}^n \rightarrow \{0, 1\}$ that is (ℓ, n^{-c_ψ}) -incompressible for size n^{c_ψ} circuits, where $\ell = n - d \cdot \log n$. Furthermore, ψ is computable in time $\text{poly}(n^{c_\psi}) \in O(n^{c_{\text{com}}})$.*

Setting parameters n, c_ψ, d as above, we let D_b be the uniform distribution over $\mathbf{x} \leftarrow \{0, 1\}^n$, conditioned on $\psi(\mathbf{x}) = b$. The theorem above immediately implies the following:

Claim. Let n, c_ψ, d, ψ be as above, let \tilde{F} be any Boolean function over $(m+1) \cdot \lambda \leq n - d \cdot \log n < (1-\alpha)n$ variables, and let C be a size n^{c_ψ} circuit with input length n and output length m . Then, over random choice of $x \leftarrow \{0, 1\}^n$, $\tilde{F} \circ C(x)$ has correlation at most $1/n^{-c_\psi}$ with $\psi(x)$.

4.2 Commitment scheme $\mathcal{C} = (\text{Com}, \text{Open})$ (length $\lambda \ll n$, hiding for poly-time adversaries, and quasi non-malleable against adversaries in $\text{Bptime}(T_{\text{com}}(n)) \cap \text{SIZE}(t_{\text{com}}(n))$)

We instantiate the commitment scheme $\mathcal{C} = (\text{Com}, \text{Open})$ with the scheme presented in [4]. Recall that the scheme has the following properties:

- Non-interactive with no-CRS.
- Perfectly binding,
- Quasi-non-malleable against in $\text{Bptime}(T_{\text{com}}(n)) \cap \text{SIZE}(t_{\text{com}}(n))$.

4.3 Non-Interactive Simulatable Proof System (Sound against uniform ppt adversaries, ZK against adversaries running in time $T_{\text{ZK}}(n)$)

Let $\Pi = (\text{P}, \text{V}, \text{Sim})$ be a NIZK proof system for NP with no CRS (Construction given in [4]) with soundness against uniform adversaries running in time $T_{\text{NIZK}}(n)$. We additionally require that the trapdoor can be extracted by uniform adversaries running in time $T'_{\text{NIZK}}(n)$.

Let $\mathcal{C}' = (\text{Com}', \text{Open}')$ be a non-interactive, perfectly binding, commitment scheme with no CRS that can be extracted in time $T_{\text{NIZK}}(n)$ and is hiding against adversaries running in time $T_{\text{ZK}}(n)$.

We also assume the existence of \mathbf{P} -certificates with soundness against adversaries running in time $T_{\text{NIZK}}(n)$.

We define the proof system $\Pi^{\text{NI}} = (\text{P}^{\text{NI}}, \text{V}^{\text{NI}}, \text{Sim}^{\text{NI}})$ for language \mathcal{L} defined in Section 3 as follows:

\mathbf{P}^{NI} : Recall that a witness w for statement $s := ([\mathbf{x}^i]_{i \in [m]}, \mathbf{c}, \text{tag}) \in \mathcal{L}$ consists of a string $\mathbf{b} = b^1, \dots, b^m$ and an opening \mathbf{d} such that (1) $\text{Open}(\mathbf{c}, \mathbf{b}, \text{tag}) = 1$ and (2) for all $i \in [m]$, $\psi(\mathbf{x}^i) = b^i$. Given a statement s and witness w , let P be a \mathbf{P} -certificate that (1) and (2) hold.

Invoke \mathbf{P} from proof system Π with the statement $s' = (s, \text{com}) \in \mathcal{L}'$ using proof system Π , where \mathcal{L}' is the language consisting of strings (s, com) such that com is a commitment to (w, P) and P is a \mathbf{P} -certificate that (1) and (2) hold for (s, w) . \mathbf{P} outputs a proof π' . \mathbf{P}^{NI} outputs proof $\pi = \text{com} \parallel \pi'$.

\mathbf{V}^{NI} : On input statement s , proof π and language \mathcal{L} : Parse $\pi := \text{com} \parallel \pi'$. Run the underlying verifier \mathbf{V} on π' for statement (s, com) and language \mathcal{L}' and output whatever it does.

\mathbf{Sim}^{NI} : On input (td, x) , and language \mathcal{L} : Set com to a commitment to 0 and invoke the underlying \mathbf{Sim} for Π with input $(\text{td}, (s, \text{com}))$ and language \mathcal{L}' .

Note that given the \mathbf{P} -certificate P , computing the NIZK proof using Π^{NI} can be done in fixed polynomial time in the length of the statement (s, com) . Moreover, given the trapdoor td , a simulated proof can also be computed in fixed polynomial time. The following claim is straightforward.

Claim. Given the above assumptions, $\Pi^{\text{NI}} = (\mathbf{P}^{\text{NI}}, \mathbf{V}^{\text{NI}}, \mathbf{Sim}^{\text{NI}})$ is a NIZK argument system for language \mathcal{L} with zero knowledge against adversaries running in time $T_{\text{ZK}}(n)$ and trapdoor that can be extracted in time $T'_{\text{NIZK}}(n)$.

4.4 Main Theorem

Theorem 7. *For any constant $c_A > 1$, $\Pi = (\mathbf{E}, \mathbf{D})$ (presented in Figure 2) is a multi-bit, non-malleable code against (uniform) tampering functions that run in time $O(n^{c_A})$, if parameters $c_\psi, c_{\text{com}}, T_{\text{com}}(n), T'_{\text{NIZK}}(n), T_{\text{ZK}}(n)$ are chosen as described above and the underlying components are instantiated in the following way:*

- For $b \in \{0, 1\}$, D_b is the distribution from Section 4.1.
- $\mathcal{C} := (\text{Com}, \text{Open})$ is the commitment scheme from Section 4.2.
- $\Pi^{\text{NI}} := (\mathbf{P}^{\text{NI}}, \mathbf{V}^{\text{NI}}, \mathbf{Sim}^{\text{NI}})$ the simulatable proof system from Section 4.3.
- $\mathcal{S} := (\text{Gen}, \text{Sign}, \text{Ver})$ is any one-time signature scheme secure against PPT adversaries.

Proof. To prove the theorem, we need to show that the necessary properties from Theorem 5 hold. We next go through these one by one.

Simulation of proofs.

1. $\Pr[g(\text{CW}_0, f(\text{CW}_0)) = 1] \stackrel{\text{negl}(n)}{\approx} \Pr[g(\text{CW}_1, f(\text{CW}_1)) = 1],$
2. $\Psi(g, \text{CW}_0, f(\text{CW}_0), \mathbf{D}(f(\text{CW}_0))) \stackrel{\text{PPT, negl}(n)}{\approx} \Psi(g, \text{CW}_1, f(\text{CW}_1), \mathbf{D}(f(\text{CW}_1))),$

where $f \in \mathcal{F}$, $\text{CW}_0 \leftarrow \mathbf{E}(\mathbf{b}_0)$ and $\text{CW}_1 \leftarrow \mathbf{E}_1(\text{td}, \mathbf{b}_0)$.

This follows by ZK property of Π^{NI} .

Simulation of Commitment.

1. $\Pr[g(\text{CW}_1, f(\text{CW}_1)) = 1] \stackrel{\text{negl}(n)}{\approx} \Pr[g(\text{CW}_2, f(\text{CW}_2)) = 1],$
2. $\Psi(g, \text{CW}_1, f(\text{CW}_1), \text{D}(f(\text{CW}_1))) \stackrel{PPT, \text{negl}(n)}{\approx} \Psi(g, \text{CW}_2, f(\text{CW}_2), \text{D}(f(\text{CW}_2))),$

where $f \in \mathcal{F}$, $\text{CW}_1 \leftarrow \text{E}_1(\text{td}, \mathbf{b}_0)$ and $\text{CW}_2 \leftarrow \text{E}_2(\text{td}, \mathbf{b}_0)$.

This follows from hiding property of the commitment scheme \mathcal{C} .

Simulation Soundness.

$$\Pr_r[\text{D}(f(\text{CW}_2)) \neq \text{D}'(f(\text{CW}_2)) \wedge g(\text{CW}_2, f(\text{CW}_2)) = 0] \in O(1/n^{c_{\text{com}}}),$$

where $f \in \mathcal{F}$, $\text{CW}_2 \leftarrow \text{E}_2(\text{td}, \mathbf{b}_0)$.

We begin by defining the following:

$$P_0(n) := \Pr[\text{D}(f(\text{CW}_0)) \neq \text{D}'(f(\text{CW}_0)) \wedge g(\text{CW}_0, f(\text{CW}_0)) = 0],$$

where $f \in \mathcal{F}$, $\text{CW}_0 \leftarrow \text{E}(\mathbf{b}_0)$

$$P_1(n) := \Pr_r[\text{D}(f(\text{CW}_1)) \neq \text{D}'(f(\text{CW}_1)) \wedge g(\text{CW}_1, f(\text{CW}_1)) = 0],$$

where $f \in \mathcal{F}$, $\text{CW}_1 \leftarrow \text{E}_1(\text{td}, \mathbf{b}_0)$

$$P_2(n) := \Pr_r[\text{D}(f(\text{CW}_2)) \neq \text{D}'(f(\text{CW}_2)) \wedge g(\text{CW}_2, f(\text{CW}_2)) = 0],$$

where $f \in \mathcal{F}$, $\text{CW}_2 \leftarrow \text{E}_2(\text{td}, \mathbf{b}_0)$.

We prove the following sequence of claims, which immediately imply the simulation soundness property.

Claim. $P_0(n) \in \text{negl}(n)$.

Since $\text{D}(f(\text{CW}_1)) \neq \text{D}'(f(\text{CW}_1))$ can only occur if the NIZK proof verifies, but the statement being proved is false, this follows from the soundness of the NIZK proof system Π^{Nl} .

Claim. $|P_1(n) - P_0(n)| \in \text{negl}(n)$.

This holds due to complexity leveraging—i.e. by appropriately setting parameters, one can check whether the statement being proved is true or false (by deciding whether \mathbf{x} is in the support of D_0 or D_1 and by extracting from the commitment scheme) without distinguishing a real from simulated proof since $T_{ZK}(n)$ is subexponential.

Claim. $|P_2(n) - P_1(n)| \in O(1/n^{c_{\text{com}}})$.

Proof. Assume $|P_2(n) - P_1(n)| \notin O(1/n^{c_{\text{com}}})$, we will construct an adversary/distinguisher (A, D) in $\mathbf{Bptime}(T_{\text{com}}(n)) \cap \mathbf{SIZE}(t_{\text{com}}(n))$ that breaks the $O(1/n^{c_{\text{com}}})$ -non-malleability of commitment scheme \mathcal{C} . Specifically, we must show an adversary A , distinguisher D in $\mathbf{Bptime}(T_{\text{com}}(n)) \cap \mathbf{SIZE}(t_{\text{com}}(n))$ such that D distinguishes the output of $\text{mim}_{\mathcal{C}}^A(\mathbf{b}_0)$ from $\text{mim}_{\mathcal{C}}^A(\mathbf{0})$ with advantage $a(n) \notin O(1/n^{c_{\text{com}}})$.

$A = (A_1, A_2)$ is specified as follows:

On input security parameter 1^n , A_1 does as follows:

- A_1 generates keys $(\text{vk}, \text{sk}) \leftarrow \text{Gen}(1^n)$
- A_1 runs in uniform time $T'_{\text{NIZK}}(n) \leq T_{\text{com}}(n)$ to recover the trapdoor td of the NIZK.
- A_1 outputs $\text{tag} := \text{vk}$ to its challenger as the desired tag and outputs td, sk to A_2 .

On input $\text{td}, \text{sk}, \text{vk}, \mathbf{c}$, A_2 does as follows:

- For $i \in [m]$, sample $\mathbf{x}^i \sim D_{b^i}$ (in time $m \cdot \text{poly}(n^{c_\psi}) \in O(n^{c_{\text{com}}})$, where poly is a fixed polynomial).
- Use td to generate a simulated proof T in fixed polynomial time and compute $\sigma \leftarrow \text{Sign}(\text{sk}, ([\mathbf{x}^i]_{i \in [m]}, \mathbf{c}, T))$ in fixed polynomial time.
- Compute $f(\text{vk}, [\mathbf{x}^i]_{i \in [m]}, \mathbf{c}, T, \sigma) = [\text{vk}', \mathbf{x}'^i]_{i \in [m]}, \mathbf{c}', T', \sigma'$.
- If the predicate g evaluates to 1, the signature σ' or proof T does not verify, output \perp (this computation takes fixed polynomial time).
- Otherwise, output $(\mathbf{c}', \text{out} := [\mathbf{x}'^i]_{i \in [m]})$. Note that in this case, $\text{vk}' \neq \text{vk}$ (corresponding to the tag of the commitment) since g evaluates to 0 and σ verifies.

Distinguisher D receives the committed value $\mathbf{v}' = v'_1, \dots, v'_m$ underlying \mathbf{c}' (or receives \perp) as well as out (the additional output of adversary A). D outputs 0 if for all $i \in [m]$, $v'_i = \psi(\mathbf{x}^i)$ (or if its input is \perp) and outputs 1 otherwise (computed in time $m \cdot \text{poly}(n^{c_\psi}) \in O(n^{c_{\text{com}}})$).

Clearly,

$$\Pr_{\mathbf{c} \leftarrow \text{Com}(\mathbf{b}_0, \text{vk})} [D(\mathbf{v}', \text{out}) = 1] = P_2(n), \quad \text{and}$$

$$\Pr_{\mathbf{c} \leftarrow \text{Com}(\mathbf{0}, \text{vk})} [D(\mathbf{v}', \text{out}) = 1] = P_1(n)$$

Thus, we have that

$$\left| \Pr_{\mathbf{c} \leftarrow \text{Com}(\mathbf{b}_0, \text{vk})} [D(\mathbf{v}', \text{out}) = 1] - \Pr_{\mathbf{c} \leftarrow \text{Com}(\mathbf{0}, \text{vk})} [D(\mathbf{v}', \text{out}) = 1] \right| \notin O(1/n^{c_{\text{com}}}).$$

Moreover, A, D are in $\mathbf{Bptime}(T_{\text{com}}(n)) \cap \mathbf{SIZE}(t_{\text{com}}(n))$. Thus, we obtain a contradiction to the $O(1/n^{c_{\text{com}}})$ non-malleability of the commitment scheme against adversaries, distinguishers in $\mathbf{Bptime}(T_{\text{com}}(n)) \cap \mathbf{SIZE}(t_{\text{com}}(n))$.

Hardness of D_b relative to Alternate Decoding.

1. $\Pr[g(\text{CW}_2, f(\text{CW}_2)) = 1] \stackrel{O(1/n^{c_\psi})}{\approx} \Pr[g(\text{CW}_3, f(\text{CW}_3)) = 1]$,
2. For every Boolean function, represented by a circuit F over m variables,

$$F \circ D'(f(\text{CW}_2)) \stackrel{\text{stat}, O(1/n^{c_\psi})}{\approx} F \circ D'(f(\text{CW}_3)),$$

where $f \in \mathcal{F}$, $\text{CW}_2 \leftarrow \text{E}_2(\text{td}, \mathbf{b}_0)$ and $\text{CW}_3 \leftarrow \text{E}_2(\text{td}, \mathbf{b}_1)$.

We consider a sequence of distributions where we switch the internal random variables of E_2 from from $\mathbf{x}^i \leftarrow D_{b_0^i}$, for all $i \in [m]$ to $\mathbf{x}^i \leftarrow D_{b_1^i}$, for all $i \in [m]$. Namely, for each $i \in \{0, \dots, m\}$ we consider a distribution where for $j \leq i$, $\mathbf{x}^j \leftarrow D_{b_1^j}$ and for $j > i$, $\mathbf{x}^j \leftarrow D_{b_0^j}$.

We must show that (1) and (2) hold for each consecutive pair of distributions. When considering the i -th consecutive pair, fix all random variables except the i -th variable \mathbf{X}^i to values $\mathbf{x}^1, \dots, \mathbf{x}^{i-1}, \mathbf{x}^{i+1}, \dots, \mathbf{x}^m$. Let \mathbf{X}^i be a random variable such that with probability $1/2$, $\mathbf{X}^i \leftarrow D_{b_0^i}$ and with probability $1/2$, $\mathbf{X}^i \leftarrow D_{b_1^i}$. $\mathbf{X}^i = \mathbf{X}^{i,\gamma}$ where $\gamma \leftarrow \{0, 1\}$, and let random variable CW^i denote the output of E_2 when using random variables $\mathbf{x}^1, \dots, \mathbf{x}^{i-1}, \mathbf{X}^i, \mathbf{x}^{i+1}, \dots, \mathbf{x}^m$.

Since proving (1) is similar, but more straightforward than proving (2), we defer the proof of (1) to [4] and proceed to prove (2) next.

To show (2), assume $\text{D}'(f(\text{CW}_2))$ and $\text{D}'(f(\text{CW}_3))$ have greater than $1/n^{c_\psi}$ statistical distance. This implies that there exists a distinguisher F (represented by an m -bit Boolean function) such that $F \circ \text{D}'(f(\text{CW}_2))$ is more than $1/n^{c_\psi}$ -far from $F \circ \text{D}'(f(\text{CW}_3))$. This implies that, for some $i \in [m]$, the output of $F \circ \text{D}'(f(\text{CW}^i))$ is $a(n) \notin O(1/n^{c_\psi})$ -correlated with $\psi(\mathbf{X}^i)$. Note that, by definition, $F \circ \text{D}'(f(\text{CW}^i)) = F \circ \text{D}'_2 \circ \text{D}'_1(f(\text{CW}^i))$, where D'_1 has output length $(m+1) \cdot \lambda$ ($m \cdot \lambda$ for the size of the non-malleable commitment and λ for the length of the tag of the non-malleable commitment). We will show that $\text{D}'_1(f(\text{CW}^i))$ can be computed by a circuit C of size $O(n^{c_\psi})$ (drawn from some distribution \mathcal{C} over circuits) with input \mathbf{X}^i . We then use Claim 4.1, which says that if C is a size $O(n^{c_\psi})$ circuit taking inputs of length n bits and producing outputs of length $(m+1) \cdot \lambda < (1-\alpha)n$ -bits and \tilde{F} is any $(m+1) \cdot \lambda < (1-\alpha)n$ -bit input Boolean function then the output of $\tilde{F}(C(\mathbf{X}^i))$ is at most $O(1/n^{c_\psi})$ -correlated with $\psi(\mathbf{X}^i)$, instantiating $\tilde{F} := F \circ \text{D}'_2$. This yields a contradiction. Details follow.

Given non-uniform advice td, f , we construct the distribution of circuits $\mathcal{C}_{f,\text{td}}^2$. A draw $C \sim \mathcal{C}_{f,\text{td}}^2$ as follows:

1. Sample signature keys $(\text{vk}, \text{SK}) \leftarrow \text{Gen}(1^n)$,
2. Sample random commitment to 0^m : $(\mathbf{c}', \mathbf{d}') \leftarrow \text{Com}(0^m, \text{tag} := \text{vk})$,
3. Sample $\mathbf{x}^1, \dots, \mathbf{x}^{i-1}$ from $D_{b_0^i}$, and $\mathbf{x}^{i+1}, \dots, \mathbf{x}^m$ from $D_{b_1^i}$.
4. Output the following circuit C that has the following structure:
 - **hardcoded variables:** $f, \mathbf{x}^1, \dots, \mathbf{x}^{i-1}, \mathbf{c}', [T_j^{\beta,i}]_{\beta \in \{0,1\}, i \in [m], j \in [n]}, \mathbf{x}^1, \dots, \mathbf{x}^{i-1}, \mathbf{x}^{i+1}, \dots, \mathbf{x}^m$.
 - **input:** \mathbf{X}^i .
 - **computes and outputs:** $\text{D}'_1(f(\text{CW}^i))$.

Given all the hardwired variables, computing CW^i can be done in time $O(n^{c_\psi})$ since it only requires computing the simulated proof T and signature σ , which can both be done in fixed polynomial time less than n^{c_ψ} . Additionally, f can be computed in time $n^{c_A} < n^{c_\psi}$, and D'_1 can be computed in fixed polynomial time less than n^{c_ψ} , since it only involves verifying the signature σ and proof T , which both take fixed polynomial time.

Acknowledgments

The first and fifth authors are supported in part by NSF grant #CCF1423306 and the Leona M. & Harry B. Helmsley Charitable Trust. The first author is additionally supported in part by an IBM Research PhD Fellowship. The second and third authors are supported in part by NSF grants #CNS-1840893, #CNS-1453045 (CAREER), by a research partnership award from Cisco and by financial assistance award 70NANB15H328 from the U.S. Department of Commerce, National Institute of Standards and Technology. The fourth author is 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. This work was performed, in part, while the first author was visiting IDC Herzliya’s FACT center and supported in part by ISF grant no. 1790/13 and the Check Point Institute for Information Security.

References

1. Applebaum, B., Artemenko, S., Shaltiel, R., Yang, G.: Incompressible functions, relative-error extractors, and the power of nondeterministic reductions. *computational complexity* 25(2), 349–418 (Jun 2016), <https://doi.org/10.1007/s00037-016-0128-9>
2. Babai, L., Fortnow, L., Nisan, N., Wigderson, A.: Bpp has subexponential time simulations unless EXP has publishable proofs. *computational complexity* 3(4), 307–318 (Dec 1993), <https://doi.org/10.1007/BF01275486>
3. Ball, M., Dachman-Soled, D., Guo, S., Malkin, T., Tan, L.Y.: Non-malleable codes for small-depth circuits. FOCS IEEE Computer Society Press (October 2018), <https://eprint.iacr.org/2018/207>, (To appear) Available at <https://eprint.iacr.org/2018/207>
4. Ball, M., Dachman-Soled, D., Kulkarni, M., Lin, H., Malkin, T.: Non-malleable codes against bounded polynomial time tampering. *Cryptology ePrint Archive, Report 2018/1015* (2018), <https://eprint.iacr.org/2018/1015>
5. Ball, M., Dachman-Soled, D., Kulkarni, M., Malkin, T.: Non-malleable codes for bounded depth, bounded fan-in circuits. In: Fischlin and Coron [30], pp. 881–908
6. Ball, M., Dachman-Soled, D., Kulkarni, M., Malkin, T.: Non-malleable codes from average-case hardness: AC^0 , decision trees, and streaming space-bounded tampering. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018, Part III. LNCS, vol. 10822, pp. 618–650. Springer, Heidelberg (Apr / May 2018)
7. Barak, B.: Constant-round coin-tossing with a man in the middle or realizing the shared random string model. In: 43rd FOCS. pp. 345–355. IEEE Computer Society Press (Nov 2002)
8. Barak, B., Ong, S.J., Vadhan, S.: Derandomization in cryptography. *SIAM J. Comput.* 37(2), 380–400 (May 2007), <http://dx.doi.org/10.1137/050641958>
9. Barak, B., Pass, R.: On the possibility of one-message weak zero-knowledge. In: Naor, M. (ed.) TCC 2004. LNCS, vol. 2951, pp. 121–132. Springer, Heidelberg (Feb 2004)

10. Bitansky, N., Lin, H.: One-message zero knowledge and non-malleable commitments. *Cryptology ePrint Archive*, Report 2018/613 (2018), <https://eprint.iacr.org/2018/613>
11. Chandran, N., Goyal, V., Mukherjee, P., Pandey, O., Upadhyay, J.: Block-wise non-malleable codes. In: Chatzigiannakis, I., Mitzenmacher, M., Rabani, Y., Sangiorgi, D. (eds.) *ICALP 2016. LIPIcs*, vol. 55, pp. 31:1–31:14. Schloss Dagstuhl (Jul 2016)
12. Chattopadhyay, E., Goyal, V., Li, X.: Non-malleable extractors and codes, with their many tampered extensions. In: Wichs and Mansour [69], pp. 285–298
13. Chattopadhyay, E., Li, X.: Non-malleable codes and extractors for small-depth circuits, and affine functions. In: Hatami, H., McKenzie, P., King, V. (eds.) *49th ACM STOC*. pp. 1171–1184. ACM Press (Jun 2017)
14. Cheraghchi, M., Guruswami, V.: Capacity of non-malleable codes. In: Naor, M. (ed.) *ITCS 2014*. pp. 155–168. ACM (Jan 2014)
15. Chung, K.M., Lin, H., Pass, R.: Constant-round concurrent zero knowledge from P-certificates. In: *FOCS 2013* [32], pp. 50–59
16. Ciampi, M., Ostrovsky, R., Siniscalchi, L., Visconti, I.: Concurrent non-malleable commitments (and more) in 3 rounds. In: Robshaw, M., Katz, J. (eds.) *CRYPTO 2016, Part III. LNCS*, vol. 9816, pp. 270–299. Springer, Heidelberg (Aug 2016)
17. Ciampi, M., Ostrovsky, R., Siniscalchi, L., Visconti, I.: Four-round concurrent non-malleable commitments from one-way functions. In: Katz and Shacham [44], pp. 127–157
18. Coron, J.S., Holenstein, T., Künzler, R., Patarin, J., Seurin, Y., Tessaro, S.: How to build an ideal cipher: The indistinguishability of the Feistel construction. *Journal of Cryptology* 29(1), 61–114 (Jan 2016)
19. Dachman-Soled, D., Katz, J., Thiruvengadam, A.: 10-round Feistel is indistinguishable from an ideal cipher. In: Fischlin and Coron [30], pp. 649–678
20. Dai, Y., Steinberger, J.P.: Indistinguishability of 8-round Feistel networks. In: Robshaw, M., Katz, J. (eds.) *CRYPTO 2016, Part I. LNCS*, vol. 9814, pp. 95–120. Springer, Heidelberg (Aug 2016)
21. Dolev, D., Dwork, C., Naor, M.: Nonmalleable cryptography. *SIAM review* 45(4), 727–784 (2003)
22. Drucker, A.: Nondeterministic direct product reductions and the success probability of SAT solvers. In: *FOCS 2013* [32], pp. 736–745
23. Dubrov, B., Ishai, Y.: On the randomness complexity of efficient sampling. In: Kleinberg, J.M. (ed.) *38th ACM STOC*. pp. 711–720. ACM Press (May 2006)
24. Dwork, C., Naor, M.: Zaps and their applications. In: *FOCS 2000* [31], pp. 283–293
25. Dziembowski, S., Pietrzak, K., Wichs, D.: Non-malleable codes. In: Yao, A.C.C. (ed.) *ICS 2010*. pp. 434–452. Tsinghua University Press (Jan 2010)
26. Faust, S., Hostáková, K., Mukherjee, P., Venturi, D.: Non-malleable codes for space-bounded tampering. In: Katz and Shacham [44], pp. 95–126
27. Faust, S., Mukherjee, P., Venturi, D., Wichs, D.: Efficient non-malleable codes and key-derivation for poly-size tampering circuits. In: Nguyen, P.Q., Oswald, E. (eds.) *EUROCRYPT 2014. LNCS*, vol. 8441, pp. 111–128. Springer, Heidelberg (May 2014)
28. Feige, U., Lapidot, D., Shamir, A.: Multiple noninteractive zero knowledge proofs under general assumptions. *SIAM Journal on Computing* 29(1), 1–28 (1999)
29. Feige, U., Lund, C.: On the hardness of computing the permanent of random matrices. *Computational Complexity* 6(2), 101–132 (1997)
30. Fischlin, M., Coron, J.S. (eds.): *EUROCRYPT 2016, Part II, LNCS*, vol. 9666. Springer, Heidelberg (May 2016)

31. 41st FOCS. IEEE Computer Society Press (Nov 2000)
32. 54th FOCS. IEEE Computer Society Press (Oct 2013)
33. 58th FOCS. IEEE Computer Society Press (2017)
34. Fortnow, L., Vadhan, S.P. (eds.): 43rd ACM STOC. ACM Press (Jun 2011)
35. Goldreich, O., Wigderson, A.: Derandomization that is rarely wrong from short advice that is typically good. In: International Workshop on Randomization and Approximation Techniques in Computer Science. pp. 209–223. Springer (2002)
36. Goyal, V.: Constant round non-malleable protocols using one way functions. In: Fortnow and Vadhan [34], pp. 695–704
37. Goyal, V., Pandey, O., Richelson, S.: Textbook non-malleable commitments. In: Wichs and Mansour [69], pp. 1128–1141
38. Goyal, V., Richelson, S., Rosen, A., Vald, M.: An algebraic approach to non-malleability. In: 55th FOCS. pp. 41–50. IEEE Computer Society Press (Oct 2014)
39. Groth, J., Ostrovsky, R., Sahai, A.: Non-interactive zaps and new techniques for NIZK. In: Dwork, C. (ed.) CRYPTO 2006. LNCS, vol. 4117, pp. 97–111. Springer, Heidelberg (Aug 2006)
40. Gutfreund, D., Shaltiel, R., Ta-Shma, A.: Uniform hardness versus randomness tradeoffs for arthur-merlin games. *computational complexity* 12(3-4), 85–130 (2003)
41. Harnik, D., Naor, M.: On the compressibility of \mathcal{NP} instances and cryptographic applications. *SIAM Journal on Computing* 39(5), 1667–1713 (2010)
42. Håstad, J., Impagliazzo, R., Levin, L.A., Luby, M.: A pseudorandom generator from any one-way function. *SIAM Journal on Computing* 28(4), 1364–1396 (1999)
43. Impagliazzo, R., Wigderson, A.: $P = BPP$ if E requires exponential circuits: Derandomizing the XOR lemma. In: 29th ACM STOC. pp. 220–229. ACM Press (May 1997)
44. Katz, J., Shacham, H. (eds.): CRYPTO 2017, Part II, LNCS, vol. 10402. Springer, Heidelberg (Aug 2017)
45. Khurana, D.: Round optimal concurrent non-malleability from polynomial hardness. In: Kalai, Y., Reyzin, L. (eds.) TCC 2017, Part II. LNCS, vol. 10678, pp. 139–171. Springer, Heidelberg (Nov 2017)
46. Khurana, D., Sahai, A.: How to achieve non-malleability in one or two rounds. In: FOCS 2017 [33], pp. 564–575
47. Klivans, A.R., Van Melkebeek, D.: Graph nonisomorphism has subexponential size proofs unless the polynomial-time hierarchy collapses. *SIAM Journal on Computing* 31(5), 1501–1526 (2002)
48. Lin, H., Pass, R.: Non-malleability amplification. In: Mitzenmacher, M. (ed.) 41st ACM STOC. pp. 189–198. ACM Press (May / Jun 2009)
49. Lin, H., Pass, R.: Constant-round non-malleable commitments from any one-way function. In: Fortnow and Vadhan [34], pp. 705–714
50. Lin, H., Pass, R., Soni, P.: Two-round and non-interactive concurrent non-malleable commitments from time-lock puzzles. In: FOCS 2017 [33], pp. 576–587
51. Lin, H., Pass, R., Venkatasubramanian, M.: Concurrent non-malleable commitments from any one-way function. In: Canetti, R. (ed.) TCC 2008. LNCS, vol. 4948, pp. 571–588. Springer, Heidelberg (Mar 2008)
52. Lindell, Y.: A simpler construction of cca2-secure public-key encryption under general assumptions. In: Biham, E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, pp. 241–254. Springer, Heidelberg (May 2003)
53. Lipton, R.J.: New directions in testing. In: Feigenbaum, J., Merritt, M. (eds.) Distributed Computing And Cryptography, Proceedings of a DIMACS Workshop, Princeton, New Jersey, USA, October 4-6, 1989. pp. 191–202 (1989)

54. Micali, S.: Computationally sound proofs. *SIAM Journal on Computing* 30(4), 1253–1298 (2000)
55. Miltersen, P.B., Vinodchandran, N.V.: Derandomizing arthur–merlin games using hitting sets. *Computational Complexity* 14(3), 256–279 (2005)
56. Naor, M., Yung, M.: Public-key cryptosystems provably secure against chosen ciphertext attacks. In: 22nd ACM STOC. pp. 427–437. ACM Press (May 1990)
57. Nisan, N., Wigderson, A.: Hardness vs randomness. *J. Comput. Syst. Sci.* 49(2), 149–167 (Oct 1994), [http://dx.doi.org/10.1016/S0022-0000\(05\)80043-1](http://dx.doi.org/10.1016/S0022-0000(05)80043-1)
58. Ostrovsky, R., Persiano, G., Venturi, D., Visconti, I.: Continuously non-malleable codes in the split-state model from minimal assumptions. *Cryptology ePrint Archive, Report 2018/542* (2018), <https://eprint.iacr.org/2018/542>
59. Pass, R., Rosen, A.: Concurrent non-malleable commitments. In: 46th FOCS. pp. 563–572. IEEE Computer Society Press (Oct 2005)
60. Pass, R., Rosen, A.: New and improved constructions of non-malleable cryptographic protocols. In: Gabow, H.N., Fagin, R. (eds.) 37th ACM STOC. pp. 533–542. ACM Press (May 2005)
61. Pass, R., Wee, H.: Constant-round non-malleable commitments from sub-exponential one-way functions. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 638–655. Springer, Heidelberg (May / Jun 2010)
62. Rivest, R.L., Shamir, A., Wagner, D.A.: Time-lock puzzles and timed-release crypto (1996)
63. Sahai, A.: Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In: 40th FOCS. pp. 543–553. IEEE Computer Society Press (Oct 1999)
64. Shaltiel, R., Umans, C.: Simple extractors for all min-entropies and a new pseudorandom generator. *Journal of the ACM (JACM)* 52(2), 172–216 (2005)
65. Shaltiel, R., Umans, C.: Pseudorandomness for approximate counting and sampling. *computational complexity* 15(4), 298–341 (2006)
66. Shaltiel, R., Umans, C.: Low-end uniform hardness versus randomness tradeoffs for am. *SIAM Journal on Computing* 39(3), 1006–1037 (2009)
67. Sudan, M., Trevisan, L., Vadhan, S.: Pseudorandom generators without the xor lemma. *Journal of Computer and System Sciences* 62(2), 236 – 266 (2001), <http://www.sciencedirect.com/science/article/pii/S002200000917306>
68. Trevisan, L., Vadhan, S.P.: Extracting randomness from samplable distributions. In: FOCS 2000 [31], pp. 32–42
69. Wichs, D., Mansour, Y. (eds.): 48th ACM STOC. ACM Press (Jun 2016)
70. Yao, A.C.: Theory and applications of trapdoor functions (extended abstract). In: 23rd Annual Symposium on Foundations of Computer Science, Chicago, Illinois, USA, 3-5 November 1982. pp. 80–91. IEEE Computer Society (1982), <https://doi.org/10.1109/SFCS.1982.45>