Electronic Colloquium on Computational Complexity, Report No. 88 (2016)

Explicit two-source extractors for near-logarithmic min-entropy

Avraham Ben-Aroya* Dean Doron† Amnon Ta-Shma‡

Abstract

We explicitly construct extractors for two independent n-bit sources of $(\log n)^{1+o(1)}$ min-entropy. Previous constructions required either $\mathrm{polylog}(n)$ min-entropy [CZ15, Mek15] or five sources [Coh16].

Our result extends the breakthrough result of Chattopadhyay and Zuckerman [CZ15] and uses the non-malleable extractor of Cohen [Coh16]. The main new ingredient in our construction is a somewhere-random condenser with a small entropy gap, used as a sampler. We construct such somewhere-random condensers using the error reduction mechanism of Raz et al. [RRV99] together with the high-error, constant degree dispersers of Zuckerman [Zuc06].

*The Blavatnik School of Computer Science, Tel-Aviv University, Tel Aviv 69978, Israel. Supported by the Israel science Foundation grant no. 994/14 and by the United States – Israel Binational Science Foundation grant no. 2010120.

†The Blavatnik School of Computer Science, Tel-Aviv University, Tel Aviv 69978, Israel. Email: deandoron@mail.tau.ac.il. Supported by the Israel science Foundation grant no. 994/14 and by the United States – Israel Binational Science Foundation grant no. 2010120.

‡The Blavatnik School of Computer Science, Tel-Aviv University, Tel Aviv 69978, Israel. Email: amnon@tau.ac.il. Supported by the Israel science Foundation grant no. 994/14 and by the United States – Israel Binational Science Foundation grant no. 2010120.

ISSN 1433-8092

1 Introduction

A graph G is *K-Ramsey* if it contains no clique or independent set of size K. In one of the first applications of the probabilistic method, Erdös showed [Erd47] that there are $2\log N$-Ramsey graphs over N vertices. Erdös raised the challenge of giving an explicit description of such a graph. A related challenge is that of constructing a *K-bipartite Ramsey graph*, i.e., a bipartite graph with no bipartite clique or bipartite independent set of size K. Any explicit bipartite Ramsey graph can be translated into an explicit (non-bipartite) Ramsey graph with about the same parameters. Erdös' probabilistic argument also shows that there are $2\log N$-bipartite Ramsey graphs (where N is the number of vertices on each side) and the problem is constructing such graphs explicitly.

From a computer science point of view, bipartite Ramsey graphs are equivalent to two-source dispersers outputting one bit. More formally, a function $\mathsf{Disp} : [N] \times [N] \to \{0,1\}$ is a (zero-error) *two-source K-disperser* if for every two sets $A, B \subseteq [N]$ of cardinality at least K, $\mathsf{Disp}(A,B) = \{0,1\}$. Such a disperser gives rise to a bipartite K-Ramsey graph with N vertices on each side.

A stronger notion is that of two-source extractors. A function $\mathsf{Ext} : [N] \times [N] \to \{0,1\}$ is a *two-source k-extractor*[1] if for every two independent distributions A, B over $[N]$ with min-entropy at least k, the bit $\mathsf{Ext}(A,B)$ has a small bias. One can see that every two-source k-extractor with any nontrivial bias readily implies a two-source 2^k-disperser, and thus also a bipartite 2^k-Ramsey graph.

Early research [Abb72, Nag75, Fra77, Chu81] culminated in the construction of 2^k-Ramsey graphs over 2^n vertices, for $k \approx 2^{\sqrt{\log n}}$ [FW81] (see also [Nao92, Alo98, Gro01, Bar06] and [Gop06]). Explictly constructing good two-source extractors was evidently more challenging. The inner product function gives a simple (and powerful) solution when $k > n/2$ [CG88]. Bourgain [Bou05, Rao07] gave a two-source extractor construction for $k = \frac{1}{2} - \alpha$ for some small constant $\alpha > 0$. Raz [Raz05] constructed a two-source extractor that has an unbalanced entropy requirement; the first source should have more than $n/2$ min-entropy, while the second source's min-entropy can be as low as $c \cdot \log n$ (for some constant c).

In a different line of research [BKS+10, BRSW12], the challenge-response mechanism was used for the construction of K-Ramsey graphs for smaller K, culminating in explicitly constructing 2^k-Ramsey graphs for $k = \mathrm{polylog}(n)$ [Coh15].

Due to the difficulty of constructing good two-source extractors, another research line focused on extracting from multiple sources having low min-entropy, trying to minimize the number of sources needed. This includes [BIW06, Rao09, Li11, Li13a, Li13b], with the later papers using alternating extraction. Eventually, Chattopadhyay and Zuckerman [CZ15] used non-malleable extractors to give a two-source extractor for $k = \mathrm{polylog}(n)$. We note that the main tool in constructing non-malleable extractors is alternating extraction.

Several improvements on the [CZ15] construction followed, including [Mek15, Li15]. Cohen and Schulman [CS15] observed that all the above constructions assume poly-logarithmic min-entropy, and managed to get the first multi-source construction for $k = (\log n)^{1+o(1)}$. Chattopadhyay and Li [CL16] reduced the number of sources in such a construction to a constant and Cohen [Coh16] put it on five. The main result in this paper is such a construction with only two sources:

[1] Throughout the paper we use uppercase letters, as K, to denote the set's cardinality, and lowercase letters, as k, to denote the corresponding min-entropy ($K = 2^k$). Under this convention, if an extractor operates on n-bit sources, it corresponds to a graph over $N = 2^n$ vertices.

Theorem 1.1. *For every large enough n, there exists an explicit, constant-error, two-source extractor $2Ext : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ for min-entropy $(\log n)^{1+o(1)}$.*

This immediately implies an explicit construction of $2^{(\log n)^{1+o(1)}}$-bipartite Ramsey graphs with 2^n vertices on each side.

1.1 An overview of the construction

Let us first recall the Chattopadhyay and Zuckerman [CZ15] construction. We are given $x_1, x_2 \in [N]$ sampled from two independent distributions X_1, X_2 with min-entropy k_1, k_2, respectively. We take a t-non-malleable (k_2, ε_2) extractor $nmEXT : [N] \times [D] \to \{0,1\}$ and write a $D \times 1$ table NM, where the rows are indexed by seeds $y \in [D]$, and in row y we write $NM[y] = nmEXT(x_2, y)$. By the properties of non-malleable extractors, we can divide the rows to good and bad such that there are few bad rows and the good rows are close to being t-wise independent.

At this stage we would have liked to output $f(NM[1], \ldots, NM[D])$ for some resilient function f that is willing to accept few bad players, and good players that are only close to being t-wise independent. Of course, since no one-source extractor exists – there is no such a function f. Nevertheless, Chattopadhyay and Zuckerman explore why this approach fails.

First, let us check the number of bad players, or equivalently, the number of bad rows in the table. The number of rows is the number of possible seeds D, and the number of bad rows is bounded by $\varepsilon_2 D$. Since we use a non-malleable extractor (or even if we just take an extractor) we must have $D \geq (\frac{1}{\varepsilon})^2$ (or in the more familiar form: $d = \log D \geq 2\log(\frac{1}{\varepsilon})$) and therefore the number of bad rows is necessarily higher than $D^{\frac{1}{2}}$. However, if we take ε_2 to be a small enough polynomial in n, and if the non-malleable extractor has seed length dependence $O(\log \frac{n}{\varepsilon_2})$, then the number of bad players is $D^{1-\alpha}$ for some constant $\alpha > 0$, and there are resilient functions that are willing to accept many bad players, as long as the good players are truly uniform.

Let us say f is (q,t) resilient if it is resilient even when there are q bad players and the good players are only t-wise independent. Chattopadhyay and Zuckerman construct such a function for $q = R^{1-\alpha}$ bad players out of the R players and $t = \mathrm{polylog}(n)$; Viola [Vio14] showed that the majority function can handle $q = R^{\frac{1}{2}-\alpha}$ and $t = O(1)$.

Yet, there are no one-source extractors, since in the table NM the good players are only (t,γ) wise independent in the sense that every t good rows are γ-close to uniform. One may conclude ([AGM03], see Lemma 2.14) that NM is $D^t\gamma$-close to a distribution that is truly t-wise independent, but the cost $D^t\gamma$ is too large, in particular prohibiting any one source extractor.

Chattopadhyay and Zuckerman use the other source to bypass the problem. They use X_1 to sample rows from the table NM, i.e., they take a (k_1, ε_1) strong extractor $Ext : [N] \times [R] \to [D]$ and output

$$\begin{aligned} 2EXT(x_1, x_2) &= f(NM[Ext(x_1, 1)], \ldots, NM[Ext(x_1, R)]) \\ &= f(nmEXT(x_2, Ext(x_1, 1)), \ldots, nmEXT(x_2, Ext(x_1, R))). \end{aligned}$$

In other words, x_1 samples R rows from the table NM, and we apply the resilient function on the sample. Because extractors are good samplers, if k_1 is large enough, almost all x_1 sample well, and the fraction of the bad players in the sampled rows will be about $\sqrt{\varepsilon_2} + \varepsilon_1$ and each t good players are $t\sqrt{\varepsilon_2}$-close to uniform (the $\sqrt{\varepsilon_2}$ appears because of averaging). We take $\varepsilon_2 \ll \varepsilon_1$, so we can just think of $\sqrt{\varepsilon_2} + \varepsilon_1$ as ε_1 fraction of bad rows. If we take a small enough ε_1 and an

extractor with seed length $O(\log(\frac{n}{\varepsilon_1}))$ we again get that, with high probability, the sample contains at most $R^{1-\alpha}$ bad players out of the R players. Also, the good players are still close to being t-wise independent. The same argument works as before, and we get that the interesting factor in the bias of the resulting output bit is $R^t\gamma = tR^t\sqrt{\varepsilon_2}$. If we use only one source, R is a function of ε_2, thus making ε_2 smaller does not help. Now, however, with two sources, R is a function of ε_1 alone and ε_2 may be chosen after R is determined, so taking ε_2 small enough, this error term becomes small and the construction magically works!

The [CZ15] construction achieves $k_1 = k_2 = \text{polylog}(n)$. This is because, as noted earlier, there are at least $R^{\frac{1}{2}}$ bad players out of the R players. All the currently known (q, t) resilient functions that handle $q \geq R^{\frac{1}{2}}$ require t which is poly-logarithmic in R.

Cohen and Schulman [CS15] note that all previous explicit multi-source extractors (or dispersers) work with min-entropy at least $\log^2 n$. They were able to construct a multi-source extractor requiring only $(\log n)^{1+o(1)}$ min-entropy using a new primitive called independence-preserving merger. In a subsequent paper, Cohen [Coh16] shows a five-source extractor for min-entropy $(\log n)^{1+o(1)}$.[2] This construction uses Majority as the resilient function and, informally speaking, uses the other four sources to guarantee that the number of bad rows is at most $R^{0.4}$. The main advantage gained by using the majority function is that it is (q, t) resilient for some *constant* t.

The starting point of this paper is the observation that condensers with a small entropy gap (that we soon define) are good samplers and their dependence on ε can get as small as $1 \cdot \log(\frac{1}{\varepsilon})$. Let us first discuss the entropy gap notion.

Definition 1.2. *A function* $Cond : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ *is a* $(n, k) \to_\varepsilon (m, k')$ *condenser if for every* (n, k)*-source* X, $Cond(X, U_d)$ *is* ε*-close to a distribution with min-entropy* k'*. The* entropy loss *of the condenser is* $k + d - k'$ *and the* entropy gap *is* $m - k'$*.*

Dodis et al. [DPW14] observe that for entropy gap 1, non-explicit constructions achieve entropy loss which is only $\log\log(\frac{1}{\varepsilon}) + O(1)$ (compare with entropy loss $2\log(\frac{1}{\varepsilon})$ when there is no entropy gap). They show that condensers with small entropy gap are still good samplers when the test set is small (see Lemma 4.4). Dodis et al. use this property for key derivation without entropy waste. They show that in non-explicit constructions the seed length dependence on the error is $1 \cdot \log(\frac{1}{\varepsilon})$ (rather than the $2\log(\frac{1}{\varepsilon})$ in extractors).

As condensers with a small entropy gap are good samplers, and their seed length dependence on ε is $\log(\frac{1}{\varepsilon})$, it is not difficult to see that if one replaces the extractor in [CZ15] with such a condenser, and uses Majority, the construction works with $(\log n)^{1+o(1)}$ min-entropy.[3]

Unfortunately, we are not aware of an *explicit* construction achieving a small entropy gap and seed length less than $2\log(\frac{1}{\varepsilon})$. We remark that the GUV condenser [GUV09] has dependence $\log(\frac{1}{\varepsilon})$ on the error but a very large entropy gap.

We bypass this problem by constructing a *somewhere-random* condenser with a constant number of blocks, very small entropy gap, and dependence $(1 + \alpha)\log(\frac{1}{\varepsilon})$ on the error, for some constant $\alpha < 1$. The idea is to start with an extractor that has the wrong dependence on the error, and amplify its error to a very small ε, in an efficient way that gives dependence smaller than $2\log(\frac{1}{\varepsilon})$. We use the error reduction scheme suggested by Raz et al. [RRV99] that works by sampling (with a sampler) a constant number of seeds and outputting all the corresponding outputs of E. They

[2] In fact, four sources suffice for the construction, because the fourth source is redundant as the advice correlation breaker works also with a weak dense seed.

[3] The min-entropy becomes logarithmic when using non-malleable extractors with the "correct" seed length.

show the obtained structure is a somewhere-random source with a small entropy gap and very low error. The sampler used in [RRV99] is obtained by taking a random walk (of constant length) on an expander. The analysis shows that such a reduction has dependence at least $2\log(\frac{1}{\varepsilon})$ on the error. Instead, we observe that what is needed in the reduction is a disperser with a very large error (i.e., the image of any large enough set is not contained in any small set), and for the parameters we need, Zuckerman [Zuc06] already constructed such dispersers having a constant degree![4] Using his dispersers in the error reduction scheme, we get the desired somewhere random condensers.

Having that, we go back to the [CZ15] construction, and replace the extractor EXT with a somewhere-random condenser. As we use a somewhere-random condenser rather than a condenser, when x_1 samples the rows of NM, we get an $R \times A$ table (rather than an $R \times 1$ table as before) where A is the number of blocks and is a constant. The property of this table is that the number of bad rows is at most $R^{0.4}$, and the good rows are, informally speaking, t-wise somewhere random. Here we apply another trick from [Coh16]: We take the parity of each row and apply the resilient function on it. The result is an almost balanced bit and we are done.

2 Definitions and preliminaries

Throughout the paper we have the convention that lowercase variables are the logarithm (in base-2) of their corresponding uppercase variables, e.g., $n = \log N$, $d = \log D$, $a = \log A$, $r = \log R$, $r' = \log R'$, etc. The density of a set $B \subseteq [D]$ is $\rho(B) = \frac{|B|}{D}$.

2.1 Random variables, min-entropy

The *statistical distance* between two distributions X and Y on the same domain D is defined as $|X - Y| = \max_{A \subseteq D}(\Pr[X \in A] - \Pr[Y \in A])$. If $|X - Y| \leq \varepsilon$ we say that X is ε-close to Y and denote it by $X \approx_\varepsilon Y$. We will denote by U_n a random variable distributed uniformly over $\{0,1\}^n$ and which is independent of all other variables. We also say that a random variable is *flat* if it is uniform over its support.

For a function $f : D_1 \to D_2$ and a random variable X distributed over D_1, $f(X)$ is the random variable, distributed over D_2, which is obtained by choosing x according to X and computing $f(x)$. For a set $A \subseteq D_1$, we simply denote $f(A) = \{f(x) \mid x \in A\}$. It is well-known that for every $f : D_1 \to D_2$ and two random variables X and Y, distributed over D_1, it holds that $|f(X) - f(Y)| \leq |X - Y|$.

The *min-entropy* of a random variable X is defined by

$$H_\infty(X) = \min_{x \in \text{Supp}(X)} \log \frac{1}{\Pr[X = x]}.$$

For $\varepsilon \geq 0$, the *smooth* min-entropy $H^\varepsilon_\infty(X)$ is the supremum of $H_\infty(X')$ over all distributions X' such that $|X - X'| \leq \varepsilon$.

We will call a random variable X distributed over $\{0,1\}^n$ with min-entropy at least k an (n,k)-source. Every distribution X with $H_\infty(X) \geq k$ can be expressed as a convex combination of flat distributions, each with min-entropy at least k. We have the following easy claim:

Claim 2.1. *If $H^{1/2}_\infty(X) \geq k$ then the support of X is at least 2^{k-1}.*

[4] Usually, the degree has to be logarithmic, but for the special parameters we need, which are also the parameters needed for [Zuc06], the degree may be constant.

2.2 Somewhere-random sources

Definition 2.2 (somewhere-random source). *A source $X = X_1 \circ \ldots \circ X_A$ is a $(k, (\alpha, \beta))$ somewhere-random (s.r.) source if there is a random variable $I \in \{0, \ldots, A\}$ such that for every $i \in [A]$, $H^{\alpha}_{\infty}(X_i | I = i) \geq k$ and $\Pr[I = 0] \leq \beta$. If $\alpha = \beta = 0$ we say X is a k s.r. source.*

Claim 2.3. *Let X be a (k, α, β) s.r. source. Then, X is $(\alpha + \beta)$-close to a k s.r. source.*

Definition 2.4. *We say X is an (n, k, ζ) s.r. source if X is ζ-close to a k s.r. source over $\{0,1\}^n$.*

Intuitively, it is often convenient to think of a k s.r. source $X = X_1 \circ \ldots \circ X_A$ as if one of the blocks X_i is having k min-entropy, and the other blocks are arbitrarily correlated with it. Formally, X is a k s.r. source if it is a convex combination of such sources.

2.3 Dispersers, extractors and s.r. condensers

Definition 2.5 (disperser). *A function $\Gamma : [N] \times [D] \to [M]$ is a (k, ε) disperser if for every $A \subseteq [N]$ with $|A| \geq 2^k$, $\left|\bigcup_{i \in [D]} \Gamma(A, i)\right| \geq (1 - \varepsilon)M$.*

Definition 2.6 (extractor). *A function $Ext : [N] \times [D] \to [M]$ is a (k, ε) strong extractor if for every (n, k)-source X, and for Y that is uniform over $[D]$, $Y \circ Ext(X, Y) \approx_{\varepsilon} Y \times U$.*

Theorem 2.7 (The GUV extractor, [GUV09]). *There exists a universal constant $c_{GUV} > 0$ such that the following holds. For all positive integers n, k and $\varepsilon > 0$ there exists an efficiently-computable (k, ε) strong seeded extractor $Ext : [N] \times [D] \to [M]$ having seed length $d = c_{GUV} \log \frac{n}{\varepsilon}$ and $m = \frac{k}{2}$ output bits.*

Definition 2.8 (two-source extractor). *A function $2Ext : [N_1] \times [N_2] \to [M]$ is an $((n_1, k_1), (n_2, k_2), \varepsilon)$ two-source extractor if for every two independent sources X_1 and X_2 where X_1 is an (n_1, k_1)-source and X_2 is an (n_2, k_2)-source, it holds that $2Ext(X_1, X_2) \approx_{\varepsilon} U_m$.*

Definition 2.9 (s.r. condenser, condenser). *A function $SRC : [N] \times [D] \times [A] \to [M]$ is a $(k \to k', \zeta)$ s.r. condenser if for every (n, k)-source X it holds that $SRC(X, U) = SRC(X, U, 1) \circ \ldots \circ SRC(X, U, A)$ is ζ-close to a k' s.r. source. D is the* degree *of the s.r. condenser, and A is its number of blocks. If $D = 1$ (i.e., $d = 0$) we say SRC is seedless. A condenser is a s.r. condenser with just one block.*

2.4 Extremely large error s.r condensers and dispersers

A s.r. condenser implies a disperser with a large error (see, e.g., [Zuc06]). Concretely,

Lemma 2.10. *Suppose $SRC : [N] \times [D] \times [A] \to [M]$ is a $(k \to k', \zeta = \frac{1}{2})$ s.r. condenser. Define $\Gamma : [N] \times [D \cdot A] \to [M]$ by $\Gamma(x; (y, a)) = SRC(x, y, a)$. Then, Γ is a $(k, 1 - \alpha)$ disperser for $\alpha = 2^{k' - 1 - m}$.*

Proof: Let $B \subseteq [N]$ be an arbitrary set such that $|B| \geq 2^k$. Then, $SRC(B, U_d) = SRC(B, U_d, 1) \circ \ldots \circ SRC(B, U_d, A)$ is $(k', \zeta = \frac{1}{2})$ s.r., with some indicator random variable I. Pick any index $i \neq 0$ in the support of I. Then, conditioned on $I = i$, we have that $C(B, U_d, i)$ is $1/2$-close to a k' source. Thus, by Claim 2.1, conditioned on $I = i$, $C(B, U_d, i)$ covers at least $2^{k' - 1}$ vertices from $[M]$. But then, even without the conditioning, $C(B, U_d, i)$ covers at least $2^{k' - 1}$ vertices from $[M]$. Therefore, $|\Gamma(B, [D \cdot A])| \geq 2^{k' - 1} = \alpha M$, as required. ∎

Zuckerman [Zuc06], using additive number theory and extending earlier results, showed:

Theorem 2.11 ([Zuc06], Theorem 8.3). *There exist two constants $0 < c_1 < c_2 < 1$, a constant $\gamma > 0$ and a $(c_1 n \to c_2 m, N^{-\gamma})$ s.r. seedless condenser with just two blocks*

$$SRC : [N] \times [1] \times [A = 2] \to [M = N^{2/3}].$$

With that Zuckerman constructs dispersers with very large error, but constant degree. Specifically,

Theorem 2.12 ([Zuc06], Theorem 1.9). *For all constant $c_1, c_2 > 0$ and $\zeta = \zeta(n) > 0$ there exists an efficient family of $(k = c_1 n, 1 - \zeta)$ dispersers*

$$\Gamma : [N] \times [D] \to [M = K^{1-c_2}]$$

with degree $D = O(\frac{n}{\log \frac{1}{\zeta}})$.

2.5 Limited independence and non-oblivious bit-fixing sources

Definition 2.13. *A distribution X over Σ^n is called (t, γ)-wise independent if the restriction of X to every t coordinates is γ-close to U_{Σ^t}.*

Alon et al. proved:

Lemma 2.14 ([AGM03]). *Let X be a distribution over $\{0,1\}^n$ that is (t, γ)-wise independent. Then, X is $(n^t \gamma)$-close to a t-wise independent distribution.*

Definition 2.15. *A source X over $\{0,1\}^n$ is called a (q, t, γ) non-oblivious bit-fixing source if there exists a subset $Q \subseteq [n]$ of size at most q such that the joint distribution of the bits in $[n] \setminus Q$ is (t, γ)-wise independent. The bits in Q are allowed to arbitrarily depend on the bits in $[n] \setminus Q$.*

Definition 2.16. *Let $f : \{0,1\}^n \to \{0,1\}$, $\mathcal{D}$ a distribution over $\{0,1\}^n$ and $Q \subseteq [n]$. Let $I_{Q,\mathcal{D}}(f)$ denote the probability that f is undetermined when the variables outside Q are sampled from $\mathcal{D}$. We define $I_{q,t,\gamma}(f)$ to be the maximum of $I_{Q,\mathcal{D}}(f)$ over all $Q \subseteq [n]$ of size q and all $\mathcal{D}$ that is a (t, γ) independent distribution.*

We say that f is (t, γ)-independent (q, ε)-resilient if $I_{q,t,\gamma}(f) \leq \varepsilon$.

Resilient functions are exactly deterministic extractors against non-oblivious bit-fixing sources outputting one bit. Chattopadhyay and Zuckerman [CZ15] derandomized the Ajtai-Linial function [AL93] and constructed a (monotone) resilient function that handles $q = n^{1-\alpha}$ for any constant α. Their construction was later improved in [Mek15]. Following [CS15, Coh16] we work hard to be able to use the majority function as the resilient function. The work of Viola [Vio14] shows that for every $\alpha > 0$, the majority function on n bits is $(t, 0)$-independent $(n^{1/2-\alpha}, O(\frac{\log t}{t} + n^{-\alpha}))$-resilient. Combining this with Lemma 2.14, we conclude:

Lemma 2.17. *There exists a constant c_{MAJ} such that for every $\alpha > 0$ and a $(q = n^{\frac{1}{2}-\alpha}, t, \gamma)$ non-oblivious bit-fixing source X on n bits,*

$$\left| \Pr[\mathrm{Maj}(X_1, \ldots, X_n) = 1] - \frac{1}{2} \right| \leq c_{MAJ} \cdot \left(\frac{\log t}{t} + n^{-\alpha} + \gamma n^t \right).$$

2.6 Non-malleable extractors

Definition 2.18. *A function $nmExt : [N] \times [D] \to [M]$ is a (k, ε) t-non-malleable extractor, if for every (n, k)-source X, Y that is uniform over $[D]$ and every deterministic functions $f_1, \ldots, f_t : [D] \to [D]$ with no fixed-point it holds that:*

$$(nmExt(X, Y), nmExt(X, f_1(Y)), \ldots, nmExt(X, f_t(Y), Y) \approx_\varepsilon \\ (U_m, nmExt(X, f_1(Y)), \ldots, nmExt(X, f_t(Y), Y).$$

Cohen [Coh16] proved:

Theorem 2.19 ([Coh16], Theorem 12.1)**.** *There exist constants $c_1, c_2 \geq 1$ such that the following holds. Given n and $\varepsilon > 0$, set $d = c_1 \log n + \log(1/\varepsilon) \cdot c_1^{\sqrt{\log\log(1/\varepsilon)}}$. For every n and $\varepsilon > 0$ there exists an efficiently-computable $(k = c_2 d, \varepsilon)$ 1–non-malleable extractor $nmExt : [N] \times [D] \to [K^{1/4}]$.*

The construction of Cohen's non-malleable extractor heavily relies on his construction of a *correlation-breaker with advice* – an object we shall not introduce here. Using a t-correlation breaker with advice gives rise to a non-malleable extractor with t adversarial functions. Taking t to be a constant and outputting only one bit, we obtain the corollary we shall require:

Corollary 2.20. *For any constant t there exist constants $c_1, c_2 \geq 1$ such that the following holds. For any integer n and for any $\varepsilon > 0$, there exists an efficiently-computable $(k = c_2 d, \varepsilon)$ t-non-malleable extractor $nmExt : [N] \times [D] \to \{0,1\}$ with seed length $d = c_1 \log n + \log(1/\varepsilon) \cdot c_1^{\sqrt{\log\log(1/\varepsilon)}} = O(\log n) + (\log(1/\varepsilon))^{1+o(1)}$.*

We will need the following lemma concerning the existence of a set of good seeds of a non-malleable extractor, given in [CZ15].

Lemma 2.21 ([CZ15], Lemma 3.5)**.** *Let $nmExt : [N] \times [D] \to \{0,1\}$ be a (k, ε) t–non-malleable extractor. Let X be any (n, k)-source. There exists a subset $R \subseteq [D]$, $|R| \geq (1 - \sqrt{\varepsilon})D$ such that for any distinct $r_1, \ldots, r_t \in R$,*

$$(nmExt(X, r_1), \ldots, nmExt(X, r_t)) \approx_{5t\sqrt{\varepsilon}} U_t.$$

Specifically, $R = [D] \setminus BAD$ where BAD is the set defined by

$$\begin{aligned} BAD \quad = \quad & \{r \in [D] \mid \exists \text{ distinct } r_1, \ldots, r_t \in [D], \forall i \in [t]\ r_i \neq r,\ |(nmExt(X, r), \\ & nmExt(X, r_1), \ldots, nmExt(X, r_t)) - (U_1, nmExt(X, r_1), \ldots, nmExt(X, r_t))| > \sqrt{\varepsilon}\}. \end{aligned}$$

We will also use the following lemma, which is a simple generalization of a one given in [CZ15].

Lemma 2.22. *Let $X_1, \ldots, X_t, Y_1, \ldots, Y_k$ be Boolean random variables. Further suppose that for any $i \in [t]$,*

$$\left(X_i, \{X_j\}_{j \neq i}, Y_1, \ldots, Y_k\right) \approx_\varepsilon \left(U_1, \{X_j\}_{j \neq i}, Y_1, \ldots, Y_k\right).$$

Then, $(X_1, \ldots, X_t, Y_1, \ldots, Y_k) \approx_{5t\varepsilon} (U_t, Y_1, \ldots, Y_k)$.

3 From large-error dispersers to low-error s.r. condensers

In this section we construct a (very) low-error s.r. condenser with seed length that is almost optimal in the small error. Specifically, the dependence of the seed length on the error is $(1+\alpha)\cdot\log(\frac{1}{\varepsilon})$ for arbitrarily small $\alpha > 0$. Furthermore, the entropy gap of the s.r. condenser is small ($O(\log(\frac{1}{\varepsilon}))$).

Theorem 3.1. *For every constant $0 < \alpha < 1$, there exists a constant A such that for every $m = m(n) \leq n$ and $0 < \varepsilon = \varepsilon(n) \leq (\frac{1}{n})^{4c_{GUV}}$ there exists an explicit $SRC : [N] \times [R'] \times [A] \to [M]$ that is a $(k = 2m + \log(\frac{1}{\varepsilon}) \to m - 2\log(\frac{1}{\varepsilon}) - O(a), \varepsilon)$ s.r. condenser with $r' = (1+\alpha)\log(\frac{1}{\varepsilon})$.*

Notice that the s.r. condenser achieves the small error ε using only a constant number of blocks, a small entropy gap (i.e., the min-entropy in the s.r. source is close to the block length) and a seed length close to $\log(\frac{1}{\varepsilon})$.

We first give an intuitive sketch of the proof. We start with an extractor $Ext : [N] \times [R] \to [M]$ that has a relatively small seed length (say, $0.5\log(\frac{1}{\varepsilon})$) and "too high" error $\varepsilon_0 = \varepsilon^{1/c}$ for some constant c. Such extractors exist, e.g., by [GUV09]. We use the [RRV99] framework to reduce the error to ε. [RRV99] argue that when we fix some high-entropy input distribution X there are two possible reasons for an error:

- Bad inputs $x \in \mathrm{Supp}(X)$ for which many seeds $y \in [R]$ fall to high-weight elements (where high-weight is defined with respect to ε), and,
- Inputs $x \in \mathrm{Supp}(X)$ that fall to high-weight elements with about the right frequency that is expected for error ε_0.

It is easy to deal with the first case: Just start with a distribution with $\log(\frac{1}{\varepsilon})$ additional min-entropy, and then the bad elements are "drawn" among all the elements in X . One way of dealing with the second case is using $O(1)$ independent seeds and arguing that one of them is good except for error ε. Indeed, [RRV99] prove the resulting distribution is a s.r. source, and then use a merger to merge it. Further, [RRV99] show that independent samples are not necessary, and any good sampler with a constant size sample set would do. [RRV99] use a short random walk on expanders. They then merge the resulting s.r. source, and get an extractor.

We are interested in getting a s.r. condenser (because we want to avoid the extra seed length required by the merger) and so we skip the merging step. Yet, the [RRV99] reduction with the expander random walk is too expansive for us, because it necessarily requires seed length that is at least $2\log(\frac{1}{\varepsilon})$. The key observation is that our construction requires dispersers with error very close to 1 (or put differently: the tests are extremely small) and using the low-degree, large-error dispersers constructed by Zuckerman [Zuc06] gives a much better dependence on the error. We now proceed with a rigorous proof.

Proof: For the proof we use the following ingredients:

- We use the GUV extractor with $\varepsilon_0 = \varepsilon^{\frac{1}{4c_{GUV}}}$ and notice that $\varepsilon_0 \leq \frac{1}{n}$. Specifically,

  $$Ext : [N] \times [R] \to [M]$$

  is an explicit $(2m, \varepsilon_0)$ strong extractor with $r = c_{GUV}\log\frac{n}{\varepsilon_0} \leq 2c_{GUV}\log\frac{1}{\varepsilon_0} = \frac{1}{2}\log(\frac{1}{\varepsilon})$. See Theorem 2.7.

- Set $r' = (1+\alpha)\log(\frac{1}{\varepsilon})$ and take

$$\Gamma : [R'] \times [A] \to [R = (R')^{\frac{1}{4}}]$$

to be the $\left(\frac{\alpha}{1+\alpha}r', \delta = 1 - 2\varepsilon_0\right)$ disperser guaranteed by Theorem 2.12, when plugging-in $c_1 = \frac{\alpha}{1+\alpha}$, $\zeta = 2\varepsilon_0$ and $c_2 = \frac{1}{2}$. Notice that $K^{\frac{1}{2}} = (R')^{\frac{1}{2}\frac{\alpha}{1+\alpha}} \geq (R')^{\frac{1}{4}}$. The degree is then $A = O\left(\frac{r'}{\log\frac{1}{2\varepsilon_0}}\right) = O\left(\frac{(1+\alpha)\log(\frac{1}{\varepsilon})}{\frac{1}{4c_{GUV}}\log(\frac{1}{\varepsilon})}\right) = O(1)$.

We define $SRC : [N] \times [R'] \times [A] \to [M]$ by:

$$SRC(x, y', z) \;=\; Ext(x, \Gamma(y', z)).$$

To prove correctness, assume X is an $(n, k = 2m + \log\frac{1}{\varepsilon})$-source. W.l.o.g., X is flat (because otherwise it is a convex combination of such sources). Let $\Delta = \frac{4A}{\varepsilon}$. We define the set of Δ-*heavy elements* $w \in [M]$ by:

$$\mathbf{H} \;=\; \left\{ w \in [M] \mid \Pr[Ext(X, U_r) = w] \geq \frac{\Delta}{M} \right\}.$$

We claim:

Claim 3.2. $|\mathbf{H}| < \frac{2\varepsilon_0}{\Delta} M$.

Proof: Notice that $|\mathbf{H}| \cdot \frac{\Delta}{M} \leq \Pr[Ext(X, U_r) \in \mathbf{H}] \leq \frac{|\mathbf{H}|}{M} + \varepsilon_0$, where the upper bound follows because $Ext(X, U_r)$ is ε_0-close to uniform. Thus $\frac{|\mathbf{H}|}{M}(\Delta - 1) \leq \varepsilon_0$ and $|\mathbf{H}| \leq \frac{\varepsilon_0}{\Delta - 1} M < \frac{2\varepsilon_0}{\Delta} M$. ∎

Next we define the set of bad inputs $x \in \mathrm{Supp}(X)$ by:

$$BadX \;=\; \left\{ x \in [N] \mid \Pr_{y \sim U_r}[Ext(x, y) \in \mathbf{H}] \geq \left(1 + \frac{2}{\Delta}\right)\varepsilon_0 \right\}.$$

Claim 3.3. $|BadX| < 2^{2m}$.

Proof: Suppose $|BadX| \geq 2^{2m}$. Let B be uniformly distributed over $BadX$. Then $Ext(B, U_r)$ is ε_0-close to uniform, and therefore

$$\left(1 + \frac{2}{\Delta}\right)\varepsilon_0 \;\leq\; \Pr_{x \sim B, y \sim U_r}[Ext(x, y) \in \mathbf{H}] \;\leq\; \frac{|\mathbf{H}|}{M} + \varepsilon_0.$$

But then $\frac{2\varepsilon_0}{\Delta} \leq \frac{|\mathbf{H}|}{M}$ – a contradiction to the previous claim. ∎

Let I be the following random variable: For $x \in [N]$ and $y' \in [R']$, $I(x, y')$ is an arbitrary $z \in [A]$ such that $Ext(x, \Gamma(y', z)) \notin \mathbf{H}$ if such a z exists, and 0 otherwise. Let I' be the same as I except that all z with $\Pr[I = z] \leq \frac{4}{\Delta}$ will be declared zero in I'.

Claim 3.4.

- $\Pr[I = 0] \leq 2\varepsilon$.

- $\Pr[I' = 0] \leq 3\varepsilon$.
- *For every* $z \in [A]$, $H_\infty(Ext(X, \Gamma(U_{r'}, I'))|I' = z) \geq m - 2\log\Delta + 2$.

Proof: For the first item, $\Pr[I = 0] \leq \Pr[X \in BadX] + \Pr[I = 0|X \notin BadX]$.

- Clearly, $\Pr[X \in BadX] \leq |BadX| \cdot 2^{-(2m+\log(\frac{1}{\varepsilon}))} \leq \varepsilon$.
- Fix an element $x \notin BadX$ and call $y \in [R]$ *bad for* x if $Ext(x, y) \in \mathbf{H}$. As $x \notin BadX$, the number of seeds y that are bad for x is at most $(1 + \frac{2}{\Delta})\varepsilon_0 R \leq 2\varepsilon_0 R$. Notice that $(I = 0|X = x)$ if some $y' \in [R']$ was chosen such that $\Gamma(y', 1), \ldots, \Gamma(y', A)$ are all bad for x. Since Γ is a $(\frac{\alpha}{1+\alpha}r', 1 - 2\varepsilon_0)$ disperser (i.e., Γ is a disperser for very large error), the number of such y'-s is at most $(R')^{\frac{\alpha}{1+\alpha}}$. Hence, $\Pr[I = 0|X \notin BadX] \leq \frac{(R')^{\frac{\alpha}{1+\alpha}}}{R'} = (R')^{-(1-\frac{\alpha}{1+\alpha})} = (\varepsilon^{1+\alpha})^{\frac{1}{1+\alpha}} = \varepsilon$, as desired.

For the second item, $\Pr[I' = 0] \leq \Pr[I = 0] + \frac{4A}{\Delta} = 3\varepsilon$.
For the third item, let w be in the support of $Ext(X, \Gamma(U_{r'}, z))$. Then,

$$\Pr[Ext(X, \Gamma(U_{r'}, z)) = w|I' = z] \quad \leq \quad \frac{\Delta}{M} \cdot \frac{1}{\Pr[I' = z]}.$$

Thus, and for every $z \in [A]$ in the support of I',

$$H_\infty(Ext(X, \Gamma(U_{r'}, I')) \mid I' = z) \geq m - 2\log\Delta + 2,$$

concluding our proof. ■

■

4 Form s.r. condensers to s.r. samplers

Extractors are good samplers, in the sense that if E is a (k, ε) extractor, then for every test S, we have that $\Pr(E(X, U) \in S)$ deviates from the density of S by at most ε. However, extractors are quite limited in the parameters they can achieve and in particular require seed length that is at least $2\log(\frac{1}{\varepsilon})$ [RTS00]. Dodis, Pietrzak and Wichs [DPW14] observed that if we are only interested in fooling sparse tests, it suffices to use condensers with a small entropy gap. Moreover, Dodis et al. note that such condensers can bypass the severe limitations that confine extractors. In this section, we show that the Dodis et al. result carries over to s.r. condensers and we obtain explicit s.r. samplers with the parameters we need, and in particular seed length that is smaller than $2\log(\frac{1}{\varepsilon})$.

Definition 4.1 (sampler)**.** *Fix* $S : [N] \times [R'] \to [D]$.

- *We say* $x \in [N]$ *is* (c, ε) *bad for* $B \subseteq [D]$ *if* $\Pr_{y' \in [R']}[S(x, y') \in B] > c\rho(B) + \varepsilon$.
- *We say* S *is a* $(K; c, \varepsilon)$ *sampler if for every* $B \subseteq [D]$,

$$|\{x \in [N] \mid x \text{ is } (c, \varepsilon) \text{ bad for } B\}| < K.$$

Definition 4.2 (s.r. sampler)**.** *Fix* $S : [N] \times [R'] \times [A] \to [D]$.

- *We say $x \in [N]$ is (c, ε) bad for $B \subseteq [D]$ if* $\Pr_{y \in [R']}[\forall_{z \in A} S(x, y', z) \in B] > c\rho(B) + \varepsilon$.
- *We say S is a $(K; c, \varepsilon)$ s.r. sampler if for every $B \subseteq [D]$,*

$$|\{x \in [N] \mid x \text{ is } (c, \varepsilon) \text{ bad for } B\}| < K.$$

Definition 4.3. *We say $S : [N] \times [R'] \times [A] \to [D]$ is* simple *if for every $x \in [N]$, and every $y'_1, y'_2 \in [R']$, $z_1, z_2 \in [A]$, if $(y'_1, z_1) \neq (y'_2, z_2)$ then $S(x, y'_1, z_1) \neq S(x, y'_2, z_2)$.*

The following lemma is based on [DPW14].

Lemma 4.4. *If Z is a $(d, d-g)$-source then for every set $S \subseteq [D]$,*

$$\Pr[Z \in S] \leq 2^g \cdot \rho(S)$$

Proof: If Z is flat, then the probability that Z is in S is bounded by the density of S inside the support of Z, i.e., it is at most $\frac{|S|}{\text{Supp}(Z)|} = \frac{|S|}{2^{d-g}} = 2^g \cdot \rho(S)$. Since every $(d, d-g)$-source is a convex combination of such flat sources, the lemma follows. ■

Lemma 4.5. *If $X = X_1 \circ \ldots \circ X_A$ is a $(d, d-g, \zeta)$-s.r. source then for every set $B \subseteq [D]$,*

$$\Pr_{x \sim X}[\forall_{z \in [A]} X_z \in B] \leq 2^g \cdot \rho(B) + \zeta$$

Proof: X is ζ-close to a $(d, d-g)$ s.r. source X'. Let I be an indicator of X'.

$$\begin{aligned}
\Pr[\forall_{z \in [A]} X'_z \in B] &\leq \sum_{i=1}^{A} \Pr[I = i] \cdot \Pr_{x \sim X'}[\forall_{z \in [A]} X'_z \in B \mid I = i] \\
&\leq \sum_{i=1}^{A} \Pr[I = i] \cdot \Pr[X'_z \in B \mid I = z] \\
&\leq \sum_{i=1}^{A} \Pr[I = i] \cdot 2^g \cdot \rho(B) \leq 2^g \cdot \rho(B),
\end{aligned}$$

where the third inequality follows from Lemma 4.4. ■

Theorem 4.6. *If $C : [N] \times [R'] \times [A] \to [D]$ is a $(k \to d-g, \varepsilon)$ s.r. condenser then C is a $(2^k; 2^g, \varepsilon)$ s.r. sampler.*

Proof: Let $B \subseteq [D]$, and let BAD denote the set of elements in $[N]$ that are $(2^g, \varepsilon)$ bad for B. If $|BAD| \geq K$, then $C(BAD, U)$ is a $(d, d-g, \varepsilon)$ s.r. source. By Lemma 4.5,

$$\Pr_{x \in BAD, y' \in [R']}[\forall_{z \in [A]} C(x, y')_z \in B] \leq 2^g \cdot \rho(B) + \varepsilon.$$

Therefore, there must exist at least one $x \in BAD$ such that $\Pr_{y' \in [R']}[\forall_{z \in [A]} C(x, y')_z \in B] \leq 2^g \cdot \rho(B) + \varepsilon$, in contradiction to the definition of BAD. Thus, $|BAD| < K$, as required. ■

We now instantiate Theorem 4.6 with the s.r. condenser from Theorem 3.1 to obtain:

Theorem 4.7. *For every constant $0 < \alpha < 1$ and $d = d(n) \leq n$ there exist constants $A, c_{mult}, c_\varepsilon$ and an explicit*

$$S : [N] \times [R'] \times [A] \to [D]$$

that is a $(K = D^2; c \leq \frac{1}{\varepsilon^{c_{mult}}}, \varepsilon = (\frac{1}{n})^{c_\varepsilon})$ simple s.r. sampler with $r' = (1+\alpha)\log(\frac{1}{\varepsilon})$.

Proof: We are given α. Set A to be the constant $A = A(\alpha)$ given in Theorem 3.1. Set $c_\varepsilon = 4c_{GUV}$, $\varepsilon = (\frac{1}{n})^{c_\varepsilon}$ and let $r' = (1+\alpha)\log(\frac{1}{\varepsilon})$. Given d, set $m = d - A - r'$. Let

$$SRC : [N] \times [R'] \times [A] \to [M]$$

denote the $(2m + \log(\frac{1}{\varepsilon}) \to m - 2\log(\frac{1}{\varepsilon}) - O(\log A), \varepsilon)$ s.r. condenser from Theorem 3.1. Define a new condenser

$$S : [N] \times [R'] \times [A] \to [R'] \times [A] \times [M]$$

by

$$S(x, y', z) = (y', z, SRC(x, y', z)).$$

It is immediate that S is simple. Notice that $M \cdot R' \cdot A = D$. By Theorem 4.6, S is a $(2^{2m+\log(\frac{1}{\varepsilon})}; c = 2^{m+r'+a-(m-2\log(\frac{1}{\varepsilon})-O(a))}, \varepsilon)$ s.r sampler, and:

- $2^{2m+\log(\frac{1}{\varepsilon})} = \frac{M^2}{\varepsilon} \leq (MR')^2 \leq D^2 = K$.
- $c = 2^{m+r'+a-(m-2\log(\frac{1}{\varepsilon})-O(a))} \leq 2^{r'+2\log(\frac{1}{\varepsilon})+O(a)} = 2^{(3+\alpha)\log(\frac{1}{\varepsilon})+O(a)} \leq \varepsilon^{-c_{mult}}$ for some constant c_{mult} (for example, one may take $c_{mult} = 4$).

■

5 From s.r. samplers to two-source extractors: Extending the CZ approach

In this section we prove:

Theorem 5.1 (Theorem 1.1 restated)**.** *For every* constant $\varepsilon > 0$ *there exists a constant c such that for every large enough integer n, there exists an explicit $((n, k_1), (n, k_2), \varepsilon)$ two-source extractor $2Ext : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ for any $k_1, k_2 \geq c \cdot \log(1/n) \cdot c^{\sqrt{\log\log n}}$.*

The proof closely follows the intuition given in the introduction, but makes it rigorous and as a result many parameters enter the discussion. We encourage the reader to read Section 1.1 before reading this section.

Proof: We are given n and a constant ε.

- Set t large enough so that $c_{MAJ}\frac{\log t}{t} \leq \frac{\varepsilon}{6}$, where c_{MAJ} is the constant from Lemma 2.17. A calculation shows that $t = \lceil \frac{36c_{MAJ}^2}{\varepsilon^2} \rceil$ suffices, so $t = O(1)$.
- Fix $\zeta = (\frac{1}{n})^{c_\varepsilon}$, $c = (\frac{1}{\zeta})^{c_{mult}}$, $R' = (\frac{1}{2\zeta})^{5/3}$, where c_ε, c_{mult} are the constants from Theorem 4.7.

- Set ε_2 small enough so that $5t\sqrt{\varepsilon_2}(R')^t \leq \frac{\varepsilon}{6}$ and $c\sqrt{\varepsilon_2} \leq \frac{1}{2}(R')^{-0.6}$. Doing the calculation one may take $\varepsilon_2 = \frac{\varepsilon^2}{30^2 t^2 c^2 (R')^{2t}}$ and so $\varepsilon_2^{-1} = (\frac{1}{cR'})^{O(1)} = n^{O(1)}$.
- Set $d = c_1 \log n + \log(1/\varepsilon_2) \cdot c_1^{\sqrt{\log\log(1/\varepsilon_2)}}$ where c_1 is the constant from Theorem 2.19.
- Define a function
$$S : [N] \times [R'] \times [A] \to [D],$$
which is a $(K = D^{c_{loss}}; c, \zeta)$ s.r. simple sampler guaranteed by Theorem 4.7 where N, R', D, c and ζ were defined before, c_{loss} is the constant 2 from Theorem 4.7, and $A = O(1)$ is the constant number of blocks guaranteed by the theorem.
- Define a function
$$nmExt : [N] \times [D] \to \{0,1\}^{},$$
which is a $(t' = tA, k_2 = c_2 d, \varepsilon_2)$ non-malleable extractor, where c_2 is the constant from Theorem 2.21. Notice that d was chosen to be sufficiently large for $nmExt$.

After fixing the above, given $x_1, x_2 \in [N]$, the construction is as follows:

1. For every $y' \in [R']$ and $z \in [A]$, compute $NM(x_1, x_2; y', z) = nmExt(x_2, S(x_1, y', z))$.
2. For every $y' \in [R']$, compute $\oplus NM(x_1, x_2; y') = \bigoplus_{z=1}^{A} NM(x_1, x_2; y', z)$.
3. Output $2Ext(x_1, x_2) = \text{Maj}(\oplus NM(x_1, x_2; 1), \ldots, \oplus NM(x_1, x_2; R'))$.

We now prove correctness. Let X_1 be an (n, k_1)-source for $k_1 = k + \log\frac{2}{\varepsilon}$ and X_2 be an (n, k_2)-source independent from X_1. Let $BAD \subseteq [D]$ be the set of density at most $\sqrt{\varepsilon_2}$ guaranteed to us by Lemma 2.21. Note that BAD depends only on X_2. We say $x_1 \in [N]$ is *bad* if x is (c, ζ) bad for BAD (see Definition 4.2) and good otherwise.

Claim 5.2. $\Pr_{x_1 \sim X_1}[x_1 \text{ is bad}] \leq \frac{\varepsilon}{2}$.

Proof: The number of bad elements is at most K, and $H_\infty(X_1) \geq k + \log\frac{2}{\varepsilon}$ so we can conclude that $\Pr_{x_1 \sim X_1}[x_1 \text{ is bad}] \leq \frac{K}{2^{k_1}} = \frac{\varepsilon}{2}$. ■

Now, fix any good $x_1 \in \text{Supp}(X_1)$. Since x_1 is good,

$$\Pr_{y' \in [R']}[\forall_{z \in A}\, S(x, y', z) \in BAD] \;\leq\; c \cdot \rho(BAD) + \zeta.$$

Call $y' \in [R']$ *a bad row* if $\forall_{z \in A}\, S(x, y', z) \in BAD$.

Lemma 5.3. *For every good $x_1 \in \text{Supp}(X_1)$,*

$$\oplus NM(x_1, X_2) = (\oplus NM(x_1, X_2, 1), \ldots, \oplus NM(x_1, X_2, R'))$$

is a (q, t, γ) non-oblivious bit-fixing source for $q = (R')^{0.4}$ and $\gamma = 5t\sqrt{\varepsilon}_2$.

Proof: First, the number of bad rows is at most $(c\rho(BAD)+\zeta)R' \leq q = (R')^{0.4}$. Next, fix t distinct good rows $y'_1,\ldots,y'_t$. Let $z_1,\ldots,z_t \in [A]$ be s.t. $S(x_1,y'_i,z_i) \notin BAD$. Then, for every $i \in [t]$,

$$\left(NM(x_1,X_2;y'_i,z_i), \left\{NM(x_1,X_2;y'_i,z)\right\}_{z\neq z_i}, \left\{NM(x_1,X_2;y'_j,z)\right\}_{j\neq i,z\in[A]}\right) \approx_{\sqrt{\varepsilon_2}} \left(U_1, \left\{NM(x_1,X_2;y'_i,z)\right\}_{z\neq z_i}, \left\{NM(x_1,X_2;y'_j,z)\right\}_{j\neq i,z\in[A]}\right),$$

where we have used Lemma 2.21 and the fact that S is simple. By Lemma 2.22,

$$\left(NM(x_1,X_2;y'_1,z_1),\ldots,NM(x_1,X_2;y'_t,z_t), \left\{NM(x_1,X_2;y'_i,z)\right\}_{(y'_i,z)\notin\left\{(y'_1,z_1),\ldots,(y'_t,z_t)\right\}}\right) \approx_{5t\sqrt{\varepsilon_2}} \left(U_t, \left\{NM(x_1,X_2;y'_i,z)\right\}_{(y'_i,z)\notin\left\{(y'_1,z_1),\ldots,(y'_t,z_t)\right\}}\right).$$

This gives $(\oplus NM(x_1,X_2,y'_1),\ldots,\oplus NM(x_1,X_2,y'_t)) \approx_{5t\sqrt{\varepsilon_2}} U_t$, as desired. ■

Therefore, by Lemma 2.17, for any good x_1,

$$\begin{aligned}\left|\Pr[\mathrm{Maj}(\oplus NM(x_1,X_2,1),\ldots,\oplus NM(x_1,X_2,R'))=1]-\frac{1}{2}\right| &\leq c_{MAJ}\left(\frac{\log t}{t}+(R')^{-0.1}+5t\sqrt{\varepsilon}_2(R')^t\right)\\ &\leq 3\cdot\frac{\varepsilon}{6} = \frac{\varepsilon}{2},\end{aligned}$$

where the probability is over X_2. Overall, we have:

$$|2Ext(X_1,X_2)-U_1| \leq \frac{\varepsilon}{2}+\frac{\varepsilon}{2} \leq \varepsilon,$$

as desired. The requested entropies are $k_1,k_2 = O(d)$. The explicitness follows from the fact that $R' = \mathrm{poly}(n)$ and the explicitness of the other ingredients. ■

We conclude with two remarks:

1. Any improvement in the seed length of explicit non-malleable extractors would reflect in the construction, and, in particular, seed length $O(\log n + \log(\frac{1}{\varepsilon}))$ would give two-source extractors for logarithmic min-entropy.
2. It is plausible that the techniques of Li [Li15] can be applied here as well to extract more bits. However, since we require $\sqrt{\varepsilon_2}(R')^t < 1$ and $R' = \mathrm{poly}(n)$, and since $\varepsilon_2 = 2^{-O(k_2)}$ it follows that $t\log R = O(k_2)$. Since the good rows are only guaranteed to be t-wise independent, we can extract at most $t = O(\frac{k_2}{\log n})$ bits from the table. This means that if we ever get to employ the scheme for $k_2 = O(\log n)$ (e.g., if the seed length for explicit non-malleable extractor improves) then the number of output bits we can extract is only a constant.

References

[Abb72] HL Abbott. Lower bounds for some ramsey numbers. *Discrete Mathematics*, 2(4):289–293, 1972.

[AGM03] Noga Alon, Oded Goldreich, and Yishay Mansour. Almost k-wise independence versus k-wise independence. *Information Processing Letters*, 88(3):107–110, 2003.

[AL93] Miklós Ajtai and Nathan Linial. The influence of large coalitions. *Combinatorica*, 13(2):129–145, 1993.

[Alo98] Noga Alon. The shannon capacity of a union. *Combinatorica*, 18(3):301–310, 1998.

[Bar06] Boaz Barak. A simple explicit construction of an $n^{\tilde{o}(\log n)}$-ramsey graph. *arXiv preprint math/0601651*, 2006.

[BIW06] Boaz Barak, Russell Impagliazzo, and Avi Wigderson. Extracting randomness using few independent sources. *SIAM Journal on Computing*, 36(4):1095–1118, 2006.

[BKS+10] Boaz Barak, Guy Kindler, Ronen Shaltiel, Benny Sudakov, and Avi Wigderson. Simulating independence: New constructions of condensers, ramsey graphs, dispersers, and extractors. *Journal of the ACM (JACM)*, 57(4):20, 2010.

[Bou05] Jean Bourgain. More on the sum-product phenomenon in prime fields and its applications. *International Journal of Number Theory*, 1(01):1–32, 2005.

[BRSW12] Boaz Barak, Anup Rao, Ronen Shaltiel, and Avi Wigderson. 2-source dispersers for $n^{o(1)}$ entropy, and ramsey graphs beating the frankl-wilson construction. *Annals of Mathematics*, 176(3):1483–1544, 2012.

[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.

[Chu81] Fan RK Chung. A note on constructive methods for ramsey numbers. *Journal of Graph Theory*, 5(1):109–113, 1981.

[CL16] Eshan Chattopadhyay and Xin Li. Explicit non-malleable extractors, multi-source extractors and almost optimal privacy amplification protocols. *ECCC*, 2016.

[Coh15] Gil Cohen. Two-source dispersers for polylogarithmic entropy and improved ramsey graphs. *arXiv preprint arXiv:1506.04428*, 2015.

[Coh16] Gil Cohen. Making the most of advice: New correlation breakers and their applications. *ECCC*, 2016.

[CS15] Gil Cohen and Leonard Schulman. Extractors for near logarithmic min-entropy. *ECCC*, 2015.

[CZ15] Eshan Chattopadhyay and David Zuckerman. Explicit two-source extractors and resilient functions. In *Electronic Colloquium on Computational Complexity (ECCC)*, volume 22, page 119, 2015.

[DPW14] Yevgeniy Dodis, Krzysztof Pietrzak, and Daniel Wichs. Key derivation without entropy waste. In *Advances in Cryptology–EUROCRYPT 2014*, pages 93–110. Springer, 2014.

[Erd47] Paul Erdös. Some remarks on the theory of graphs. *Bulletin of the American Mathematical Society*, 53(4):292–294, 1947.

[Fra77] Peter Frankl. A constructive lower bound for ramsey numbers. *Ars Combinatoria*, 3(297-302):28, 1977.

[FW81] Peter Frankl and Richard M. Wilson. Intersection theorems with geometric consequences. *Combinatorica*, 1(4):357–368, 1981.

[Gop06] Parikshit Gopalan. Constructing ramsey graphs from boolean function representations. In *Computational Complexity, 2006. CCC 2006. Twenty-First Annual IEEE Conference on*, pages 14–pp. IEEE, 2006.

[Gro01] Vince Grolmusz. Low rank co-diagonal matrices and ramsey graphs. *Journal of combinatorics*, 7(1):R15–R15, 2001.

[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.

[Li11] Xin Li. Improved constructions of three source extractors. In *Computational Complexity (CCC), 2011 IEEE 26th Annual Conference on*, pages 126–136. IEEE, 2011.

[Li13a] Xin Li. Extractors for a constant number of independent sources with polylogarithmic min-entropy. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, pages 100–109. IEEE, 2013.

[Li13b] Xin Li. New independent source extractors with exponential improvement. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 783–792. ACM, 2013.

[Li15] Xin Li. Improved constructions of two-source extractors. *arXiv preprint arXiv:1508.01115*, 2015.

[Mek15] Raghu Meka. Explicit resilient functions matching Ajtai-Linial. *CoRR*, abs/1509.00092, 2015.

[Nag75] Zs Nagy. A constructive estimation of the ramsey numbers. *Mat. Lapok*, 23:301–302, 1975.

[Nao92] Moni Naor. Constructing ramsey graphs from small probability spaces. *IBM Research Report RJ*, 8810, 1992.

[Rao07] Anup Rao. An exposition of bourgains 2-source extractor. In *Electronic Colloquium on Computational Complexity (ECCC)*, volume 14, 2007.

[Rao09] Anup Rao. Extractors for a constant number of polynomially small min-entropy independent sources. *SIAM Journal on Computing*, 39(1):168–194, 2009.

[Raz05] Ran Raz. Extractors with weak random seeds. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 11–20. ACM, 2005.

[RRV99] Ran Raz, Omer Reingold, and Salil Vadhan. Error reduction for extractors. In *Foundations of Computer Science, 1999. 40th Annual Symposium on*, pages 191–201. IEEE, 1999.

[RTS00] Jaikumar Radhakrishnan and Amnon Ta-Shma. Bounds for dispersers, extractors, and depth-two superconcentrators. *SIAM Journal on Discrete Mathematics*, 13(1):2–24, 2000.

[Vio14] Emanuele Viola. Extractors for circuit sources. *SIAM Journal on Computing*, 43(2):655–672, 2014.

[Zuc06] David Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 681–690. ACM, 2006.

ECCC ISSN 1433-8092

http://eccc.hpi-web.de