

Capacity of Non-Malleable Codes

Mahdi Cheraghchi*
CSAIL
MIT
Cambridge, MA 02139

Venkatesan Guruswami†
Computer Science Department
CMU
Pittsburgh, PA 15213

Abstract

Non-malleable codes, introduced by Dziembowski, Pietrzak and Wichs (ICS 2010), encode messages s in a manner so that tampering the codeword causes the decoder to either output s or a message that is independent of s . While this is an impossible goal to achieve against unrestricted tampering functions, rather surprisingly non-malleable coding becomes possible against every fixed family \mathcal{F} of tampering functions that is not too large (for instance, when $|\mathcal{F}| \leq 2^{2^{\alpha n}}$ for some $\alpha < 1$ where n is the number of bits in a codeword).

In this work, we study the “capacity of non-malleable coding,” and establish optimal bounds on the achievable rate as a function of the family size, answering an open problem from Dziembowski et al. (ICS 2010). Specifically,

- We prove that for every family \mathcal{F} with $|\mathcal{F}| \leq 2^{2^{\alpha n}}$, there exist non-malleable codes against \mathcal{F} with rate arbitrarily close to $1 - \alpha$ (this is achieved w.h.p. by a randomized construction).
- We show the existence of families of size $\exp(n^{O(1)}2^{\alpha n})$ against which there is no non-malleable code of rate $1 - \alpha$ (in fact this is the case w.h.p for a random family of this size).
- We also show that $1 - \alpha$ is the best achievable rate for the family of functions which are only allowed to tamper the first αn bits of the codeword, which is of special interest.

As a corollary, this implies that the capacity of non-malleable coding in the split-state model (where the tampering function acts independently but arbitrarily on the two halves of the codeword, a model which has received some attention recently) equals $1/2$.

We also give an efficient Monte Carlo construction of codes of rate close to 1 with polynomial time encoding and decoding that is non-malleable against any fixed $c > 0$ and family \mathcal{F} of size 2^{n^c} , in particular tampering functions with say cubic size circuits.

*Email: mahdi@csail.mit.edu. Research supported in part by V. Guruswami’s Packard Fellowship, MSR-CMU Center for Computational Thinking, and the Swiss National Science Foundation research grant PA00P2-141980.

†Email: guruswami@cmu.edu. Research supported in part by the National Science Foundation under Grant No. CCF-0963975. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Contents

1	Introduction	3
1.1	Our results	4
1.2	Proof ideas	5
2	Preliminaries	6
2.1	Notation	6
2.2	Definitions	6
3	Probabilistic construction of non-malleable codes	8
3.1	Proof of Theorem 3.1 for bijective adversaries	9
3.2	Proof of Theorem 3.1 for general adversaries	10
3.3	Efficiency in the random oracle model	20
4	A Monte Carlo construction for computationally bounded adversaries	21
5	Impossibility bounds	25
5.1	General adversaries	25
5.2	Random adversaries	26
5.3	General adversaries acting on a subset of positions	27
A	Rate $1/2$ barrier for the uniform coding scheme.	32
B	Useful tools	33

1 Introduction

Non-malleable codes are a fascinating new concept put forth in [8], following the program on non-malleable cryptography which was introduced by the seminal work of Dolev, Dwork and Naor [6]. Non-malleable codes are aimed at protecting the integrity of data in situations where it might be corrupted in ways that precludes error-correction or even error-detection. Informally, a code is non-malleable if the corrupted codeword either encodes the original message, or a completely unrelated value. This is akin to the notion of non-malleable encryption in cryptography which requires the intractability of, given a ciphertext, producing a different ciphertext so that the corresponding plaintexts are related to each other.

A non-malleable (binary¹) code against a family \mathcal{F} of tampering functions each mapping $\{0, 1\}^n$ to $\{0, 1\}^n$, consists of a randomized encoding function $\text{Enc} : \{0, 1\}^k \rightarrow \{0, 1\}^n$ and a deterministic decoding function $\text{Dec} : \{0, 1\}^n \rightarrow \{0, 1\}^k \cup \{\perp\}$ (where \perp denotes error-detection) which satisfy $\text{Dec}(\text{Enc}(s)) = s$ always, and the following non-malleability property with error ϵ : For every message $s \in \{0, 1\}^k$ and every function $f \in \mathcal{F}$, the distribution of $\text{Dec}(f(\text{Enc}(s)))$ is ϵ -close to a distribution \mathcal{D}_f that depends only on f and is independent² of s . In other words, if some adversary (who has full knowledge of the code and the message s , but not the internal randomness of the encoder) tampers with the codeword $\text{Enc}(s)$ corrupting it to $f(\text{Enc}(s))$, he cannot control the relationship between s and the message the corrupted codeword $f(\text{Enc}(s))$ encodes.

In general, it is impossible to achieve non-malleability against arbitrary tampering functions. Indeed, the tampering function can decode the codeword to compute the original message s , flip the last bit of s to obtain a related message \tilde{s} , and then reencode \tilde{s} . This clearly violates non-malleability as the tampered codeword encodes the message \tilde{s} which is closely related to s . Therefore, in order to construct non-malleable codes, one focuses on a restricted class of tampering functions. For example, the body of work on error-correcting codes consists of functions which can flip an arbitrary subset of bits up to a prescribed limit on the total number of bit flips.

The notion of non-malleable coding becomes more interesting for families against which error-correction is not possible. A simple and natural such family is the set of functions causing arbitrary “additive errors,” namely $\mathcal{F}_{\text{add}} = \{f_{\Delta} \mid \Delta \in \{0, 1\}^n\}$ where $f_{\Delta}(x) := x + \Delta$. Note that there is no restriction on the Hamming weight of Δ as in the case of channels causing bounded number of bit flips. While error-correction is impossible against \mathcal{F}_{add} , *error-detection* is still possible — the work of Cramer et al. [5] constructed codes of rate approaching 1 (which they called “Algebraic Manipulation Detection” (AMD) codes) such that offset by an arbitrary $\Delta \neq 0$ will be detected with high probability. AMD codes give a construction of non-malleable codes against the family \mathcal{F}_{add} .

Even error-detection becomes impossible against many other natural families of tampering functions. A particularly simple such class consists of all constant functions $f_c(x) := c$ for $c \in \{0, 1\}^n$. This family includes some function that maps all inputs to a valid codeword c^* , and hence one cannot detect tampering. Note, however, that non-malleability is trivial to achieve against this family — the rate 1 code with identity encoding function is itself non-malleable as the output distribution of a constant function is trivially independent of the message. A natural function family for which non-malleability is non-trivial to achieve consists of *bit-tampering functions* f in which the different of bits of the codewords are tampered indepen-

¹Throughout this paper we deal only with binary codes. We point out that non-malleability is mainly interesting over small alphabets, since when the alphabet size is large enough, even error-detection (e.g., via MDS codes) is possible at rates achieved by non-malleable codes.

²The formal definition (see Definition 2.3) has to accommodate the possibility that Dec error-corrects the tampered codeword to the original message s ; and this is handled in a manner independent of s by including a special element same in the support of \mathcal{D}_f .

dently (i.e., either flipped, set to 0/1, or left unchanged); formally $f(x) = (f_1(x_1), f_2(x_2), \dots, f_n(x_n))$ for arbitrary 1-bit functions f_1, f_2, \dots, f_n [8].

The family \mathcal{F}_{all} of all functions $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ has size given by $\log \log |\mathcal{F}_{\text{all}}| = n + \log n$. The authors of [8] show the existence of a non-malleable code against *any* small enough family \mathcal{F} (for which $\log \log |\mathcal{F}| < n$). The rate of the code is constant if $\log \log |\mathcal{F}| \leq \alpha n$ for some constant $\alpha \in (0, 1)$. The question of figuring out the optimal rates of non-malleable codes for various families of tampering functions was left as an open problem in [8]. In this work we give a satisfactory answer to this question, pinning down the rate for many natural function families. We describe our results next.

1.1 Our results

Our results include improvements to the rate achievable as a function of the size of the family of tampering functions, as well as limitations of non-malleable codes demonstrating that the achieved rate cannot be improved for natural families of the stipulated size. Specifically, we establish the following results concerning the possible rates for non-malleable coding as a function of the size of the family of tampering functions:

1. (Rate lower bound) We prove in Section 3 that if $|\mathcal{F}| \leq 2^{2^{\alpha n}}$, then there exists a (strong) non-malleable code of rate arbitrarily close to $1 - \alpha$ which is non-malleable w.r.t \mathcal{F} with error $\exp(-\Omega(n))$. This significantly improves the probabilistic construction of [8], which achieves a rate close to $(1 - \alpha)/3$ using a delicate Martingale argument. In particular, for arbitrary small families, of size $2^{2^{o(n)}}$, our result shows that the rate can be made arbitrarily close to 1. This was not known to be possible even for the family of bit-tampering functions (which has size 4^n), for which $1/3$ was the best known rate³ [8]. In fact, we note (in Appendix A) why the proof strategy of [8] is limited to a rate of $1/2$ even for a very simple tampering function such as the one that flips the first bit. As discussed in Section 3.3, our probabilistic construction is equipped with an encoder and decoder that can be efficiently and exactly implemented with access to a uniformly random permutation oracle and its inverse (corresponding to the ideal-cipher model in cryptography). This is a slight additional advantage over [8], where only an approximation of the encoder and decoder is shown to be efficiently computable.
2. (Upper bound/limitations on rate) The above coding theorem shows that the “capacity” of a function family $|\mathcal{F}|$ for non-malleable coding is at least $1 - (\log \log |\mathcal{F}|)/n$. We also address the natural “converse coding question” of whether this rate bound is the best achievable (Section 5). This turns out to be false in general due to the existence of uninteresting large families for which non-malleable coding with rate close to 1 is easy. But we do prove that the $1 - \alpha$ rate is best achievable in “virtually all” situations:
 - (a) We prove that for *random* families of size $2^{2^{\alpha n}}$, with high probability it is not possible to exceed a rate of $1 - \alpha$ for non-malleable coding with small error.
 - (b) For the family of tampering functions which leave the last $(1 - \alpha)n$ bits intact and act arbitrarily on the first αn bits, we prove that $1 - \alpha$ is the best achievable rate for non-malleable coding. (Note that a rate of $1 - \alpha$ is trivial to achieve for this family, by placing the message bits in the last $(1 - \alpha)n$ bits of the codeword, and setting the first αn bits of the codeword to all 0s.)

³Assuming the existence of one-way functions, an explicit construction of non-malleable codes of rate close to 1 was proposed in [8]. This construction, however, only satisfies a weaker definition of non-malleability that considers computational indistinguishability rather than statistical security.

The result 2b, together with the existential result 1 above, pins down the optimal rate for non-malleable codes in the *split-state model* to $1/2$. In the split-state model, which was the focus of a couple of recent works [7, 1], the tampering function operates independently (but in otherwise arbitrary ways) on the two halves of the codeword, i.e., $f(x) = ((f_1(x_1), f_2(x_2)))$ where x_1, x_2 are the two halves of x and f_1, f_2 are functions mapping $n/2$ bits to $n/2$ bits. The recent work [1] gave an explicit construction in this model with polynomially small rate. Our work shows that the capacity of the split-state model is $1/2$, but we do not offer any explicit construction. For the more restrictive class of bit-tampering functions (where each bit is tampered independently), in a follow-up work [3] we give an explicit construction with rate approaching $1/2$ [3]. We also present in that work a reduction of non-malleable coding for the split-state model to a new notion of non-malleable two-source extraction.

Monte Carlo construction for small families. Our result 1 above is based on a random construction which takes exponential time (and space). Derandomizing this construction, in Section 4 we are able to obtain an efficient Monte Carlo construction of non-malleable codes of rate close to $1/2$ (with polynomial time encoding and decoding, and inverse polynomially small error) for an *arbitrary* family of size $\exp(n^c)$ for any fixed $c > 0$. Note that in particular this includes tampering functions that can be implemented by circuits of any fixed polynomial size, or simpler families such as bit-tampering adversaries. The construction does not rely on any computational hardness assumptions, at the cost of using a small amount of randomness.

1.2 Proof ideas

Rate lower bound. Our construction of rate $\approx 1 - (\log \log |\mathcal{F}|)/n$ codes is obtained by picking for each message, a random blob of t codewords, such that blobs corresponding to distinct messages are disjoint. For each tampering function f , our proof analyzes the distribution of $\text{Dec}(f(\text{Enc}(s)))$ for each message s separately, and shows that w.h.p. they are essentially close to the same distribution \mathcal{D}_f . In order to achieve sufficiently small error probability allowing for a union bound, the proof uses a number of additional ideas, including a randomized process that gradually reveals information about the code while examining the t codewords in each blob in sequence. The analysis ensures that as little information is revealed in each step as possible, so that enough independence remains in the conditional joint distribution of the codewords throughout the analysis. Finally, strong concentration bounds are used to derive the desired bound on the failure probability. The proof for the special case of bijective tampering functions turns out to be quite straightforward, and as a warm-up we present this special case first in Section 3.1.

Monte Carlo construction. Since the analysis of the probabilistic code construction considers each message s separately, we observe that it only needs limited (t -wise) independence of the codewords. On the other hand, the code construction is designed to be sparse, namely taking $t = \text{poly}(n, \log |\mathcal{F}|, 1/\epsilon)$ suffices for the analysis. This is the key idea behind our efficient Monte Carlo construction for small families with $\log |\mathcal{F}| \leq \text{poly}(n)$.

The birthday paradox implies that picking the blob of codewords encoding each message independently of other messages, while maintaining disjointness of the various blobs, limits the rate to $1/2$. Therefore, we construct the code by means of a t -wise independent *decoding* function implemented via a random low-degree polynomial. After overcoming some complications to ensure an efficient encoding function, we get our efficient randomized construction for small families of tampering functions.

Rate upper bounds. Our main impossibility result for the family of adversaries that only tamper the first αn bits of the codeword uses an information theoretic argument. We argue that if the rate of the code is

sufficiently large, one can always find messages s_0 and s_1 and a set $X_\eta \subseteq \{0, 1\}^{\alpha n}$ such that the following holds: The first αn bits of the encoding of s_0 has a noticeable chance of being in X_η , whereas this chance for s_1 is quite small. Using this property, we design an adversary that maps the first αn bits of the encoding to a dummy string if they belong to X_η and leaves the codeword intact otherwise. This suffices to violate non-malleability of the code.

2 Preliminaries

2.1 Notation

We use \mathcal{U}_n for the uniform distribution on $\{0, 1\}^n$ and U_n for the random variable sampled from \mathcal{U}_n and independently of any existing randomness. For a random variable X , we denote by $\mathcal{D}(X)$ the probability distribution that X is sampled from. Moreover, for an event \mathcal{E} , we use $\mathcal{D}(X|\mathcal{E})$ to denote the conditional distribution on the random variable X on the event \mathcal{E} . Generally, we will use calligraphic symbols (such as \mathcal{X}) for probability distributions and the corresponding capital letters (such as X) for related random variables. For a discrete distribution \mathcal{X} , we denote by $\mathcal{X}(x)$ the probability mass assigned to x by \mathcal{X} . Two distributions \mathcal{X} and \mathcal{Y} being ϵ -close in statistical distance is denoted by $\mathcal{X} \approx_\epsilon \mathcal{Y}$. We will use $(\mathcal{X}, \mathcal{Y})$ for the product distribution with the two coordinates independently sampled from \mathcal{X} and \mathcal{Y} . All unsubscripted logarithms are taken to the base 2. Support of a discrete random variable (or distribution) X is denoted by $\text{supp}(X)$. With a slight abuse of notation, for various bounds we condition probabilities and expectations on random variables rather than events (e.g., $\mathbb{E}[X|Y]$, or $\Pr[\mathcal{E}|Y]$). In such instances, the notation means that the statement holds for *every* possible realization of the random variables that we condition on.

2.2 Definitions

In this section, we review the formal definition of non-malleable codes as introduced in [8]. First, we recall the notion of *coding schemes*.

Definition 2.1 (Coding schemes). A pair of functions $\text{Enc}: \{0, 1\}^k \rightarrow \{0, 1\}^n$ and $\text{Dec}: \{0, 1\}^n \rightarrow \{0, 1\}^k \cup \{\perp\}$ where $k \leq n$ is said to be a coding scheme with block length n and message length k if the following conditions hold.

1. The encoder Enc is a randomized function; i.e., at each call it receives a uniformly random sequence of coin flips that the output may depend on. This random input is usually omitted from the notation and taken to be implicit. Thus for any $s \in \{0, 1\}^k$, $\text{Enc}(s)$ is a random variable over $\{0, 1\}^n$. The decoder Dec is; however, deterministic.
2. For every $s \in \{0, 1\}^k$, we have $\text{Dec}(\text{Enc}(s)) = s$ with probability 1.

The *rate* of the coding scheme is the ratio k/n . A coding scheme is said to have relative distance δ , for some $\delta \in [0, 1)$, if for every $s \in \{0, 1\}^k$ the following holds. Let $X := \text{Enc}(s)$. Then, for any $\Delta \in \{0, 1\}^n$ of Hamming weight at most δn , $\text{Dec}(X + \Delta) = \perp$ with probability 1. \square

Before defining non-malleable coding schemes, we find it convenient to define the following notation.

Definition 2.2. For a finite set Γ , the function $\text{copy}: (\Gamma \cup \{\text{same}\}) \times \Gamma \rightarrow \Gamma$ is defined as follows:

$$\text{copy}(x, y) := \begin{cases} x & x \neq \text{same}, \\ y & x = \text{same}. \end{cases} \quad \square$$

The notion of non-malleable coding schemes from [8] can now be rephrased as follows.

Definition 2.3 (Non-malleability). A coding scheme (Enc, Dec) with message length k and block length n is said to be non-malleable with error ϵ (also called *exact security*) with respect to a family \mathcal{F} of tampering functions acting on $\{0, 1\}^n$ (i.e., each $f \in \mathcal{F}$ maps $\{0, 1\}^n$ to $\{0, 1\}^n$) if for every $f \in \mathcal{F}$ there is a distribution \mathcal{D}_f over $\{0, 1\}^k \cup \{\perp, \text{same}\}$ such that the following holds. Let $s \in \{0, 1\}^k$ and define the random variable $S := \text{Dec}(f(\text{Enc}(s)))$. Let S' be independently sampled from \mathcal{D}_f . Then,

$$\mathcal{D}(S) \approx_\epsilon \mathcal{D}(\text{copy}(S', s)). \quad \square$$

Remark 2.4. The above definition allows the decoder to output a special symbol \perp that corresponds to error detection. It is easy to note that any such code can be transformed to one where the decoder never outputs \perp without affecting the parameters (e.g., the new decoder may simply output 0^k whenever the original decoder outputs \perp).

Dziembowski et al. [8] also consider the following stronger variation of non-malleable codes.

Definition 2.5 (Strong non-malleability). A pair of functions as in Definition 2.3 is said to be a *strong* non-malleable coding scheme with error ϵ with respect to a family \mathcal{F} of tampering functions acting on $\{0, 1\}^n$ if the conditions (1) and (2) of Definition 2.3 is satisfied, and additionally, the following holds. For any message $s \in \{0, 1\}^k$, let $E_s := \text{Enc}(s)$, consider the random variable

$$D_s := \begin{cases} \text{same} & \text{if } f(E_s) = E_s, \\ \text{Dec}(f(E_s)) & \text{otherwise,} \end{cases}$$

and let $\mathcal{D}_{f,s} := \mathcal{D}(D_s)$. It must be the case that for every pair of distinct messages $s_1, s_2 \in \{0, 1\}^k$, $\mathcal{D}_{f,s_1} \approx_\epsilon \mathcal{D}_{f,s_2}$. \square

Remark 2.6 (Computational security). Dziembowski et al. also consider the case where statistical distance is replaced with computational indistinguishability with respect to a bounded computational model. As our goal is to understand information-theoretic limitations of non-malleable codes, we do not consider this variation in this work. It is clear, however, that our negative results in Section 5 apply to this model as well. A related (but incomparable) model that we consider in Section 4 is when the distinguishability criterion is still statistical; however the adversary is computationally bounded (e.g., one may consider the family of polynomial sized Boolean circuits). For this case, we construct an efficient Monte Carlo coding scheme that achieves any rate arbitrarily close to 1.

Remark 2.7 (Efficiency of sampling \mathcal{D}_f). The original definition of non-malleable codes in [8] also requires the distribution \mathcal{D}_f to be efficiently samplable given oracle access to the tampering function f . We find it more natural to remove this requirement from the definition since even combinatorial non-malleable codes that are not necessarily equipped with efficient components (such as the encoder, decoder, and sampler for \mathcal{D}_f) are interesting and highly non-trivial to construct. It should be noted; however, that for any non-malleable coding scheme equipped with an efficient encoder and decoder, the following is a valid and efficiently samplable choice for the distribution \mathcal{D}_f (possibly incurring a constant factor increase in the error parameter):

1. Let $S \sim \mathcal{U}_k$, and $X := \text{Enc}(S)$.
2. If $\text{Dec}(X) = S$, output same. Otherwise, output $\text{Dec}(X)$.

Our Monte Carlo construction in Section 4 is equipped with a polynomial-time encoder and decoder. So is the case for our probabilistic construction in Section 3 in the random oracle model.

3 Probabilistic construction of non-malleable codes

In this section, we introduce our probabilistic construction of non-malleable codes. Contrary to the original construction of Dziembowski et al. [8], where they pick a uniformly random truth table for the decoder and do not allow the \perp symbol. Our code, on the other hand, is quite sparse. In fact, in our construction $\text{Dec}(U_n) = \perp$ with high probability. As we observe in Section A, this is the key to our improvement, since uniformly random decoders cannot achieve non-malleability even against extremely simple adversaries at rates better than $1/2$. Moreover, our sparse construction offers the added feature of having a large minimum distance in the standard coding sense; any tampering scheme that perturbs the codeword in a fraction of the positions bounded by a prescribed limit will be detected by the decoder with probability 1. Another advantage of sparsity is allowing a compact representation for the code. We exploit this feature in our Monte Carlo construction of Section 4. Our probabilistic coding scheme is described in Construction 1.

We remark that Construction 1 can be efficiently implemented in the ideal-cipher model, which in turn implies an efficient approximate implementation in the random oracle model (see the discussion following the proof of Theorem 3.1 in Section 3.3). In turn, this implies that the distribution \mathcal{D}_f in Definition 2.3 for this construction can be efficiently sampled in both models (see Remark 2.7).

- *Given:* Integer parameters $0 < k \leq n$ and integer $t > 0$ such that $t2^k \leq 2^n$, and a relative distance parameter δ , $0 \leq \delta < 1/2$.
- *Output:* A pair of functions $\text{Enc}: \{0, 1\}^k \times \{0, 1\}^n$ and $\text{Dec}: \{0, 1\}^n \rightarrow \{0, 1\}^k$, where Enc may also use a uniformly random seed which is hidden from that notation, but Dec is deterministic.
- *Construction:*
 1. Let $\mathcal{N} := \{0, 1\}^n$.
 2. For each $s \in \{0, 1\}^k$, in an arbitrary order,
 - Let $E(s) := \emptyset$.
 - For $i \in \{1, \dots, t\}$:
 - (a) Pick a uniformly random vector $w \in \mathcal{N}$.
 - (b) Add w to $E(s)$.
 - (c) Let $\Gamma(w)$ be the Hamming ball of radius δn centered at w . Remove $\Gamma(w)$ from \mathcal{N} (note that when $\delta = 0$, we have $\Gamma(w) = \{w\}$).
 3. Given $s \in \{0, 1\}^k$, $\text{Enc}(s)$ outputs an element of $E(s)$ uniformly at random.
 4. Given $w \in \{0, 1\}^n$, $\text{Dec}(w)$ outputs the unique s such that $w \in E(s)$, or \perp if no such s exists.

Construction 1: Probabilistic construction of non-malleable codes.

The main theorem of this section is the result below that proves non-malleability of the coding scheme in Construction 1.

Theorem 3.1. *Let $\mathcal{F}: \{0, 1\}^n \rightarrow \{0, 1\}^n$ be any family of tampering functions. For any $\epsilon, \eta > 0$, with probability at least $1 - \eta$, the coding scheme (Enc, Dec) of Construction 1 is a strong non-malleable code with respect to \mathcal{F} and with error ϵ and relative distance δ , provided that both of the following conditions are satisfied.*

1. $t \geq t_0$, for some

$$t_0 = O\left(\frac{1}{\epsilon^6} \left(\log \frac{|\mathcal{F}|N}{\eta}\right)\right). \quad (1)$$

2. $k \leq k_0$, for some

$$k_0 \geq n(1 - h(\delta)) - \log t - 3 \log(1/\epsilon) - O(1), \quad (2)$$

where $h(\cdot)$ denotes the binary entropy function.

Thus by choosing $t = t_0$ and $k = k_0$, the construction satisfies

$$k \geq n(1 - h(\delta)) - \log \log(|\mathcal{F}|/\eta) - \log n - 9 \log(1/\epsilon) - O(1).$$

In particular, if $|\mathcal{F}| \leq 2^{2^{\alpha n}}$ for any constant $\alpha \in (0, 1)$, the rate of the code can be made arbitrarily close to $1 - h(\delta) - \alpha$ while allowing $\epsilon = 2^{-\Omega(n)}$.

Remark 3.2. (Error detection) An added feature of our sparse coding scheme is the error-detection capability. However, observe that any probabilistic coding scheme that is non-malleable against all families of adversaries of bounded size over $\{0, 1\}^n$ (such as Construction 1, Construction 2, and the probabilistic construction of [8]) can be turned into one having relative distance δ (and satisfying the same non-malleability guarantees) by composing the construction with a fixed outer code \mathcal{C} of block length n and relative distance δ . Indeed, any class \mathcal{F} of tampering functions for the composed code corresponds to a class \mathcal{F}' of the same size or less for the original construction. Namely, each function $f' \in \mathcal{F}'$ equals $\text{Dec}_{\mathcal{C}} \circ f$ ($\text{Dec}_{\mathcal{C}}$ being the decoder of \mathcal{C}) for some $f \in \mathcal{F}$. The caveat with this approach (rather than directly addressing distance as in Construction 1) is that the composition may lose strong non-malleability even if the original code is strongly non-malleable. Indeed, it may be the case that f is a sophisticated tampering function whereas its projection f' becomes as simple as the identity function. If so, non-malleability may be satisfied by choosing $\mathcal{D}_f := \mathcal{D}(\text{same})$ whereas strong non-malleability does not hold.

3.1 Proof of Theorem 3.1 for bijective adversaries

We first prove the theorem for adversaries that are bijective and have no fixed points. This case is still broad enough to contain interesting families of adversaries such as additive error adversaries \mathcal{F}_{add} mentioned in the introduction, for which case we reconstruct the existence proof of AMD codes (although optimal explicit constructions of AMD codes are already known [5]).

As it turns out, the analysis for this case is quite straightforward, and significantly simpler than the general case that we will address in Section 3.2.

Let $N := 2^n$, $K := 2^k$, and consider a fixed message $s \in \{0, 1\}^k$ and a fixed bijective tampering function $f: \{0, 1\}^n \rightarrow \{0, 1\}^n$ such that for all $x \in \{0, 1\}^n$, $f(x) \neq x$. We show that the non-malleability requirement of Definition 2.3 holds with respect to the distribution \mathcal{D}_f that is entirely supported on $\{\perp\}$. That is, we wish to show that with high probability, the coding scheme (Enc, Dec) of Construction 1 is so that

$$\Pr[\text{Dec}(f(\text{Enc}(s))) \neq \perp] \leq \epsilon. \quad (3)$$

By taking a union bound over all choices of f and s , this would imply that with high probability, the code is non-malleable (in fact, strongly non-malleable) for the entire family \mathcal{F} .

Let $E(s) := \text{supp}(\text{Enc}(s))$ be the set of the t codewords that are mapped to s by the decoder. Let E_1, \dots, E_t be the codewords in this set in the order they are picked by the code construction. For any $x \in \{0, 1\}^n \setminus E(s)$, we know that

$$\Pr[\text{Dec}(x) \neq \perp] \leq t(K-1)/(N-t) \leq \frac{\gamma}{1-\gamma},$$

where $\gamma := tK/N$. This can be seen by observing that each codeword in $E(s')$ for $s' \neq s$ is uniformly distributed on the set $\{0, 1\}^n \setminus E(s)$, and taking a union bound. Thus, in particular since $\{f(E_1), \dots, f(E_t)\}$ is a set of size t outside $E(s)$, we see that $\Pr[\text{Dec}(f(E_1)) \neq \perp] \leq \frac{\gamma}{1-\gamma}$. In fact, the same argument holds for $\text{Dec}(E_2)$ conditioned on any realization of $f(E_1)$, and more generally, one can derive for each $i \in [t]$,

$$\Pr[\text{Dec}(f(E_i)) \neq \perp \mid f(E_1), \dots, f(E_{i-1})] \leq \frac{\gamma}{1-\gamma}. \quad (4)$$

Define indicator random variables $0 = X_0, X_1, \dots, X_t \in \{0, 1\}$, where $X_i = 1$ iff $\text{Dec}(f(E_i)) \neq \perp$. From (4) and using Proposition B.1, we can deduce that for all $i \in [t]$, $\Pr[X_i = 1 \mid X_0, \dots, X_{i-1}] \leq \frac{\gamma}{1-\gamma}$. Now, using Proposition B.5, letting $X := X_1 + \dots + X_t$,

$$\Pr[X > \epsilon t] \leq \left(\frac{e\gamma}{\epsilon(1-\gamma)} \right)^{\epsilon t}.$$

Assuming $\gamma \leq \epsilon/4$, the above upper bound simplifies to $\exp(-\Omega(\epsilon t))$. By taking a union bound over all possible choices of s and f (that we trivially upper bound by $N|\mathcal{F}|$), it can be seen that, as long as $t \geq t_0$ for some choice of $t_0 = O\left(\frac{1}{\epsilon} \log\left(\frac{N|\mathcal{F}|}{\eta}\right)\right)$, the probability that (Enc, Dec) fails to satisfy (3) for some choice of s and f is at most η .

Finally, observe that the assumption $\gamma \leq \epsilon/4$ can be satisfied provided that $K \leq K_0$ for some choice of $K_0 = \Omega(\epsilon N/t)$, or equivalently, when $k \leq k_0$ for some choice of $k \geq n - \log t - \log(1/\epsilon)$. Note that for this case the proof obtains a better dependence on ϵ compared to (1) and (2).

3.2 Proof of Theorem 3.1 for general adversaries

First, we present a proof sketch describing the ideas and intuitions behind the general proof, and then proceed with a full proof of the theorem.

- **Proof sketch**

In the proof for bijective adversaries, we heavily used the fact that the tampering of each set $E(s)$ of codewords is a disjoint set of the same size. For general adversaries; however, this may not be true. Intuitively, since the codewords in $E(s)$ are chosen uniformly and almost independently at random (ignoring the distinctness dependencies), the tampered distribution $f(E(s))$ should look similar to $f(\mathcal{U}_n)$ for all s , if $|E(s)|$ is sufficiently large. Indeed, this is what shown in the proof. The proof also adjusts the probability mass of same according to the fraction of the fixed points of f , but we ignore this technicality for the proof sketch.

Note that the distribution $f(\mathcal{U}_n)$ may be arbitrary, and may assign a large probability mass to a small set of the probability space. For example, f may assign half of the probability mass to a single point. We call the points in $\{0, 1\}^n$ such that receive a noticeable share of the probability mass in $f(\mathcal{U}_n)$ the *heavy elements* of $\{0, 1\}^n$, and fix the randomness of the code construction so that the decoder's values at heavy elements

are revealed before analyzing each individual message s . Doing so allows us to analyze each message s separately and take a union bound on various choices of s as in the case of bijective adversaries. Contrary to the bijective case; however, the distribution \mathcal{D}_f is no longer entirely supported on \perp ; but we show that it still can be made to have a fairly small support; roughly $\text{poly}(n, \log |\mathcal{F}|)$. More precisely, the proof shows non-malleability with respect to the choice of \mathcal{D}_f which is explicitly defined to be the distribution of the following random variable:

$$D := \begin{cases} \text{same} & \text{if } f(U_n) = U_n, \\ \text{Dec}(f(U_n)) & \text{if } f(U_n) \neq U_n \text{ and } f(U_n) \in H, \\ \perp & \text{otherwise,} \end{cases}$$

where $H \subseteq \{0, 1\}^n$ is the set of heavy elements formally defined as

$$H := \{x \in \{0, 1\}^n : \Pr[f(U_n) = x] > 1/r\},$$

for an appropriately chosen $r = \Theta(\epsilon^2 t)$.

Although the above intuition is natural, turning it into a rigorous proof requires substantially more work than the bijective case, and the final proof turns out to be rather delicate even though it only uses elementary probability tools. The first subtlety is that revealing the decoder at the heavy elements creates dependencies between various random variables used in the analysis. In order to make the proof more intuitive, we introduce a random process, described as an algorithm *Reveal*, that gradually reveals information about the code as the proof considers the codewords E_1, \dots, E_t corresponding to the picked message s . The process outputs a list of elements in $\{0, 1\}^k$, and we show that the empirical distribution of this list is close to the desired \mathcal{D}_f for all messages s .

Roughly speaking, at each step $i \in [t]$ the analysis estimates the distribution of $\text{Dec}(f(E_i))$ conditioned on the particular realizations of the previous codewords. There are three subtleties that we need to handle to make this work:

1. The randomness corresponding to some of the E_i is previously revealed by the analysis and thus such codewords cannot be assumed to be uniformly distributed any more. This issue may arise due to the revealing of the decoder's values at heavy elements in the beginning of analysis, or existence of cycles in the evaluation graph of the tampering function f . Fortunately, it is straightforward to show that the number of such codewords remain much smaller than t with high probability, and thus they may simply be ignored.
2. At each step of the analysis, the revealed information make the distribution of $\text{Dec}(f(E_i))$ gradually farther from the desired \mathcal{D}_f . The proof ensures that the expected increase at each step is small, and using standard Martingale concentration bounds the total deviation from \mathcal{D}_f remains sufficiently small with high probability at the end of the analysis.
3. Obtaining small upper bounds (e.g., $\exp(-cn)$ for some $c < 1$) on the probability of various bad events in the analysis (e.g., $\text{Dec}(f(\text{Enc}(s)))$ significantly deviating from \mathcal{D}_f) is not difficult to achieve. However, extra care is needed to ensure that the probabilities are much smaller than $1/(2^k |\mathcal{F}|)$ (to accommodate the final union bound), where the latter may easily be doubly-exponentially small in n . An exponential upper bound of $\exp(-cn)$ does not even suffice for moderately large families of adversaries such as bit-tampering adversaries, for which we have $|\mathcal{F}| = 4^n$.

• **Complete proof of Theorem 3.1**

First, observe that by construction, the minimum distance of the final code is always greater than δn ; that is, whenever $\text{Dec}(w_1) \neq \perp$ and $\text{Dec}(w_2) \neq \perp$ for any pair of vectors $w_1 \neq w_2$, we have

$$\text{dist}_h(w_1, w_2) > \delta n,$$

where $\text{dist}_h(\cdot)$ denotes the Hamming distance. This is because whenever a codeword is picked, its δn neighborhood is removed from the sample space for the future codewords. Let V denote the volume of a Hamming ball of radius δn . It is well known that $V \leq 2^{nh(\delta)}$, where $h(\cdot)$ is the binary entropy function.

Fix an adversary $f \in \mathcal{F}$. We wish to show that the coding scheme (Enc, Dec) defined by Construction 1 is non-malleable with high probability for the chosen f .

Define $p_0 := \Pr[f(U_n) = U_n]$. In the sequel, assume that $p_0 < 1$ (otherwise, there is nothing to prove). For every $x \in \{0, 1\}^n$, define $p(x) := \Pr[f(U_n) = x \wedge x \neq U_n]$. Observe that

$$\sum_x p(x) = 1 - p_0.$$

We say that a string $x \in \{0, 1\}^n$ is *heavy* if

$$p(x) > 1/r,$$

for a parameter $r \leq t$ to be determined later. Note that the number of heavy strings must be less than r . Define

$$\begin{aligned} H &:= \{x \in \{0, 1\}^n : p(x) > 1/r\}, \\ \gamma &:= t/N, \\ \gamma' &:= tK/N. \end{aligned}$$

Fix the randomness of the code construction so that $\text{Dec}(x)$ is revealed for every heavy x . We will argue that no matter how the decoder's outcome on heavy elements is decided by the randomness of the code construction, the construction is non-malleable for every message s and the chosen function f with overwhelming probability. We will then finish the proof with a union bound over all choices of s and f .

Consider a random variable D defined over $\{0, 1\}^k \cup \{\perp, \text{same}\}$ in the following way:

$$D := \begin{cases} \text{same} & \text{if } f(U_n) = U_n, \\ \text{Dec}(f(U_n)) & \text{if } f(U_n) \neq U_n \text{ and } f(U_n) \in H, \\ \perp & \text{otherwise.} \end{cases} \quad (5)$$

For the chosen f , we explicitly define the distribution \mathcal{D}_f as $\mathcal{D}_f := \mathcal{D}(D)$.

Now, consider a fixed message $s \in \{0, 1\}^k$, and define the random variable $E_s := \text{Enc}(s)$. That is, E_s is uniformly supported on the set $E(s)$ (this holds by the way that the encoder is defined). Observe that the marginal distribution of each individual set $E(s)$ (with respect to the randomness of the code construction) is the same for all choices of s , regardless of the ordering assumed by Construction 1 on the message space $\{0, 1\}^k$.

Furthermore, define the random variable D_s as follows.

$$D_s := \begin{cases} \text{same} & \text{if } f(E_s) = E_s, \\ \text{Dec}(f(E_s)) & \text{otherwise.} \end{cases} \quad (6)$$

Our goal is to show that the distribution of D_s (for the final realization of the code) is ϵ -close to \mathcal{D}_f with high probability over the randomness of the code construction. Such assertion is quite intuitive by comparing the way the two distributions D_s and \mathcal{D}_f are defined. In fact, it is not hard to show that the assertion holds with probability $1 - \exp(-\Omega(n))$. However, such a bound would be insufficient to accommodate a union bound of even moderate sizes such as 2^n , which is needed for relatively simple classes such as bit-tampering adversaries. More work needs to be done to ensure that it is possible to achieve a high probability statement with failure probability much smaller than $1/|\mathcal{F}|$, which may in general be doubly exponentially small in n .

The claim below shows that closeness of $\mathcal{D}(D_s)$ to \mathcal{D}_f would imply non-malleability of the code.

Claim 3.3. *Suppose that for every $s \in \{0, 1\}^k$, we have $\mathcal{D}(D_s) \approx_\epsilon \mathcal{D}_f$ for the choice of \mathcal{D}_f defined in (5). Then, (Enc, Dec) is a non-malleable coding scheme with error ϵ and a strong non-malleable coding scheme with error 2ϵ .*

Proof. In order to verify Definition 2.5, we need to verify that for every distinct pair of messages $s_1, s_2 \in \{0, 1\}^k$, $\mathcal{D}(D_{s_1}) \approx_{2\epsilon} \mathcal{D}(D_{s_2})$. But from the assumption, we know that $\mathcal{D}(D_{s_1})$ and $\mathcal{D}(D_{s_2})$ are both ϵ -close to \mathcal{D}_f . Thus the result follows by the triangle inequality.

It is of course possible now to use [8, Theorem 3.1] to deduce that Definition 2.3 is also satisfied. However, for the clarity of presentation, here we give a direct argument that shows that non-malleability is satisfied with the precise choice of \mathcal{D}_f defined in (5) and error ϵ . Let $s \in \{0, 1\}^k$, and let $E_s := \text{Enc}(s)$ and $S := \text{Dec}(f(E_s))$. Let $S' \sim \mathcal{D}_f$ and $S'' \sim \mathcal{D}(D_s)$ be sampled independently. We need to show that

$$\mathcal{D}(S) \approx_\epsilon \mathcal{D}(\text{copy}(S', s)). \quad (7)$$

From the definition of D_s in (6), since $\text{Dec}(f(E_s)) = s$ when $f(E_s) = E_s$, we see that $\mathcal{D}(\text{copy}(S'', s)) = \mathcal{D}(\text{Dec}(f(E_s))) = \mathcal{D}(S)$. Now, since by assumption $\mathcal{D}(S') \approx_\epsilon \mathcal{D}(S'')$, it follows that $\mathcal{D}(\text{copy}(S', s)) \approx_\epsilon \mathcal{D}(\text{copy}(S'', s))$ which proves (7). \square

Let the random variables E_1, \dots, E_t be the elements of $E(s)$, in the order they are sampled by Construction 1.

Define, for $i \in [t]$,

$$S_i := \begin{cases} \text{same} & \text{if } f(E_i) = E_i, \\ \text{Dec}(f(E_i)) & \text{otherwise.} \end{cases}$$

We note that, no matter how the final code is realized by the randomness of the construction, the distribution D_s is precisely the empirical distribution of S_1, \dots, S_t as determined by the code construction.

In the sequel, for each $i \in [t]$, we analyze the distribution of the variable S_i conditioned on the values of S_1, \dots, S_{i-1} and use this analysis to prove that the empirical distribution of the sequence (S_1, \dots, S_t) is close to \mathcal{D}_f .

In order to understand the empirical distribution of the S_i , we consider the following process Reveal that considers the picked codewords E_1, \dots, E_t in order, gradually reveals information about the code construction, and outputs a subset of the S_i . We will ensure that

1. The process outputs a large subset of $\{S_1, \dots, S_t\}$, and,
2. The empirical distribution of the sequence output by the process is close to \mathcal{D}_f with high probability.

The above guarantees would in turn imply that the empirical distribution of the entire sequence S_i is also close to \mathcal{D}_f with high probability. We define the process as follows.

Process Reveal:

1. Initialize the set $\text{Skip} \subseteq [t]$ with the empty set. Recall that the values of $\text{Dec}(w)$ for all $w \in H$ are already revealed in the analysis, as well as $\text{Dec}(\Gamma(w))$ for those for which $\text{Dec}(w) \neq \perp$.
2. For each heavy element $w \in H$, if $\text{Dec}(w) = s$, consider the unique $j \in [t]$ such that $E_j = w$. Reveal⁴ j and E_j , and add j to Skip .
3. For i from 1 to t , define *the i th stage* as follows:
 - 3.1. If $i \in \text{Skip}$, declare a *skip* and continue the loop with the next i . Otherwise, follow the remaining steps.
 - 3.2. Reveal $\Gamma(E_i)$. Note that revealing E_i implies that $\text{Dec}(E_i)$ is revealed as well, since $\text{Dec}(E_i) = s$. Moreover, recall that for any $x \in \Gamma(E_i) \setminus E_i$, $\text{Dec}(x) = \perp$ by the code construction.
 - 3.3. If $\text{Dec}(f(E_i))$ is not already revealed:
 - 3.3.1. Reveal $\text{Dec}(f(E_i))$.
 - 3.3.2. If $\text{Dec}(f(E_i)) = s$, consider the unique $j \in [t]$ such that $E_j = f(E_i)$. It must be that $j > i$, since $\text{Dec}(E_j)$ has not been revealed before. Reveal j and add it to Skip .
 - 3.3.3. Declare that an *unveil* has happened if $\text{Dec}(f(E_i)) \neq \perp$. If so, reveal $\text{Dec}(f(x))$ for all $x \in \Gamma(f(E_i)) \setminus E_i$ to equal \perp .
 - 3.4. Reveal and output S_i .

For $i \in [t]$, we use the notation Reveal_i to refer to all the information revealed from the beginning of the process up to the time the i th stage begins. We also denote by $\text{Next}(i)$ the least $j > i$ such that a skip does not occur at stage j ; define $\text{Next}(i) := t + 1$ if no such j exists, and define $\text{Next}(0)$ to be the index of the first stage that is not skipped. Moreover, for $w \in \{0, 1\}^n$, we use the notation $w \in \text{Reveal}_i$ as a shorthand to denote the event that the process Reveal has revealed the value of $\text{Dec}(w)$ at the time the i th stage begins.

By the way the code is constructed, the decoder's value at each given point is most likely \perp . We make this intuition more rigorous and show that the same holds even conditioned on the information reveal by the process Reveal .

Claim 3.4. *For all $i \in [t]$ and any $a \in \text{supp}(\text{Reveal}_i)$,*

$$\Pr[\text{Dec}(x) \neq \perp \mid (\text{Reveal}_i = a) \wedge (x \notin \text{Reveal}_i)] \leq \gamma' / (1 - 3\gamma V).$$

Proof. Suppose $x \notin \text{Reveal}_i$, and observe that Reveal_i at each step reveals at most the values of the decoder at $2V$ points; namely, $\Gamma(E_i)$ and $\Gamma(f(E_i))$. Moreover, before the first stage, decoder's value is revealed at up to r heavy points and its Hamming neighborhood at radius δn . In total, the total number of points at which decoder's value is revealed by the information in Reveal_i is at most

$$(|H| + 2(i - 1))|V| \leq (2t + r)V \leq 3\gamma VN.$$

⁴In a rigorous sense, by revealing a random variable we mean that we condition the probability space on the event that a particular value is assumed by the variable. For example, revealing E_i means that the analysis branches to a conditional world where the value of E_i is fixed to the revealed value. In an intuitive way, one may think of a reveal as writing constraints on the realization of the code construction on a blackboard, which is subsequently consulted by the analysis (in form of the random variable Reveal_i that the analysis defines to denote the information revealed by the process before stage i).

Let

$$\mathcal{C} := \bigcup_s E(s)$$

be the set of all codewords of the coding scheme. Some of the elements of \mathcal{C} are already included in Reveal_i , and by assumption we know that none of these is equal to x .

The distribution of each unrevealed codeword, seen in isolation, is uniform over the $N(1 - 3\gamma V)$ remaining vectors in $\{0, 1\}^n$. Thus by taking a union bound on the probability of each such codeword hitting the point x (which is the only way to make $\text{Dec}(x) \neq \perp$, we deduce that

$$\Pr[\text{Dec}(x) \neq \perp \mid \text{Reveal}_i = a] \leq \frac{tK}{N(1 - 3\gamma V)} = \gamma'/(1 - 3\gamma V). \quad \square$$

Ideally, for each $i \in [t]$ we desire to have E_i almost uniformly distributed, conditioned on the revealed information, so that the distribution of $\text{Dec}(f(E_i))$ (which is described by S_i when E_i does not hit a fixed point of f) becomes close to $\text{Dec}(f(U_n))$. However, this is not necessarily true; for example, when the process Reveal determines the decoder's value on the heavy elements, the value of, say, E_1 may be revealed, at which point there is no hope to ensure that E_1 is nearly uniform. This is exactly what the set Skip is designed for, to isolate the instances when the value of E_i is already determined by the prior information. More precisely, we have the following.

Claim 3.5. *Suppose that $i \notin \text{Skip}$ when the i th stage of Reveal begins. Then, for any $a \in \text{supp}(\text{Reveal}_i)$,*

$$\mathcal{D}(E_i \mid \text{Reveal}_i = a) \approx_\nu \mathcal{U}_n,$$

where $\nu := (3\gamma V)/(1 - 3\gamma V)$.

Proof. Note that, without any conditioning, the distribution of E_i is exactly uniform on $\{0, 1\}^n$. If at any point prior to reaching the i th stage it is revealed that $\text{Dec}(E_i) = s$, either line 2 or line 3.2.2 of process Reveal ensures that i is added to the set Skip .

If, on the other hand, the fact that $\text{Dec}(E_i) = s$ has not been revealed when the i th stage begins, the distribution of E_i becomes uniform on the points in $\{0, 1\}^n$ that have not been revealed yet. As in Claim 3.4, the number of revealed points is at most $(2t + r)V \leq 3\gamma V N$. Thus, the conditional distribution E_i remains $((3\gamma V)/(1 - 3\gamma V))$ -close to uniform by Proposition B.2. \square

For each $i \in [t]$, define a random variable $S'_i \in \{0, 1\}^k \cup \{\text{same}, \perp\}$ as follows (where U_n is independently sampled from \mathcal{U}_n):

$$S'_i := \begin{cases} \text{same} & \text{if } f(U_n) = U_n, \\ \text{Dec}(f(U_n)) & \text{if } f(U_n) \neq U_n \wedge f(U_n) \in \text{Reveal}_i, \\ \perp & \text{otherwise.} \end{cases} \quad (8)$$

Note that $\mathcal{D}(S'_1) = \mathcal{D}_f$.

Intuitively, S'_i is the ‘‘cleaned up’’ version of the random variable S_i that we are interested in. As defined, S'_i is an independent random variable, and as such we are more interested in its *distribution* than value. Observe that the distribution of S'_i is randomly determined according to the randomness of the code construction (in particular, the knowledge of Reveal_i completely determines $\mathcal{D}(S'_i)$). The variable S'_i is defined so that its distribution approximates the distribution of the actual S_i conditioned on the revealed information before stage i . Formally, we can show that conditional distributions of these two variables are (typically) similar. Namely,

Claim 3.6. *Suppose that $i \notin \text{Skip}$ when the i th stage of Reveal begins. Then, for any $a \in \text{supp}(\text{Reveal}_i)$,*

$$\mathcal{D}(S_i | \text{Reveal}_i = a) \approx_\nu \mathcal{D}(S'_i | \text{Reveal}_i = a),$$

where $\nu := (3\gamma V + \gamma')/(1 - 3\gamma V)$.

Proof. First, we apply Claim 3.5 to ensure that

$$\mathcal{D}(E_i | \text{Reveal}_i = a) \approx_{\nu'} \mathcal{U}_n,$$

where $\nu' = (3\gamma V)/(1 - 3\gamma V)$. Thus we can assume that the conditional distribution of E_i is exactly uniform at cost of a ν' increase in the final estimate.

Now, observe that, conditioned on the revealed information, the way S_i is sampled at stage i of Reveal can be rewritten as follows:

1. Sample $E_i \sim \mathcal{U}_n$.
2. If $f(E_i) = E_i$, set $S_i \leftarrow \text{same}$.
3. Otherwise, if $f(E_i) \in \text{Reveal}_i$, set S_i to $\text{Dec}(f(E_i))$ as determined by the revealed information.
4. Otherwise, reveal $\text{Dec}(f(E_i))$ (according to its conditional distribution on the knowledge of Reveal_i) and set S accordingly.

This procedure is exactly the same as how S'_i is sampled by (8); with the difference that at the third step, S'_i is set to \perp whereas S_i is sampled according to the conditional distribution of $\text{Dec}(f(E_i))$. However, we know by Claim 3.4 that in this case,

$$\Pr[\text{Dec}(f(E_i)) \neq \perp \mid \text{Reveal}_i = a] \leq \gamma'/(1 - 3\gamma V).$$

Thus we see that S_i changes the probability mass of \perp in $\mathcal{D}(S'_i)$ by at most $\gamma'/(1 - 3\gamma V)$. The claim follows. \square

Recall that the distribution of S'_1 is the same as \mathcal{D}_f . However, for subsequent stages this distribution may deviate from \mathcal{D}_f . We wish to ensure that by the end of process Reveal, the deviation remains sufficiently small.

For $i \in [t - 1]$, define Δ_i as

$$\Delta_i := \text{dist}(\mathcal{D}(S'_{i+1}), \mathcal{D}(S'_i)).$$

where $\text{dist}(\cdot)$ denotes statistical distance. Note that Δ_i is a random variable that is determined by the knowledge of Reveal_{i+1} (recall that Reveal_i determines the exact distribution of S'_i). We show that the conditional values attained by this random variable are small in expectation.

Claim 3.7. *For each $i \in [t - 1]$, and all $a \in \text{supp}(\text{Reveal}_i)$,*

$$\mathbb{E}[\Delta_i | \text{Reveal}_i = a] \leq \frac{2\gamma'}{r(1 - 3\gamma V)}. \tag{9}$$

Moreover, $\Pr[\Delta_i \leq 2/r \mid \text{Reveal}_i = a] = 1$.

Proof. Recall that the distribution of S'_{i+1} is different from S'_i depending on the points at which the decoder's value is revealed during stage i of Reveal. If a skip is declared at stage i , we have $\text{Reveal}_{i+1} = \text{Reveal}_i$ and thus, $\Delta_i = 0$. Thus in the following we may assume that this is not the case.

However, observe that whenever for some $x \in \{0, 1\}^n$, the decoder's value $\text{Dec}(x)$ is revealed at stage i , the new information affects the probability distribution of S'_i only if $\text{Dec}(x) \neq \perp$. This is because when $\text{Dec}(x) = \perp$, some of the probability mass assigned by S_i to \perp in (8) is removed and reassigned by S'_{i+1} to $\text{Dec}(x)$, which is still equal to \perp . Thus, changes of this type can have no effect on the distribution of S'_i . We conclude that only revealing the value of E_i and an unveil (as defined in line 3.3.3 of process Reveal) can contribute to the statistical distance between S'_i and S'_{i+1} .

Whenever an unveil occurs at stage i , say at point $x \in \{0, 1\}^n$, some of the probability mass assigned to \perp by S'_i is moved to $\text{Dec}(x)$ in the distribution of S'_{i+1} . Since we know that $x \notin H$, the resulting change in the distance between the two distributions is bounded by $1/r$, *no matter* what the realization of x and $\text{Dec}(x)$ are. Overall, using Claim 3.4, the expected change between the two distributions contributed by the occurrence of an unveil is upper bounded by the probability of an unveil occurring times $1/r$, which is at most

$$\frac{\gamma'/r}{1 - 3\gamma V}. \quad (10)$$

The only remaining factor that may contribute to an increase in the distance between distribution of S'_i and S'_{i+1} is the revealing of E_i at stage i . The effect of this reveal in the statistical distance between the two distributions is $p(E_i)$, since according to (8) the value of S'_{i+1} is determined by the outcome of $f(U_n)$, and thus the probability mass assigned to $\text{Dec}(E_i)$ by S'_{i+1} is indeed $\Pr[f(U_n) = E_i]$. Let \mathcal{D}_E be the distribution of E_i conditioned on the knowledge of Reveal_i . Observe that, since the values $\{p(x) : x \in \{0, 1\}^n\}$ defines a probability distribution on N points, we clearly have

$$\sum_{x \in \text{supp}(\mathcal{D}_E)} p(x) \leq 1. \quad (11)$$

On the other hand, by the assumption that a skip has not occurred at stage i , we can deduce using the argument in Claim 3.5 that \mathcal{D}_E is uniformly supported on a support of size at least $N(1 - 3\gamma V)$. Therefore, using (11), the expected contribution to Δ_i by the revealing of E_i is (which is the expected value of $p(E_i)$) is at most

$$\frac{1}{N(1 - 3\gamma V)} \leq \frac{\gamma'/r}{(1 - 3\gamma V)}, \quad (12)$$

where the inequality uses $r \leq \gamma'N = tK$. The desired bound follows by adding up the two perturbations (10) and (12) considered.

Finally, observe that each of the perturbations considered above cannot be more than $1/r$, since stage i never reveals the decoder's value on a heavy element (recall that all heavy elements are revealed before the first stage begins and the choices of E_i that correspond to heavy elements are added to Skip when Reveal begins). Thus, the conditional value of Δ_i is never more than $2/r$. \square

Using the above result, we can deduce a concentration bound on the summation of the differences Δ_i .

Claim 3.8. *Let $\Delta := \Delta_1 + \dots + \Delta_{t-1}$, and suppose*

$$\frac{\gamma'}{1 - 3\gamma V} \leq \frac{\epsilon r}{32t}. \quad (13)$$

Then,

$$\Pr[\Delta \geq \epsilon/8] \leq \exp(-\epsilon^2 r^2 / (2048t)) =: \eta_0. \quad (14)$$

Proof. For $i \in [t-1]$, define $\Delta'_i := \Delta_i r/2$, $\Delta'_0 := 0$, and $\Delta' := \Delta'_1 + \dots + \Delta'_{t-1}$. Since Reveal_i determines Δ'_{i-1} , by Claim 3.7 we know that

$$\mathbb{E}[\Delta'_i | \Delta'_0, \dots, \Delta'_{i-1}] \leq \nu,$$

where $\nu := \frac{\gamma'}{1-3\gamma V} \leq \epsilon r / (32t)$. In the above, conditioning on $\Delta'_0, \dots, \Delta'_{i-1}$ instead of Reveal_i (for which Claim 3.7 applies), is valid in light of Proposition B.1, since the knowledge of Reveal_i determines $\Delta'_0, \dots, \Delta'_{i-1}$.

Moreover, again by the Claim 3.7, we know that the Δ'_i are between 0 and 1. Using Proposition B.3, it follows that

$$\Pr[\Delta \geq \epsilon/8] = \Pr[\Delta' \geq \frac{\epsilon r}{16t} \cdot t] \leq \eta_0. \quad \square$$

Next, we prove a concentration bound for the total number of unveils that can occur in line 3.3.3 of process Reveal .

Claim 3.9. *Let u be the total number of unveils that occur in process Reveal . Assuming $\gamma' / (1 - 3\gamma V) \leq \epsilon/8$ (which is implied by (13)), we have*

$$\Pr[u \geq \epsilon t / 4] \leq \exp(-\epsilon^2 t / 128) \leq \eta_0.$$

Proof. Let X_1, \dots, X_t be indicator random variable such that $X_i = 1$ iff an unveil occurs at stage i , and let $X_0 := 0$. Recall that an unveil can only occur at a stage that is not skipped. Thus, if $i \in [t]$ when the i th stage begins, we can deduce that $X_i = 0$.

Consider $i \in [t]$ such that $i \notin \text{Skip}$ when the i th stage begins. An unveil occurs when $\text{Dec}(f(E_i)) \notin \text{Reveal}_i$. In this case, by Claim 3.4, we get that

$$\Pr[\text{Dec}(f(E_i)) \neq \perp \mid \text{Reveal}_i] \leq \gamma' / (1 - 3\gamma V).$$

Since Reveal_i determines all the revealed information in each prior stage, and in particular the values of X_0, \dots, X_{i-1} , we can use Proposition B.1 to deduce that

$$\Pr[X_i = 1 \mid X_0, \dots, X_{i-1}] \leq \gamma' / (1 - 3\gamma V).$$

Finally, Proposition B.3 derives the desired concentration bound on the number of unveils, which is $X_1 + \dots + X_t$. \square

We are now ready to wrap up the proof and show that with overwhelming probability, the empirical distribution of S_1, \dots, S_t is ϵ -close to \mathcal{D}_f .

Suppose that process Reveal outputs a subset of the S_i . Let $T \subseteq [t]$ be the set of indices i such Reveal outputs S_i in the end of the i th stage. Note that $T = [t] \setminus \text{Skip}$, where Skip denotes the skip set when Reveal terminates. Observe that $|\text{Skip}|$ is at most the total number of unveils occurring at line 3.3.3 of Reveal plus r (which upper bounds the number of heavy elements in H). Thus, using Claim 3.9 we see that, assuming (13),

$$\Pr[t - |T| \geq r + \epsilon t / 4] \leq \eta_0. \quad (15)$$

Let δ_i for $i \in [t]$ denote the statistical distance between S'_i and \mathcal{D}_f . We know that δ_i is a random variable depending on Reveal_i . Thus, the value of δ_i becomes known to a particular fixed value conditioned on the

outcome of every Reveal_j , $j \geq i$. Define $\delta_0 := \max_i \delta_i$, which is a random variable that becomes revealed by the knowledge of Reveal_t in the end of the process.

Using Claim 3.6, we thus know that for any $a \in \text{supp}(\text{Reveal}_i)$ and $i \in T$,

$$\mathcal{D}(S_i | \text{Reveal}_i = a) \approx_{\nu_0 + \delta_0} \mathcal{D}_f,$$

where

$$\nu_0 := (3\gamma V + \gamma') / (1 - 3\gamma V).$$

Let \mathcal{S} denote the empirical distribution of $\{S_i : i \in T\}$, and define $S_0 := \perp$. From the above conclusion, using Proposition B.1 we can now write, for $i \in T$,

$$\mathcal{D}(S_i | (S_j : j \in T \cap \{1, \dots, i-1\})) \approx_{\nu_0 + \delta_0} \mathcal{D}_f.$$

Recall that $|\text{supp}(\mathcal{D}_f)| \leq r + 2$. Assuming that

$$\nu_0 + \delta_0 \leq \epsilon/4, \tag{16}$$

Proposition B.7 implies (after simple manipulations) that with probability $1 - \eta_1$, where

$$\eta_1 \leq 2^{r+4 - \Omega(\epsilon^2 |T|)}, \tag{17}$$

\mathcal{S} is $(\epsilon/2)$ -close to \mathcal{D}_f .

Recall that $\mathcal{D}(S'_1) = \mathcal{D}_f$. Using the triangle inequality for statistical distance, for every $i \in [t]$ we can write

$$\text{dist}(S'_i, \mathcal{D}_f) = \text{dist}(S'_i, S'_1) \leq \Delta_1 + \dots + \Delta_{i-1} \leq \Delta,$$

and thus deduce that $\delta_0 \leq \Delta$. Recall that by Claim 3.8, we can ensure that, assuming (13), $\Delta \leq \epsilon/8$ (and thus, $\delta_0 \leq \epsilon/8$) with probability at least $1 - \eta_0$. Thus under the assumption that

$$\nu_0 \leq \epsilon/8, \tag{18}$$

and (13), which we recall below

$$\frac{\gamma'}{1 - 3\gamma V} \leq \frac{\epsilon r}{32t},$$

we can ensure that $\nu_0 + \delta_0 \leq \epsilon/4$ with probability at least $1 - \eta_0$. Moreover, conditioned on the event $\nu_0 + \delta_0 \leq \epsilon/4$ (recall that δ_0 is a random variable), we have already demonstrated that with probability at least $1 - \eta_1$, \mathcal{S} is $(\epsilon/2)$ -close to \mathcal{D}_f . After removing conditioning on the bound on δ_0 , we may deduce that overall (under the assumed inequalities (13) and (18)), with probability at least $1 - O(\eta_0 + \eta_1)$,

$$\mathcal{S} \approx_{\epsilon/2} \mathcal{D}_f,$$

which in turn, implies that the empirical distribution of S_1, \dots, S_t becomes ϵ' -close to uniform, where

$$\epsilon' := \epsilon/2 + (1 - |T|/t).$$

Finally, we can use (15) to ensure that (assuming (13)), $\epsilon' \leq \epsilon$ and $|T|/t \geq 1 - \epsilon/2$ with probability at least $1 - O(\eta_0 + \eta_1)$ as long as

$$r \leq \epsilon t / 4. \tag{19}$$

By comparing (17) with (14), we also deduce that $\eta_1 = O(\eta_0)$ (and also that (19) holds) as long as $r \leq r_0$ for some

$$r_0 = \Omega(\epsilon^2 t). \quad (20)$$

Altogether, we arrive at the conclusion that under assumptions (13), (18), and by taking $r := r_0$, with probability at least $1 - O(\eta_0)$,

$$(\text{empirical distribution of } (S_1, \dots, S_t)) \approx_\epsilon \mathcal{D}_f,$$

which ensures the required non-malleability condition for message s and tampering function f . By taking a union bound over all possible choices of s and f , the probability of failure becomes bounded by

$$O(\eta_0 K |\mathcal{F}|) =: \eta_2.$$

We can now ensure that $\eta_2 \leq \eta$ for the chosen value for r by taking $t \geq t_0$ for some

$$t_0 = O\left(\frac{1}{\epsilon^6} \left(\log \frac{|\mathcal{F}|N}{\eta}\right)\right). \quad (21)$$

Furthermore, in order to satisfy assumptions (13), (18), and the requirement $tKV \leq 1$ which is needed to make the construction possible, it suffices to have $K \leq K_0$ for some

$$K_0 = \Omega(\epsilon^3 N / (tV)).$$

Using the bound $V \leq 2^{nh(\delta)}$, where $h(\cdot)$ is the binary entropy function, and taking the logarithm of both sides, we see that it suffices to have $k \leq k_0$ for some

$$k_0 \geq n(1 - h(\delta)) - \log t - 3 \log(1/\epsilon) - O(1).$$

This concludes the proof of Theorem 3.1.

3.3 Efficiency in the random oracle model

One of the main motivations of the notion of non-malleable codes proposed in [8] is the application for *tamper-resilient security*. In this application, a stateful consists of a public functionality and a private state $s \in \{0, 1\}^k$. The state is stored in form of its non-malleable encoding, which is prone to tampering by a family of adversaries. It is shown in [8] that the security of the system with encoded private state can be guaranteed (in a naturally defined sense) provided that the distribution \mathcal{D}_f related to the non-malleable code is efficiently samplable. In light of Remark 2.7, efficient sampling of \mathcal{D}_f can be assured if the non-malleable code is equipped with an efficient encoder and decoder.

Although the code described by Construction 1 may require exponential time to even describe, it makes sense to consider efficiency of the encoder and the decoder in the *random oracle model*, where all involved parties have oracle access to a shared, exponentially long, random string. The uniform decoder construction of [8] is shown to be efficiently implementable in the random oracle model in an *approximate* sense (as long as all involved parties query the random oracle a polynomial number of times), assuming existence of an efficient algorithm implementing a uniformly random permutation Π and its inverse Π^{-1} .

We observe that Construction 1, for the distance parameter $\delta = 0$ (which is what needed for strong non-malleability as originally defined in [8]) can be *exactly* implemented efficiently (without any further

assumptions on boundedness of the access to the random oracle) assuming access to a uniformly random permutation and its inverse (i.e., the so-called ideal-cipher model). This is because our code is designed so that the codewords are picked uniformly at random and without replacement. More precisely, the encoder, given message $s \in \{0, 1\}^k$, can sample a uniformly random $i \in [t]$, and output $\Pi(s, i)$, where (s, i) is interpreted as an element of $\{0, 1\}^n$ (possibly after padding).

As noted in [8], efficient approximate implementations of uniformly random permutations exist in the random oracle model. In particular, [4] show such an approximation with security $\text{poly}(q)/2^n$, where q is the number of queries to the random oracle.

4 A Monte Carlo construction for computationally bounded adversaries

An important feature of Construction 1 is that the proof of non-malleability, Theorem 3.1, only uses limited independence of the permutation defining the codewords $E(s)$ corresponding to each message. This is because the proof analyzes the distribution of $\text{Dec}(f(\text{Enc}(s)))$ for each individual message separately, and then takes a union bound on all choice of s .

More formally, below we show that Theorem 3.1 holds for a broader range of code constructions than the exact Construction 1.

Definition 4.1 (ℓ -wise independent schemes). Let (Enc, Dec) be any randomized construction of a coding scheme with block length n and message length k . For each $s \in \{0, 1\}^k$, define $E(s) := \text{supp}(\text{Enc}(s))$ and let $t_s := |\text{supp}(\text{Enc}(s))|$. We say that the construction is ℓ -wise independent if the following are satisfied.

1. For any realization of (Enc, Dec) , the distribution of $\text{Enc}(s)$ (with respect to the internal randomness of Enc) is uniform on $\text{supp}(\text{Enc}(s))$.
2. The distribution of the codewords defined by the construction is ℓ -wise independent. Formally, we require the following. Let $\mathcal{C} := \bigcup_{s \in \{0, 1\}^k} \text{supp}(\text{Enc}(s))$. Suppose the construction can be described by a deterministic function⁵ $E: \{0, 1\}^k \times \mathbb{N} \times \mathbb{N} \rightarrow \{0, 1\}^n$ such that for a bounded random oracle \mathcal{O} over \mathbb{N} (describing the random bits used by the construction), the sequence

$$(E(s, i, \mathcal{O}))_{s \in \{0, 1\}^k, i \in [t_s]}$$

enumerates the set \mathcal{C} . Moreover, for any set of t indices $S = \{(s_j, i_j) : j \in [t], s_j \in \{0, 1\}^k, i_j \in [t_{s_j}]\}$, we have

$$\mathcal{D}(E(s_1, i_1, \mathcal{O}), \dots, E(s_\ell, i_\ell, \mathcal{O})) = \mathcal{D}(\Pi(1), \dots, \Pi(\ell))$$

for a uniformly random bijection $\Pi: [2^n] \rightarrow \{0, 1\}^n$.

Lemma 4.2. *Let (Enc, Dec) be any randomized construction of a coding scheme with block length n and message length k . For each $s \in \{0, 1\}^k$, define $E(s) := \text{supp}(\text{Enc}(s))$. Suppose that for any realization of (Enc, Dec) , and for every $s_1, s_2 \in \{0, 1\}^k$, we have*

1. $|E(s_1)| \geq t_0$, where t_0 is the parameter defined in Theorem 3.1.
2. $|E(s_2)| = O(|E(s_1)|)$.

⁵As an example, in Construction 1, all the values t_s are equal to the chosen t , and moreover, one can take $E(s, i, \mathcal{O}) = \Pi(s, i)$, where $\Pi: \{0, 1\}^k \times [2^{n-k}] \rightarrow \{0, 1\}^n$ is a uniformly random bijection defined by the randomness of \mathcal{O} .

Moreover, suppose that $k \leq k_0$, for k_0 as in Theorem 3.1. Let $t := \max_s |E(s)|$. Then, assuming that the construction is $(3t)$ -wise independent, the conclusion of Theorem 3.1 for distance parameter $\delta = 0$ holds for the coding scheme (Enc, Dec) .

Proof. We argue that the proof of Theorem 3.1 holds without any technical change if

1. The codewords in $\text{supp}(\text{Enc}(\mathcal{U}_k))$ are chosen not fully independently but $(3t)$ -wise independently, and
2. Each set $E(s)$ is not necessarily of exact size t but of size at least t_0 and $\Theta(t)$.

The key observation to be made is that the proof analyzes each individual message $s \in \{0, 1\}^k$ separately, and then applies a union bound on all choices of s . Thus we only need sufficient independence to ensure that the view of the analysis on each individual choice of the message is statistically the same as the case where the codewords are chosen fully independently.

Observe that the bulk of the information about the code looked up by the analysis for analyzing each individual message is contained in the random variable Reveal_{t+1} defined in the proof of Theorem 3.1, that is defined according to how the process Reveal evolves. Namely, Reveal_{t+1} summarizes all the information revealed about the code by the end of the process Reveal .

For a fixed message $s \in \{0, 1\}^n$ the process Reveal iterates for $|E(s)| \leq t$ step. At each step, the location of at most two codewords in $\text{supp}(\text{Enc}(\mathcal{U}_k))$ is revealed. Moreover, before the process starts, the values of the decoder on the heavy elements in H , which can correspond to less than t codewords, are revealed by the process. The only other place in the proof where an independent codeword is required is the union bound in the proof of Claim 3.4, which needs another degree of independence. Altogether, we conclude that the proof of Theorem 3.1 only uses at most $3t$ degrees of independence in the distribution of the codewords picked by the construction.

Moreover, for each message s , the analysis uses the fact that $|E(s)| \geq t_0$ to ensure that the code does not satisfy non-malleability for the given choice of s and tampering function remains below the desired level. Since $|E(s)|$ for different values of s are assumed to be within a constant factor of each other, the requirement (20) may also be satisfied by an appropriate choice of the hidden constant. Finally, using the fact that $\max_s |E(s)| = O(\min_s |E(s)|)$, we can also ensure that assumptions (13), and (18) can be satisfied for appropriate choices of the hidden constants in asymptotic bounds. \square

In order to implement an efficient ℓ -wise independent coding scheme, we use the bounded independence property of polynomial evaluations over finite fields. More precisely, we consider the coding scheme given in Construction 2.

The advantage of using the derandomized Monte Carlo construction is that the number of random bits required to describe the code is dramatically reduced from $O(tnK)$ bits (which can be exponentially large if the rate of the code is $\Omega(1)$) to only $O(tn)$ bits, which is only polynomially large if $t = \text{poly}(n)$. In order to efficiently implement the derandomized construction, we use bounded independence properties of polynomial evaluation. Using known algorithms for finite field operations and root finding, the implementation can be done in polynomial time.

Lemma 4.3. *Consider the pair $(\text{EncMC}, \text{DecMC})$ defined in Construction 2. For every $\eta > 0$, there is a $t_0 = O(n + \log(1/\eta))$ such that for every $t \geq t_0$ (where t is a power of two), with probability at least $1 - \eta$ the following hold.*

- *Given:* Integer parameters $0 < k \leq n$ and integer $t > 1$ which is a power of two. Let $b := \log(2t)$ and $m := n - k - b$.
- *Output:* A coding scheme (EncMC, DecMC) of block length n and message length k .
- *Randomness of the construction:* A uniformly random polynomial $P \in \mathbb{F}_{2^n}[9t - 1]$.
- *Construction of EncMC:* Given $s \in \{0, 1\}^k$,
 1. Initialize a set $E \subseteq \{0, 1\}^n$ to the empty set.
 2. For every $z \in \{0, 1\}^b$,
 - (a) Construct a vector $y := (s, 0^m, z) \in \{0, 1\}^n$ and regard it as an element of \mathbb{F}_{2^n} .
 - (b) Solve $P(X) = y$, and add the set of solutions (which is of size at most $9t - 1$) to E .
 3. Output a uniformly random element of E .
- *Construction of DecMC:* Given $x \in \{0, 1\}^n$, interpret x as an element of \mathbb{F}_{2^n} , and let $y := P(x)$, interpreted as a vector $(y_1, \dots, y_n) \in \{0, 1\}^n$. If $(y_{k+1}, y_{k+2}, \dots, y_{k+m}) = 0^m$, output (y_1, \dots, y_k) . Otherwise, output \perp .

Construction 2: The Monte Carlo Construction.

1. (EncMC, DecMC) is a $(9t)$ -wise independent coding scheme.
2. For all $s \in \{0, 1\}^k$, $|\text{supp}(\text{EncMC}(s))| \in [t, 3t]$.

Proof. Let $N := 2^n$ and $K := 2^k$. Consider the vector $X := (X_1, \dots, X_N) \in \mathbb{F}_{2^n}^N$, where $X_i := P(i)$ and each i is interpreted as an element of \mathbb{F}_{2^n} . Since the polynomial P is of degree $9t - 1$, the distribution of X_1, \dots, X_N over the randomness of the polynomial P is $(9t)$ -wise independent with each individual X_i being uniformly distributed on \mathbb{F}_{2^n} . This standard linear-algebraic fact easily follows from invertibility of square Vandermonde matrices.

Note that the decoder function DecMC in Construction 2 is defined so that

$$\text{DecMC}(U_n) = \begin{cases} \perp & \text{with probability } 1 - 2tK/N \\ s \in \{0, 1\}^k & \text{with probability } 2t/N. \end{cases} \quad (22)$$

For $s \in \{0, 1\}^k$, let $E(s) := \text{supp}(\text{EncMC}(s))$. Note that the encoder, given s , is designed to output a uniformly random element of $E(s)$. Since the definition of the EncMC(s) is so that it exhausts the list of all possible words in $\{0, 1\}^n$ that can lie in $\text{DecMC}^{-1}(s)$, it trivially follows that (EncMC, DecMC) is always a valid coding scheme; that is, for any realization of the code and for all $s \in \{0, 1\}^k$, we have $\text{DecMC}(\text{EncMC}(s)) = s$ subject to the guarantee that $|E(s)| > 0$.

Fix some $s \in \{0, 1\}^k$. Let $Z_1, \dots, Z_N \in \{0, 1\}$ be indicator random variable such that $Z_i = 1$ iff $\text{DecMC}(i) = s$ (when i is interpreted as an n -bit string). Recall that (Z_1, \dots, Z_N) is a $(9t)$ -wise independent random vector with respect to the randomness of the code construction. Let $Z := Z_1 + \dots + Z_N$, and note that $Z = |E(s)|$. From (22), we see that

$$\mathbb{E}[Z] = \mathbb{E}[|E(s)|] = 2t.$$

Using Theorem B.6 with $\ell := t/4$ and $A := \mathbb{E}[Z]/2 = t$, we see that

$$\Pr[|Z - 2t| \geq t] \leq 8(3/4)^{t/4}.$$

By taking a union bound over all choices of $s \in \{0, 1\}^k$, we conclude that with probability at least $1 - \eta_0$, where we define $\eta_0 := 8N(3/4)^{t/4}$, the realization of (EncMC, DecMC) is so that

$$(\forall s \in \{0, 1\}^k): |E(s)| \in [t, 3t].$$

This bound suffices to show the desired conclusion. \square

By combining the above tools with Theorem 3.1, we can derive the following result on the performance of Construction 2.

Theorem 4.4. *Let $\mathcal{F}: \{0, 1\}^n \rightarrow \{0, 1\}^n$ be any family of tampering functions. For any $\epsilon, \eta > 0$, with probability at least $1 - \eta$, the pair (EncMC, DecMC) in Construction 2 can be set up so achieve a non-malleable coding scheme with respect to \mathcal{F} and with error ϵ . Moreover, the scheme satisfies the following.*

1. *The code achieves $k \geq n - \log \log(|\mathcal{F}|/\eta) - \log n - 9 \log(1/\epsilon) - O(1)$.*
2. *The number of random bits needed to specify the code is $O\left((n + \log(|\mathcal{F}|/\eta))n/\epsilon^6\right)$.*
3. *The encoder and the decoder run in worst case time $\text{poly}(\log(|\mathcal{F}|/\eta)n/\epsilon)$.*

Proof. Let t_0 and k_0 be the parameters promised by Theorem 3.1. We instantiate Construction 2 with parameter $t := t_0$ and $k := k_0$. Observe that this choice of t is large enough to allow Lemma 4.3 to hold. Thus we can ensure that, with probability at least $1 - \eta$, (EncMC, DecMC) is a $(9t)$ -wise independent coding scheme where, for every $s \in \{0, 1\}^k$, $|E(s)| \in [t_0, 3t_0]$. Thus we can now apply Lemma 4.2 to conclude that with probability at least $1 - 2\eta$, (EncMC, DecMC) is a strong non-malleable code with the desired parameters.

The number of random bits required to represent the code is the bit length of the polynomial $P(X)$ in Construction 2, which is $9tn$. Plugging in the value of t from (21) gives the desired estimate.

The running time of the decoder is dominated by evaluation of the polynomial $P(X)$ at a given point. Since the underlying field is of characteristic two, a representation of the field as well as basic field operations can be computed in deterministic polynomial time in the degree n of the extension using Shoup's algorithm [10].

The encoder is, however, slightly more complicated as it needs to iterate through $O(t)$ steps, and at each iteration compute all roots of a given degree $9t - 1$ polynomial. Again, since characteristic of the underlying field is small, this task can be performed in deterministic polynomial time in the degree $9t - 1$ of the polynomial and the degree n of the extension (e.g., using [11]). After plugging in the bound on t from (21), we obtain the desired bound on the running time. \square

As a corollary, we observe that the rate of the Monte Carlo construction can be made arbitrarily close to 1 while keeping the bit-representation of the code as well as the running time of the encoder and decoder at $\text{poly}(n)$ provided that $\epsilon = 1/\text{poly}(n)$ and $|\mathcal{F}| = 2^{\text{poly}(n)}$. In particular, we see that the Monte Carlo construction achieves strong non-malleability even with respect to such powerful classes of adversaries as polynomial-sized Boolean circuits (with n outputs bits) and virtually any interesting computationally bounded model.

Remark 4.5. Since in this construction the error ϵ is only polynomially small, for cryptographic applications such as tamper-resilient security it is important to set up the code so as to ensure that $1/\epsilon$ is significantly larger than the total number of tampering attempts made by the adversary.

Caveat. We point out that any explicit coding scheme for computationally bounded models (such as polynomial-sized Boolean circuits) necessarily implies an explicit lower bound for the respective computational model. This is because a function in the restricted model cannot be powerful enough to compute the decoder function, as otherwise, the following adversary would violate non-malleability:

Consider fixed tuples $(s_1, x_1), (s_2, x_2) \in \{0, 1\}^k \times \{0, 1\}^n$, where $s_1 \neq s_2$, $\text{Dec}(x_1) = s_1$ and $\text{Dec}(x_2) = s_2$. Given a codeword $x \in \{0, 1\}^n$, compute $s := \text{Dec}(x)$. If $s = s_1$, output x_2 . If $s = s_2$, output x_1 . Otherwise, output x .

Remark 4.6. (Alternative Monte Carlo construction) In addition to Construction 2, it is possible to consider a related Monte Carlo construction when polynomial evaluation is performed at the encoder and root finding is done by the encoder. More precisely, the encoder, given $s \in \{0, 1\}^k$, may sample $i \in [t]$ uniformly at random, and output $P(s, i)$ where (s, i) is interpreted as an element of \mathbb{F}_{2^n} (possibly after padding). The drawback with this approach is that the rate of the code would be limited by $1/2$, since for larger rates there is a noticeable chance that the encoder maps different messages to the same codeword.

5 Impossibility bounds

In this section, we show that the bounds obtained by Theorem 3.1 are essentially optimal. In order to do so, we consider three families of adversaries. Throughout the section, we use k and n for the message length and block length of coding schemes and define $N := 2^n$ and $K := 2^k$.

5.1 General adversaries

The first hope is to demonstrate that Theorem 3.1 is the best possible for *every* family of the tampering functions of a prescribed size. We rule out this possibility and demonstrate a family \mathcal{F} of tampering functions achieving $\log \log |\mathcal{F}| \approx n$ for which there is a non-malleable code achieving rate $1 - \gamma$ for arbitrarily small $\gamma > 0$.

Let $S \subseteq \{0, 1\}^n$ be any set of size at least $N^{1-\alpha}$ and at most $N/2$. Consider the family \mathcal{F} of functions satisfying the property that

$$(\forall f \in \mathcal{F})(\forall x \in S): f(x) = x.$$

We can take the union of such families over all choices of S ; however, for our purposes it suffices to define \mathcal{F} with respect to a single choice of S . Observe that

$$|\mathcal{F}|N^{N-|S|} \geq N^{N/2},$$

which implies

$$\log \log |\mathcal{F}| \geq n - 1.$$

However, there is a trivial coding scheme that is non-malleable with zero error for all functions in \mathcal{F} . Namely, the encoder Enc is a deterministic function that maps messages to distinct elements of S , whereas

the decoder Dec inverts the encoder and furthermore, maps any string outside S to \perp . In this construction, we see that

$$(\forall f \in \mathcal{F})(\forall x \in \{0, 1\}^k): \text{Dec}(f(\text{Enc}(x))) = x,$$

since f necessarily fixes all the points in S (in particular, in Definition 2.3 one can take $\mathcal{D}_f := \mathcal{D}(\text{same})$). Finally, observe that the rate of this coding scheme is at least $1 - \gamma$. In fact, this result holds for any $\gamma \geq 1/n$, implying that the rate of the code can be made $1 - o(1)$.

5.2 Random adversaries

The observation in Section 5.1 rules out the hope for a general lower bound that only depends on the size of the adversarial family. However, in this section we show that for “virtually all” families of tampering functions of a certain size, Theorem 3.1 gives the best possible bound. More precisely, we construct a family \mathcal{F} of a designed size M as follows: For each $i \in [M]$, sample a uniformly random function $f_i: \{0, 1\}^n \rightarrow \{0, 1\}^n$ and add f_i to the family. Since some of the f_i may turn out to be the same (albeit with negligible probability), $|\mathcal{F}|$ may in general be lower than M (which can only make a lower bound stronger).

We prove the following.

Theorem 5.1. *For any $\alpha > 0$, there is an M_0 satisfying*

$$\log \log M_0 \leq \alpha n + O(\log n)$$

such that with probability $1 - \exp(-n)$, a random family \mathcal{F} with designed size $M \geq M_0$ satisfies the following: There is no coding scheme achieving rate at least $1 - \alpha$ and error $\epsilon < 1$ that is non-malleable with respect to the tampering family \mathcal{F} .

Proof. We begin with the following simple probabilistic argument:

Claim 5.2. *Let $\mathcal{C} \subseteq [q]^N$ be a multi-set of vectors each chosen uniformly and independently at random. For any integer $\ell \in [N]$ and parameter $\gamma > 0$, there is an $M_0 = O(\ell q^\ell \log(qN/\gamma))$ such that as long as $|\mathcal{C}| \geq M_0$, the following holds with probability at least $1 - \gamma$: For every $S \subseteq [N]$ with $|S| \leq \ell$, the set of vectors in \mathcal{C} restricted to the positions picked by S is equal to $[q]^{|S|}$.*

Proof. Fix any choice of the set S (where, without loss of generality, $|S| = \ell$) and let \mathcal{C}_S be the set of vectors in \mathcal{C} restricted to the positions in S . For any $w \in [q]^{|S|}$, we have

$$\Pr[w \notin \mathcal{C}_S] = \left(1 - \frac{1}{q^\ell}\right)^{|\mathcal{C}|} \leq \exp(-\Omega(|\mathcal{C}|/q^\ell)).$$

By taking a union bound on all the choices of w and S , the probability that \mathcal{C} does not satisfy the desired property can be seen to be at most

$$(qN)^\ell \exp(-\Omega(|\mathcal{C}|/q^\ell)),$$

which can be made no more than γ for some

$$|\mathcal{C}| = O\left(q^\ell(\ell \log(qN) + \log(1/\gamma))\right). \quad \square$$

Let $\gamma > 0$ be a parameter to be determined later. By Claim 5.2, with probability at least $1 - \gamma$ over the randomness of the family \mathcal{F} , we can ensure that for all sets $S \subseteq \{0, 1\}^n$ of size at most $2N^\alpha$, and for all functions $f_S: S \rightarrow \{0, 1\}^n$, there is a function $f \in \mathcal{F}$ that agrees with f_S on all points in S . This guarantee holds if we take $\mathcal{F} \geq M_0$ for some

$$M_0 = O\left(N^{(4N^\alpha)}(8N^\alpha \log(N/\gamma))\right).$$

Overestimating the above bound yields

$$\log \log M_0 \leq \alpha n + \log \log(N/\gamma) + O(1)$$

which is at most $\alpha n + O(\log n)$ for $\gamma = \exp(-n)$. Assuming that the family \mathcal{F} attains the above-mentioned property, we now proceed as follows.

Consider any coding scheme (Enc, Dec) with block length n and message length k which is non-malleable for the family \mathcal{F} randomly constructed as above and achieving rate at least $1 - \alpha$ for some $\alpha > 0$ and any non-trivial error $\epsilon < 1$. For any message $s \in \{0, 1\}^k$, let

$$E(s) := \text{supp}(\text{Enc}(s)) \subseteq [N]$$

and observe that $E(s) \cap E(s') = \emptyset$ for all $s \neq s'$. Observe that

$$\mathbb{E}[|E(U_k)|] \leq N^\alpha$$

by the disjointness property of the $E(s)$ and the assumption on the rate of the code. By Markov's bound,

$$\Pr[|E(U_k)| \geq 2N^\alpha] < 1/2$$

implying that for at least half of the choices of $s \in \{0, 1\}^k$, we can assume $|E(s)| \leq 4N^\alpha$. Take two distinct vectors $s_1, s_2 \in \{0, 1\}^k$ satisfying this bound.

Now, let $S := E(s_1) \cup E(s_2)$, where $|S| \leq 2N^\alpha$ as above. Consider any $c_1 \in E(s_1)$ and $c_2 \in E(s_2)$ and define $f_S: S \rightarrow \{0, 1\}^n$ such that

$$(\forall x \in E(s_1)): f_S(x) = c_2 \quad \text{and} \quad (\forall x \in E(s_2)): f_S(x) = c_1.$$

By the choice of \mathcal{F} , we know that there is $f \in \mathcal{F}$ that agrees with f_S on all the points in S . This choice of the adversary ensures that

$$\Pr[\text{Dec}(f(\text{Enc}(s_1))) = s_2] = 1 \quad \text{and} \quad \Pr[\text{Dec}(f(\text{Enc}(s_2))) = s_1] = 1$$

with respect to the randomness of the encoder. Since the two distributions $\text{Dec}(f(\text{Enc}(s_1)))$ and $\text{Dec}(f(\text{Enc}(s_2)))$ are maximally far from each other and moreover, the adversary f always tampers codewords in $E(s_1)$ and $E(s_2)$ to a codeword corresponding to a different message, we conclude that there is no choice of \mathcal{D}_f in Definition 2.3 that ensures non-malleability with any error less than 1. \square

5.3 General adversaries acting on a subset of positions

An important family of adversaries is the one that is only restricted by the subset of bits it acts upon. More precisely, let $T \subseteq [n]$ be a fixed set of size αn , for a parameter $\alpha \in (0, 1)$. For $x \in \{0, 1\}^n$, we use the

notation $x_T \in \{0, 1\}^{|T|}$ for the restriction of x to the positions in T . Without loss of generality, assume that T contains the first $|T|$ coordinate positions so that $x = (x_T, x_{\bar{T}})$, where $\bar{T} := [n] \setminus T$. We consider the family \mathcal{F}_T of all functions $f: \{0, 1\}^n \rightarrow \{0, 1\}^n$ such that

$$f(x) = (g(x_T), x_{\bar{T}})$$

for some $g: \{0, 1\}^{|T|} \rightarrow \{0, 1\}^{|T|}$. Observe that $|\mathcal{F}_T| \leq N^{(\alpha N^\alpha)}$ which implies $\log \log |\mathcal{F}_T| \leq \alpha n$.

We prove the following lower bound, which is a variation of the classical Singleton bound for non-malleable codes. What makes this variation much more challenging to prove is the fact that 1) non-malleable codes allow a randomized encoder, and 2) non-malleability is a more relaxed requirement than error detection, and hence the proof must rule out the case where the decoder does not detect errors (i.e., outputs a wrong message) while still satisfies non-malleability.

Theorem 5.3. *Let $T \subseteq [n]$ be of size αn and consider the family \mathcal{F}_T of the tampering functions that only act on the coordinate positions in T (as defined above). Then, there is a $\delta_0 = O((\log n)/n)$ such that the following holds. Let (Enc, Dec) be any coding scheme which is non-malleable for the family \mathcal{F}_T and achieves rate $1 - \alpha + \delta$, for any $\delta \in [\delta_0, \alpha]$ and error ϵ . Then, $\epsilon \geq \delta/(16\alpha)$. In particular, when α and δ are absolute constants, $\epsilon = \Omega(1)$.*

Before proving the theorem, we state the following immediate corollary.

Corollary 5.4. *Let \mathcal{F} be the family of split-state adversaries acting on n bits. That is, each $f \in \mathcal{F}$ interprets the input as a pair (x_1, x_2) where $x_1 \in \{0, 1\}^{\lfloor n/2 \rfloor}$ and $x_2 \in \{0, 1\}^{\lceil n/2 \rceil}$, and outputs $(f_1(x_1), f_2(x_2))$ for arbitrary tampering functions f_1 and f_2 (acting on their respective input lengths).*

Moreover, for a fixed constant $\delta \in (0, 1)$, let \mathcal{F}_δ be the class of tampering functions where $f \in \mathcal{F}_\delta$ iff every bit of $f(x)$ depends on at most $\lfloor \delta n \rfloor$ of the bits of x .

Let $(\text{Enc}_1, \text{Dec}_1)$ (resp., $(\text{Enc}_\delta, \text{Dec}_\delta)$) be any coding scheme which is non-malleable for the class \mathcal{F} (resp., \mathcal{F}_δ) achieving error at most ϵ and rate R (resp., R_δ). Then, for every fixed constant $\gamma > 0$, there is a fixed constant $\epsilon_0 > 0$ such that if $\epsilon \leq \epsilon_0$, the following bounds hold.

- (i) $R \leq 1/2 - \gamma$,
- (ii) $R_\delta \leq 1 - \delta - \gamma$.

The proof of Theorem 5.3 uses basic tools from information theory, and the core ideas can be described as follows. Assume that the codeword is (X_1, X_2) where the adversary acts on X_1 , which is of length αn . We show that for any coding scheme with rate slightly larger than $(1 - \alpha)n$, there is a set $X_\eta \subseteq \{0, 1\}^{\alpha n}$ such that

1. For some message s_0 , X_1 lies in X_η with noticeable probability.
2. For a ‘‘typical’’ message s_1 , X_1 is unlikely to land in X_η .
3. There is a vector $w \in \{0, 1\}^{\alpha n}$ that cannot be extended to a codeword (w, w') that maps to either s_0 or s_1 by the decoder.

We then use the above properties to design the following strategy that violates non-malleability of the code: Given (X_1, X_2) , if $X_1 \in X_\eta$, the adversary tampers the codeword to (w, X_2) , which decodes to a message outside $\{s_0, s_1\}$. This ensures that $\text{Dec}(f(\text{Enc}(s_0)))$ has a noticeable chance of being tampered to an incorrect message. Otherwise, the adversary leaves the codeword unchanged, ensuring that $\text{Dec}(f(\text{Enc}(s_1)))$ has little chance of being tampered at all. Thus there is no choice for a distribution \mathcal{D}_f that sufficiently matches both $\text{Dec}(f(\text{Enc}(s_0)))$ and $\text{Dec}(f(\text{Enc}(s_1)))$.

Proof of Theorem 5.3

Throughout the proof, we use standard information theoretic tools, such as the notation $H(X)$ for the Shannon entropy of a discrete random variable X and $I(X; Y)$ for the mutual information between discrete random variables X and Y . We will need the following standard information-theoretic fact.

Claim 5.5. *Suppose $H(X) \leq r$ and let $p(x) := \Pr[X = x]$. For any $\eta > 0$, and define*

$$X_\eta := \{x \in \text{supp}(X) : p(x) > \frac{1}{2^{r/(1-\eta)}}\}.$$

Then, $\Pr[X \in X_\eta] \geq \eta$ and $|X| < 2^{r/(1-\eta)}$.

Proof. The upper bound on $|X_\eta|$ is immediate from the definition of X_η . Let $\bar{X}_\eta := \text{supp}(X) \setminus X_\eta$. We need to show that $\Pr[X \in \bar{X}_\eta] \leq 1 - \eta$. If this is not the case, we can write

$$\begin{aligned} H(X) &\geq \sum_{x \in \bar{X}_\eta} p(x) \log(1/p(x)) \\ &\geq \sum_{x \in \bar{X}_\eta} \frac{rp(x)}{1-\eta} \\ &= \Pr[x \in \bar{X}_\eta] r / (1-\eta) > r, \end{aligned}$$

a contradiction. □

Suppose there is a coding scheme (Enc, Dec) that is non-malleable for the family \mathcal{F}_T and achieving rate at least $1 - \alpha + \delta$, for an arbitrarily small parameter $\delta \in (0, \alpha]$. Let $S \sim \mathcal{U}_k$, $X := \text{Enc}(S)$ and suppose $X = (X_1, X_2)$ where $X_1 := X_T$ and $X_2 := X_{\bar{T}}$.

For any $s \in \{0, 1\}^k$, define $E(s) := \text{supp}(\text{Enc}(s))$. Observe that

$$\mathbb{E}_S |E(S)| \leq N/N^{1-\alpha+\delta} = N^{\alpha-\delta}$$

By Markov's bound, for any $\gamma \in (0, 1]$,

$$\Pr[|E(S)| > N^{\alpha-\delta}/\gamma] < \gamma. \tag{23}$$

By the assumption on rate, $H(S) \geq n(1 - \alpha + \delta)$. Also, $H(X_2|S) \leq H(X_2) \leq n - |T| = n(1 - \alpha)$. Thus,

$$I(X_2; S) = H(S) - H(S|X_2)$$

Using the chain rule for mutual information,

$$\begin{aligned} I(X_1; S) &= I(X_1, X_2; S) - I(X_2; S|X_1) \\ &= (H(S) - H(S|X_1, X_2)) - (H(X_2|X_1) - H(X_2|S, X_1)) \\ &\geq H(S) - H(X_2|X_1) & (24) \\ &\geq H(S) - H(X_2) & (25) \\ &\geq (1 - \alpha + \delta)n - (1 - \alpha)n = \delta n, & (26) \end{aligned}$$

where (24) holds because $S = \text{Dec}(X_1, X_2)$ and thus $H(S|X_1, X_2) = 0$, in addition to non-negativity of entropy; (25) uses the fact that conditioning does not increase entropy; and (26) holds because of the assumption on the rate of the code and the length of X_2 . From this, we can deduce that

$$H(X_1|S) = H(X_1) - I(X_1; S) \leq H(X_1) - \delta n.$$

Note that the latter inequality in particular implies that $H(X_1) \geq \delta n$, and that $\text{supp}(X_1) \geq 2^{\delta n}$. By Markov's bound,

$$|\{s \in \{0, 1\}^k : H(X_1|S = s) > (H(X_1) - \delta n)(1 + 4\gamma)\}| < \frac{2^k}{1 + 4\gamma} \leq (1 - 2\gamma)2^k. \quad (27)$$

By combining (23) and (27) using a union bound, there is a choice of $s_0 \in \{0, 1\}^k$ such that

$$|E(s_0)| \leq N^{\alpha - \delta}/\gamma, \text{ and, } H(X_1|S = s_0) \leq (H(X_1) - \delta n)(1 + 4\gamma).$$

We can take $\gamma := \delta/(8\alpha)$ so that the above becomes

$$|E(s_0)| \leq 8\alpha N^{\alpha - \delta}/\delta, \text{ and, } H(X_1|S = s_0) \leq H(X_1) - \delta n/2. \quad (28)$$

For a parameter $\eta > 0$, to be determined later, we can now apply Claim 5.5 to the conditional distribution of X_1 subject to $S = s_0$ and construct a set $X_\eta \subseteq \{0, 1\}^{\alpha n}$ such that

$$\begin{aligned} \Pr[X_1 \in X_\eta | S = s_0] &\geq \eta, \\ |X_\eta| &\leq 2^{(H(X_1) - \delta n/2)/(1 - \eta)}. \end{aligned} \quad (29)$$

Let $\eta' := \Pr[X_1 \in X_\eta]$, and let $h(\cdot)$ denote the binary entropy function. Using a simple information-theoretic rule that follows from the definition of Shannon entropy, we can write

$$\begin{aligned} H(X_1) &= h(\eta') + \eta' H(X_1|X_1 \in X_\eta) + (1 - \eta') H(X_1|X_1 \notin X_\eta) \\ &\leq h(\eta') + \eta' \cdot \frac{H(X_1) - (\delta/2)n}{1 - \eta} + (1 - \eta') H(X_1|X_1 \notin X_\eta) \end{aligned} \quad (30)$$

$$\leq h(\eta') + \eta' \cdot \frac{H(X_1) - (\delta/2)n}{1 - \eta} + (1 - \eta') H(X_1), \quad (31)$$

where (30) is due to the upper bound on the support size of X_η and (31) holds since conditioning does not increase entropy. After simple manipulations, (31) simplifies to

$$\eta' \leq \frac{2h(\eta')(1 - \eta)}{\delta n - 2\eta H(X_1)} \leq \frac{2h(\eta')}{n(\delta - 2\eta\alpha)}. \quad (32)$$

Now, we take $\eta := \delta/(4\alpha)$, so that the above inequalities, combined with the estimate $h(\eta') = O(\eta' \log(1/\eta'))$ yields

$$h(\eta')/\eta' \geq \delta n/4 \Rightarrow \log(1/\eta') = \Omega(\delta n) \Rightarrow \eta' \leq \exp(-\Omega(\delta n)).$$

From the above inequality, straightforward calculations ensure that

$$\eta' \leq \eta/4 = \delta/(16\alpha), \quad (33)$$

as long as $\delta \geq \delta_0 = O((\log n)/n)$.

From (33), recalling that $\eta' = \Pr[X_1 \in X_\eta]$ and using Markov's bound,

$$|\{s: \Pr[X_1 \in X_\eta | S = s] > \eta/2\}|/2^k < 1/2.$$

Combined with (23) and a union bound, there is a fixed $s_1 \in \{0, 1\}^k$ such that

$$|E(s_1)| \leq 8\alpha N^{\alpha-\delta}/\delta, \text{ and, } \Pr[X_1 \in X_\eta | S = s_1] \leq \eta/2. \quad (34)$$

Assuming the chosen lower bound for δ , we can also ensure that, using (28), that $|E(s_0) \cup E(s_1)| < N^\alpha$. Thus, there is a fixed string $w \in \{0, 1\}^{\alpha n}$ that cannot be extended to any codeword in $E(s_0)$ or in $E(s_1)$; i.e.,

$$\Pr[X_1 = w | (S = s_0) \vee (S = s_1)] = 0,$$

which in turn implies

$$(\forall x_2 \in \{0, 1\}^{n(1-\alpha)}): \text{Dec}(w, x_2) \notin \{s_0, s_1\}. \quad (35)$$

Now, we consider the following tampering strategy $f: \{0, 1\}^{|T|} \times \{0, 1\}^{n-|T|} \rightarrow \{0, 1\}^{|T|} \times \{0, 1\}^{n-|T|}$ acting on the coordinate positions in T :

- Given $(x_1, x_2) \in \{0, 1\}^{|T|} \times \{0, 1\}^{n-|T|}$, if $x_1 \in X_\eta$, output (w, x_2) .
- Otherwise, output (x_1, x_2) .

Suppose the coding scheme (Enc, Dec) satisfied Definition 2.3 for a particular distribution \mathcal{D}_f over $\{0, 1\}^n \cup \{\underline{\text{same}}, \perp\}$ for the tampering function f .

Since f does not alter any string with the first component outside X_η , (34) implies that

$$\Pr[f(X_1, X_2) = (X_1, X_2) | S = s_1] \geq 1 - \eta/2. \quad (36)$$

On the other hand, by (29) and (35),

$$\Pr[\text{Dec}(f(X_1, X_2)) \notin \{s_0, s_1\} | S = s_0] \geq \eta. \quad (37)$$

By (37) and Definition 2.3, \mathcal{D}_f must be ϵ -close to a distribution D_0 that assigns at most $1 - \eta$ of the probability mass to $\{\underline{\text{same}}, s_0, s_1\}$. On the other hand, by (36), \mathcal{D}_f must be ϵ -close to a distribution D_1 that assigns at least $1 - \eta/2$ of the probability mass to $\{\underline{\text{same}}, s_1\}$. Thus, the statistical distance between D_0 and D_1 is at least $\eta/2$ (from the distinguisher corresponding to the event $\{\underline{\text{same}}, s_1\}$). By triangle inequality, however, D_0 and D_1 are (2ϵ) -close. Therefore, $\epsilon \geq \eta/4$ and the result follows.

References

- [1] D. Aggarwal, Y. Dodis, and S. Lovett. Non-malleable codes from additive combinatorics. *ECCC Technical Report TR13-081*, 2013. 5
- [2] M. Bellare and J. Rompel. Randomness efficient oblivious sampling. In *Proceedings of the 35th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 1994. 35

- [3] M. Cheraghchi and V. Guruswami. Non-malleable coding against bit-wise and split-state tampering. Manuscript in preparation, 2013. 5
- [4] J.-S. Coron, J. Patarin, and Y. Seurin. The random oracle model and the ideal cipher model are equivalent. In *Proceedings of CRYPTO 2008*, 2008. 21
- [5] R. Cramer, Y. Dodis, S. Fehr, C. Padró, and D. Wichs. Detection of algebraic manipulation with applications to robust secret sharing and fuzzy extractors. In *Proceedings of EUROCRYPT 2008*, pages 471–488, 2008. 3, 9
- [6] D. Dolev, C. Dwork, and M. Naor. Nonmalleable cryptography. *SIAM Journal on Computing*, 30(2):391–437, 2000. 3
- [7] S. Dziembowski, T. Kazana, and M. Obremski. Non-malleable codes from two-source extractors. In *Proceedings of CRYPTO*, 2013. 5
- [8] S. Dziembowski, K. Pietrzak, and D. Wichs. Non-malleable codes. In *Proceedings of Innovations in Computer Science (ICS 2010)*, 2010. 3, 4, 6, 7, 8, 9, 13, 20, 21, 32
- [9] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995. 35
- [10] V. Shoup. New algorithms for finding irreducible polynomials over finite fields. *Mathematics of Computation*, 54:435–447, 1990. 24
- [11] J. von zur Gathen and V. Shoup. Computing frobenius maps and factoring polynomials. *Computational Complexity* 2:187–224, 2:187–224, 1992. 24

A Rate $1/2$ barrier for the uniform coding scheme.

Dziembowski et al. [8] consider the uniformly random coding scheme (Enc, Dec) in which the decoder Dec maps any given input $x \in \{0, 1\}^n$ to a uniform and independent random string in $\{0, 1\}^k$. Moreover, the encoder, given $s \in \{0, 1\}^k$, outputs a uniformly random element of $\text{Dec}^{-1}(s)$. In this section, we argue that the uniform coding scheme cannot achieve a rate better than $1/2$ even with respect to very simple tampering functions.

Suppose that the scheme is indeed non-malleable with error upper bounded by a small constant (say $1/4$), and consider any bijective tampering function $f: \{0, 1\}^n \rightarrow \{0, 1\}^n$. For example, one may think of f as the function that flips the first bit of the input. For simplicity, we assume that the coding scheme achieves strong non-malleability (as proved by Dziembowski et al. [8]). Since the chosen tampering function does not have any fixed points (i.e., $f(x) \neq x$), Definition 2.5 implies that there is a choice of \mathcal{D}_f that has no support on $\{\text{same}\}$, and we can restrict to such a distribution. However, it can be shown that the argument extends to the weaker definition of non-malleability as well.

Let $X := \text{Enc}(U_k)$ and observe that $\mathcal{D}(X) = U_n$, which in turn implies that $\mathcal{D}(f(X)) = U_n$. Consider $S := \text{Dec}(f(X))$. Note that $\mathcal{D}(S)$ is a random variable depending on the randomness of the code construction (namely, it is the empirical distribution of the truth table of the decoder). With respect to this randomness, we have

$$\mathbb{E}[\mathcal{D}(S)] = U_k.$$

Moreover, with overwhelming probability, the realization of the code is so that

$$\mathcal{D}(S) \approx_{o(1)} \mathcal{U}_k.$$

Suppose this is the case and fix the randomness of the code construction accordingly.

Since for every $s \in \{0, 1\}^k$, we know that $\mathcal{D}(\text{Dec}(f(\text{Enc}(s))))$ is close (in the sense described by Definition 2.3) to \mathcal{D}_f , it follows that the convex combination

$$\sum_{s \in \{0, 1\}^k} \Pr[\text{Dec}(U_n) = s] \cdot \mathcal{D}(\text{Dec}(f(\text{Enc}(s))))$$

is equally close to \mathcal{D}_f . But, since $f(\text{Enc}(\text{Dec}(\mathcal{U}_n))) = f(\mathcal{U}_n) = \mathcal{U}_n$, the above convex combination is exactly $\mathcal{D}(\text{Dec}(\mathcal{U}_n)) = \mathcal{D}(S)$, which we know is close to \mathcal{U}_k .

Thus it follows that for every $s \in \{0, 1\}^k$,

$$(s, \text{Dec}(f(\text{Enc}(s)))) \approx_{o(1)} (s, \mathcal{U}_k),$$

and, for $U \sim \mathcal{U}_k$,

$$(U, \text{Dec}(f(\text{Enc}(U)))) \approx_{o(1)} \mathcal{U}_{2k}. \quad (38)$$

Since $(U, \text{Dec}(f(\text{Enc}(U))))$ is a function of $\text{Enc}(U)$, we get

$$H(U, \text{Dec}(f(\text{Enc}(U)))) \leq n.$$

On the other hand, 38 implies that the above entropy is close to $2k$. Thus, $k \leq (n/2)(1 + o(1))$.

B Useful tools

In many occasions in the paper, we deal with a chain of correlated random variables $0 = X_0, X_1, \dots, X_n$ where we wish to understand an event depending on X_i conditioned on the knowledge of the previous variables. That is, we wish to understand

$$\mathbb{E}[f(X_i) | X_0, \dots, X_{i-1}].$$

The following proposition shows that in order to understand the above quantity, it suffices to have an estimate with respect to a more restricted event than the knowledge of X_0, \dots, X_{i-1} . Formally, we can state the following, where X stands for X_i in the above example and Y stands for (X_0, \dots, X_{i-1}) .

Proposition B.1. *Let X and Y be possibly correlated random variables and let Z be a random variable such that the knowledge of Z determines Y ; that is, $Y = f(Z)$ for some function f . Suppose that for every possible outcome of the random variable Z , namely, for every $z \in \text{supp}(Z)$, and for some real-valued function g , we have*

$$\mathbb{E}[g(X) | Z = z] \in I. \quad (39)$$

for a particular interval I . Then, for every $y \in \text{supp}(Y)$,

$$\mathbb{E}[g(X) | Y = y] \in I.$$

Similarly, suppose for some distribution \mathcal{D} , and all $z \in \text{supp}(Z)$,

$$\mathcal{D}(X | Z = z) \approx_\epsilon \mathcal{D}.$$

Then, for all $y \in \text{supp}(Y)$,

$$\mathcal{D}(X | Y = y) \approx_\epsilon \mathcal{D}.$$

Proof. Let $T = \{z \in \text{supp}(Z) : f(z) = y\}$, and let $p(z) := \Pr[Z = z | Y = y]$. Then,

$$\mathbb{E}[g(X)|Y = y] = \sum_{z \in T} p(z) \mathbb{E}[g(X)|Z = z].$$

Since by (39), each $\mathbb{E}[g(X)|Z = z]$ lies in I and $\sum_{z \in T} p(z) = 1$, we deduce that

$$\mathbb{E}[g(X)|Y = y] \in I.$$

Proof of the second part is similar, by observing that if a collection of distributions is statistically close to a particular distribution \mathcal{D} , any convex combination of them is equally close to \mathcal{D} as well. \square

Proposition B.2. *Let the random variable $X \in \{0, 1\}^n$ be uniform on a set of size at least $(1 - \epsilon)2^n$. Then, $\mathcal{D}(X)$ is $(\epsilon/(1 - \epsilon))$ -close to \mathcal{U}_n .*

We will use the following tail bounds on summation of possibly dependent random variables, which are direct consequences of Azuma's inequality.

Proposition B.3. *Let $0 = X_0, X_1, \dots, X_n$ be possibly correlated random variables in $[0, 1]$ such that for every $i \in [n]$ and for some $\gamma \geq 0$,*

$$\mathbb{E}[X_i | X_0, \dots, X_{i-1}] \leq \gamma.$$

Then, for every $c \geq 1$,

$$\Pr\left[\sum_{i=1}^n X_i \geq cn\gamma\right] \leq \exp(-n\gamma^2(c-1)^2/2),$$

or equivalently, for every $\delta > \gamma$,

$$\Pr\left[\sum_{i=1}^n X_i \geq n\delta\right] \leq \exp(-n(\delta - \gamma)^2/2).$$

Proof. The proof is a standard Martingale argument. For $i \in [n]$, define

$$X'_i := X_i - \gamma,$$

and

$$S_i := \sum_{j=1}^i X'_j = \sum_{j=1}^i X_j - i\gamma.$$

By assumption, S_i is a super-martingale, that is, assuming $S_0 := 0$,

$$\mathbb{E}[S_{i+1} | S_0, \dots, S_i] \leq S_i.$$

Thus, by Azuma's inequality, for all $t \geq 0$,

$$\Pr[S_n \geq t] \leq \exp(-t^2/(2n)).$$

Substituting $t := (c - 1)n\gamma$ proves the claim. \square

In a similar fashion (using Azuma's inequality for sub-martingales rather than super-martingales in the proof), we may obtain a tail bound when we have a lower bound on conditional expectations.

Proposition B.4. Let $0 = X_0, X_1, \dots, X_n$ be possibly correlated random variables in $[0, 1]$ such that for every $i \in [n]$ and for some $\gamma \geq 0$,

$$\mathbb{E}[X_i | X_0, \dots, X_{i-1}] \geq \gamma.$$

Then, for every $\delta < \gamma$,

$$\Pr\left[\sum_{i=1}^n X_i \leq n\delta\right] \leq \exp(-n(\delta - \gamma)^2/2).$$

The following tail bound is similar in flavor to the one given by Proposition B.3, but only applies to indicator random variables. However, it can be better when the individual expectations are low and the target deviation from mean is very large.

Proposition B.5. Let $0 = X_0, X_1, \dots, X_n \in \{0, 1\}$ be indicator, possibly dependent, random variables such that for every $i \in [n]$,

$$\mathbb{E}[X_i | X_1, \dots, X_{i-1}] \leq p,$$

for some $p \in [0, 1]$. Let $X := X_1 + \dots + X_n$. Then, for every $c \geq 1$,

$$\Pr[X > cnp] \leq (e/c)^{cnp}.$$

Proof. We closely follow the standard proof of Chernoff bounds for independent indicator random variables (see, e.g., [9]). Using Markov's bound on the exponential moment of X , we can write, for a parameter $t > 0$ to be determined later,

$$\Pr[X > cnp] \leq \frac{\mathbb{E}[\exp(tX)]}{\exp(tcnp)} = \frac{\mathbb{E}[\exp(tX_1) \cdots \exp(tX_n)]}{\exp(tcnp)}. \quad (40)$$

However, we can write down the expectation of product as the following chain of conditional expectations

$$\begin{aligned} \mathbb{E}_{(X_1, \dots, X_n)}[\exp(tX)] &= \mathbb{E}_{X_1} \left[e^{tX_1} \mathbb{E}_{(X_2|X_1)} [e^{tX_2} \cdots \mathbb{E}_{(X_n|X_1, \dots, X_{n-1})} e^{tX_n} \cdots] \right] \\ &\leq (p \exp(t) + 1)^n. \end{aligned}$$

where the inequality uses the fact that the X_i are Bernoulli random variables and thus

$$\mathbb{E}[\exp(tX_i) | X_1, \dots, X_{i-1}] \leq p \exp(t) + (1 - p) \exp(0) \leq p \exp(t) + 1.$$

Using the inequality $(1 + x)^n \leq \exp(nx)$ the above simplifies to

$$\mathbb{E}[\exp(tX)] \leq \exp(np \exp(t)),$$

and thus, plugging the above result into (40),

$$\Pr[X > cnp] \leq \frac{\exp(np \exp(t))}{\exp(tcnp)}.$$

Choosing $t := \ln c$ yields the desired conclusion. □

For summation of ℓ -wise independent random variables, we use the following tail bound from [2]:

Theorem B.6. Let $\ell > 1$ be an even integer, and let $X_1, \dots, X_n \in [0, 1]$ be ℓ -wise independent variables. Define $X := X_1 + \dots + X_n$ and $\mu := \mathbb{E}[X]$. Then,

$$\Pr[|X - \mu| \geq A] \leq 8 \left(\frac{\ell(\mu + \ell)}{A^2} \right)^{\ell/2}.$$

Approximating distributions by fuzzy correlated sampling

In this section, we show that it is possible to sharply approximate a distribution \mathcal{D} with finite support by sampling possibly correlated random variables X_1, \dots, X_n where the distribution of each X_i is close to \mathcal{D} conditioned on the previous outcomes, and computing the empirical distribution of the drawn samples.

Lemma B.7. *Let \mathcal{D} be a distribution over a finite set Σ such that $|\text{supp}(\mathcal{D})| \leq r$. For any $\eta, \epsilon, \gamma > 0$ such that $\gamma < \epsilon$, there is a choice of*

$$n = O((r + 2 + \log(1/\eta))/(\epsilon - \gamma)^2)$$

such that the following holds. Suppose $0 = X_0, X_1, \dots, X_n \in \Sigma$ are possibly correlated random variables such that for all $i \in [n]$ and all values $0 = x_0, x_1, \dots, x_n \in \text{supp}(\mathcal{D})$,

$$\mathcal{D}(X_i | X_0 = x_0, \dots, X_{i-1} = x_{i-1}) \approx_\gamma \mathcal{D}.$$

Then, with probability at least $1 - \eta$, the empirical distribution of the outcomes X_1, \dots, X_n is ϵ -close to \mathcal{D} .

Proof. First, we argue that without loss of generality, we can assume that $|\Sigma| \leq r + 1$. This is because if not, we can define a function $f: \Sigma \rightarrow \text{supp}(\mathcal{D}) \cup \{\star\}$ as follows:

$$f(x) := \begin{cases} x & \text{if } x \in \text{supp}(\mathcal{D}) \\ \star & \text{otherwise.} \end{cases}$$

Observe that for any distribution \mathcal{D}' over Σ , $\text{dist}(\mathcal{D}', \mathcal{D}) = \text{dist}(f(\mathcal{D}'), \mathcal{D})$, since the elements outside $\text{supp}(\mathcal{D})$ always contribute to the statistical distance and we aggregate all such mass on a single extra point \star , and by doing so do not affect the statistical distance. Thus the empirical distribution of (X_1, \dots, X_n) is ϵ -close to \mathcal{D} iff the empirical distribution of $(f(X_1), \dots, f(X_n))$ is.

Now suppose $|\Sigma| \leq r + 1$. Let $A \subseteq \Sigma$ be any non-empty event, and denote by \mathcal{D}' the empirical distribution of the outcomes X_1, \dots, X_n . Let $p := \mathcal{D}(A)$, and define indicator random variables

$$Y_i := \begin{cases} 0 & X_i \notin A, \\ 1 & X_i \in A. \end{cases}$$

for $i \in [n]$ and $Y_0 := 0$. Observe that

$$\mathcal{D}'(A) = \frac{\sum_{i=1}^n Y_i}{n},$$

and, by the assumption on the closeness of conditional distributions of the X_i to \mathcal{D} ,

$$\mathbb{E}[Y_i | Y_0, \dots, Y_{i-1}] \in [p - \gamma, p + \gamma].$$

By Propositions B.3 and B.4, we can thus obtain a concentration bound

$$\Pr[|\mathcal{D}'(A) - p| > \epsilon] \leq 2 \exp(-(\epsilon - \gamma)^2 n / 2).$$

Now we can apply a union bound on all possible choices of A and conclude that

$$\Pr[\neg(\mathcal{D}' \approx_\epsilon \mathcal{D})] \leq 2^{r+2} \exp(-(\epsilon - \gamma)^2 n / 2),$$

which can be ensured to be at most η for some choice of

$$n = O((r + 2 + \log(1/\eta))/(\epsilon - \gamma)^2). \quad \square$$