

Two-Source Condensers with Low Error and Small Entropy Gap via Entropy-Resilient Functions

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

Abstract

In their seminal work, Chattopadhyay and Zuckerman (STOC'16) constructed a two-source extractor with error ε for n -bit sources having min-entropy $\text{poly} \log(n/\varepsilon)$. Unfortunately, the construction running-time is $\text{poly}(n/\varepsilon)$, which means that with polynomial-time constructions, only polynomially-large errors are possible. Our main result is a $\text{poly}(n, \log(1/\varepsilon))$ -time computable two-source *condenser*. For any $k \geq \text{poly} \log(n/\varepsilon)$, our condenser transforms two independent (n, k) -sources to a distribution over $m = k - O(\log(1/\varepsilon))$ bits that is ε -close to having min-entropy $m - o(\log(1/\varepsilon))$. Hence, achieving entropy gap of $o(\log(1/\varepsilon))$.

The bottleneck for obtaining low error in recent constructions of two-source extractors lies in the use of resilient functions. Informally, this is a function that receives input bits from r players with the property that the function's output has small bias even if a bounded number of corrupted players feed adversarial inputs after seeing the inputs of the other players. The drawback of using resilient functions is that the error cannot be smaller than $1/r$. This, in return, forces the running time of the construction to be polynomial in $1/\varepsilon$.

A key component in our construction is a variant of resilient functions which we call *entropy-resilient functions*. This variant can be seen as playing the above game for several rounds. The goal of the corrupted players is to reduce, with as high probability as they can, the min-entropy accumulated throughout the rounds. We show that while the bias decreases only polynomially with the number of players in a one-round game, their success probability decreases *exponentially* in the entropy gap they are attempting to incur in a repeated game.

*The Blavatnik School of Computer Science, Tel-Aviv University, Tel Aviv, Israel. Supported by the Israel science Foundation grant no. 994/14 and by Len Blavatnik and the Blavatnik Family foundation.

†Department of Computer Science, Princeton University, Princeton, USA. Email: gile@princeton.edu.

‡The Blavatnik School of Computer Science, Tel-Aviv University, Tel Aviv, Israel. Supported by the Israel science Foundation grant no. 994/14, and by Len Blavatnik and the Blavatnik Family foundation. Email: deandoron@mail.tau.ac.il

§The Blavatnik School of Computer Science, Tel-Aviv University, Tel Aviv, Israel. Supported by the Israel science Foundation grant no. 994/14. Email: amnon@tau.ac.il

Contents

1	Introduction	1
1.1	Resilient functions – the barrier for obtaining extractors with low error	2
1.2	Entropy-resilient functions	3
1.3	The two-source condensers we obtain	3
1.4	Seedless condensers for a single block-source	4
1.5	A three-source extractor	5
2	Preliminaries	5
2.1	Random variables, min-entropy	5
2.2	Limited independence	6
2.3	Seeded extractors	6
2.4	Two-source condensers	7
2.5	Non-malleable extractors	7
2.6	Fooling AC circuits	8
3	Entropy-Resilient Functions	9
3.1	Functions with one output bit	10
3.2	Functions with multiple output bits	11
4	Low Error Two-Source Condensers	14
4.1	From two independent sources to a non-oblivious Σ -fixing source	15
4.2	Low error condensers with high entropy loss	16
4.3	Deterministically condensing a single block-source	17
4.4	Low error condensers with small entropy gap outputting many bits	19

1 Introduction

The problem of extracting randomness from imperfect random sources can be traced back to von Neumann [Neu51]. Ideally, and somewhat informally, a randomness extractor is an algorithm that produces, or extracts, truly random bits from an imperfect source of randomness. Going beyond that particular task, randomness extractors have found dozens of applications for error correcting codes, cryptography, combinatorics, and circuit lower bounds to name a few.

An imperfect source of randomness is modelled by a random variable X that, for convenience sake, is assumed to be supported on n -bit strings. The standard measure for the amount of randomness in X is its *min-entropy* [CG88], which is the maximum $k \geq 0$ for which one cannot guess X with probability larger than 2^{-k} . For any such k , we say that X is an (n, k) -source, or a k -source for short.

Ideally, a randomness extractor would have been defined as a function $\text{Ext}: \{0, 1\}^n \rightarrow \{0, 1\}^m$ with the property that for every random variable X with sufficiently high min-entropy, the output $\text{Ext}(X)$ is ε -close to the uniform distribution on $\{0, 1\}^m$ in the statistical distance, which we write as $\text{Ext}(X) \approx_\varepsilon U_m$. Unfortunately, such a function Ext , even for very high min-entropy $k = n - 1$ and, when set with a modest error guarantee $\varepsilon = 1/4$ and a single output bit $m = 1$, does not exist. In light of that, several types of randomness extractors, that relax in different ways the above ideal definition, have been introduced and studied in the literature. In this work, we focus on one such well-studied instantiation.

Definition 1.1 (Two-source extractors [CG88]). *A function $\text{Ext}: \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^m$ is a two-source extractor for min-entropy k with error guarantee ε if for every pair of independent (n, k) -sources X, Y , the output distribution $\text{Ext}(X, Y) \approx_\varepsilon U_m$.*

The existence of a two-source extractor for any min-entropy $k = \Omega(\log(n/\varepsilon))$ with $m = 2k - O(\log(1/\varepsilon))$ output bits was proved in [CG88]. In the same paper, an explicit construction of a two-source extractor for min-entropy $k > n/2$ was obtained. Remarkably, despite much attention [Bou05, Raz05, BSZ11] and progress on relaxed settings [BKS⁺10, Rao09, Li13b, Li13a, Li15, Coh15], the problem of constructing two-source extractors even for min-entropy as high as $k = 0.49n$ with $m = 1$ output bits remained open for 30 years. To appreciate the difficulty of constructing two-source extractors, we remark that such constructions yield explicit constructions of Ramsey graphs, a notoriously hard problem in combinatorics [BKS⁺10, BRSW12, Coh16b].

In their breakthrough result, Chattopadhyay and Zuckerman [CZ16] were finally able to obtain an explicit two-source extractor for min-entropy $k = \text{poly}(\log(n/\varepsilon))$. Partially motivated by the problem of constructing Ramsey graphs, a line of followup works [CS16, CL16, Coh16a, BADTS17, Coh16c, Li17, Li18] focused on the case of constant error ε and was devoted for reducing the min-entropy requirement as a function of n . The state of the art result in this line of work is due to Li [Li18] and requires min-entropy $\frac{\log n \cdot \log \log n}{\log \log \log n} \cdot \text{poly}(1/\varepsilon)$.

1.1 Resilient functions – the barrier for obtaining extractors with low error

Unfortunately, despite the fact that the dependence of the min-entropy of the Chattopadhyay-Zuckerman extractor on ε is polynomially-close to optimal, the running-time of their construction depends polynomially on $1/\varepsilon$ rather than the desired poly $\log(1/\varepsilon)$ dependence. The same holds for all subsequent constructions. That is, these constructions are not strongly polynomial-time and, in particular, the error guarantee cannot be taken to be sub-polynomial in n while maintaining running-time $\text{poly}(n)$. This stands in contrast to classical extractors for high min-entropy [CG88, Raz05, Bou05] that are strongly polynomial-time, and can support low error.

Informally speaking, the reason for this undesired dependence of the running-time on ε lies in the use of a so-called *resilient function* [BOL85]. A q -resilient function $f: \{0, 1\}^r \rightarrow \{0, 1\}$ can be thought of as an r -player game. If all players feed uniform and independent inputs to f , the output distribution has small bias, and, furthermore, this property is retained even if any q players decide to deviate from the rules of the game and choose their inputs as a function of all other inputs to f .

Majority on r input bits is an example of a q -resilient function with $q = O(\sqrt{r})$. Ajtai and Linial proved, using the probabilistic method, the existence of a q -resilient function for $q = O(\frac{r}{\log^2 r})$ [AL93]. The KKL Theorem [KKL88] implies that the Ajtai-Linial function is tight up to a $\log r$ factor. Chattopadhyay and Zuckerman [CZ16] constructed a derandomized version of the Ajtai-Linial function with $q = r^{1-\delta}$, for any constant $\delta > 0$. Their construction has further desirable properties. In a subsequent work, Meka obtained a derandomized version of the Ajtai-Linial function with the same parameters as the randomized construction [Mek17]. However, no matter what function is chosen, [KKL88] showed that there is always a single corrupted player that has influence $p = \Omega(\frac{\log r}{r})$, i.e., with probability p over the input fed by the other players, the single corrupted player can fully determine the result.

Almost all constructions of randomness extractors following [CZ16] can be divided into two steps. First, the two n -bit sources X, Y are “transformed” to a single r -bit source $Z = h(X, Y)$ with some structure. A resilient function $f: \{0, 1\}^r \rightarrow \{0, 1\}$ is then applied to Z so to obtain the output $\text{Ext}(X, Y) = f(h(X, Y))$. In all works, the function h is based on non-malleable extractors or on related primitives such as correlation breakers. As mentioned above, the use of the resilient function implies that even a single corrupted player has influence $\Omega(\frac{\log r}{r})$ and so to obtain an error guarantee ε , the number of players r must be taken larger than $1/\varepsilon$. This results in running-time $\Omega(1/\varepsilon)$.¹

¹There is one exception to the above scheme. In [BACD⁺18], it is shown that if very strong t -non-malleable extractors can be explicitly constructed then the function f can be replaced with the parity function (which is not resilient at all) and low error two-source extractors with low min-entropy requirement can be obtained. However, it is not known how to explicitly construct such t -non-malleable extractors.

1.2 Entropy-resilient functions

To obtain our condenser, we extend the notion of resilient functions to functions outputting many bits. Informally speaking, instead of considering an r -player game in which the bad players try to bias the output, we study a repeated game version in which the r players play for m rounds. The bad players attempt to decrease, with as high probability as they can, the min-entropy of the m -bit outcome (and we will even allow the bad players to cast their votes after the good players played all rounds).

Recall that, by [KKL88], when $m = 1$, even a single player can bias the result by $\Omega(\frac{\log r}{r})$. Put differently, viewing this bias as the error of a deterministic extractor, the error is bound to be at least polynomially-small in the number of players. Our key insight is that when m becomes large, the probability that the bad players can reduce g bits of entropy from the output (creating an “entropy gap” of g) is *exponentially small* in g . We further show that this holds for a specific function f , induced by the Ajtai-Linial function, even when the honest players are only t -wise independent (for $t = \text{poly} \log(r/\varepsilon)$). Our analysis uses and extends ideas from the work of Chattopadhyay and Zuckerman [CZ16].

1.3 The two-source condensers we obtain

The main contribution of this work is an explicit construction of a two-source *condenser* with low error and small entropy gap, outputting almost all of the entropy from one source.

Definition 1.2 (Two-source condensers). *A function $\text{Cond}: \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^m$ is a two-source condenser for min-entropy k with min-entropy gap g and error guarantee ε if for every pair of independent (n, k) -sources, $\text{Cond}(X, Y)$ is ε -close to an $(m, m - g)$ -source.*

Note that a two-source extractor is a two-source condenser with entropy gap $g = 0$. Thus, condensers can be seen as a relaxation of extractors in which some, hopefully small, “gap” of min-entropy in the output distribution is allowed. Despite having a weaker guarantee, condensers play a key role in the construction of many types of randomness extractors, including two-source extractors [BADTS17], their variants [Raz05, BKS⁺10, Zuc07, Rao09, Li13a], and seeded-extractors [GUV09]. Most related to our work is a paper by Rao [Rao08] that, for every $\delta > 0$, constructed a $\text{poly}(n, \log(1/\varepsilon))$ -time computable two-source condenser² for min-entropy $k = \delta n$ having $m = \Omega(\delta n)$ output bits with entropy gap $g = \text{poly}(1/\delta, \log(1/\varepsilon))$.

In this work, we obtain a strongly polynomial-time construction of a two-source condenser with low error and small min-entropy gap.

Theorem 1.3 (Main result). *For all integers n, k and every $\varepsilon > 0$ such that $n \geq k \geq \text{poly} \log(\frac{n}{\varepsilon})$, there exists a $\text{poly}(n, \log(1/\varepsilon))$ -time computable two-source condenser*

$$\text{Cond}: \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^m$$

²To the matter of fact, Rao entitled his construction a “two-source almost extractor” – a suitable name given its small entropy gap.

for min-entropy k , with error guarantee ε , min-entropy gap $g = o(\log \frac{1}{\varepsilon})$, and $m = k - O(\log(1/\varepsilon))$ output bits.

Note that the entropy gap g is independent of the entropy k and scales *sub*-logarithmically with $1/\varepsilon$. We prove Theorem 1.3, whose formal statement is given in Theorem 4.6, in two steps. First, we construct a two-source condenser with the same guarantees as provided by Theorem 1.3, though only with $m = k^\alpha$ output bits, where $0 < \alpha < 1$ is some small universal constant (see Theorem 4.2). This part of the construction is based on our study of entropy-resilient functions (Section 3) and on the adaptation of the Chattopadhyay-Zuckerman construction for entropy-resilient functions. To reduce the huge entropy-loss we incur (i.e., to increase the output length from k^α to $k - O(\log(1/\varepsilon))$), in the second step, we construct a seedless condenser for block-sources—a result that we believe is of independent interest on which we now elaborate.

1.4 Seedless condensers for a single block-source

A (k_1, k_2) -*block-source* is a pair of random variables X_1, X_2 that, although may be dependent, have the following guarantee. First, X_1 is a k_1 -source, and second, conditioned on any fixing of X_1 , the random variable X_2 has min-entropy k_2 . Throughout this section, we denote the length of X_1 by n_1 and the length of X_2 by n_2 . Informally, the notion of a block-source “lies between” a single source and two independent sources. Indeed, any (k_1, k_2) -block-source is a $(k_1 + k_2)$ -source. Moreover, if X_1 is a k_1 -source and X_2 is an independent k_2 -source then X_1, X_2 is a (k_1, k_2) -block-source.

Block-sources are key to almost all constructions of seeded extractors as well as to the construction of Ramsey graphs. As mentioned above, there is no one-source extractor, whereas two-source extractors exist even for very low min-entropy. Despite being more structured than a general source, it is a well-known fact that there is no extractor for a single block-source (with non-trivial parameters).

A key component that allows us to increase the output length of our condenser discussed above is a seedless condenser for a single block-source. Let X_1, X_2 be a (k_1, k_2) -block-source. Write $g = n_2 - k_2$ for the entropy gap of X_2 . For any given $\varepsilon > 0$, we show how to *deterministically* transform X_1, X_2 to a single m -bit random variable, where $m = k_1 - g - O(\log(1/\varepsilon))$, that is ε -close to having min-entropy $m - g - 1$. That is, informally, we are able to condense X_1 roughly to its entropy content k_1 using (the dependent random variable) X_2 while inheriting the entropy gap of X_2 both in the resulted entropy gap and entropy loss. We stress that this transformation is deterministic. This demonstrates that despite the well-known fact that a block-source extractor does not exist, a block-source condenser does. For a formal treatment, see Section 4.3.

1.5 A three-source extractor

An immediate implication of Theorem 1.3 are low error three-source extractors supporting min-entropies $k_1 = k_2 = \text{poly log}(n/\varepsilon)$ and $k_3 = \Omega(\log(1/\varepsilon))$. This is achieved by feeding our condenser's output $Y = \text{Cond}(X_1, X_2)$ as a seed to a seeded extractor that supports small entropies (see, e.g., Theorem 2.6), outputting $\text{Ext}(X_3, Y)$. We can compensate for the tiny entropy gap of Y by employing Ext with a slightly lower error.

Corollary 1.4. *For all integers n, k, k' and every $\varepsilon > 0$ such that $n \geq k \geq \text{poly log}(\frac{n}{\varepsilon})$ and $n \geq k' \geq \Omega(\log \frac{1}{\varepsilon})$ there exists a $\text{poly}(n, \log(1/\varepsilon))$ -time computable three-source extractor*

$$\text{Ext}: \{0, 1\}^n \times \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^m$$

for min-entropies k, k, k' and error guarantee ε , where $m = k' - O(\log(1/\varepsilon))$.

We note that when ε is taken sub-polynomial in n , in which case the two-source extractor of [CZ16] is not polynomial-time computable, Corollary 1.4 modestly improves upon known three-source extractors that either require all three-sources to have min-entropy $\text{poly log}(\frac{n}{\varepsilon})$ [Li15] or require, for any parameter of choice $\delta > 0$, min-entropies δn , $\text{poly}(\frac{1}{\delta}) \log(\frac{n}{\varepsilon})$, $\text{poly}(\frac{1}{\delta}) \log(\frac{\log n}{\varepsilon})$ [Coh15].

To conclude, we believe that the use of entropy-resilient functions as a tool to extract almost all the entropy from bit-fixing sources while suffering only a small error is both natural and interesting on its own. We hope the tools and constructions developed in this paper will be of further use, possibly for constructing low error two-source extractors.

2 Preliminaries

We use $\log(x)$ for $\log_2(x)$. For an integer n , we denote by $[n]$ the set $\{1, \dots, n\}$. The density of a subset $B \subseteq A$ is denoted by $\mu(B) = \frac{|B|}{|A|}$.

2.1 Random variables, min-entropy

The *statistical distance* between two distributions X and Y over the same domain Ω is defined by $\text{SD}(X, Y) = \max_{A \subseteq \Omega} (\Pr[X \in A] - \Pr[Y \in A])$. If $\text{SD}(X, Y) \leq \varepsilon$ we say X is ε -close to Y and denote it $X \approx_\varepsilon Y$. We denote by U_n the random variable distributed uniformly over $\{0, 1\}^n$.

For a function $f: \Omega_1 \rightarrow \Omega_2$ and a random variable X distributed over Ω_1 , $f(X)$ is the random variable distributed over Ω_2 obtained by choosing $x \sim X$ and outputting $f(x)$. For every $f: \Omega_1 \rightarrow \Omega_2$ and two random variables X, Y over Ω_1 it holds that $\text{SD}(f(X), f(Y)) \leq \text{SD}(X, Y)$.

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

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

A random variable X is an (n, k) -source if X is distributed over $\{0, 1\}^n$ and has min-entropy at least k . When n is clear from the context we sometimes omit it and simply say that X is a k -source.

2.2 Limited independence

Definition 2.1. A distribution X over $\{0, 1\}^n$ is called (t, γ) -wise independent if the restriction of X to every t coordinates is γ -close to U_t .

Lemma 2.2 ([AGM03]). Let $X = X_1, \dots, X_n$ 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.

2.3 Seeded extractors

Definition 2.3 (Seeded extractors). A function

$$\text{Ext}: \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$$

is a (k, ε) -seeded-extractor if the following holds. For every (n, k) -source X , the output $\text{Ext}(X, Y) \approx_\varepsilon U_m$, where Y is uniformly distributed over $\{0, 1\}^d$ and is independent of X . Further, Ext is a (k, ε) -strong-seeded-extractor if $(\text{Ext}(X, Y), Y) \approx_\varepsilon (U_m, Y)$.

Theorem 2.4 ([GUV09]). There exists a universal constant $c_{\text{GUV}} \geq 2$ for which the following holds. For every integers $n \geq k$ and $\varepsilon > 0$ there exists a $\text{poly}(n, \log(1/\varepsilon))$ -time computable (k, ε) -strong-seeded-extractor

$$\text{Ext}: \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$$

with seed length $d = c_{\text{GUV}} \log(n/\varepsilon)$ and $m = k/2$ output bits.

Extractors can be used for sampling using weak sources.

Theorem 2.5 ([Zuc97]). Let $\text{Ext}: \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ be a (k_1, ε) -seeded-extractor. Identify $\{0, 1\}^d$ with $[2^d]$ and let $S(X) = \{\text{Ext}(X, 1), \dots, \text{Ext}(X, 2^d)\}$. Then, for every (n, k_2) -source X and any set $T \subseteq \{0, 1\}^m$,

$$\Pr_{x \sim X} \left[\left| \frac{|S(x) \cap T|}{2^m} - \mu(T) \right| > \varepsilon \right] \leq 2^{-(k_2 - k_1)}.$$

The following extractor allows us to extract all the min-entropy, at the cost of a larger seed-length.

Theorem 2.6 ([GUV09]). *There exists a universal constant c such that the following holds. For all integers $n \geq k$ and any $\varepsilon > 0$ such that $k \geq 2 \log(1/\varepsilon) + O(1)$, there exists a poly($n, \log(1/\varepsilon)$)-time computable (k, ε) -strong-seeded-extractor*

$$\text{Ext}: \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$$

with seed length $d = c \log n \cdot \log \frac{n}{\varepsilon}$ and $m = k - 2 \log \frac{1}{\varepsilon} - O(1)$ output bits.

2.4 Two-source condensers

Definition 2.7 (Condensers). *A function*

$$\text{Cond}: \{0, 1\}^{n_1} \times \{0, 1\}^{n_2} \rightarrow \{0, 1\}^m$$

is an $((n_1, k_1), (n_2, k_2)) \rightarrow_\varepsilon (m, k' = m - g)$ condenser if the following holds. For every (n_1, k_1) -source X_1 and an independent (n_2, k_2) -source X_2 , the output $\text{Cond}(X_1, X_2)$ is ε -close to an (m, k') -source. We refer to ε as the error guarantee and to g as the entropy gap of Cond .

Definition 2.8 (Strong condensers). *A function*

$$\text{Cond}: \{0, 1\}^{n_1} \times \{0, 1\}^{n_2} \rightarrow \{0, 1\}^m$$

is a $((n_1, k_1), (n_2, k_2)) \rightarrow_{\varepsilon_1, \varepsilon_2} (m, k')$ -strong-condenser (in the first source) if the following holds. For every (n_1, k_1) -source X_1 and an independent (n_2, k_2) -source X_2 , with probability $1 - \varepsilon_1$ over $x_1 \sim X_1$, the output $\text{Cond}(x_1, X_2)$ is ε_2 -close to an (m, k') -source.

Similarly, one can define, in the natural way, a condenser that is strong in the second source.

2.5 Non-malleable extractors

Definition 2.9. *A function $\text{nmExt}: \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is a (k, ε) t -non-malleable extractor, if for every (n, k) -source X , for every independent random variable Y that is uniform over $\{0, 1\}^d$ and every functions $f_1, \dots, f_t: \{0, 1\}^d \rightarrow \{0, 1\}^d$ with no fixed-points³ it holds that:*

$$\begin{aligned} (\text{nmExt}(X, Y), \text{nmExt}(X, f_1(Y)), \dots, \text{nmExt}(X, f_t(Y), Y)) &\approx_\varepsilon \\ (U_m, \text{nmExt}(X, f_1(Y)), \dots, \text{nmExt}(X, f_t(Y), Y)). \end{aligned}$$

³That is, for every i and every x , we have $f_i(x) \neq x$.

We will need the following lemma concerning the existence of a set of good seeds of a non-malleable extractor, given in [CZ16].

Lemma 2.10 ([CZ16]). *Let $\text{nmExt}: \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ be a (k, ε) t -non-malleable extractor. Let X be any (n, k) -source. Let BAD be the set defined by*

$$\text{BAD} = \{r \in [D] \mid \exists \text{ distinct } r_1, \dots, r_t \in [D], \forall i \in [t] \ r_i \neq r, \ |(\text{nmExt}(X, r), \text{nmExt}(X, r_1), \dots, \text{nmExt}(X, r_t)) - (U_m, \text{nmExt}(X, r_1), \dots, \text{nmExt}(X, r_t))| > \sqrt{\varepsilon}\}.$$

Then, $\mu(\text{BAD}) \leq \sqrt{\varepsilon}$. We refer to the set $[D] \setminus \text{BAD}$ as the set of good seeds (with respect to the underlying distribution of X).

Lemma 2.11. *Let X_1, \dots, X_t be random variables over $\{0, 1\}^m$. Further suppose that for any $i \in [t]$,*

$$(X_i, \{X_j\}_{j \neq i}) \approx_\varepsilon (U_m, \{X_j\}_{j \neq i}).$$

Then, $(X_1, \dots, X_t) \approx_{t\varepsilon} U_{tm}$.

Finally, good explicit constructions of t -non-malleable extractors exist. The following choice of parameters will be sufficient for us.

Theorem 2.12 ([CGL16, Coh16c, Li17]). *There exists a universal constant $c_{\text{nm}} \geq 2$ such that for all integers n, k, t , and every $\varepsilon > 0$ such that $n \geq k \geq c_{\text{nm}} t^2 \log^2(n/\varepsilon)$, there exists a $\text{poly}(n, \log(1/\varepsilon))$ -time computable (k, ε) t -non-malleable extractor*

$$\text{nmExt}: \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$$

with $m = \frac{k}{3t}$ output bits and seed length $d = c_{\text{nm}} t^2 \log^2(n/\varepsilon)$.

2.6 Fooling AC circuits

A Boolean circuit is an $\text{AC}[d, s]$ circuit if it has depth d , size s and unbounded fan-in. We say that a circuit C with n input bits is ε -fooled by a distribution D if $\text{SD}(C(D), D(U_n)) \leq \varepsilon$.

Harsha and Srinivasan [HS16], improving upon Braverman's seminal result [Bra10] (see also [Tal17]) proved:

Theorem 2.13 ([HS16]). *There exists a constant $c > 0$ such that the following holds. For every integers s, d, t , any $\text{AC}[d, s]$ circuit is ε -fooled by any t -wise independent distribution, where $\varepsilon = 2^{(\log s)^{3d+c-t}}$.*

Definition 2.14 (ε -bias). *Let X be a distribution over $\{0, 1\}^n$. We say $\text{bias}(X) \leq \varepsilon$ if for every non-empty $S \subseteq [n]$, $\Pr[\bigoplus_{i \in S} X_i = 1] \in [\frac{1}{2} \pm \varepsilon]$.*

Lemma 2.15 ([Vaz86]). *Let X be a distribution over $\{0, 1\}^n$. Then, $\text{SD}(X, U_n) \leq 2^{n/2} \cdot \text{bias}(X)$.*

We need a slight generalization of Theorem 2.13:

Lemma 2.16. *There exists a constant $c > 0$ such that the following holds for every integers n, m, d, s . Let $C: \{0, 1\}^n \rightarrow \{0, 1\}^m$ be an $\text{AC}[d, s]$ circuit. Then, C is ε -fooled by any t -wise independent distribution, where $\varepsilon = 2^{(m+\log s)^{3d+c}-t}$.*

Proof. Fix $S \subseteq [m]$ and consider the circuit $C_S: \{0, 1\}^n \rightarrow \{0, 1\}$ given by

$$C_S(x) = \bigoplus_{i \in S} C(x)_i.$$

As the parity over $|S|$ bits can be expressed as a CNF (or DNF) of at most $2^{|S|}$ terms, C_S is an $\text{AC}[d+1, s']$ circuit, for $s' = O(s + m2^m)$. By Theorem 2.13, every t -wise distribution ε' -fools C_S , where $\varepsilon' = 2^{(m+\log s)^{3d+c'}-t}$ for some universal constant $c' > 0$. That is, for every t -wise distribution D and non-empty $S \subseteq [m]$, $\text{SD}(C_S(D), C_S(U_n)) \leq \varepsilon'$. By Lemma 2.15,

$$\varepsilon = \text{SD}(C(D) - C(U_n)) \leq 2^{m/2} \varepsilon' = 2^{(m+\log s)^{3d+c}-t}$$

for some universal constant $c > c' > 0$. □

3 Entropy-Resilient Functions

Definition 3.1 (Non-oblivious sources). *Let $\Sigma = \{0, 1\}^m$. A (q, t) -non-oblivious Σ -fixing source $X = (X_1, \dots, X_r)$ is a random variable over $\Sigma^r = \{0, 1\}^{rm}$ for which there exists a set $R_{\text{bad}} \subseteq [r]$ of cardinality $q' \leq q$ such that:*

- *The joint distribution of $\{(X_i)_j \mid i \in [r] \setminus R_{\text{bad}}, j \in [m]\}$, denoted by G_X , is t -wise independent over $\{0, 1\}^{(r-q')m}$; and*
- *Each of the random variables in $B_X \triangleq \{(X_i)_j\}$ with $i \in R_{\text{bad}}$ and $j \in [m]$ may depend arbitrarily on all other random variables in G_X and B_X .*

If $t = (r - q')m$ we say X is a q -non-oblivious Σ -fixing source. If $m = 1$ we say X is a bit-fixing source and the definition coincides with the standard definition of non-oblivious bit-fixing sources [BOL85]. When X is clear from context, we write G and B for G_X and B_X , respectively.

Definition 3.2 (Entropy-resilient functions). *Let $\Sigma = \{0, 1\}^m$. A function $f: \Sigma^r \rightarrow \Sigma$ is a (q, t, g, ε) -entropy-resilient function if for every (q, t) -non-oblivious Σ -fixing source X over Σ^r , the output $f(X)$ is ε -close to an $(m, m-g)$ -source. If $g = 0$ we say f is (q, t, ε) -resilient.*

3.1 Functions with one output bit

Definition 3.3. Let $f: \{0, 1\}^r \rightarrow \{0, 1\}$ be an arbitrary function. Let X be a (q, t) -non-oblivious bit-fixing source over $\{0, 1\}^r$. Define $E(f)$ to be the event in which the bits tossed by the good players do not determine the value of the function f . We define the influence of the bad players by $I(f) = \Pr[E(f)]$.

Balanced resilient functions can be seen as deterministic extractors against non-oblivious bit-fixing sources outputting one bit. Chattopadhyay and Zuckerman [CZ16], followed by an improvement by Meka [Mek17], derandomized the Ajtai-Linial function [AL93] and obtained an explicit construction of an almost-balanced resilient function which is also computable by monotone AC^0 circuits.

Theorem 3.4 ([CZ16, Mek17]). For every constant $0 < \delta < 1$, there exists a constant $c_\delta \geq 1$ such that for every constant $c \geq c_\delta$ and integer r there exists a monotone function $\text{Res}: \{0, 1\}^r \rightarrow \{0, 1\}$ such that for every $t \geq c \log^4 r$,

- For every (q, t) -non-oblivious bit-fixing source X , $I(\text{Res}) \leq c \cdot \frac{q}{r^{1-\delta}}$.
- For every t -wise independent distribution D , $\text{bias}(\text{Res}(D)) \leq r^{-1/c}$.

The function Res is computable by a uniform depth 3 monotone circuit of size r^c . Further, the function $c_\delta(\delta)$ is continuous and monotonically decreasing.

Throughout the paper we make use of the following corollary.

Corollary 3.5. For every constant $0 < \gamma < 1$ there exist constants $0 < \alpha < \beta < 1$ such that for every integer r there exists a function $\text{Res}: \{0, 1\}^r \rightarrow \{0, 1\}$ which for every $t \geq \frac{1}{\beta} \log^4 r$ satisfies: For every $(r^{1-\gamma}, t)$ -non-oblivious bit-fixing source X ,

$$\begin{aligned} I(\text{Res}) &\leq \frac{1}{\beta} \cdot r^{-\alpha}, \\ \text{bias}(\text{Res}(X) \mid \neg E(\text{Res})) &\leq \frac{3}{\beta} \cdot r^{-\alpha}. \end{aligned}$$

The function Res is computable by a uniform depth 3 monotone circuit of size $r^{\frac{1}{\beta}}$.

Proof. Using the notations of Theorem 3.4, assume that for every η , $c_\eta > \frac{1}{2\eta}$ (if not, we can always increase c_η). Given $\gamma > 0$, set δ to be the constant satisfying the equation $f(\delta) = \delta - \gamma + \frac{1}{2c_\delta} = 0$. Such a δ exists, as $f(\delta) \leq 2\delta - \gamma$ and therefore $f(\delta) < 0$ when δ approaches 0, and $f(\delta) > 0$ when δ approaches γ . Note that by our choice of δ , it holds that

$$\delta < \gamma = \delta + \frac{1}{2c_\delta} < \delta + \frac{1}{c_\delta}.$$

Set $\alpha = \gamma - \delta > 0$ and $\beta = \frac{1}{c_\delta}$. Note that indeed $\beta > \alpha$.

By Theorem 3.4, applied with the constant δ , it holds that $I(\text{Res}) \leq c_\delta \frac{r^{1-\gamma}}{r^{1-\delta}} = \frac{1}{\beta} r^{-\alpha}$. Further, $\text{bias}(\text{Res}(D)) \leq r^{-\beta}$.

Following similar arguments as in [CZ16], we have that $\text{bias}(\text{Res}(X)) \leq \frac{1}{\beta} r^{-\alpha} + r^{-\beta}$, so

$$\text{bias}(\text{Res}(X) \mid \neg E(\text{Res})) \leq \frac{\frac{1}{\beta} r^{-\alpha} + r^{-\beta}}{1 - \frac{1}{\beta} r^{-\alpha}} \leq \frac{3}{\beta} r^{-\alpha}.$$

□

3.2 Functions with multiple output bits

The output bit of a (q, t, ε) -resilient function $f: \{0, 1\}^r \rightarrow \{0, 1\}$ applied to a (q, t) -non-oblivious bit-fixing source is indeed ε -close to uniform, but, as shown by [KKL88] even when $q = 1$, ε cannot be smaller than $\frac{\ln r}{r}$ (and the simpler bound $\varepsilon \geq \frac{1}{r}$ is almost trivial). We show that when we output many bits, and allow $o(\log \frac{1}{\varepsilon})$ entropy gap, we may obtain much smaller error. We do that by exhibiting an entropy-resilient function based on a parallel application of the (derandomized version of the) Ajtai-Linial function.

A construction of an entropy-resilient function. Given a constant $0 < \gamma < 1$ and integers $r \geq m$ let $\text{Res}: \{0, 1\}^r \rightarrow \{0, 1\}^r$ be the function guaranteed by Corollary 3.5 with respect to γ . Define $\Sigma = \{0, 1\}^m$ and

$$\text{EntRes}: \Sigma^r \rightarrow \Sigma$$

as follows. On input $x \in \Sigma^r$,

$$\text{EntRes}(x) = (\text{Res}(y_1), \dots, \text{Res}(y_m)),$$

where x_i stands for the i -th column of x , when we view x as a $r \times m$ table.

Theorem 3.6. *For every constant $0 < \gamma < 1$ there exist constants $0 < \alpha < \beta < 1$ and $c' \geq 1$ such that the following holds. For every integers $r, m \leq r^{\alpha/2}$, every $\varepsilon > 0$, and for every integer $t \geq (m + \log(r/\varepsilon))^{c'}$, the function $\text{EntRes}: \Sigma^r \rightarrow \Sigma$ is $(q = r^{1-\gamma}, t, g, \varepsilon)$ -entropy-resilient with entropy gap $g = o(\log(1/\varepsilon))$.*

The proof of Theorem 3.6 is done in two steps. First, in Section 3.2.1, we analyze the theorem for the special case in which the distribution G_X of the given non-oblivious Σ -fixing source X is uniform. Then, based on that result, in Section 3.2.2 we prove Theorem 3.6.

3.2.1 The uniform case

In this section, we prove the following lemma.

Lemma 3.7. *Keeping the notations of Theorem 3.6, the function $\text{EntRes}: \Sigma^r \rightarrow \Sigma$ is $(q = r^{1-\gamma}, g, \varepsilon)$ -entropy-resilient with entropy gap*

$$g = c_{\text{ent}_1} \frac{\ln \frac{1}{\varepsilon}}{\ln \ln \frac{1}{\varepsilon} + c_{\text{ent}_2} \ln r} = o(\log(1/\varepsilon))$$

for some universal constant $c_{\text{ent}_1} > 0$ and a constant $c_{\text{ent}_2} > 0$ that depends only on γ .

Proof of Lemma 3.7. Let X be a $(q = r^{1-\gamma})$ -non-oblivious Σ -fixing source. Let $R_{\text{bad}} \subseteq [r]$ be the set of bad players, and E_i the event that the values of the good players in X_i do not determine the value of Res . Note that we shall also denote E_i as an indicator for that event. By Corollary 3.5, there exists constants $0 < \alpha < \beta < 1$ such that

$$\Pr[E_i = 1] \leq \frac{1}{\beta} \cdot r^{-\alpha}$$

for every $i \in [m]$. Observe that the random variables E_1, \dots, E_m are independent, as the value of E_i depends only on the values of the good players in the i -th column, and by assumption all these values are independent of the corresponding values in the other columns.

Write

$$\mu = m \cdot \frac{1}{\beta} \cdot r^{-\alpha}$$

and note that since $m \leq r^{\alpha/2}$, $\mu < 1$. Set

$$c = \frac{4 \ln \frac{1}{\varepsilon}}{\mu} \cdot \frac{1}{\ln \frac{1}{\mu}}$$

and observe that $c > 1$. By the Chernoff bound,

$$\Pr \left[\sum_{i=1}^m E_i > c\mu \right] \leq \left(\frac{e^{c-1}}{c^c} \right)^\mu \leq e^{-\frac{1}{2}\mu c \ln c} \leq \varepsilon,$$

where the last inequality follows from the fact that $c \ln c \geq \frac{2 \ln \frac{1}{\varepsilon}}{\mu}$.

By Corollary 3.5, for every $i \in [m]$,

$$\text{bias}(\text{Res}(X_i) \mid E_i = 0) \leq \frac{3}{\beta} \cdot r^{-\alpha}.$$

Assume that the event $\sum_{i=1}^m E_i \leq c\mu$ holds, and let $I \subseteq [m]$, $|I| \geq m - c\mu$ be the set of good columns I for which $E_i = 0$. For every $w \in \{0, 1\}^m$, we have:

$$\begin{aligned} \Pr[\text{EntRes}(X) = w] &\leq \Pr[\text{EntRes}(X)_I = w_I] \\ &\leq \left(\frac{1}{2} + \frac{3}{\beta} \cdot r^{-\alpha}\right)^{m-c\mu} \\ &\leq 2^{-m+c\mu} e^{\frac{6}{\beta} r^{-\alpha} m} \\ &\leq 2^{-m+c\mu} 2^{10\mu}. \end{aligned}$$

Now, we have

$$c\mu + 10\mu \leq 2c\mu \leq \frac{8 \ln \frac{1}{\varepsilon}}{\ln \ln \frac{1}{\varepsilon} + \ln \frac{1}{\mu}} \leq \frac{8 \ln \frac{1}{\varepsilon}}{\ln \ln \frac{1}{\varepsilon} + \frac{4}{\alpha} \ln r} = o\left(\log \frac{1}{\varepsilon}\right).$$

We have shown that except with probability ε , the output $\text{EntRes}(X)$ has min-entropy $m - o(\log(1/\varepsilon))$, as desired. More specifically, the min-entropy in the good columns alone is at least $m - o(\log(1/\varepsilon))$, and we stress that the good columns are not fixed but depend on the sample itself. \square

3.2.2 The bounded-independence case – proof of Theorem 3.6

Throughout this section, we use the same notations as in Lemma 3.7. We are given X that is a (q, t) -non-oblivious Σ -fixing source. We use a similar approach to the one taken in [CZ16]. For the sake of the proof, we:

- Let GU be the distribution in which the good players are jointly uniform, and the bad players are arbitrary.
- Define a small-depth circuit C' that is related to EntRes so that $H_\infty(\text{EntRes}(X)) \geq H_\infty(C'(X))$.

We will show that $C'(X)$ and $C'(GU)$ are statistically close to each other. Finally, the results of Section 3.2.1 proves that except for a small probability, $H_\infty(C'(GU)) \geq m - o(\log(1/\varepsilon))$.

Proof of Theorem 3.6. Fix a (q, t) -non-oblivious Σ -fixing source X . Let GU be the distribution where the good players are jointly uniform, and the bad players are arbitrary. We construct a circuit $C': \{0, 1\}^{rm} \rightarrow \{0, 1\}^m$ such that:

$$(C'(x))_i = \begin{cases} \text{EntRes}(x)_i & \text{If } E_i(x) = 0, \\ 0 & \text{Otherwise.} \end{cases}$$

Recall that E_i is *fully determined* by the good players, and so does $\text{EntRes}(X)_i$ when $E_i = 0$. Hence, C' is fully determined by the good players.

We can write a small-depth circuit computing C' . Let C be the depth-3 size $r^{1/\beta}$ circuit that computes the function $\text{Res}: \{0, 1\}^r \rightarrow \{0, 1\}$ as guaranteed by Theorem 3.4. Construct a circuit for C' as follows:

- For $i \in [m]$ and $b \in \{0, 1\}$ let $C_{i,b}$ be a copy of C where we wire $(x_i)_j$ for every good player $j \in [r]$, and the value b for every bad player.
- The top part contains m comparators, outputting the output of $C_{i,0}$ if the output of $C_{i,0}$ is the same as the output of $C_{i,1}$, and 0 otherwise.

The circuit has depth 4 and size $s'' = O(mr^{1/\beta})$ and its correctness is guaranteed by the fact that Res is monotone (so it is sufficient to consider the case where the bad players voted unanimously).

By Lemma 2.16, $\text{SD}(C'(GU), C'(X)) \leq 2^{(m+\log r)c''-t}$ for some large enough universal constant $c'' > 0$. For every $w \in \{0, 1\}^m$:

$$\begin{aligned} \Pr[\text{EntRes}(X) = w] &\leq \Pr[\text{EntRes}(X)_I = w_I] \\ &= \Pr[C'(X)_I = w_I] \\ &\leq \Pr[C'(GU)_I = w_I] + 2^{(m+\log r)c''-t} \\ &\leq 2^{-m + \frac{8 \ln \frac{1}{\varepsilon}}{\ln \ln \frac{1}{\varepsilon} + \frac{4}{\alpha} \ln r}} + 2^{(m+\log r)c''-t}, \end{aligned}$$

where in the last inequality we have used Lemma 3.7. We can set the constant c' stated in the theorem to be larger than c'' and get that

$$\Pr[\text{EntRes}(X) = w] \leq 2 \cdot 2^{-m + \frac{8 \ln \frac{1}{\varepsilon}}{\ln \ln \frac{1}{\varepsilon} + \frac{4}{\alpha} \ln r}}.$$

To conclude, note that the above holds with probability at least $1 - \varepsilon$, and then $\text{EntRes}(X)$ has min-entropy at least $-1 + m - \frac{8 \ln \frac{1}{\varepsilon}}{\ln \ln \frac{1}{\varepsilon} + \frac{4}{\alpha} \ln r} = m - o(\log(1/\varepsilon))$, as desired. \square

4 Low Error Two-Source Condensers

Chattopadhyay and Zuckerman showed a reduction from two independent sources to non-oblivious bit-fixing sources. In Section 4.1 we extend this to many output bits and show a reduction from two independent sources to non-oblivious Σ -fixing sources. In Section 4.2 we use this together with the results of Section 3 to get a low error two-source condenser with many output bits, yet still far from getting almost all of the possible entropy from the two sources. In Section 4.4 we show how the condenser obtained in Section 4.2 can be used to extract more bits and get a condenser extracting almost all the entropy from one of the sources. To this end, we use the connection between condensers with small entropy gap and samplers with multiplicative error (Section 4.3) and ideas from [RRV99].

4.1 From two independent sources to a non-oblivious Σ -fixing source

In this section, we revisit the [CZ16] transformation of two independent sources to a non-oblivious bit-fixing source, and extend it to sources with several bits. Throughout this section, we refer to $c_{\text{GUV}}, c_{\text{nm}}$ as the constants that appear in Theorem 2.4 and Theorem 2.12, respectively.

Theorem 4.1. *For every integers n, t, m, k , with $n \geq k \geq (tm \log n)^5$ and set $\Sigma = \{0, 1\}^m$, there exists a poly(n)-time computable function*

$$\text{TwoSourcesToNOF}: \{0, 1\}^n \times \{0, 1\}^n \rightarrow \Sigma^r,$$

where $r = n^{2c_{\text{GUV}}}$ such that the following holds. Let X_1, X_2 be a pair of independent (n, k) -sources. Then, with probability at least $1 - 2^{-k/2}$ over $x_2 \sim X_2$, the output

$$\text{TwoSourcesToNOF}(X_1, x_2)$$

is (n^{-mt}) -close to an $(r^{1-\frac{1}{4c_{\text{GUV}}}}, t)$ -non-oblivious Σ -fixing source.

Proof. We start by setting the following parameters:

Setting of parameters.

- Set $\varepsilon_{\text{GUV}} = \frac{1}{n}$.
- Set $d_{\text{GUV}} = c_{\text{GUV}} \log\left(\frac{n}{\varepsilon_{\text{GUV}}}\right) = 2c_{\text{GUV}} \log n$.
- Set $\varepsilon_{\text{nm}} = 2^{-4mt(d_{\text{GUV}} + \log m)}$.
- Set $d_{\text{nm}} = c_{\text{nm}} t^2 \log^2\left(\frac{n}{\varepsilon_{\text{nm}}}\right)$.

Note that $\varepsilon_{\text{nm}} = 2^{-\Theta(mt \log n)}$ and that $d_{\text{nm}} = \Theta(t^4 m^2 \log^2 n)$.

Building blocks. For the construction of TwoSourcesToNOF, we make use of the following ingredients:

- Let $\text{Ext}: \{0, 1\}^n \times \{0, 1\}^{d_{\text{GUV}}} \rightarrow \{0, 1\}^{d_{\text{nm}}}$ be the $(k/2, \varepsilon_{\text{GUV}})$ -strong-seeded-extractor, guaranteed by Theorem 2.4. One can verify that $k/2 \geq 2d_{\text{nm}}$ as required by Theorem 2.4.
- Let $\text{nmExt}: \{0, 1\}^n \times \{0, 1\}^{d_{\text{nm}}} \rightarrow \{0, 1\}^m$ be the $(k, \varepsilon_{\text{nm}})$ t -non-malleable extractor, guaranteed by Theorem 2.12. Note that $k \geq 3tm$ so the hypothesis of Theorem 2.12 is met with our choice of parameters.

The construction. We identify $[r]$ with $\{0, 1\}^{d_{\text{GUV}}}$. On inputs $x_1, x_2 \in \{0, 1\}^n$, we define $\text{TwoSourcesToNOF}(x_1, x_2)$ to be the $r \times m$ matrix whose i -th row is given by

$$\text{TwoSourcesToNOF}(x_1, x_2)_i = \text{nmExt}(x_1, \text{Ext}(x_2, i)).$$

Analysis. Write $D_{nm} = 2^{d_{nm}}$ and identify $[D_{nm}]$ with $\{0, 1\}^{d_{nm}}$. Let $G \subseteq [D_{nm}]$, $|G| \geq (1 - \sqrt{\varepsilon_{nm}})D_{nm}$, be the set of good seeds guaranteed by Lemma 2.10. By Lemma 2.11, for any distinct $r_1, \dots, r_t \in G$,

$$(\text{nmExt}(X_1, r_1), \dots, \text{nmExt}(X_1, r_t)) \approx_{t\sqrt{\varepsilon_{nm}}} U_{tm}.$$

Let $S(X_2) = \{\text{Ext}(X_2, 1), \dots, \text{Ext}(X_2, 2^{d_{nm}})\}$. By Theorem 2.5,

$$\Pr_{x_2 \sim X_2} [|S(x_2) \cap G| \leq (1 - \sqrt{\varepsilon_{nm}} - \varepsilon_{\text{GUV}}) \cdot r] \leq 2^{-k/2}.$$

We say that $x_2 \in \text{supp}(X_2)$ is good if it induces a good sample, that is if $|S(x_2) \cap G| > (1 - \sqrt{\varepsilon_{nm}} - \varepsilon_{\text{GUV}})r$. Fix a good x_2 and let $Z = \text{TwoSourcesToNOF}(X_1, x_2)$. In the good seeds, every t elements of Z are $(t\sqrt{\varepsilon_{nm}})$ -close to uniform, and there are at most $q \leq (\sqrt{\varepsilon_{nm}} + \varepsilon_{\text{GUV}})r$ bad rows. Applying Lemma 2.2, we get that Z is $\zeta = t\sqrt{\varepsilon_{nm}}(rm)^{mt}$ to a (q, t) -non-oblivious bit-fixing source. By our choice of ε_{nm} ,

$$\zeta = 2^{-2mt(d_{\text{GUV}} + \log m)} 2^{mt \log(rm)} \leq 2^{-mt \log r} \leq n^{-mt}.$$

Further,

$$q \leq (\sqrt{\varepsilon_{nm}} + \varepsilon_{\text{GUV}})r \leq 2\varepsilon_{\text{GUV}}r = 2r^{-\frac{1}{2c_{\text{GUV}}} + 1} \leq r^{1 - \frac{1}{4c_{\text{GUV}}}}.$$

We now analyse the running-time. We first apply Ext to compute $S(x_2)$, which takes time $\text{poly}(n, \log(1/\varepsilon_{\text{GUV}})) = \text{poly}(n)$. Then, applying each nmExt takes

$$\text{poly}(n, \log(1/\varepsilon_{nm})) = \text{poly}(n, m, t, d_{\text{GUV}}) = \text{poly}(n)$$

time and we do it for $r = \text{poly}(n)$ times. Overall, the running time is $\text{poly}(n)$, as required. In particular, as $n \geq k \geq m$, the running time is also poly-logarithmic in the errors of the construction, $2^{-k/2}$ and n^{-mt} . \square

4.2 Low error condensers with high entropy loss

Theorem 4.2. *There exists a universal constant $c \geq 1$ such that the following holds. For every integers n, k, m and every $\varepsilon > 0$ such that $n \geq k \geq (m \log(n/\varepsilon))^c$ there exists a $\text{poly}(n)$ -time computable $((n, k), (n, k)) \rightarrow_{\varepsilon, 2^{-k/2}} (m, m - g)$ -condenser*

$$\text{Cond}' : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^m,$$

that is strong in the second source, with entropy gap $g = o(\log(1/\varepsilon))$.

Proof. We start by describing the construction of our condenser Cond' and then turn to the analysis. As usual, we let c_{GUV} be the constant that is given by Theorem 2.4.

Setting of parameters.

- Set $\gamma = \frac{1}{4c_{\text{GUV}}}$ and let $0 < \alpha < \beta < 1$ and c' be the constants from Theorem 3.6 with respect to this γ .
- Set $r = n^{2c_{\text{GUV}}}$.
- Set $t = (m + \log(r/\varepsilon))^{c'}$.
- Set c , the constant stated in this theorem, to $c = \max(10c', 2/\alpha)$.

Building blocks.

- Let $\text{TwoSourcesToNOF}: \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^{r \times m}$ be the function that is given by Theorem 4.1. We are about to apply TwoSourcesToNOF to (n, k) -sources, and indeed k is large enough to satisfy the hypothesis of Theorem 4.1.
- Let $\text{EntRes}: \{0, 1\}^{r \times m} \rightarrow \{0, 1\}^m$ be the function from Theorem 3.6 when set with the parameter γ as defined above. Note that the hypothesis of Theorem 3.6 holds, as since $c \geq 2/\alpha$ we have that $m < r^{\alpha/2}$, and t is large enough.

The construction. On inputs $x_1, x_2 \in \{0, 1\}^n$, we define

$$\text{Cond}'(x_1, x_2) = \text{EntRes}(\text{TwoSourcesToNOF}(x_1, x_2)).$$

Analysis. Clearly, EntRes is computable in $\text{poly}(m, r) = \text{poly}(n)$ time. Let X_1, X_2 be a pair of independent (n, k) -sources. By Theorem 4.1, except with probability $2^{-k/2}$ over $x_2 \sim X_2$, the output $\text{TwoSourcesToNOF}(X_1, x_2)$ is n^{-mt} -close to an $(r^{1-\gamma}, t)$ -non-oblivious bit-fixing source. For every x_2 for which this event holds, the output $\text{EntRes}(\text{TwoSourcesToNOF}(X_1, x_2))$ is $(n^{-mt} + \varepsilon)$ -close to an $(m, m - o(\log(1/\varepsilon)))$ -source, and $n^{-mt} \leq \varepsilon$. \square

4.3 Deterministically condensing a single block-source

A distribution (X, Y) is a blockwise source if both X has sufficient min-entropy and also for every $x \in \text{supp}(X)$, $(Y | X = x)$ has sufficient min-entropy.

Lemma 4.3 (Deterministically condensing a blockwise source). *Let X be an (n, k) -source. Let Y be a d -bit random variable (that may depend on X) such that for every $x \in \text{supp}(X)$, the random variable $(Y | X = x)$ is ε_{B} -close to a $(d, d - g)$ -source.*

Let $\text{Ext}: \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ be a $(k_{\text{Ext}}, \varepsilon_{\text{Ext}})$ -seeded-extractor. Suppose $k \geq k_{\text{Ext}} + \log(1/\varepsilon_{\text{Ext}})$. Then, $\text{Ext}(X, Y)$ is $(2^{g+2}\varepsilon_{\text{Ext}} + 2\varepsilon_{\text{B}})$ -close to an $(m, m - g - 1)$ -source.

Proof. Fix any $T \subseteq \{0, 1\}^m$. Define the set

$$\text{OverHit}_T = \left\{ x \in \{0, 1\}^n : \Pr_{y \sim U_d} [\text{Ext}(x, y) \in T] > \mu(T) + \varepsilon_{\text{Ext}} \right\}.$$

Claim 4.4. $|\text{OverHit}_T| < 2^{k_{\text{Ext}}}$.

Proof. Suppose towards a contradiction that $|\text{OverHit}_T| \geq 2^{k_{\text{Ext}}}$ and let B denote the random variable that is uniform over the set OverHit_T . Since B has min-entropy at least k_{Ext} , the output $\text{Ext}(B, U_d)$ is ε_{Ext} -close to uniform, and therefore $\Pr_{x \sim B, y \sim U_d} [\text{Ext}(x, y) \in T] \leq \mu(T) + \varepsilon_{\text{Ext}}$. This stands in contradiction to the definition of B . \square

Now,

$$\Pr [\text{Ext}(X, Y) \in T] \leq \Pr [\text{Ext}(X, Y) \in T \mid X \notin \text{OverHit}_T] + \Pr [X \in \text{OverHit}_T].$$

By Claim 4.4, $\Pr [X \in \text{OverHit}_T] \leq 2^{k_{\text{Ext}}-k}$. Also, for every $x \notin \text{OverHit}_T$ let

$$GY_x = \left\{ y \in \{0, 1\}^d : E(x, y) \in T \right\}.$$

By definition, $\mu(GY_x) \leq \mu(T) + \varepsilon_{\text{Ext}}$. Therefore,

$$\Pr_{y \sim (Y|X=x)} [E(x, y) \in T] = \Pr_{y \sim (Y|X=x)} [y \in GY_x] \leq \varepsilon_B + |GY_x|2^{g-d} \leq \varepsilon_B + 2^g(\mu(T) + \varepsilon_{\text{Ext}}).$$

Thus,

$$\begin{aligned} \Pr [\text{Ext}(X, Y) \in T] &\leq \Pr [\text{Ext}(X, Y) \in T \mid X \notin \text{OverHit}_T] + \Pr [X \in \text{OverHit}_T] \\ &\leq 2^g \mu(T) + 2^g \varepsilon_{\text{Ext}} + \varepsilon_B + 2^{k_{\text{Ext}}-k} \\ &\leq 2^g \mu(T) + (2^g + 1) \varepsilon_{\text{Ext}} + \varepsilon_B. \end{aligned}$$

But,

Claim 4.5. Let Z be a random variable over n -bit strings such that for every $T \subseteq \{0, 1\}^n$, $\Pr [Z \in T] \leq 2^g \mu(T) + \varepsilon$. Then, Z is 2ε -close to an $(n, n - g - 1)$ -source.

Proof. Set $H = \{x : \Pr [X = x] > 2^{-(n-g-1)}\}$. On the one hand,

$$\Pr [Z \in H] = \sum_{x \in H} \Pr [X = x] \geq 2^{g+1} \mu(H).$$

On the other hand, by our assumption, $\Pr [X \in H] \leq 2^g \mu(H) + \varepsilon$. Together, we get that $2^g \mu(H) \leq \varepsilon$. Thus, $\Pr [X \in H] \leq 2\varepsilon$. As H are all the heavy elements, we conclude that Z is 2ε -close to a distribution with $n - g - 1$ min-entropy. \square

We can therefore summarize that $\text{Ext}(X, Y)$ is $(2^{g+2} \varepsilon_{\text{Ext}} + 2\varepsilon_B)$ -close to an $(m, m - g - 1)$ -source. \square

4.4 Low error condensers with small entropy gap outputting many bits

In this section we prove the following theorem that readily implies Theorem 1.3.

Theorem 4.6 (Main theorem). *There exists a universal constant $c \geq 1$ such that the following holds. For every integers $n \geq k$ and every $\varepsilon > 0$ such that $k \geq \log^c(n/\varepsilon)$ there exists a $\text{poly}(n, \log(1/\varepsilon))$ -time computable $((n, k), (n, k)) \rightarrow_\varepsilon (m, m - g)$ two-source condenser*

$$\text{Cond}: \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^m$$

with $m = k - 5 \log(1/\varepsilon) - O(1)$ and $g = o(\log(1/\varepsilon))$.

Proof. We start by setting some parameters.

Parameters.

- Set $\varepsilon_{\text{Cond}'} = \varepsilon/8$.
- Set $\varepsilon_{\text{Ext}} = \varepsilon^2/32$.
- Set $k_{\text{Ext}} = k - \log(2/\varepsilon)$.
- Set $d_{\text{Ext}} = c' \log n \cdot \log(n/\varepsilon_{\text{Ext}})$ where c' is the constant that is given by Theorem 2.6.

For the construction we make use of the following building blocks.

- Let $\text{Ext}: \{0, 1\}^n \times \{0, 1\}^{d_{\text{Ext}}} \rightarrow \{0, 1\}^m$ be the $(k_{\text{Ext}}, \varepsilon_{\text{Ext}})$ -strong-seeded-extractor that is given by Theorem 2.6. By that theorem, $m = k_{\text{Ext}} - 2 \log(1/\varepsilon_{\text{Ext}}) - O(1)$.
- Let $\text{Cond}': \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^{d_{\text{Ext}}}$ be the $((n, k), (n, k)) \rightarrow_{\varepsilon_{\text{Cond}'}, 2^{-k/2}} (d_{\text{Ext}}, d_{\text{Ext}} - g')$ -condenser, strong in the second source, that is given by Theorem 4.2, with $g' = o(\log(1/\varepsilon_{\text{Cond}'}))$. Note that our choice of parameters satisfies the hypothesis of Theorem 4.2 for a large enough constant c .

The construction. On inputs $x_1, x_2 \in \{0, 1\}^n$, we define

$$\text{Cond}(x_1, x_2) = \text{Ext}(x_2, \text{Cond}'(x_1, x_2)).$$

Analysis. Let X_1, X_2 be a pair of independent (n, k) -sources. By Theorem 4.2, with probability at least $1 - 2^{-k/2}$ over $x_2 \sim X_2$, the random variable $\text{Cond}'(X_1, x_2)$ is $\varepsilon_{\text{Cond}'}$ -close to a $(d, d - g')$ -source. Lemma 4.3 implies that $\text{Ext}(X_2, \text{Cond}'(X_1, X_2))$ is $2^{-k/2} + (2^{g'+2}\varepsilon_{\text{Ext}} + 2\varepsilon_{\text{Cond}'})$ -close to an $(m, m - g' - 1)$ -source.

By our choice of parameters, $2^{-k/2} + 2^{g'+1}\varepsilon_{\text{Ext}} + 2\varepsilon_{\text{Cond}'} \leq \varepsilon$. Note that $k - m = \log(2/\varepsilon) + 2 \log(1/\varepsilon_{\text{Ext}}) = 5 \log(1/\varepsilon) + O(1)$. The running-time of the construction readily follows from the running-times of Cond' and Ext . \square

References

- [AGM03] N. Alon, O. Goldreich, and Y. Mansour. Almost k -wise independence versus k -wise independence. *Information Processing Letters*, 88(3):107–110, 2003.
- [AL93] M. Ajtai and N. Linial. The influence of large coalitions. *Combinatorica*, 13(2):129–145, 1993.
- [BACD⁺18] A. Ben-Aroya, E. Chattopadhyay, D. Doron, X. Li, and A. Ta-Shma. A new approach for constructing low-error, two-source extractors. In *LIPICs-Leibniz International Proceedings in Informatics*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018. (To Appear).
- [BADTS17] A. Ben-Aroya, D. Doron, and A. Ta-Shma. An efficient reduction from two-source to non-malleable extractors: achieving near-logarithmic min-entropy. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 1185–1194. ACM, 2017.
- [BKS⁺10] B. Barak, G. Kindler, R. Shaltiel, B. Sudakov, and A. Wigderson. Simulating independence: New constructions of condensers, Ramsey graphs, dispersers, and extractors. *Journal of the ACM (JACM)*, 57(4):20, 2010.
- [BOL85] M. Ben-Or and N. Linial. Collective coin flipping, robust voting schemes and minima of banzhaf values. In *Proceedings of the 26th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 408–416. IEEE, 1985.
- [Bou05] J. Bourgain. More on the sum-product phenomenon in prime fields and its applications. *International Journal of Number Theory*, 1(01):1–32, 2005.
- [Bra10] M. Braverman. Polylogarithmic independence fools AC0 circuits. *Journal of the ACM (JACM)*, 57(5):28, 2010.
- [BRSW12] B. Barak, A. Rao, R. Shaltiel, and A. 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.
- [BSZ11] E. Ben-Sasson and N. Zewi. From affine to two-source extractors via approximate duality. In *Proceedings of the 43rd annual ACM Symposium on Theory of computing (STOC)*, pages 177–186. ACM, 2011.
- [CG88] B. Chor and O. Goldreich. Unbiased bits from sources of weak randomness and probabilistic communication complexity. *SIAM Journal on Computing*, 17(2):230–261, 1988.

- [CGL16] E. Chattopadhyay, V. Goyal, and X. Li. Non-malleable extractors and codes, with their many tampered extensions. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 285–298. ACM, 2016.
- [CL16] E. Chattopadhyay and X. Li. Explicit non-malleable extractors, multi-source extractors and almost optimal privacy amplification protocols. In *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2016.
- [Coh15] G. Cohen. Local correlation breakers and applications to three-source extractors and mergers. In *IEEE 56th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 845–862. IEEE, 2015.
- [Coh16a] G. Cohen. Making the most of advice: New correlation breakers and their applications. In *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2016.
- [Coh16b] G. Cohen. Two-source dispersers for polylogarithmic entropy and improved Ramsey graphs. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 278–284. ACM, 2016.
- [Coh16c] G. Cohen. Two-source extractors for quasi-logarithmic min-entropy and improved privacy amplification protocols. In *Electronic Colloquium on Computational Complexity (ECCC)*, number 114, 2016.
- [CS16] G. Cohen and L. Schulman. Extractors for near logarithmic min-entropy. In *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, page 14, 2016.
- [CZ16] E. Chattopadhyay and D. Zuckerman. Explicit two-source extractors and resilient functions. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC)*, pages 670–683. ACM, 2016.
- [GUV09] V. Guruswami, C. Umans, and S. Vadhan. Unbalanced expanders and randomness extractors from Parvaresh–Vardy codes. *Journal of the ACM*, 56(4):20, 2009.
- [HS16] P. Harsha and S. Srinivasan. On polynomial approximations to AC0. *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, 2016.

- [KKL88] K. Kahn, G. Kalai, and N. Linial. The influence of variables on boolean functions. In *Proceedings of the 29th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 68–80. IEEE, 1988.
- [Li13a] X. Li. Extractors for a constant number of independent sources with polylogarithmic min-entropy. In *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 100–109, 2013.
- [Li13b] X. 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] X. Li. Three-source extractors for polylogarithmic min-entropy. In *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 863–882. IEEE, 2015.
- [Li17] X. 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 (STOC)*, pages 1144–1156. ACM, 2017.
- [Li18] X. Li. Pseudorandom correlation breakers, independence preserving mergers and their applications. In *Electronic Colloquium on Computational Complexity (ECCC)*, number 28, 2018.
- [Mek17] R. Meka. Explicit resilient functions matching ajtai-linial. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1132–1148. SIAM, 2017.
- [Neu51] J. V. Neumann. Various techniques used in connection with random digits. 12(1):36–38, 1951.
- [Rao08] A. Rao. A 2-source almost-extractor for linear entropy. In *Approximation, Randomization and Combinatorial Optimization. Algorithms and Techniques*, pages 549–556. Springer, 2008.
- [Rao09] A. Rao. Extractors for a constant number of polynomially small min-entropy independent sources. *SIAM Journal on Computing*, 39(1):168–194, 2009.
- [Raz05] R. Raz. Extractors with weak random seeds. In *Proceedings of the 37th annual ACM Symposium on Theory of Computing (STOC)*, pages 11–20, 2005.
- [RRV99] R. Raz, O. Reingold, and S. Vadhan. Error reduction for extractors. In *40th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 191–201. IEEE, 1999.

- [Tal17] A. Tal. Tight bounds on the fourier spectrum of AC0. In *LIPICs-Leibniz International Proceedings in Informatics*, volume 79. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.
- [Vaz86] U. V. Vazirani. *Randomness, Adversaries and Computation*. University of California, Berkeley, 1986.
- [Zuc97] D. Zuckerman. Randomness-optimal oblivious sampling. *Random Structures and Algorithms*, 11(4):345–367, 1997.
- [Zuc07] D. Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. *Theory of Computing*, 3:103–128, 2007.