

An Entropy Lower Bound for Non-Malleable Extractors

Tom Gur ^{*} Igor Shinkar [†]

September 22, 2019

Abstract

A (k, ε) -non-malleable extractor is a function $\text{nmExt} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}$ that takes two inputs, a weak source $X \sim \{0, 1\}^n$ of min-entropy k and an independent uniform seed $s \in \{0, 1\}^d$, and outputs a bit $\text{nmExt}(X, s)$ that is ε -close to uniform, even given the seed s and the value $\text{nmExt}(X, s')$ for an adversarially chosen seed $s' \neq s$. Dodis and Wichs (STOC 2009) showed the existence of (k, ε) -non-malleable extractors with seed length $d = \log(n - k - 1) + 2 \log(1/\varepsilon) + 6$ that support sources of min-entropy $k > \log(d) + 2 \log(1/\varepsilon) + 8$.

We show that the foregoing bound is essentially tight, by proving that any (k, ε) -non-malleable extractor must satisfy the min-entropy bound $k > \log(d) + 2 \log(1/\varepsilon) - \log \log(1/\varepsilon) - C$ for an absolute constant C . In particular, this implies that non-malleable extractors require min-entropy at least $\Omega(\log \log(n))$. This is in stark contrast to the existence of strong seeded extractors that support sources of min-entropy $k = O(\log(1/\varepsilon))$.

Our techniques strongly rely on coding theory. In particular, we reveal an inherent connection between non-malleable extractors and error correcting codes, by proving a new lemma which shows that any (k, ε) -non-malleable extractor with seed length d induces a code $\mathcal{C} \subseteq \{0, 1\}^{2^k}$ with relative distance $\frac{1}{2} - 2\varepsilon$ and rate $\frac{d-1}{2^k}$.

1 Introduction

Randomness extractors are central objects in the theory of computation. Loosely speaking, a *seeded extractor* [NZ96] is a randomized algorithm that extracts nearly uniform bits from biased random sources, using a short seed of randomness. A *non-malleable extractor* [DW09] is a seeded extractor that satisfies a very strong requirement regarding the lack of correlations of the output of the extractor with respect to different seeds.

More accurately, a (k, ε) -non-malleable extractor is a function $\text{nmExt} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}$ such that for every (weak) source X of min-entropy k and a random variable s uniformly distributed on $\{0, 1\}^d$ it holds that $\text{nmExt}(X, s)$ is ε -close to uniform, even given the seed $s \in \{0, 1\}^d$ and the value $\text{nmExt}(X, s')$ for any seed $s' \neq s$ that is determined as an arbitrary function of s . More generally, if $\text{nmExt}(X, s)$ is ε -close to uniform, even given $\text{nmExt}(X, s'_1), \dots, \text{nmExt}(X, s'_t)$ for t adversarially chosen seeds such that $s'_i \neq s$ for all $i \in [t]$, we say it is a (k, ε) - t -non-malleable extractor [CRS14].

The notion of non-malleable extractors is strongly motivated by applications to privacy amplification protocols in presence of an active adversary. This problem considers a setting in which two

^{*}Centre for Discrete Mathematics and its Applications (DIMAP), University of Warwick, United Kingdom. Email: tom.gur@warwick.ac.uk.

[†]School of Computing Science, Simon Fraser University, Burnaby, BC, Canada. Email: ishinkar@sfu.ca.

parties begin by sharing a secret whose distribution may be far from uniform. The parties interact over a public communication channel in the presence of an active adversary (who can arbitrarily change the messages), and would like to securely agree on a nearly uniform secret using small number of rounds of communication. Dodis and Wichs [DW09] proposed an elegant approach for solving this problem by relying on the notion of non-malleable extractors, significantly strengthening the notion of a strong extractors.

Non-malleable extractors have also proven to be a fundamental notion in the theory of pseudo-randomness, as has been recently exemplified by the key role it played in the breakthrough construction of explicit two-source extractors by Chattopadhyay and Zuckerman [CZ16], and the related notions of to Ramsey theory (via two-source extractors and two-source dispersers) [PR04, BKS⁺05].

Non-malleable extractors are a strengthening of the notion of *strong seeded extractors*. These are functions $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}$ such that for a weak source X and seed s it holds that $\text{Ext}(X, s)$ is ε -close to uniform, *even given the seed* $s \in \{0, 1\}^d$. We stress that this is a much weaker guarantee than that of non-malleable extractors. In particular, there exist a blackbox transformation of seeded extractors into strong seeded extractors with roughly the same parameters [RSW06], whereas no such transformation is known for non-malleable extractors.

By a simple probabilistic argument (see, e.g., [Vad12]), there exists a (strong) seeded extractor $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}$ for sources of seed length $d = \log(n) + 2\log(1/\varepsilon) + O(1)$ and min-entropy $k = 2\log(1/\varepsilon) + O(1)$. Moreover, by a long line of research, starting with the seminal work of Nisan and Zuckerman [NZ96], and culminating with [GUV09, DKSS13, TU12] we now know of *explicit* constructions that nearly achieve the optimal parameters.

For non-malleable extractors the parameters achievable by current constructions are weaker. Dodis and Wichs [DW09] showed the existence of (k, ε) -non-malleable extractors with seed length $d = \log(n - k - 1) + 2\log(1/\varepsilon) + 6$, and min-entropy $k > \log(d) + 2\log(1/\varepsilon) + 8$; and in particular, for $k \geq \log \log(n) + 2\log(1/\varepsilon)$. Following a long line of research [Li12a, Li12b, DLWZ14, CRS14, CGL16, Coh16, Coh17, Li17]. The current best explicit construction, due to [Li19] achieves seed length $d = O(\log n) + \tilde{O}(\log(1/\varepsilon))$ for min-entropy $k = \Omega(\log \log n + \log(1/\varepsilon))$.

Note that while for (strong) seeded extractors there are constructions that support sources of min-entropy $k = 2\log(1/\varepsilon) + O(1)$, without any dependence on n , all known constructions of non-malleable extractors require the min-entropy of the source to be at least doubly-logarithmic in n . This naturally raises the question of whether the dependence on n is indeed necessary for non-malleable extractors.

Question: Is it true that in any (k, ε) -non-malleable extractor the min-entropy k must grow with n ?

In this paper we give a positive answer to this question, as well as reveal a simple yet fundamental connection between non-malleable extractors and error-correcting codes, which we believe to be of independent interest.

1.1 Our results

Our main result is a lower bound on the min-entropy required by non-malleable extractors, which essentially matches the one obtained by the probabilistic construction. In particular, we show that any (k, ε) -non-malleable extractor requires the source min-entropy k to be at least $\log \log(n) + (2 - o_\varepsilon(1)) \log(1/\varepsilon)$. In fact, we prove the min-entropy lower bound for the more general notion of t -non-malleable extractors.

Theorem 1 (Main result). *Let $n, k, d, t \in \mathbb{N}$ be parameters such that $t \leq 2^{d/2}$, and let $\varepsilon \in (0, c_0)$ for some absolute constant c_0 . If $\text{nmExt} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}$ is a (k, ε) - t -non-malleable extractor, then $d > \log(n - k) + 2 \log(1/\varepsilon) - C$ and $k \geq \log(d) + 2 \log(1/\varepsilon) - \log \log(1/\varepsilon) + \log(t) - C$ for an absolute constant C .*

We remark that by recent results of [BCD⁺17] and [Li19] the lower bound on d in the theorem is tight up to an additive factor of $O(\log(t))$, and our lower bound on k is almost tight in ε , up to an additive factor of $\log \log(1/\varepsilon)$. Furthermore, since as we mentioned above, there exist (strong) seeded extractors for sources of min-entropy $k = 2 \log(1/\varepsilon) + O(1)$, Theorem 1 implies a chasm between non-malleable extractors and (strong) seeded extractors; in particular, it rules out the possibility of transforming seeded extractors into non-malleable extractors, while preserving the parameters.

Note also that for (k, ε) - t -non-malleable extractor it holds that $k \geq t - O_\varepsilon(1)$. Although the lower bound on k we prove is only logarithmic in t , we find it interesting as this bound is *in addition* to the $\log(d) + (2 - o_\varepsilon(1)) \log(1/\varepsilon)$ lower bound obtained for the case of $t = 1$. We believe that k must be at least $\log(d) + 2 \log(1/\varepsilon) + t$, and we leave this as an open problem.

A key technical tool that we use to prove Theorem 1 is a lemma, which shows that any non-malleable extractor induces an error correcting code with a good distance. We believe this lemma is of independent interest.

Lemma 2. *If there exists a (k, ε) -non-malleable extractor $\text{nmExt} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}$, then there exists an error correcting code $\mathcal{C} \subseteq \{0, 1\}^{2^k}$ with relative distance $\frac{1}{2} - 2\varepsilon$ and rate $\frac{d-1}{2^k}$.*

In fact, we actually prove a more general lemma, which shows that t -non-malleable extractors induce codes with rate that grows with t . See Section 4 for details.

The idea underlying the proof of Lemma 2 has appeared several times in the past, e.g., in [?]. However, to the best of our knowledge, the results were stated in the language of (almost) pair-wise independence, and not expressed in terms of distance of an underlying code, which allows us to apply the known results on the tradeoff between rate and distance of binary codes.

1.2 Technical overview

We provide a high-level overview of the proof of our main result, the min-entropy lower bound in Theorem 1, for the simple case of $t = 1$ (i.e., for standard non-malleable extractors). See Section 4 for the complete details of the proof for the general case. We assume basic familiarity with coding theory and extractors (see Section 2 for the necessary preliminaries).

Consider a non-malleable extractor nmExt . Our strategy for showing a lower bound on the source min-entropy of nmExt consists of the following two steps.

1. Derive a binary code \mathcal{C} with high distance and rate from nmExt , as captured by Lemma 2.
2. Show bounds on the rate of binary codes with a given minimum distance, and apply them to \mathcal{C} to obtain a min-entropy lower bound.

That is, we show that if the parameters of nmExt were too good, then the implied code \mathcal{C} would have parameters that would violate the rate bounds in the second step. Below, we elaborate on each of the steps.

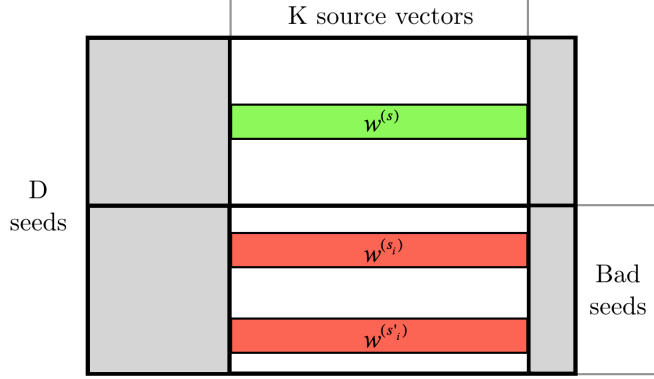


Figure 1: Truth table of a (k, ε) -non-malleable extractor $\text{nmExt} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}$. Rows correspond to the $D = 2^d$ seeds. Columns correspond to all n -bit vectors, out of which we highlight the $K = 2^k$ vectors of the flat source X . Each vector $w^{(s)} = (\text{nmExt}(x, s))_{x \in X}$ consists of the values corresponding to seed s and all vectors of X . The vectors $w^{(s_i)}$ and $w^{(s'_i)}$ correspond to a pair of “bad” seeds $s_i, s'_i \in B$, and hence they are close to each other.

Deriving codes with high distance from non-malleable extractors. We start with a (k, ε) -non-malleable extractor $\text{nmExt} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}$. Denote $K = 2^k$, and consider a (flat) source X , which we view as a collection of K vectors $X \subseteq \{0, 1\}^n$. We show that there is a large subset S of the seeds such that the evaluations of nmExt , with respect to X and S , constitute a code with high distance and rate.

More accurately, denote by $w^{(s)}$ the *evaluation vector* of nmExt on the source X and seed $s \in \{0, 1\}^d$; that is, $w^{(s)} = (\text{nmExt}(x, s))_{x \in X}$. We show that there exists a large subset of seeds $S \subseteq \{0, 1\}^d$ such that

$$\mathcal{C} \stackrel{\text{def}}{=} \{w^{(s)} : s \in S\}$$

is a code with distance $\frac{1}{2} - 2\varepsilon$ and rate $(d - 1)/K$.

As a warmup, it is instructive to note that the definition of (standard) *seeded extractors* only requires that a random coordinate of a random $w^{(s)}$ is nearly uniformly distributed. *Strong seeded extractors* also imply that most evaluation vectors are roughly balanced (i.e., contain a similar number of zeros and ones),¹ as a strong seeded extractor needs to output a nearly uniform bit, even given the seed (i.e., even when the identity of $w^{(s)}$ is known).

An important observation is that the structure of *non-malleable extractors* asserts that there exists a large subset of seeds whose corresponding evaluation vectors are (close to) *pairwise uncorrelated*, and hence constitute a code with large distance. Details follow.

Denote the number of seeds by $D = 2^d$. We wish to show that there exists a subset $S \subset \{0, 1\}^d$ of $D/2$ seeds whose corresponding evaluation vectors are pairwise $(\frac{1}{2} - 2\varepsilon)$ -far. Suppose the contrary, i.e., that *every* set S of $D/2$ seeds contains at least two distinct seeds s, s' such that $w^{(s)}$ is $(\frac{1}{2} - 2\varepsilon)$ -close to $w^{(s')}$. This means that we can iteratively select a set of $D/2$ “bad” seeds $B \stackrel{\text{def}}{=} \{s_1, \dots, s_{D/4}, s'_1, \dots, s'_{D/4}\}$ such that $w^{(s_i)}$ and $w^{(s'_i)}$ are $(\frac{1}{2} - 2\varepsilon)$ -close in Hamming distance, for every $i \in [D/4]$. (See Fig. 1.)

¹ We stress that elements of a set of nearly-balanced vectors are not necessarily pairwise-far, unless this set is a *linear space*. Hence, the foregoing property of strong seeded extractors does *not* imply a good code in general.

The crux is that having many pairs of correlated evaluation vectors violates the assumption that nmExt is a non-malleable extractor. Intuitively, this holds because for each $w^{(s_i)}$ corresponding to a bad seed $s_i \in B$, the output of $\text{nmExt}(X, s_i)$ is biased given $\text{nmExt}(X, s'_i)$. Hence, a non-malleable extractor cannot have a large set of bad seeds.

In Section 4.1 we make this intuition precise by exhibiting an adversarial function $\mathcal{A}: \{0, 1\}^d \rightarrow \{0, 1\}^d$ (with no fixed points) that matches pairs of bad seeds such that we can construct a distinguisher that, for a random variable U_d uniformly distributed on the seeds $\{0, 1\}^d$, can tell apart with confidence ε between $\text{nmExt}(X, U_d)$ and a uniform bit, even when given $\text{nmExt}(X, \mathcal{A}(U_d))$ and U_d .

Rate bounds for binary codes. After deriving a binary code \mathcal{C} with distance $\frac{1}{2} - 2\varepsilon$ and rate $(d-1)/K$ from a (k, ε) -non-malleable extractor nmExt , we wish to apply upper bounds on the rate of binary codes, which will imply a lower bound on the min-entropy required by nmExt .

Our starting point is the state-of-the-art upper bound of McEliece, Rodemich, Rumsey and Welch [MRR⁺77], which, loosely speaking, states that any family of binary codes with relative distance $\frac{1}{2} - \varepsilon$, for sufficiently small $\varepsilon > 0$, has rate $O(\varepsilon^2 \log(1/\varepsilon))$ as the blocklength n tends to infinity.

Unfortunately, the aforementioned bound does not suffice for the min-entropy lower bound, as we need an explicit quantitative bound on the required *blocklength* n of the code (rather than a statement that holds as n tends to infinity). We prove the following theorem, which provides the refined bound that we need.

Theorem 3. *Fix a constant $c \in (0, 1/20)$, and let $\varepsilon \in (0, c)$. For $K > \frac{c}{\varepsilon^2}$ let $\mathcal{C} \subseteq \{0, 1\}^K$ be a code with relative distance $\delta = \frac{1}{2} - \varepsilon$. Then $|\mathcal{C}| < 2^{\frac{23}{e} \varepsilon^2 \log(1/\varepsilon) K}$.*

The proof of Theorem 3 is implied by a result on the high rank of perturbed identity matrices due to Alon [Alo09]. For completeness, we provide in Section 3 an alternative proof, which relies on the spectral approach of Navon and Samorodnitsky [NS09].

To conclude the proof of the min-entropy lower bound, we argue that if the non-malleable extractor nmExt could support min-entropy that is smaller than stated in Theorem 1, then the code \mathcal{C} we derive via Lemma 2 would have rate that would violate the lower bound in Theorem 3.

Relation to t -wise δ -dependent variables. The focus of this paper is concerned with (k, ε) - t -non-malleable extractors. However, we remark that our techniques also apply to t -wise δ -dependent variables.² We posit that stronger bounds might be obtained via an argument that holds for (k, ε) - t -non-malleable extractors but *not* for t -wise δ -dependent variables, and leave this question to future work.

1.3 Organization

In Section 2 we present the required preliminaries. In Section 3 we prove the refined bounds on the rate of binary codes. Finally, in Section 4 we prove our main result, Theorem 1, as well as Lemma 2, which captures the connection between non-malleable extractors and error correcting codes.

²This is because non-malleable extractors imply that for any flat source X there exists a large set of seeds S such that the collection of random variables $\{\text{nmExt}(X, s)\}_{s \in S}$ is t -wise δ -dependent. Note that this holds for every flat source with large min-entropy, where our argument only relies on this being true for some source.

2 Preliminaries

We cover the notation and basic definitions used in this paper.

2.1 Notation

For $n \in \mathbb{N}$, we denote by $[n]$ the set $\{1, \dots, n\}$, and by U_n the random variable that is uniformly distributed over $\{0, 1\}^n$. Throughout, $\log(x)$ is defined as $\log_2(x)$. The *binary min-entropy function* $H: [0, 1] \rightarrow [0, 1]$ is given by $H(x) = -x \log(x) - (1-x) \log(1-x)$. We denote by $\mathbf{1}_E$ the *indicator* of an event E . For a finite set X , we denote by $\Pr_{x \in X}[\cdot]$ the probability over an element x that is chosen *uniformly at random* from X .

Distance. The *relative Hamming distance* (or just *distance*), over alphabet Σ , between two vectors $x, y \in \Sigma^n$ is denoted $\text{dist}(x, y) \stackrel{\text{def}}{=} \frac{|\{i \in [n] : x_i \neq y_i\}|}{n}$. If $\text{dist}(x, y) \leq \varepsilon$, we say that x is ε -close to y , and otherwise we say that x is ε -far from y . Similarly, the *relative distance* of $x \in \Sigma^n$ from a non-empty set $S \subseteq \Sigma^n$ is denoted $\text{dist}(x, S) \stackrel{\text{def}}{=} \min_{y \in S} \text{dist}(x, y)$. If $\text{dist}(x, S) \leq \varepsilon$, we say that x is ε -close to S , and otherwise we say that x is ε -far from S .

The *total variation distance* between two random variables X_1, X_2 over domain Ω is denoted by $\text{dist}_{\text{TV}}(X_1, X_2) \stackrel{\text{def}}{=} \sup_{S \subseteq \Omega} |\Pr[X_1 \in S] - \Pr[X_2 \in S]|$, and is equivalent, up to a factor 2, to their ℓ_1 distance $\|X_1 - X_2\|_1 \stackrel{\text{def}}{=} \sum_{\omega \in \Omega} |\Pr[X_1 = \omega] - \Pr[X_2 = \omega]|$. We say that X_1 is ε -close to X_2 if $\text{dist}_{\text{TV}}(X_1, X_2) \leq \varepsilon$, and otherwise we say that X_1 is ε -far from X_2 .

Remark. In order to show that X_1 is ε -far from X_2 it suffices to show a randomized distinguisher $\mathcal{D}: \Omega \rightarrow \{0, 1\}$ such that $|\Pr[\mathcal{D}(X_1) = 1] - \Pr[\mathcal{D}(X_2) = 1]| > \varepsilon$, where the probabilities are over the random variables X_1, X_2 and the randomness of \mathcal{D} . Note that if such randomized distinguisher exists, then, by averaging, there is also a *deterministic* distinguisher with the same property. This, naturally, defines the event $S_{\mathcal{D}} = \{\omega \in \Omega : \mathcal{D}(\omega) = 1\} \subseteq \Omega$. for which we have $\text{dist}_{\text{TV}}(X_1, X_2) = \sup_{S \subseteq \Omega} |\Pr[X_1 \in S] - \Pr[X_2 \in S]| \geq |\Pr[X_1 \in S_{\mathcal{D}}] - \Pr[X_2 \in S_{\mathcal{D}}]| > \varepsilon$, and hence X_1 is ε -far from X_2 .

2.2 Error correcting codes

Let $k, n \in \mathbb{N}$, and let Σ be a finite alphabet. An *error correcting code* is a set $\mathcal{C} \subseteq \Sigma^n$, and the elements of \mathcal{C} are called its *codewords*. The parameter n is called the *blocklength* of \mathcal{C} , and $k = \log_{|\Sigma|}(|\mathcal{C}|)$ is the *dimension* of \mathcal{C} . The *relative distance* of a code \mathcal{C} is the minimal relative Hamming distance between its codewords, and is denoted by $\delta = \min_{c \neq c' \in \mathcal{C}} \{\text{dist}(c, c')\}$. The *rate* of the code, measuring the redundancy of the encoding, is the ratio of its dimension and blocklength, and is denoted by $\rho = k/n$. If the alphabet is binary, i.e., $\Sigma = \{0, 1\}$, we say that \mathcal{C} is a *binary code*.

2.3 Randomness extractors

We recall the standard definitions of random sources and several types of extractors, as well as state known bounds that we will need.

Weak sources. For integers $n > k$, an (n, k) -random source X of min-entropy k is a random variable taking values in $\{0, 1\}^n$ such that for every $x \in \{0, 1\}^n$ it holds that $\Pr[X = x] \leq 2^{-k}$. An (n, k) -random source X is *flat* if it is uniformly distributed over some subset $S \subseteq \{0, 1\}^n$ of size 2^k .

It is well known [CG88] that the distribution of any (n, k) -random source is a convex combination of distributions of flat (n, k) -random sources, and thus it typically suffices to consider flat sources. We follow the literature, restrict our attention to flat (n, k) -random sources, and refer to them simply as (n, k) -sources.

Seeded extractors. A function $\text{Ext}: \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}$ is a (k, ε) -seeded extractor if for any (n, k) -source X , the distribution of $\text{Ext}(X, U_d)$ is ε -close to U_1 , i.e., $\text{dist}_{TV}(\text{Ext}(X, U_d), U_1) \leq \varepsilon$. (Recall that U_m denotes the random variable that is uniformly distributed on $\{0, 1\}^m$.)

A function $\text{Ext}: \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}$ is a (k, ε) -strong seeded extractor if for any (n, k) -source X the distribution of $(\text{Ext}(X, U_d), U_d)$ is ε -close to U_{d+1} . We will need the following lower bound on the source min-entropy required by strong seeded extractors, due to Radhakrishnan and Ta-Shma [RT00] (see also [NZ96]).

Theorem 2.1 ([RT00] Theorem 1.9). *Let $\text{Ext}: \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}$ be a (k, ε) -strong seeded extractor. Then, it holds that*

$$d > \log(n - k) + 2 \log(1/\varepsilon) - c \text{ and } k \geq 2 \log(1/\varepsilon) - c,$$

for some absolute constant $c \in \mathbb{R}$.

Non-malleable extractors. Informally, a *non-malleable extractor* nmExt is a seeded extractor that for any source X and seed s outputs a bit $\text{nmExt}(X, s)$ that is nearly uniform even if given the seed s and value $\text{nmExt}(X, s')$ for an adversarially selected seed s' .

Formally, we say that a function $\mathcal{A}: \{0, 1\}^d \rightarrow \{0, 1\}^d$ is an *adversarial function* if it has no fixed points, i.e., if $\mathcal{A}(s) \neq s$ for all $s \in \{0, 1\}^d$. Non-malleable extractors are defined as follows.

Definition 2.2. *A function $\text{nmExt}: \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}$ is a (k, ε) -non-malleable extractor if for any (n, k) -source X , and for any adversarial function $\mathcal{A}: \{0, 1\}^d \rightarrow \{0, 1\}^d$, it holds that the distribution of the 3-tuple $(\text{nmExt}(X, U_d), \text{nmExt}(X, \mathcal{A}(U_d)), U_d)$ is ε -close to $(U_1, \text{nmExt}(X, \mathcal{A}(U_d)), U_d)$; that is,*

$$\text{dist}_{TV} \left((\text{nmExt}(X, U_d), \text{nmExt}(X, \mathcal{A}(U_d)), U_d), (U_1, \text{nmExt}(X, \mathcal{A}(U_d)), U_d) \right) \leq \varepsilon.$$

We will also consider the more general notion of *t-non-malleable extractors*, in which it is possible to extract randomness even given *multiple* (namely, t) outputs of the extractor with respect to adversarially chosen seeds.

Definition 2.3. *A function $\text{nmExt}: \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}$ is a (k, ε) -t-non-malleable extractor if for any (n, k) -source X and for any t adversarial functions $\mathcal{A}_1, \dots, \mathcal{A}_t: \{0, 1\}^d \rightarrow \{0, 1\}^d$ it holds that*

$$\text{dist}_{TV} \left((\text{nmExt}(X, U_d), (\text{nmExt}(X, \mathcal{A}_i(U_d)))_{i=1}^t, U_d), (U_1, (\text{nmExt}(X, \mathcal{A}_i(U_d)))_{i=1}^t, U_d) \right) \leq \varepsilon.$$

3 Refined coding bounds

As we mentioned in the technical overview (Section 1.2), we prove our min-entropy lower bound for non-malleable extractors by deriving codes from extractors and bounding the rate of these codes. To this end, in this section we prove refined bounds on the rate of binary codes with a given minimum distance. Our starting point is the seminal result of McEliece, Rodemich, Rumsey and Welch [MRR⁺77].

Theorem 3.1 ([MRR⁺77]). *Any code $\mathcal{C} \subseteq \{0, 1\}^n$ with relative distance $\delta \in (0, \frac{1}{2})$ has rate at most $H\left(\frac{1}{2} - \sqrt{\delta(1-\delta)}\right) + o(1)$, where $o(1)$ is some function that tends to zero as n grows to infinity.*

Observe that in particular, by plugging in $\delta = \frac{1}{2} - \varepsilon$ for sufficiently small $\varepsilon > 0$, and letting n be sufficiently large Theorem 3.1 implies that any family of binary codes with blocklength n and relative distance $\frac{1}{2} - \varepsilon$ has rate $\rho = O(\varepsilon^2 \log(1/\varepsilon))$.

However, the above does not suffice for our needs, as to prove our main result (Theorem 1) we need a quantitative bound on n . Specifically, we need the bound to hold for $n > \frac{c}{\varepsilon^2}$ for some constant $c > 0$. We prove the following theorem, which provides the refined bound that we seek.

Theorem 3.2. *Fix some constant $c \in (0, 1/20)$, and let $\varepsilon \in (0, c)$. For $n > \frac{c}{\varepsilon^2}$, let $\mathcal{C} \subseteq \{0, 1\}^n$ be a code with relative distance $\delta = \frac{1}{2} - \varepsilon$. Then $|\mathcal{C}| < 2^{\frac{23}{c}\varepsilon^2 \log(1/\varepsilon)n}$.*

Proof. The proof follows the general approach of Navon and Samorodnitsky [NS09], who provide a spectral graph theoretic framework to prove upper bounds on the rate of binary codes.

We will need the following definition, which generalizes the notion of a *maximal eigenvalue* to subsets of the hypercube.

Definition 3.3. *Let $A \in \{0, 1\}^{2^n \times 2^n}$ be the adjacency matrix of the hypercube graph; that is, $A_{x,y} = 1$ if and only if $x \in \{0, 1\}^n$ and $y \in \{0, 1\}^n$ differ in exactly one coordinate. Given a set $B \subseteq \{0, 1\}^n$, we define*

$$\lambda_B = \max_{\substack{f: \{0,1\}^n \rightarrow \mathbb{R} \\ \text{supp}(f) \subseteq B}} \frac{\langle Af, f \rangle}{\langle f, f \rangle}.$$

To better understand the definition of λ_B , it is convenient to consider the subgraph H_B of the hypercube graph $\{0, 1\}^n$ induced by the vertices in B , and observe that λ_B is the maximal eigenvalue of the adjacency matrix of H_B . Navon and Samorodnitsky [NS09] prove the following result.

Proposition 3.4 ([NS09, Proposition 1.1]). *Let $\mathcal{C} \subseteq \{0, 1\}^n$ be a code with relative distance $\delta > 0$, and let $\varepsilon > 0$. Suppose that for a subset $B \subseteq \{0, 1\}^n$ it holds that $\lambda_B \geq (1 - 2\delta + \varepsilon)n$. Then $|\mathcal{C}| \leq |B|/\varepsilon$.*

The foregoing theorem naturally suggest the following proof strategy: to upper bound the rate of a binary code \mathcal{C} with relative distance $\delta = \frac{1}{2} - \varepsilon$, it suffices to exhibit a (small as possible) set $B \subseteq \{0, 1\}^n$ whose corresponding maximal eigenvalue satisfies $\lambda_B \geq 3\varepsilon n$; note that the smaller B is, the better upper bound we get on the rate of \mathcal{C} .

Towards this end, let $r \in [n]$ be a parameter to be chosen later, and let

$$B = \{x \in \{0, 1\}^n : \text{weight}(x) \in \{r, r + 1\}\},$$

where $\text{weight}(x)$ denotes the (absolute) Hamming weight of x . We lower bound the maximal eigenvalue λ_B by showing a particular function f that is supported on B , such that $\frac{\langle Af, f \rangle}{\langle f, f \rangle} \geq 3\epsilon n$. Specifically, for some $a, b \in \mathbb{R}$ to be chosen later, we define $f : \{0, 1\}^n \rightarrow \mathbb{R}$ as

$$f(x) = \begin{cases} a & \text{if } \text{weight}(x) = r \\ b & \text{if } \text{weight}(x) = r + 1 \\ 0 & \text{otherwise} . \end{cases}$$

Clearly $\text{supp}(f) \subseteq B$. Observe that

$$\frac{\langle Af, f \rangle}{\langle f, f \rangle} = \frac{ab \binom{n}{r} \cdot (n - r) + ab \binom{n}{r+1} (r + 1)}{a^2 \binom{n}{r} + b^2 \binom{n}{r+1}} = \frac{2ab \binom{n}{r} \cdot (n - r)}{a^2 \binom{n}{r} + b^2 \binom{n}{r} \cdot \frac{n-r}{r+1}} > \frac{2ab \cdot r(n - r)}{a^2 \cdot r + b^2 \cdot (n - r)} .$$

By choosing r to be an integer in the interval $\left[\frac{9\epsilon^2}{c}n, \frac{10\epsilon^2}{c}n \right]$ and letting $b = a\sqrt{\frac{r}{n}}$ we get that³

$$\frac{\langle Af, f \rangle}{\langle f, f \rangle} > \frac{2a^2 \sqrt{r/n} \cdot r(n - r)}{a^2 r + a^2 \cdot (r/n) \cdot (n - r)} = \frac{2\sqrt{rn}(n - r)}{2n - r} > 6\epsilon n ,$$

where the last inequality uses the assumptions that $\epsilon < c < 1/20$, which implies that $r \leq \frac{10\epsilon^2}{c}n < \frac{n}{2}$. Therefore, by applying Proposition 3.4 we get that

$$|\mathcal{C}| \leq \frac{|B|}{4\epsilon} = \frac{\binom{n}{r} + \binom{n}{r+1}}{4\epsilon} \leq \binom{n}{r} \cdot \frac{n}{4r\epsilon} \leq \frac{c}{9\epsilon^3} \binom{n}{\frac{10\epsilon^2}{c}n} \leq \frac{c}{9\epsilon^3} \left(\frac{ce}{10\epsilon^2} \right)^{\frac{10\epsilon^2}{c}n} < 2^{\frac{23\epsilon^2 \log(1/\epsilon)}{c}n} ,$$

which concludes the proof of Theorem 3.2. □

4 Proof of Theorem 1

In this section we prove Theorem 1, which we restate here with slightly more specific parameters than those stated above.

Theorem 1 (restated): *Let $n, k, d, t \in \mathbb{N}$ be parameters such that $t \leq 2^{d/2}$, and let $\epsilon \in (0, c_0/2)$ for $c_0 = \min\{1/2^c, 1/20\}$, where $c > 0$ is the constant from Theorem 2.1. If $\text{nmExt} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}$ is a (k, ϵ) - t -non-malleable extractor, then*

$$d > \log(n - k) + 2 \log(1/\epsilon) - O(1) \text{ and } k \geq \log(d) + 2 \log(1/\epsilon) - \log \log(1/\epsilon) + \log(t) - O(1).$$

We start, in Section 4.1, with the proof of Theorem 1 for the special case where $t = 1$ (i.e., for standard non-malleable extractors). Then, in Section 4.2, we provide the full proof for general values of t .

³Note that by the assumption in the theorem we have $1 < \frac{\epsilon^2}{c}n < n$. In particular, the interval $\left[\frac{9\epsilon^2}{c}n, \frac{10\epsilon^2}{c}n \right]$ contains an integer.

4.1 Proof of Theorem 1 for $t = 1$

Following the outline provided in Section 1.2, we start the proof with the following lemma, showing that any non-malleable extractor induces an error correcting code with good distance.

Lemma 4.1 (Lemma 2, restated). *If there exists a (k, ε) -non-malleable extractor $\text{nmExt} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}$, then there exists an error correcting code $\mathcal{C} \subseteq \{0, 1\}^{2^k}$ with relative distance $\frac{1}{2} - 2\varepsilon$ and rate $\frac{d-1}{2^k}$.*

Proof. Let $\text{nmExt} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}$ be a (k, ε) -non-malleable extractor, and let X be an (n, k) -source. That is, $X \subseteq \{0, 1\}^n$ is a collection of $K = 2^k$ vectors, which we denote by $X = \{x_1, \dots, x_K\} \subseteq \{0, 1\}^n$. For each seed $s \in \{0, 1\}^d$, let $w^{(s)} \in \{0, 1\}^K$ be the K -bit *evaluation vector* defined as

$$w^{(s)} = (\text{nmExt}(x_i, s))_{i \in \{1, \dots, K\}} .$$

We claim that the (multi-)set $\{w^{(s)} : s \in \{0, 1\}^d\} \subseteq \{0, 1\}^K$ contains an error correcting code $\mathcal{C} \subseteq \{0, 1\}^K$ with relative distance $\frac{1}{2} - 2\varepsilon$ and rate $\frac{d-1}{K}$.

Claim 4.2. *There exists a subset $S \subseteq \{0, 1\}^d$ of size 2^{d-1} such that for every two distinct $s, s' \in S$ it holds that $\text{dist}(w^{(s)}, w^{(s')}) \geq \frac{1}{2} - 2\varepsilon$.*

Proof. Suppose towards contradiction that for every subset $S' \subseteq \{0, 1\}^d$ of size at least 2^{d-1} there exist distinct seeds $s, s' \in S'$ such that $\text{dist}(w^{(s)}, w^{(s')}) < \frac{1}{2} - 2\varepsilon$. We show below that this contradicts the assumption that nmExt is a (k, ε) -non-malleable extractor.

Indeed, by the assumption, we can find $s_1, s'_1 \in \{0, 1\}^d$ such that $\text{dist}(w^{(s_1)}, w^{(s'_1)}) < \frac{1}{2} - 2\varepsilon$. Then, we can remove s_1, s'_1 from $\{0, 1\}^d$, and apply the assumption again, to obtain $s_2, s'_2 \in \{0, 1\}^d \setminus \{s_1, s'_1\}$ such that $\text{dist}(w^{(s_2)}, w^{(s'_2)}) < \frac{1}{2} - 2\varepsilon$. By iteratively repeating this argument $D/4$ times, where $D = 2^d$, we obtain $D/4$ pairs of distinct elements $(s_1, s'_1), \dots, (s_{D/4}, s'_{D/4})$ such that

$$\forall j \in [D/4] \quad \text{dist}(w^{(s_j)}, w^{(s'_j)}) < \frac{1}{2} - 2\varepsilon . \quad (1)$$

Let $B = \{s_j, s'_j : j \in [D/4]\} \subseteq \{0, 1\}^d$ denote the set of all such “bad” seeds, and define an adversarial function $\mathcal{A} : \{0, 1\}^d \rightarrow \{0, 1\}^d$ that matches each pair of bad seeds by mapping $\mathcal{A}(s_j) = s'_j$ and $\mathcal{A}(s'_j) = s_j$ for all $j \in [D/4]$, and defining $\mathcal{A}(s)$ arbitrarily for all other seeds $s \notin B$.

Next we prove that nmExt is not a (k, ε) -non-malleable extractor by arguing that the distribution of the random variable consisting of the 3-tuple $(\text{nmExt}(X, U_d), \text{nmExt}(X, \mathcal{A}(U_d)), U_d)$ is ε -far from $(U_1, \text{nmExt}(X, \mathcal{A}(U_d)), U_d)$, where recall that U_m denotes the random variable that is uniformly distributed over $\{0, 1\}^m$. Indeed, consider the following distinguisher $\mathcal{D} : \{0, 1\} \times \{0, 1\} \times \{0, 1\}^d \rightarrow \{0, 1\}$, defined as

$$\mathcal{D}(b, b', s) = \begin{cases} \mathbf{1}_{b=b'}, & \text{if } s \in B \\ U_1, & \text{otherwise} . \end{cases}$$

Clearly $\Pr[\mathcal{D}(U_1, \text{nmExt}(X, \mathcal{A}(U_d)), U_d) = 1] = \frac{1}{2}$. On the other hand, by Eq. (1), for s sampled from U_d we have

$$\Pr[\mathcal{D}(\text{nmExt}(X, s), \text{nmExt}(X, \mathcal{A}(s)), s) = 1] \geq \left(\frac{1}{2} + 2\varepsilon\right) \Pr[s \in B] + \frac{1}{2} \Pr[s \notin B] \geq \frac{1}{2} + \varepsilon ,$$

thus contradicting the assumption that nmExt is a (k, ε) -non-malleable extractor. This concludes the proof of Claim 4.2. \square

Therefore, by Claim 4.2 there exists a set $\mathcal{C} = \{w^{(s)} : s \in S\} \subseteq \{0, 1\}^K$ of size 2^{d-1} such that for every $x, y \in \mathcal{C}$ it holds that $\text{dist}(x, y) \geq \frac{1}{2} - 2\varepsilon$, i.e., \mathcal{C} is an error correcting code with relative distance $\frac{1}{2} - 2\varepsilon$ and rate $\frac{d-1}{2^k}$, which completes the proof of Lemma 4.1. \square

By applying the bound from Theorem 3.2 to the code obtained in Lemma 4.1, we prove Theorem 1 for the case of $t = 1$.

Proof of Theorem 1 for $t = 1$. Since every non-malleable extractor is, in particular, a strong seeded extractor, then by Theorem 2.1 it holds that the seed length is $d > \log(n - k) + 2 \log(1/\varepsilon) - c$, as required. Furthermore, Theorem 2.1 also implies that

$$k \geq 2 \log(1/\varepsilon) - c. \quad (2)$$

By Lemma 4.1, if $\text{nmExt} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}$ is a (k, ε) -non-malleable extractor, then there exists an error correcting code $\mathcal{C} \subseteq \{0, 1\}^{2^k}$ with relative distance $\frac{1}{2} - 2\varepsilon$ and rate $\frac{d-1}{2^k}$.

Next, we wish to apply Theorem 3.2 to the code \mathcal{C} . Recall that by the assumption it holds that $\varepsilon < c_0$ and $c_0 < 1/2^c$, and observe that by Eq. (2) we have $2^k \geq \frac{2^{-c}}{\varepsilon^2} > \frac{c_0}{\varepsilon^2}$. Therefore, by applying Theorem 3.2, with respect to c_0 (recall that $c_0 < 1/20$) and $2\varepsilon < c_0$ we get that

$$2^{d-1} \leq |\mathcal{C}| < 2^{\frac{23}{c_0} \cdot (2\varepsilon)^2 \log(1/2\varepsilon) 2^k},$$

and thus $k \geq \log(d) + 2 \log(1/\varepsilon) - \log \log(1/\varepsilon) - O(1)$, as required. \square

4.2 Proof of Theorem 1 for general t

Next, we extend the idea presented in Section 4.1 to larger values of t . The key step is the following lemma.

Lemma 4.3. *If there exists a (k, ε) - t -non-malleable extractor $\text{nmExt} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}$, then, there exists an error correcting code $\mathcal{C} \subseteq \{0, 1\}^{2^k}$ with relative distance $\frac{1}{2} - 2\varepsilon$ such that $|\mathcal{C}| \geq (2^{d-1}/t)^{\lfloor t/2 \rfloor}$.*

Proof. Let $\text{nmExt} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}$ be a (k, ε) - t -non-malleable extractor. Similarly to the proof of Lemma 4.1, we set $K = 2^k$, and let X be an (n, k) -source, which we view as a collection of vectors $X = \{x_1, \dots, x_K\} \subseteq \{0, 1\}^n$. For each seed $s \in \{0, 1\}^d$, let $w^{(s)} \in \{0, 1\}^K$ be the K -bit *evaluation vector*, defined as $w^{(s)} = (\text{nmExt}(x_i, s))_{i \in \{1, \dots, K\}}$. Hereafter, all sums involving binary vectors are summations over $\text{GF}(2)$. For $x \in \{0, 1\}^n$, we denote by $\text{weight}(x)$ the (absolute) Hamming weight of x .

Whereas before, in the proof of Lemma 4.1, we showed that the multi-set of evaluation vectors $\{w^{(s)} : s \in \{0, 1\}^d\} \subseteq \{0, 1\}^K$ simply contains an error correcting code with good parameters, here we will derive our code by considering all $\text{GF}(2)$ -linear combinations of $\lfloor t/2 \rfloor$ elements of a carefully selected subset of the evaluation vectors.

Towards that end, the next claim shows that there exists a large subset of seeds such that any linear combination of $t + 1$ of the evaluation vectors that corresponds to these seeds has large Hamming weight.

Claim 4.4. *There is a subset $S \subseteq \{0, 1\}^d$ of size 2^{d-1} such that for every subset $I \subseteq S$ of size $|I| \leq t + 1$ it holds that $\text{weight}(\sum_{s \in I} w^{(s)}) \geq (\frac{1}{2} - 2\varepsilon)K$.*

Proof. Assume towards contradiction that for every subset $S' \subseteq \{0, 1\}^d$ of size at least 2^{d-1} there are $t' \leq t + 1$ distinct seeds $s_1, \dots, s_{t'} \in S'$ such that

$$\Pr_{x \in X} \left[\sum_{j=1}^{t'} \text{nmExt}(x, s_j) = 1 \right] < \frac{1}{2} - 2\varepsilon .$$

We show below that this contradicts the assumption that nmExt is a (k, ε) - t -non-malleable extractor.

By our assumption, there is a subset of seeds $S_1 \subseteq \{0, 1\}^d$ for which there exists $I_1 \subseteq S_1$ of size $|I_1| = t'_1 \leq t + 1$ such that $\text{weight}(\sum_{s \in I_1} w^{(s)}) < (\frac{1}{2} - 2\varepsilon)K$. We remove I_1 from $\{0, 1\}^d$, and apply the assumption again to obtain $I_2 \subseteq \{0, 1\}^d$ of size $|I_2| = t'_2 \leq t + 1$ such that $\text{weight}(\sum_{s \in I_2} w^{(s)}) < (\frac{1}{2} - 2\varepsilon)K$. We then remove I_2 from $\{0, 1\}^d \setminus I_1$, and apply the assumption again with respect to $\{0, 1\}^d \setminus (I_1 \cup I_2)$. By repeating this argument as long as $|\cup_j I_j| < 2^{d-1}$, we obtain R disjoint subsets I_1, \dots, I_R , where the size of each I_j is $t'_j \leq t + 1$, such that $\sum_{j=1}^R |I_j| \geq 2^{d-1}$ and

$$\text{weight} \left(\sum_{s \in I_j} w^{(s)} \right) < \left(\frac{1}{2} - 2\varepsilon \right) K , \quad (3)$$

for all $j \in [R]$. Analogously to the proof of Lemma 4.1, the set $I_1 \cup \dots \cup I_R$ consists of the “bad seeds” that correspond to evaluation vectors whose $(t + 1)$ -element linear combinations are of low weight.

To prove that the foregoing collection of “bad seeds” violates the assumption that nmExt is a (k, ε) - t -non-malleable extractor, we exhibit t adversarial functions $\mathcal{A}_1, \dots, \mathcal{A}_t: \{0, 1\}^d \rightarrow \{0, 1\}^d$ (with no fixed points) for which there exists a function that distinguishes between the random variables consisting of the $(t + 2)$ -tuples

$$\left(\text{nmExt}(X, U_d), \left(\text{nmExt}(X, \mathcal{A}_\ell(U_d)) \right)_{\ell \in [t]}, U_d \right) \text{ and } \left(U_1, \left(\text{nmExt}(X, \mathcal{A}_\ell(U_d)) \right)_{\ell \in [t]}, U_d \right)$$

with confidence ε , where recall that U_m denotes the random variable that is uniformly distributed over $\{0, 1\}^m$.

We define the family $\{\mathcal{A}_\ell\}_{\ell \in [t]}$ in the natural way, by mapping each of the bad seeds to the set of seeds with which its linear combination is a low weight vector. That is, for each $j \in [R]$ let $I_j = \{s_1, \dots, s_{t'_j}\}$, where $t'_j \leq t + 1$. Then, for all $\ell \in [t]$ we define

$$\mathcal{A}_\ell(s_i) = \begin{cases} s_{i+\ell \pmod{t'_j}}, & \text{for } s_i \in I_j, j \in [R] \\ \text{arbitrary,} & \text{for } s \in \{0, 1\}^d \setminus (\cup_{j \in [R]} I_j) . \end{cases}$$

Note that by definition of the \mathcal{A}_ℓ 's, for all $j \in [R]$ and $s \in I_j$ it holds that $\{s\} \cup \{\mathcal{A}_\ell(s)\}_{\ell \in [t]} = I_j$, and so, by Eq. (3) we have that

$$\Pr_{x \in X} \left[\text{nmExt}(x, s) = \sum_{i=1}^{t'_j-1} \text{nmExt}(x, \mathcal{A}_i(s)) \right] = \frac{\text{weight}(\sum_{s \in I_j} w^{(s)})}{K} < \left(\frac{1}{2} - 2\varepsilon \right) K .$$

Next, we define the distinguisher $\mathcal{D}: \{0, 1\} \times \{0, 1\}^t \times \{0, 1\}^d \rightarrow \{0, 1\}$ as

$$\mathcal{D}(b, b_1, \dots, b_t, s) = \begin{cases} \mathbf{1}_{b = \sum_{i \in [t'_j-1]} b_i}, & \text{if } s \in I_j \text{ for some } j \in [R] \\ U_1, & \text{otherwise .} \end{cases}$$

Clearly $\Pr \left[\mathcal{D} \left(U_1, (\text{nmExt}(X, \mathcal{A}_\ell(U_d)))_{\ell \in [t]}, U_d \right) = 1 \right] = \frac{1}{2}$. On the other hand, for s sampled from U_d we have

$$\begin{aligned} & \Pr \left[\mathcal{D} \left(\text{nmExt}(X, s), (\text{nmExt}(X, \mathcal{A}_\ell(s)))_{\ell \in [t]}, s \right) = 1 \right] \\ & \geq \left(\frac{1}{2} + 2\varepsilon \right) \Pr \left[s \in \cup_{j \in [R]} I_j \right] + \frac{1}{2} \Pr \left[s \in \{0, 1\}^d \setminus \cup_{j \in [R]} I_j \right] \geq \frac{1}{2} + \varepsilon , \end{aligned}$$

thus contradicting the assumption that nmExt is a (k, ε) - t -non-malleable extractor. This concludes the proof of Claim 4.4. \square

Let $S \subseteq \{0, 1\}^d$ be the set guaranteed by Claim 4.4, and consider the code

$$\mathcal{C} \stackrel{\text{def}}{=} \left\{ \sum_{s \in I} w^{(s)} : I \subseteq S, |I| \leq \lfloor t/2 \rfloor \subseteq \{0, 1\}^K \right\} .$$

Note that for $D = 2^d$ we have $|\mathcal{C}| \geq \binom{D/4}{\lfloor t/2 \rfloor} \geq (D/2t)^{\lfloor t/2 \rfloor}$. By the guarantee of Claim 4.4, for every distinct $x, y \in \mathcal{C}$ it holds that $\text{dist}(x, y) \geq \frac{1}{2} - 2\varepsilon$; that is $\mathcal{C} \subseteq \{0, 1\}^K$ is an error correcting code with relative distance $\frac{1}{2} - 2\varepsilon$, which completes the proof of Lemma 4.3. \square

We prove Theorem 1 by applying the bound from Theorem 3.2 to the code obtained in Lemma 4.3, analogously to the way we proved the theorem for the restricted case of $t = 1$ before.

Proof of Theorem 1 (general case). Since every t -non-malleable extractor is, in particular, a strong seeded extractor, then by Theorem 2.1 it holds that the seed length is $d > \log(n - k) + 2 \log(1/\varepsilon) - c$, as required. Furthermore, Theorem 2.1 also implies that $k \geq 2 \log(1/\varepsilon) - c$.

By Lemma 4.3, if $\text{nmExt} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}$ is a (k, ε) -non-malleable extractor, then there exists an error correcting code $\mathcal{C} \subseteq \{0, 1\}^{2^k}$ with relative distance $\frac{1}{2} - 2\varepsilon$ such that $|\mathcal{C}| \geq (2^{d-1}/t)^{\lfloor t/2 \rfloor}$.

We wish to apply Theorem 3.2 to the code \mathcal{C} . Recall that by the assumption it holds that $\varepsilon < c_0$ and $c_0 < 1/2^c$, and observe that according to the bound on k given by Theorem 2.1, we have that $2^k \geq \frac{2^{-c}}{\varepsilon^2} > \frac{c_0}{\varepsilon^2}$. Therefore, by applying Theorem 3.2, with respect to c_0 (recall that $c_0 < 1/20$) and $2\varepsilon < c_0$, we get that

$$(2^{d-1}/t)^{\lfloor t/2 \rfloor} \leq |\mathcal{C}| < 2^{\frac{23}{c_0} \cdot (2\varepsilon)^2 \log(1/2\varepsilon) 2^k} ,$$

and by the assumption that $\log(t) < d/2$ we get that

$$\frac{23}{c_0} \cdot (2\varepsilon)^2 \log(1/2\varepsilon) 2^k \geq (d - 1 - \log(t)) \cdot \lfloor t/2 \rfloor \geq \Omega(d \cdot t) .$$

This implies that $k \geq \log(d) + \log(t) + 2 \log(1/\varepsilon) - \log \log(1/\varepsilon) - O(1)$, as required. \square

Acknowledgements

We are thankful to Gil Cohen for helpful discussions, and to Venkatesan Guruswami for a discussion regarding the MRRW bounds. We also thank the anonymous reviewers for comments that improved the presentation of this work.

References

- [Alo09] Noga Alon. Perturbed identity matrices have high rank: Proof and applications. *Combinatorics, Probability and Computing*, 18(1-2):3–15, 2009.
- [BCD⁺17] Avraham Ben-Aroya, Eshan Chattopadhyay, Dean Doron, Xin Li, and Amnon Ta-Shma. A reduction from efficient non-malleable extractors to low-error two-source extractors with arbitrary constant rate. *ECCC TR17-027*, 2017. Manuscript.
- [BKS⁺05] Boaz Barak, Guy Kindler, Ronen Shaltiel, Benny Sudakov, and Avi Wigderson. Simulating independence: New constructions of condensers, ramsey graphs, dispersers, and extractors. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 1–10, 2005.
- [CG88] Benny Chor and Oded Goldreich. Unbiased bits from sources of weak randomness and probabilistic communication complexity. *SIAM Journal on Computing*, 17(2):230–261, 1988.
- [CGL16] Eshan Chattopadhyay, Vipul Goyal, and Xin Li. Non-malleable extractors and codes, with their many tampered extensions. In *Proceedings of the Forty-eighth Annual ACM Symposium on Theory of Computing*, STOC '16, pages 285–298, 2016.
- [Coh16] Gil Cohen. Non-malleable extractors - new tools and improved constructions. In *31st Conference on Computational Complexity (CCC 2016)*, 2016.
- [Coh17] Gil Cohen. Towards optimal two-source extractors and ramsey graphs. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1157–1170, 2017.
- [CRS14] Gil Cohen, Ran Raz, and Gil Segev. Nonmalleable extractors with short seeds and applications to privacy amplification. *SIAM Journal on Computing*, 43(2):450–476, 2014.
- [CZ16] Eshan Chattopadhyay and David Zuckerman. Explicit two-source extractors and resilient functions. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*, pages 670–683, 2016.
- [DKSS13] Zeev Dvir, Swastik Kopparty, Shubhangi Saraf, and Madhu Sudan. Extensions to the method of multiplicities, with applications to key sets and mergers. *SIAM Journal on Computing*, 42(6):2305–2328, 2013.
- [DLWZ14] Yevgeniy Dodis, Xin Li, Trevor D. Wooley, and David Zuckerman. Privacy amplification and non-malleable extractors via character sums. *SIAM Journal on Computing*, 43(2):800–830, 2014.
- [DW09] Yevgeniy Dodis and Daniel Wichs. Non-malleable extractors and symmetric key cryptography from weak secrets. In *Proceedings of the Forty-first Annual ACM Symposium on Theory of Computing*, STOC '09, pages 601–610, New York, NY, USA, 2009. ACM.

- [GUV09] Venkatesan Guruswami, Christopher Umans, and Salil Vadhan. Unbalanced expanders and randomness extractors from parvaresh–vardy codes. *Journal of the ACM (JACM)*, 56(4):20, 2009.
- [Li12a] Xin Li. Design extractors, non-malleable condensers and privacy amplification. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing*, pages 837–854, 2012.
- [Li12b] Xin Li. Non-malleable extractors, two-source extractors and privacy amplification. In *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science*, pages 688–697, 2012.
- [Li17] Xin Li. Improved non-malleable extractors, non-malleable codes and independent source extractors. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1144–1156. ACM, 2017.
- [Li19] Xin Li. Non-malleable extractors and non-malleable codes: Partially optimal constructions. In *Proceedings of the 34th Conference on Computational Complexity*, 2019.
- [MRR⁺77] Robert J. McEliece, Eugene R. Rodemich, Howard Rumsey, Lloyd, and R. Welch. New upper bounds on the rate of a code via the Delsarte-MacWilliams inequalities. *IEEE Transactions on Information Theory*, IT-23(2):157–166, 1977.
- [NS09] Michael Navon and Alex Samorodnitsky. Linear programming bounds for codes via a covering argument. *Discrete & Computational Geometry*, 41(2):199–207, 2009.
- [NZ96] Noam Nisan and David Zuckerman. Randomness is linear in space. *Journal of Computer and System Sciences*, 52(1):43–52, 1996.
- [PR04] P. Pudlak and V. Rödl. Pseudorandom sets and explicit constructions of ramsey graphs. *Quad Mat*, 13, 2004.
- [RSW06] Omer Reingold, Ronen Shaltiel, and Avi Wigderson. Extracting randomness via repeated condensing. *SIAM Journal on Computing*, 35(5):1185–1209, 2006.
- [RT00] Jaikumar Radhakrishnan and Amnon Ta-Shma. Bounds for dispersers extractors and depth-two super concentrators. *SIAM J. Discrete Math.*, 13(1):2–24, 2000.
- [TU12] Amnon Ta-Shma and Christopher Umans. Better condensers and new extractors from parvaresh–vardy codes. In *Computational Complexity (CCC), 2012 IEEE 27th Annual Conference on*, pages 309–315, 2012.
- [Vad12] Salil Vadhan. Pseudorandomness. *Foundations and Trends in Theoretical Computer Science*, 7(1–3):1–336, 2012.