

Extractors for Adversarial Sources via Extremal Hypergraphs

Eshan Chattopadhyay*
Cornell University
eshanc@cornell.edu

Jesse Goodman*
Cornell University
jpmgoodman@cs.cornell.edu

Vipul Goyal†
Carnegie Mellon University
vipul@cmu.edu

Xin Li‡
Johns Hopkins University
lixints@cs.jhu.edu

December 13, 2019

Abstract

Randomness extraction is a fundamental problem that has been studied for over three decades. A well-studied setting assumes that one has access to multiple independent weak random sources, each with some entropy. However, this assumption is often unrealistic in practice. In real life, natural sources of randomness can produce samples with no entropy at all or with unwanted dependence. Motivated by this and applications from cryptography, we initiate a systematic study of randomness extraction for the class of *adversarial sources* defined as follows.

A weak source \mathbf{X} of the form $\mathbf{X}_1, \dots, \mathbf{X}_N$, where each \mathbf{X}_i is on n bits, is an (N, K, n, k) -source of locality d if the following hold:

1. **Somewhere good sources:** at least K of the \mathbf{X}_i 's are independent, and each contains min-entropy at least k . We call these \mathbf{X}_i 's *good sources*, and their locations are unknown.
2. **Bounded dependence:** each remaining (bad) source can depend arbitrarily on at most d good sources.

We focus on constructing extractors with negligible error, in the regime where most of the entropy is contained *within* a few sources instead of *across* many (i.e., k is at least polynomial in K). In this setting, even for the case of 0-locality, very little is known prior to our work. For $d \geq 1$, essentially no previous results are known. We present various new extractors for adversarial sources in a wide range of parameters, and some of our constructions work for locality $d = K^{\Omega(1)}$. As an application, we also give improved extractors for small-space sources.

The class of adversarial sources generalizes several previously studied classes of sources, and our explicit extractor constructions exploit tools from recent advances in extractor machinery, such as two-source non-malleable extractors and low-error condensers. Thus, our constructions can be viewed as a new application of non-malleable extractors. In addition, our constructions combine the tools from extractor theory in a novel way through various sorts of explicit extremal hypergraphs. These connections leverage recent progress in combinatorics, such as improved bounds on cap sets and explicit constructions of Ramsey graphs, and may be of independent interest.

*Supported by NSF grant CCF-1849899.

†Supported in part by NSF grant 1916939, a gift from Ripple, a JP Morgan Faculty Fellowship, and a CyLab seed funding award.

‡Supported by NSF Award CCF-1617713 and NSF CAREER Award CCF-1845349.

1 Introduction

The use of randomness is widespread in computer science, particularly in areas such as cryptography, algorithm design, and distributed computing. Randomness is also useful in running Monte Carlo simulations of complex systems and in various sampling tasks. It is often the case that these applications crucially need access to high-quality randomness, i.e., a stream of uniform and independent bits. For instance, it was shown [DOPS04] that it is impossible to do basic cryptographic tasks such as bit commitment schemes and secret sharing schemes without access to high-quality random bits. This poses a challenging problem since most sources of randomness in nature are typically far from producing pure random bits, and in fact produce a stream of correlated bits containing little or no entropy. In addition, even originally high quality random bits can be compromised adversarially by side channel attacks.

The area of randomness extraction is motivated by the above problem. Informally, a randomness extractor is a *deterministic algorithm* that purifies a weak random source to produce a distribution that is close to uniform. As is standard in this area, we measure the randomness of a weak source \mathbf{X} using min-entropy, defined as:

$$H_\infty(\mathbf{X}) := \min_x \{-\log(\Pr[\mathbf{X} = x])\}.$$

Define an (n, k) -source to be a distribution on $\{0, 1\}^n$ with min-entropy at least k , and the entropy rate to be k/n . Thus, if \mathbf{X} is an (n, k) -source, then for any $x \in \{0, 1\}^n$, we have $\Pr[\mathbf{X} = x] \leq 2^{-k}$.

Definition 1.1. Let \mathcal{X} be a family of distributions over $\{0, 1\}^n$. We say that a function $\text{Ext} : \{0, 1\}^n \rightarrow \{0, 1\}^m$ is an extractor for \mathcal{X} with error ϵ if, for all $\mathbf{X} \in \mathcal{X}$,

$$|\text{Ext}(\mathbf{X}) - \mathbf{U}_m| \leq \epsilon.$$

Here $|\cdot|$ refers to the standard statistical distance, \mathbf{U}_m denotes the uniform distribution on m bits, and ϵ is known as the error of the extractor. A folklore result shows that it is impossible to extract even one random bit from a single (n, k) -source. More precisely, there cannot exist an extractor $\text{Ext} : \{0, 1\}^n \rightarrow \{0, 1\}$ such that for any $(n, n-1)$ -source \mathbf{X} , $|\text{Ext}(\mathbf{X}) - \mathbf{U}_1| < 1/2$.

Given the above bottleneck, there are two major directions that researchers have explored in randomness extraction over the last 3 decades. The first is to assume access to a short independent uniform *seed* \mathbf{U}_d to extract randomness out of a single (n, k) -source \mathbf{X} . Such extractors are called seeded extractors, and from a beautiful line of work we now have constructions with near optimal parameters [LRVW03, GUV09, DKSS13].

The second direction, which is more relevant to this paper, assumes special structures in the weak source \mathbf{X} . In particular, the most well studied model assumes that \mathbf{X} is of the form $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_C$, where each \mathbf{X}_i is an independent (n, k) -source. Indeed, recently there has been an exciting line of work on extracting randomness from independent sources, which we discuss in more details in Section 1.3. However, these works typically assume that all the sources are independent and have sufficient min-entropy, which is often unrealistic in practice. In real life, computers generate random numbers by combining various “unpredictable” sources such as keystrokes, mouse movements, timestamps, processor temperatures, and so on. It is quite possible that some of these sources are “bad” in the following senses. First, some of them may be predictable and thus contain no entropy. Second, while it is reasonable to assume *some* independence across the sources, there can also certainly be some degree of (adversarial) dependence between them. Developing a theory

of randomness extraction in the presence of adversarial sources is thus a natural generalization of the well-studied model of independent sources, and may eventually help us build better random number generators for computers. To the best of our knowledge, little work has been done in this setting, and in this paper we seek to initialize a systematic study of this natural question.

1.1 Adversarial sources

To capture the setting discussed above, we generalize the model of independent sources in two non-trivial ways and introduce the class of adversarial sources.

Definition 1.2. *Let N, K, n, k, d be nonnegative integers. A distribution $\mathbf{X} = \mathbf{X}_1, \dots, \mathbf{X}_N$, where each \mathbf{X}_i is on n bits, is called an (N, K, n, k) -source of locality d , if the following conditions hold:*

1. **Somewhere good sources:** *There is a set $S \subseteq [N]$, $|S| \geq K$ such that for any $i \in S$, $H_\infty(\mathbf{X}_i) \geq k$. We call the sources \mathbf{X}_i , $i \in S$ good sources and the remaining bad sources.*
2. **Bounded dependence:** *The set of good sources are independent, and each bad source can depend arbitrarily on at most d good sources.*

As discussed before, the *somewhere good sources* condition captures the natural setting where a physical source of randomness (e.g., a Zener diode) outputs a stream of bits, where entropy is localized in certain unknown chunks. The bounded dependence condition captures possible troublesome dependence between chunks of different bits. As it turns out, our model also has natural motivations from cryptography.

In the domain of cryptography, extractors for adversarial sources may allow us to generate a uniform random string with the help of several parties each having an imperfect random source, even if some of these parties are adversarial. As a simple example, consider coin flipping protocols with synchronous channels. If all parties simply broadcast their strings, we get several strings which are good (but imperfect) and some other strings which can be adversarially chosen (though independent of the good strings). By applying an extractor for adversarial sources with 0-locality, one can then obtain a uniform random string. Going to asynchronous channels, the strings of adversarial parties may depend on a set of good strings due to the order of messages in the protocol, and hence extractors for adversarial sources with larger locality can be useful.

As another example, several primitives in cryptography such as non-interactive zero knowledge (NIZK) require a random “common reference string” (CRS). A number of works have investigated the setting where the CRS might be imperfect [CPS07, LPV09] and even the setting where there are multiple CRS and some of them may be adversarially chosen [GK08, GGJS11, GO14] (but the good ones are uniform). Extractors for adversarial sources may allow us to handle the second setting.

1.2 Summary of our results

We will be mainly interested in extracting from (N, K, n, k) -sources of locality d in the negligible error setting, motivated by applications in cryptography. Further, we will focus on the setting $k \geq K^\gamma$, for any constant $\gamma > 0$ (i.e., entropy is more concentrated *within* a few sources, rather than spread across them; or, roughly, there are a few long sources). Here, our goal is to construct extractors with output and error of the form $m = k^{\Omega(1)}, \epsilon = 2^{-k^{\Omega(1)}}$. In Section 6, we motivate our study of this regime and show that in the complementary regime, there is a relatively simple construction based on prior work. In our setting of interest, the only known result is in the

case of 0-locality, where the work of Kamp et al. [KRVZ06] implies negligible error extractors for (N, K, n, k) -sources, as long as $Kk \geq (Nn)^{1-\gamma}$, for some tiny constant $\gamma > 0$ arising from estimates in additive combinatorics. For the case of d -locality with $d \geq 1$, to the best of our knowledge there are no known previous results. We discuss other related prior work in [Section 1.3](#).

In our first three main theorems, we construct an explicit extractor for adversarial sources that produces *polynomially many bits* with *negligible error*, even if the good sources have just *poly-logarithmic entropy*. Several of our extractors use the small parameter \mathcal{R}_N , which we define below.

Definition 1.3. *We let \mathcal{R}_N denote the smallest number such that there exists an explicit construction of bipartite Ramsey graphs over $2N$ vertices with no bipartite clique nor independent set of size $2\mathcal{R}_N$. Currently, $\mathcal{R}_N = (\log N)^{o(\log \log N)}$, and this also holds for non-bipartite Ramsey graphs [Li19].*

In our first main theorem, we extract from (N, K, n, k) -sources of locality 0, given just $K \geq \mathcal{R}_N^2$ good sources, as long as one extra condition holds:

Theorem 1. *There exist universal constants $C, \gamma > 0$ such that for all $N, K, n, k \in \mathbb{N}$ satisfying $k \geq \log^C n$ and $K \geq \mathcal{R}_N^2$, there exists an explicit extractor $\text{Ext} : (\{0, 1\}^n)^N \rightarrow \{0, 1\}^m$ for (N, K, n, k) -sources of locality 0, with output length $m = k^{\Omega(1)}$ and error $\epsilon = 2^{-k^{\Omega(1)}}$, provided $N \leq k^\gamma$.*

Thus, for 0-local sources, we obtain extractors for extremely small k and K , under the condition that the number of sources is not too large compared to the entropy in the good sources, i.e., $N \leq k^\gamma$. It is natural to ask if we can completely remove this restriction. Our second main theorem does exactly this.

Theorem 2. *There exists a universal constant $C > 0$ such that for all $N, K, n, k \in \mathbb{N}$ satisfying $k \geq \log^C n$ and $K \geq \sqrt{N} \cdot \mathcal{R}_N$, there exists an explicit extractor $\text{Ext} : (\{0, 1\}^n)^N \rightarrow \{0, 1\}^m$ for (N, K, n, k) -sources of locality 0, with output length $m = k^{\Omega(1)}$ and error $\epsilon = 2^{-k^{\Omega(1)}}$.*

Thus, we see that if we increase the number of good sources from $K \geq N^{o(1)}$ to $K \geq N^{0.5+o(1)}$, we are able to remove any restriction between N and k . Our third main theorem shows that, in fact, we can extend our constructions to handle *polynomial locality*. We state an interesting parameter setting of our more general theorem ([Theorem 5.1](#)) here:

Theorem 3. *There exist universal constants $C, \gamma > 0$ such that for all $N, K, n, k, d \in \mathbb{N}$ satisfying $k \geq \log^C n$ and $K \geq N^{1-\gamma}$, and $d \leq K^\gamma$, there exists an explicit extractor $\text{Ext} : (\{0, 1\}^n)^N \rightarrow \{0, 1\}^m$ for (N, K, n, k) -sources of locality d , with output length $m = k^{\Omega(1)}$ and error $\epsilon = 2^{-k^{\Omega(1)}}$, provided $N \leq k^\gamma$.*

We also show (non-explicitly) that extractors with negligible error exist for adversarial sources that contain just $K = N^\alpha$ good sources and have locality $d = K^{1-\alpha}$, for any constant $\alpha > 0$.

Theorem 4. *For any constant $0 < \gamma < 1$ there exists a constant $\alpha > 0$ such that for all $N, K, n, k, d \in \mathbb{N}$ satisfying $k \geq (1 + \gamma) \log n$ and $K \geq N^\gamma$, and $d \leq K^{1-\gamma}$, there exists a (possibly non-explicit) extractor for (N, K, n, k) -sources of locality d with output length $m = k^{\Omega(1)}$ and error $\epsilon = 2^{-\Omega(k)}$, provided $N \leq k^\alpha$.*

The theorem that we prove is actually more general (see [Theorem 7.2](#)), and the proof makes use of a more robust variant of seedless non-malleable extractors that we introduce. We also show that it is impossible to construct an extractor for adversarial sources if half of the sources are good, but each bad source can depend on all the good sources. For more details, we refer the reader to [Section 7](#).

Finally, we show that our constructions also give improved extractors for sources sampled by algorithms that have limited memory, in the negligible error regime. These sources were initially studied by [\[KRVZ06\]](#), and fit into the line of work initiated by [\[TV00\]](#) on extracting from sources that are samplable using limited resources.

Theorem 5. *For any fixed $\gamma > 0$ and all $n, k, s \in \mathbb{N}$ satisfying $k \geq n^{2/3+\gamma}$ and $s \leq (k/n)^{3+\gamma} \cdot n$, there exists an explicit extractor $\text{Ext} : \{0, 1\}^n \rightarrow \{0, 1\}^m$ for space s sources of min-entropy k , with output length $m = n^{\Omega(1)}$ and error $\epsilon = 2^{-n^{\Omega(1)}}$.*

Previously, the best extractor for s -space sources [\[KRVZ06\]](#) with negligible error required min-entropy $k \geq n^{1-\gamma}$ (for some tiny constant $\gamma > 0$) for about the same space $s \leq (k/n)^3 n$, and had error $2^{-n^{\Omega(1)}}$. In the same paper, Kamp et al. reduce the entropy requirement to $k > n^{0.81}$ for space $s = 1$ sources with an extra restriction. We note that [Theorem 5](#) reduces the entropy requirement to $k > n^{0.67}$, and works for large space with no such restrictions.

We present an overview of our techniques and explicit constructions in [Section 2](#), but first we review some relevant prior work.

1.3 Relevant prior work

Relation of adversarial sources to other structured sources Special cases of adversarial sources are studied by works on randomness extraction for some other kinds of sources in prior work. Hence our model of adversarial sources can also be viewed as a generalization of several previous models. We discuss some details below.

- *Bit-fixing sources:* Oblivious bit-fixing sources correspond to (N, K, n, k) -sources of locality 0, with $n = k = 1$. Thus, they are distributions on $\{0, 1\}^N$, with some unknown K coordinates being uniform and independent, while the rest of the bits are fixed and do not depend on the random bits. They are studied in the works [\[CGH+85, KZ06, GRS06\]](#). The best known extractors in different regimens of error are the following: (i) Kamp and Zuckerman [\[KZ06\]](#) constructed an extractor that works for any $K > 0$ with error $1/\text{poly}(K)$, and (ii) Rao [\[Rao09b\]](#) constructed an extractor that works for any $K \geq \text{poly}(\log N)$ with error $2^{-K^{\Omega(1)}}$.

Non-oblivious bit-fixing sources allow the non-random bits to arbitrarily depend on the random bits. Thus, they correspond to $(N, K, 1, 1)$ -sources of locality K . The best known results [\[Mek17, CZ19\]](#) can handle $K \geq N - O\left(\frac{N}{\log^2 N}\right)$, with error $1/N^{\Omega(1)}$. The KKL theorem [\[KKL88\]](#) implies that the best K one could hope for in this setting is $N - O\left(\frac{N}{\log N}\right)$.

- *Symbol-fixing sources:* Kamp and Zuckerman [\[KZ06\]](#) introduced the class of symbol fixing sources, generalizing bit-fixing sources. Symbol-fixing sources correspond to (N, K, n, k) -sources with $k = n$. The locality is 0 for oblivious symbol-fixing sources and is K for non-oblivious symbol fixing sources. The results mentioned above on total entropy sources capture the best known extractors for oblivious symbol-fixing sources. To the best of our knowledge,

there is no non-trivial construction of extractors for non-oblivious symbol-fixing sources other than using known extractors for non-oblivious bit-fixing sources.

- *Independent sources*: The most well-studied model of seedless extraction assumes that the weak source \mathbf{X} is of the form $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_C$, where each \mathbf{X}_i is an *independent* (n, k) -source. Thus, these sources correspond to (C, C, n, k) -sources of locality 0. The probabilistic method provides existential proof of extractors for such sources, called C -source extractors, with strong parameters. In particular, it can be shown that there exists a 2-source extractor with error ϵ for $k \geq \log n + 2 \log(1/\epsilon) + O(1)$.

An explicit construction of a 2-source extractor was given by Chor and Goldreich [CG88], but they required min-entropy $k > n/2$ for both of the sources. The entropy requirement was marginally improved by Bourgain [Bou05] to $k > 0.499n$, and Raz [Raz05] improved the entropy requirement of one of the sources to $O(\log n)$ (but required the other source to have entropy $> n/2$). Recently, an impressive line of work [Coh16b, CL16a, Li16, BADTS17, Coh17, Li17, Mek17, BDT18, CZ19, Li19] improved the entropy requirement to $(\log n)^{1+o(1)}$. However, the recent progress has a major drawback in terms of the error parameter, and in particular, the best known 2-source extractor construction for error $\epsilon = 1/n^{\omega(1)}$ requires min-entropy $(1/2 - \delta)n$, for some small constant δ [Bou05, Lew19].

Assuming access to 3 or more independent sources, a long line of work [BKS⁺05, BIW06, Rao09a, Li11a, Li13a, Li13b, Li15b, Coh16a] explicitly constructed excellent extractors. In particular, Li [Li15b] constructed an explicit 3-source extractors with $k \geq \text{poly}(\log n)$ and error $2^{-k^{\Omega(1)}}$.

Also closely related to adversarial sources are *total entropy sources*. Introduced by Koenig and Maurer [KM05], an (N, n, Γ) -total entropy source consists of N independent sources of length n such that the *sum* of min-entropies across the sources is at least Γ . Thus, an (N, K, n, k) -source of locality 0 is an (N, n, Kk) -total entropy source. Plugging in the best known extractor for total entropy sources in the regime of negligible error [KRVZ06] implies an explicit extractor for (N, K, n, k) -sources of 0-locality with error $2^{-n^{\Omega(1)}}$ as long as $Kk \geq (Nn)^{1-\gamma}$, for some tiny constant γ that arises from sum-product estimates in additive combinatorics.

Kamp et al. [KRVZ06] constructed total entropy extractors in another extreme setting of parameters, where there are a large number of short sources. Their results imply explicit extractors for (N, K, n, k) -sources of 0-locality, as long as $Kk \geq \omega(2^n \sqrt{Nn})$. The error of the extractor is $2^{\Omega(-(Kk)^2/(Nn2^{2n}))}$, and the extractor runs in time $\text{poly}(N, 2^n)$. Thus, this gives an explicit construction with negligible error as long as $n = o(\log N)$ (i.e., the number of sources is exponential in the length of the sources).

Finally, in the regime of larger error, Chattopadhyay and Li [CL16b] constructed an explicit extractor for $(N, 2, n, \text{poly}(\log n))$ -sources of locality 0. They refer to these sources as $(n, \text{poly}(\log n), N)$ -somewhere-2 sources, and the error of the extractor in their construction is $\epsilon = 1/n^{\Omega(1)}$.

Other models of seedless extraction Apart from the models discussed above, other examples of structured sources that have been studied by researchers include affine sources [Bou07, Li11b, Yeh11, Li16], polynomial and variety sources [DGW09, Dvi12], sources sampled by small-space algorithms [KRVZ06, CL16b], and sources sampled by small circuits [TV00, Vio14, Li16].

Comparison to SHELA sources Very recently, a work by Aggarwal et. al. [AOR⁺19] introduced another model that generalizes independent sources by allowing dependence, which they call SHELA (Somewhere Honest Entropic Look Ahead) sources, and studied randomness extraction in this model. Although similar in spirit to our model, there are several important differences.

First, both models can be viewed as a stream of ℓ sources, where some unknown t sources are good, and the positions of the good sources are selected by an adversary. The rest of the sources are bad and can depend on the good sources. However, in the model of SHELA sources, the dependence is only *one-way*: any bad source can only depend on good sources that come before it, and hence the name “look-ahead”. In contrast, in our model the dependence of bad sources on good sources can go *both ways*: a bad source can depend on good sources both before it and after it. In this sense our model is more general than theirs.

Second, [AOR⁺19] shows that if the fraction of good sources is any *constant* γ , then it is impossible to achieve randomness extraction from SHELA sources if the number of blocks ℓ is a large enough constant depending on γ . Thus, the authors turn to a less ambitious goal: to obtain from such sources a convex combination of *somewhere random sources*. A “ T -out-of- L ” somewhere random source is a sequence of L sources, where some fixed but unknown T sources are jointly independent and uniformly distributed, while the other sources can have arbitrary dependence on the T sources. Aggarwal et al. show that they can construct efficient *somewhere extractors* which output a convex combination of T -out-of- L somewhere random sources with a small L , when the fraction of good sources γ is any constant. They then show that these somewhere random sources can still be used in several cryptographic applications. In contrast, our results build real randomness extractors for adversarial sources, and thus the outputs in our constructions are truly close to uniform, instead of being only somewhere random. Hence, they can be used in all cryptographic applications universally. Moreover, we give explicit constructions even when the fraction of good sources is sub-constant (e.g., $K = N^{1-\gamma}$ for some constant γ), as compared to a constant fraction in [AOR⁺19]. However, in order to circumvent the impossibility result in [AOR⁺19], we have to limit our locality d to be not too large. In this sense, our model and the model in [AOR⁺19] are incomparable.

Finally, the entropy requirement on good sources in our work is better than that in [AOR⁺19]. In [AOR⁺19] the constructions require the good sources to have linear min-entropy, while all our constructions work for the case of $k \geq \text{poly}(\log n)$.

Organization First, we provide an overview of our constructions in [Section 2](#). Then, in [Section 3](#), we define several preliminaries that will be important for formalizing these ideas. We present our main extractors for adversarial sources of locality 0 in [Section 4](#). Then, in [Section 5](#), we present our extractors for adversarial sources of polynomial locality. These extractors work best when most of the entropy is contained *within* the sources, rather than *across them*; we motivate this setting by giving a simple explicit construction for the complementary setting in [Section 6](#). Next, we put our results in context by providing existential and impossibility results in [Section 7](#). (Our existential results rely on a new generalized seedless non-malleable extractor; in [Appendix A](#), we introduce these objects and prove their existence.) Finally, we show that our explicit constructions give improved extractors for total entropy and small-space sources in [Section 8](#), and in [Section 9](#) we suggest future directions of research.

2 Overview of our constructions

At a high level, all our constructions use two key ideas. The first idea is to design a well-structured hypergraph around the N sources (represented as vertices), and try to extract separately from each hyperedge. While it is easy to guarantee that some (unknown) hyperedge produces uniform bits, we must produce a *single* uniform string. Thus, we must combine the output from each hyperedge and hope that the uniform bits are not destroyed in the process. In all our constructions this is done by computing the XOR of the outputs.

This brings us to our second key idea. In order for the XOR to work, we need to break the correlations between the uniform output bits from some hyperedge and the outputs from the other hyperedges. For this purpose we crucially rely on recent constructions of *non-malleable extractors*. We identify and explicitly construct certain classes of extremal hypergraphs with the following goals: to minimize the size of their largest independent set (for some general notion of independent set), while maintaining some sort of limited interaction between their hyperedges. The size of largest independent set controls the number of good sources we need, while the limited interaction will make it easier to break correlations between the random variables produced by the hyperedges, using the property of non-malleable extractors.

2.1 Extracting from 0-locality

Our first goal is to construct negligible-error extractors for (N, K, n, k) -sources of locality 0. As shown in [CL16b], for $K = 2$ and $k = 0.51n$, this is straightforward: we may simply call the 2-source Hadamard extractor (Theorem 3.2), **Had**, over all pairs of sources, and take the bitwise XOR of the results. This works because some call to **Had** must use the two good sources (call them **X** and **Y**), and the remaining calls use at most one of **X**, **Y**. If we fix the XOR of the calls that use **X** but not **Y**, we introduce no correlation between them, and Lemma 3.9 tells us that the entropy of **X** drops by very little. We can do the same for the calls that use **Y** but not **X**. This shows that with high probability, the last remaining call to **Had** outputs near-uniform bits, and they remain uniform after taking their bitwise XOR with the fixed bits.

It is natural to ask if we can extract with negligible error from much smaller k , if we allow larger K . Because there exist explicit constructions of negligible-error three-source extractors for polylogarithmic entropy (Theorem 3.3), the naive idea would be to alter the above construction to call a three-source extractor **3Ext** over all triples of sources, and XOR the results. It is true that for just $K = 3$, some call to **3Ext** in this construction is guaranteed to use three good sources. However, it will also be the case that there are other calls that use two of the good sources, and we cannot fix these outputs without introducing correlation between them. Thus, this idea fails.

In order to replace **Had** in the above construction with a different extractor (say, **3Ext**) that can handle lower entropy, we must do something more clever than just applying **3Ext** over all triples of sources. The main idea behind our 0-local extractors is that we must *carefully select* triples over which to call **3Ext**, in order to ensure two properties:

1. **Activation**: given K good sources, some call to **3Ext** is guaranteed to use three good sources.
2. **Fragile correlation**: all other calls to **3Ext** can be fixed without ruining the near-uniform output of the good call (i.e., without destroying the entropy of its inputs or introducing correlation between them).

If we can accomplish this, then we can reduce the entropy requirement of the good sources from $k = 0.51n$ to $k = \log^C n$, for some universal constant $C \geq 1$. This can be easily achieved if we have $K > 2N/3$ good sources by simply calling `3Ext` over disjoint sets of sources. However, we want to accomplish the above using as few good sources, K , as possible. To do this, we will design a hypergraph over N vertices whose hyperedges will be used to select triples of sources (vertices) on which to call `3Ext`.¹ The hypergraph will have a structural constraint that will guarantee *fragile correlation*, and we seek such a hypergraph with the smallest possible max independent set (for some generalized notion), which roughly corresponds to the number of sources needed for *activation*.

The STS-extractor To be more concrete, we must answer the following question: what structure must a 3-uniform hypergraph have such that if some hyperedge is *activated* (contains three good sources), then every other hyperedge makes a call to `3Ext` that can be safely fixed without ruining the output of the call to `3Ext` from the activated hyperedge? One answer is to enforce that each pair of hyperedges share at most one source. In particular, if the activated hyperedge contains good sources $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$, then every other hyperedge contains at most one of these. Thus, fixing the outputs of the other hyperedges does not introduce correlation between $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$, and we can again use [Lemma 3.9](#) to show that such fixings only decrease their entropy by just a little.

Thus, we can ensure fragile correlation by selecting sources using a hypergraph with the following property: no two hyperedges share more than one vertex. Such hypergraphs are well-studied in combinatorial design theory, and are known as *partial Steiner triple systems* (STS's). Furthermore, recalling that an independent set in a hypergraph is a set of vertices that contains no hyperedge, we see that we can guarantee *activation* using just K sources if the partial Steiner triple system contains no independent set of size K (equivalently, it should have *independence number* $\alpha < K$).

We therefore construct a so-called STS-extractor for (N, K, n, k) -sources of locality 0 as follows. Let $H = (V, \mathcal{E})$ be an STS over N vertices, and define $\text{stsExt}_H : (\{0, 1\}^n)^N \rightarrow \{0, 1\}^m$ as:

$$\text{stsExt}_H(\mathbf{X}) := \bigoplus_{(h,i,j) \in \mathcal{E}} \text{3Ext}(\mathbf{X}_h, \mathbf{X}_i, \mathbf{X}_j).$$

For a more precise definition, see [Definition 4.6](#), and for an illustration, see [Figure 1a](#).

As per our discussion, this will successfully extract uniform bits as long as K exceeds the size of the largest independent set in H . Furthermore, it inherits the *polylogarithmic entropy* requirement of `3Ext` ([Theorem 3.3](#)), along with its *polynomially large output length* and *negligible error*. Thus, the challenge is to explicitly construct an STS $H = (V, \mathcal{E})$ over N with small α . We achieve such an explicit construction by identifying V with $\mathbb{F}_3^{\log N}$, identifying \mathcal{E} with the lines in $\mathbb{F}_3^{\log N}$, and showing that recent bounds on the cap set problem [[CLP17](#), [EG17](#)] immediately imply $\alpha \leq O(N^{0.923})$. As a result, instantiating `stsExt` with H yields an explicit extractor for polynomially few good sources (see [Theorem 4.1](#)).

It would be nice if we could extract from even fewer good sources. However, lower bounds on the cap set problem [[Ede04](#)] show that we cannot use lines in $\mathbb{F}_3^{\log N}$ to achieve better than $K \geq N^{0.724}$, and impossibility results on Steiner systems [[RŠ94](#)] show that one cannot hope to achieve $K \ll \sqrt{N \log N}$ using these objects. Thus, we need new ideas if we want to drastically decrease K .

¹Each call to `3Ext` will need the hyperedge to order its vertices, but the ordering will not be important, so we induce one by simply assuming the vertices are identified with $[N]$.

The wedge-extractor Towards this end, we show that STS’s actually have more structure than is required for fragile correlation. Indeed, we show that if we replace 3Ext with a more *robust* three-source extractor 3Ext^+ , we can extract using a much larger class of hypergraphs, and thereby reduce the size K needed for activation. In particular, in order to construct 3Ext^+ , we make use of two recent advances in extractor theory. First, we will use a *two-source non-malleable extractor*, 2nmExt , which is a robust variant of a two-source extractor that, given two independent sources \mathbf{X}, \mathbf{Y} , outputs bits $2\text{nmExt}(\mathbf{X}, \mathbf{Y})$ that look uniform *even conditioned on* knowing the value of $2\text{nmExt}(f(\mathbf{X}), g(\mathbf{Y}))$ or $2\text{nmExt}(g(\mathbf{Y}), f(\mathbf{X}))$, where f, g are so-called *tampering functions* that have no fixed points (see [Definition 3.6](#) for a formal definition). If the output of 2nmExt looks uniform even conditioned on its output under up to t pairs of tampering functions, we say 2nmExt has *degree* t . The motivation behind using these objects is as follows: previously, if we fixed any random variables that depended on two of the good sources, we would introduce correlation between them. This will no longer be the case, and thus we have more power to ensure fragile correlation.


Second, we will use a *two-source condenser*, 2Cond , which is a weaker version of a two-source extractor that only guarantees its output to have high entropy rate. 2Cond will also be *strong*, in the sense that it will work even conditioned on fixing its second source, with high probability ([Theorem 3.4](#)). The motivation here is that 2nmExt only works for sources with high entropy, and 2Cond is able to condense a source with just polylogarithmic entropy into one (on fewer bits) with almost full entropy. Thus, we can maintain our requirement that $k = \log^C n$. Our new robust three-source extractor is defined as $3\text{Ext}^+(\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3) := 2\text{nmExt}(2\text{Cond}(\mathbf{X}_1, \mathbf{X}_3), 2\text{Cond}(\mathbf{X}_2, \mathbf{X}_3))$.²

We again consider the following question, with respect to our robust three-source extractor: what structure must a 3-uniform hypergraph have such that if some hyperedge is *activated*, then every other hyperedge makes a call to 3Ext^+ that can be safely fixed without ruining the output of the call to 3Ext^+ from the activated hyperedge? We notice that here, each call to 3Ext^+ requires us to specify three sources, and indicate one of these to be *special*, in that it will be reused in both calls to 2Cond . One way to encode this information is as a hyperedge A of size 3, containing a hyperedge B of size 2 (which leaves out the special source; we call B the *representative edge* of A).


Thus, we consider using hypergraphs that have hyperedges of the above form to make calls to 3Ext^+ . We now argue the following: if we construct such a hypergraph such that the representative edge B of a hyperedge A is also the representative edge of *any other hyperedge containing both its vertices* (call this *representative edge agreement*), then we can satisfy fragile correlation. Consider such a hypergraph, and suppose it has an activated hyperedge A that contains three good sources, $\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3$, and a representative edge that holds $\mathbf{X}_1, \mathbf{X}_2$. If we fix all sources excluding $\mathbf{X}_1, \mathbf{X}_2$, we can note a few things: first, by the strength of 2Cond , each of $\mathbf{Y}_1 := 2\text{Cond}(\mathbf{X}_1, \mathbf{X}_3), \mathbf{Y}_2 := 2\text{Cond}(\mathbf{X}_2, \mathbf{X}_3)$ are now independent and have high entropy, with high probability. Next, because of our *representative edge agreement* property, we know that $\mathbf{X}_1, \mathbf{X}_2$ will never show up together in a single 2Cond in *any* call to 3Ext^+ .

Thus, any call to 3Ext^+ made from a hyperedge outside of the activated hyperedge can fall into one of four categories: (1) it involves neither source $\mathbf{X}_1, \mathbf{X}_2$; (2) it involves \mathbf{X}_1 but not \mathbf{X}_2 ; (3) it involves \mathbf{X}_2 but not \mathbf{X}_1 ; or (4) it involves both $\mathbf{X}_1, \mathbf{X}_2$, but by the representative edge agreement property, they are guaranteed to be in different 2Cond calls. To ensure fragile correlation, we want to fix the calls to 3Ext^+ from each category without destroying the uniform bits produced by the activated hyperedge. Note that the calls in (1) are already fixed. If we fix the calls to (2), (3),

²There are some minor technical details to ensure that 2nmExt will work, such as *tagging* each of its inputs uniquely.



(a) The *Fano plane*, a 3-uniform hypergraph $H = (V, \mathcal{E})$ that is a Steiner triple system.



(b) A graph $G = (V, E)$ that contains many wedges, inspired by the Fano plane.

Figure 1: Extracting from 0-local adversarial sources using Steiner triple systems and wedges. In Figures 1a and 1b, sources are represented as nodes of a (hyper)graph. Figure 1a represents our STS-extractor, stsExt_H , and Figure 1b our wedge-extractor, wExt_G . In Figure 1a, some hyperedge is guaranteed to be activated iff at least 5 sources are good (green), while in Figure 1b, some wedge is guaranteed to be activated iff at least 4 sources are good (green).

we know that no correlation is introduced between $\mathbf{X}_1, \mathbf{X}_2$, and each loses just a little entropy, by Lemmas 3.9 and 3.10. Finally, we know that if (4) has no more calls than the degree of 2nmExt , we can fix these calls and use the non-malleability of 2nmExt to ensure that the bits from our activated hyperedge still look uniform. Observe that the number of calls in (4) is at most the number of hyperedges that share the same representative edge. Thus, because hyperedges have size at most 3, and we assume no hyperedge has more than one copy, we know that the number of calls in (4) is at most $N - 2$, and thus we can perform these fixings as long as $N \leq k^\gamma$, for a small constant γ , by the parameters in Lemma 3.8. Thus, we can ensure fragile correlation.


Is there a nicer way to describe such hypergraphs with hyperedges of size 3, and representative edges of size 2, such that the *representative edge agreement* property holds? In fact, there is a very natural way to do so: these are exactly the hypergraphs that can be constructed via taking a standard graph G , and selecting some wedges (sets of size 3 that induce a 2-hop-path in G) to turn into hyperedges, where the two non-adjacent vertices of each wedge (the *terminals*) make up the representative edge. Thus, we are motivated to define a new extractor over the wedges of graphs.

In particular, we construct a so-called wedge-extractor for (N, K, n, k) -sources of locality 0 as follows. Let G be a graph over N vertices, and let \mathcal{W} be the collection of sets of size 3 in G that induce a wedge. We order each $W \in \mathcal{W}$ as a triple (h, i, j) so that h, i are the terminals of W , and define $\text{wExt}_G : (\{0, 1\}^n)^N \rightarrow \{0, 1\}^m$ as:


$$\text{wExt}_G(\mathbf{X}) := \bigoplus_{(h,i,j) \in \mathcal{W}} 3\text{Ext}^+(\mathbf{X}_h, \mathbf{X}_i, \mathbf{X}_j).$$

For a more precise definition, see Definition 4.15, and for an illustration, see Figure 1b.

As per our discussion, this will successfully extract uniform bits (provided $N \leq k^\gamma$) as long as K good sources are guaranteed to activate some hyperedge; note that here, this simply means that any subset of size K in $V(G)$ covers some wedge in G , or that the size of the largest so-called *wedge-independent set*, $\alpha_{\mathcal{W}}$, is less than K . Furthermore, note that wExt_G inherits the *polylogarithmic*



(a) A graph $G = (V, E)$ that contains many wedges.



(b) A fragile set system $H = (G, \mathcal{S})$, with $\mathcal{S} = \{\{1, 2, 4, 5, 6\}, \{2, 3, 4, 5, 6\}, \{1, 3, 4, 5, 6\}\}$.

Figure 2: Extracting from 0-local adversarial sources using wedges and fragile set systems. In Figures 2a and 2b, sources are represented as nodes of a (hyper)graph. Figure 2a represents our wedge-extractor, wExt_G , and Figure 2b our FSS-extractor, fssExt_H . The activating sets in Figures 2a and 2b are exactly the same, but each representative edge in Figure 2a appears in 3 wedges, while each representative edge in Figure 2b appears in just 1 fragile set.

entropy requirement of 2Cond (Theorem 3.4), and the *polynomially large output length* and *negligible error* of both 2Cond , 2nmExt (Theorems 3.4 and 3.7). Thus, the challenge is to explicitly construct a graph $G = (V, E)$ such that its largest wedge-independence set has a small size α_W .

We achieve such a construction by showing that a Ramsey graph with no clique nor independent set of size ℓ actually also has no wedge independent set of size ℓ^2 (see Lemma 4.17 for more details). To see this, we observe that a set of vertices that covers no wedge must be a disjoint collection of cliques (with no crossing edges), and thus a wedge independent set of size ℓ^2 would imply a clique or independent set of size ℓ (by taking the largest clique, or a single vertex from each clique, in the wedge-independent set). This immediately yields Theorem 1.

The FSS-extractor We note that Theorem 1 extracts from very few good sources with very little entropy, under the condition that $N \leq k^\gamma$. While this condition is reasonable in many settings, it would be nice to get rid of it completely. We construct a new extractor that succeeds in doing so, and in fact generalizes all of the constructions we have seen so far. The main idea is the same as with the wedge-extractor, with one small but powerful twist. In particular, recall that our restriction $N \leq k^\gamma$ arises from the observation that up to $N - 2$ hyperedges may share the same representative edge. If we can reduce this number, then we can relax and even remove this restriction. We achieve this by coming up with a more general hypergraph structure.

In particular, we generalize the previous hypergraph to allow hyperedges of any size greater than 2, such that each hyperedge still contains a representative edge (hyperedge of size 2). Again, we enforce the *representative edge agreement* property that the representative edge B of a hyperedge A is also the representative edge of any other hyperedge containing it. Note that our three-source extractor is no longer well-defined, since each hyperedge could indicate more than three sources over which to attempt extraction. Indeed, we extend our extractor as follows. Each hyperedge

$A \subseteq [N]$ with representative edge $B = \{h, i\}$ identifies a call of the form:

$$2\text{nmExt}(2\text{Cond}(\mathbf{X}_h, \oplus_{j \in A \setminus B} \mathbf{X}_j), 2\text{Cond}(\mathbf{X}_i, \oplus_{j \in A \setminus B} \mathbf{X}_j)),$$

and our extractor will take the XOR over all hyperedges of these calls. Furthermore, we can redefine *activation* to a much more relaxed notion: instead of requiring that some hyperedge contains *all* good sources, we simply require that some hyperedge has good sources on the endpoints of its representative edge, and one more good source outside of its representative edge. Then, if some hyperedge is activated by good sources $\mathbf{X}_h, \mathbf{X}_i, \mathbf{X}_j$, the representative edge agreement property guarantees that our extractor will work for the same reasons in our analysis of the wedge-extractor.

Is there a nicer way to describe our new, more general, hypergraphs? The answer is yes: these are exactly the hypergraphs that can be constructed via taking a standard graph G , selecting some of its so-called *fragile sets* (sets in G that contain exactly one edge), turning each fragile set into a hyperedge, and turning the edge in the fragile set into its representative edge. We call such a hypergraph (consisting of G and a collection \mathcal{S} of *some* of its fragile sets) a *fragile set system*. We say the *degree* of a fragile set system H , denoted $\deg(H)$ is the max number of fragile sets $S \in \mathcal{S}$ that contain the same edge. Together with the generality of this new structure, this new parameter will give us fine control over removing the restriction $N \leq k^\gamma$, by replacing it with $\deg(H) - 1 \leq k^\gamma$. Thus, we are motivated to define a new extractor over fragile set systems.

In particular, we construct a so-called FSS-extractor for (N, K, n, k) -sources of locality 0 as follows. Let G be a graph over N vertices, and \mathcal{S} be a collection of fragile sets in G , thus creating the fragile set system $H = (G, \mathcal{S})$. We write each $S \in \mathcal{S}$ as a triple (u, v, S') where u, v are the endpoints of the edge in S , and S' are the remaining vertices. We define $\text{fssExt}_H : (\{0, 1\}^n)^N \rightarrow \{0, 1\}^m$ as:

$$\text{fssExt}_H(\mathbf{X}) := \bigoplus_{(u, v, S') \in \mathcal{S}} 2\text{nmExt}(2\text{Cond}(\mathbf{X}_u, \oplus_{j \in S'} \mathbf{X}_j), 2\text{Cond}(\mathbf{X}_v, \oplus_{j \in S'} \mathbf{X}_j)).$$

For a more precise definition, see [Definition 4.23](#), and for an illustration, see [Figure 2](#).

As per our discussion, this will successfully extract uniform bits (provided $\deg(H) - 1 \leq k^\gamma$) as long as K good sources are guaranteed to activate some hyperedge; here, this simply means that some fragile set contains three good sources, two of which lie on the endpoints of its edge. Equivalently, we need $\alpha_{\text{FSS}} < K$, where α_{FSS} denotes the *FSS-independence number*, or the size of the largest set that activates *no* hyperedge. Thus, the challenge is to explicitly construct a fragile set system $H = (G, \mathcal{S})$ with small $\deg(H)$ and small α_{FSS} .

We achieve such a construction for $\deg(H) \leq 1$ and $\alpha_{\text{FSS}} < \sqrt{N \cdot \mathcal{R}_N}$, which therefore extracts from $K \geq \sqrt{N \cdot \mathcal{R}_N} = N^{0.5+o(1)}$ good sources, while completely removing any restriction between N and k^γ , thereby yielding [Theorem 2](#). It is worth noting that given optimal Ramsey graphs, this would exactly match (up to constant factors) the best result that is existentially possible with partial Steiner triple systems. The construction of such a fragile set system works by placing N vertices into roughly \sqrt{N} clouds $C_1, C_2, \dots, C_{\sqrt{N}}$ of size \sqrt{N} , drawing a bipartite Ramsey graph between each pair of clouds, and adding *one* fragile set for each edge (and thus, the degree is 1). The fragile set simply includes that edge, considers the endpoint in the smaller-labeled cloud, and adds all non-neighbors of this endpoint that are in the higher-labeled cloud. It is then straightforward to show that given $K = N^{0.5+o(1)}$ vertices, two clouds must have enough vertices so that if some fragile set were not activated, a large bipartite clique or independent set must exist (see [Lemma 4.25](#) for more details).

2.2 Extracting from polynomial locality

Thus far, we have constructed explicit extractors for the 0-local setting that are quite general, in the sense that each of our extractors can take any hypergraph from a certain class (STS's, wedges in graphs, and fragile set systems) to instantiate the extractor, and the parameters that can be achieved by that extractor are directly related to the parameters of the hypergraph used to instantiate it. We show that, in fact, we can find hypergraphs to instantiate our extractors so that they succeed in extracting from up to polynomial locality.

Because our FSS-extractor generalizes the other constructions, we show that it can extract from polynomial locality. Indeed, we prove an even stronger result that its specialization as the wedge-extractor can also succeed in doing so. In particular, recall that our wedge-extractor works by explicitly constructing a graph G over N vertices, identifying the sources with the N vertices, calling 3Ext^+ over all triples that identify a wedge in G , and taking the XOR of the results.

We use the exact same ideas in the (≥ 1) -locality setting, except there are additional complications, in particular when establishing fragile correlation. Recall that previously, if some hyperedge (wedge) W was activated by good sources, then we could fix every source but the two sources $\mathbf{X}_1, \mathbf{X}_2$ in the representative edge of the wedge (i.e., its non-edge), and use [Lemmas 3.9 and 3.10](#) and the non-malleability of 2nmExt to fix the output of every other 3Ext^+ call, while keeping the output of the 3Ext^+ call over W near-uniform. But we could only do this because we were in the 0-local setting, since using wedges to select sources guaranteed that $\mathbf{X}_1, \mathbf{X}_2$ would never show up together as the arguments to a single 2Cond call.

While it is still true in the (≥ 1) -local setting that $\mathbf{X}_1, \mathbf{X}_2$ never show up in a single 2Cond call, it might be the case that random variables (bad sources) correlated to $\mathbf{X}_1, \mathbf{X}_2$ show up together in a 2Cond call. In this case, we cannot hope to fix the output of the call to 3Ext^+ involving this 2Cond call without introducing correlation between $\mathbf{X}_1, \mathbf{X}_2$.

In order to fix this issue, we must prevent this from happening. One way to do this is to note that when using our wedge-extractor, two sources (good or bad) show up together in a call to 2Cond *only if* their corresponding vertices are connected by an edge. Thus, consider the case that some wedge W is covered by good sources, and the sources on its terminals are $\mathbf{X}_1, \mathbf{X}_2$. Let $\text{cloud}(\mathbf{X}_1)$ denote the vertices corresponding to sources correlated with \mathbf{X}_1 , and $\text{cloud}(\mathbf{X}_2)$ denote those corresponding to sources correlated with \mathbf{X}_2 . Observe that if there are no edges between $\text{cloud}(\mathbf{X}_1)$ and $\text{cloud}(\mathbf{X}_2)$, and they are disjoint, then we can perform fixings as usual, and guarantee that our extractor works.

Using this idea, we tackle the 1-local setting as follows. Analogously to the 0-local setting, given a graph G that will be used to instantiate the wedge-extractor, we define a new flavor of *activating set* of vertices. As in the 0-local setting, we want this activating set to have the property that if the good sources land on it, then the wedge-extractor is guaranteed to extract uniform bits from the 1-local source.

As hinted above, we will define an activating set to be any set of vertices S in G such that no matter how we draw a separate cloud around each $s \in S$ (making sure that no two clouds intersect), there will be three clouds such that the three vertices from S they contain cover a wedge in G , such that the terminals of that wedge lie in two distinct clouds with no edges between them. We call this structure a *cloud-wedge* (refer to [Definition 5.3](#) for a formal definition). Thus, the goal is to construct a graph G such that no matter how one selects K vertices and draws K disjoint clouds around them, a cloud-wedge is guaranteed to appear (for the smallest K possible). The selection of K vertices represents the placement of K good sources among the N total sources in our adversarial source, and the drawing of clouds indicates which bad sources will be dependent on




Figure 3: Extracting from 1-local adversarial sources using wedges. As before, a green node represents a good source and a red node represents a bad source. The red *clouds* D_2, D_7, D_8, D_{11} represent *dependencies*: cloud D_i contains all sources correlated with good source i . This placement of good sources and dependencies over our graph induces a cloud-wedge: $(\{2, 7, 8\}, \{D_2, D_7, D_8\})$. The other good wedges do not induce cloud-wedges due to crossing edges in their terminal clouds.

which good sources. If one can always find a cloud-wedge for a given K , then the wedge extractor is guaranteed to work for just K good sources (see Lemma 5.6). We refer the reader to Figure 3 for an illustration.

We show that one family of graphs that exhibits the above-desired property are graphs with no cycle of length 4, and with no big independent set. Through some structural lemmas, we show that these two properties ensure that any relatively large set of vertices in such a graph must cover a large *star* (complete bipartite graph with 1 vertex on the left), and any big collection of nonempty disjoint subsets in such a graph must have two subsets with no edges crossing between them. It is straightforward to show that, together, these so-called “star-dense” and “anti-cloud-clique” properties ensure that in the aforementioned process, we will always be able to find a cloud-wedge (see Lemma 5.8).

Thus, we reduce the question of constructing extractors for 1-local adversarial sources to that of explicitly constructing C_4 -free graphs with no big independent set. Fortunately, explicit constructions of such objects are known [Alo86], and so we are able to successfully extract from 1-local sources. In order to extract from higher locality, we provide a reduction from d -local sources to 1-local sources, in the spirit of Viola’s reduction from samplable sources to affine sources [Vio14]. By combining this reduction with our extractors for 1-local adversarial sources, we are able to yield Theorem 3.

2.3 Non-explicit results

As discussed above, the reason behind our choice of hypergraphs is that we need them to carefully control the dependence between different random variables produced in the computation of our extractors. However, this task would become easier if we had sufficiently strong non-malleable extractors to break the dependence. To this end, we introduce the notion of a *generalized s -source non-malleable extractor with tampering degree t* . This is a non-malleable extractor that takes as input s independent weak sources, and is secure against t tampering outputs. In each tampering, the adversary can produce s tampered sources, where each tampered source depends on *at most* $s - 1$ of the original s sources. As long as each tampering has no fixed point, we show that such

generalized non-malleable extractors exist with excellent parameters. Given such extractors, it is simple to extract from adversarial sources with high locality: just apply the non-malleable extractor on every s -tuple of sources and compute the XOR. As long as there is a subset S of s good sources such that no bad source depends on all good sources in S , we can fix all good sources outside S , and all calls to the non-malleable extractor over tuples not equal to S , and the property of the non-malleable extractor guarantees that the output will be close to uniform. By taking s to be a large enough constant, we can handle arbitrary polynomially few good sources and $K^{0.99}$ locality, proving [Theorem 4](#). We refer the reader to [Section 7](#) and [Appendix A](#) for details.

We note that the study of non-malleable extractors where several sources may be tampered together was recently undertaken by Goyal et al. [[GSZ19](#)] in the context of designing better non-malleable secret sharing schemes. However, their work only provides a construction for the so-called *cover-free* tampering function family, which does not include our setting where any tampered source may be a result of tampering any $s - 1$ (out of s) sources jointly.

3 Preliminaries

Throughout, we use \circ to denote string concatenation. For two strings $x, y \in \{0, 1\}^n$, we let $x \oplus y$ denote bitwise XOR. Given a graph $G = (V, E)$ and set $S \subseteq V$, we let $G[S]$ denote the subgraph induced by S .

3.1 Extractors and condensers for independent sources

First, we recall that the *statistical distance* of two distributions \mathbf{D}_1 and \mathbf{D}_2 (on the same support) is given by

$$|\mathbf{D}_1 - \mathbf{D}_2| := \frac{1}{2} \sum_x |\Pr[\mathbf{D}_1 = x] - \Pr[\mathbf{D}_2 = x]|,$$

and \mathbf{D}_1 is ϵ -close to \mathbf{D}_2 if $|\mathbf{D}_1 - \mathbf{D}_2| \leq \epsilon$. Next, we recall the definition of a multi-source extractor:

Definition 3.1. *Let $C \in \mathbb{N}$. We call a function $\text{Ext} : (\{0, 1\}^n)^C \rightarrow \{0, 1\}^m$ a C -source extractor for entropy k , output length m , and error ϵ if, given any C independent (n, k) -sources $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_C$,*

$$|\text{Ext}(\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_C) - \mathbf{U}_m| \leq \epsilon.$$

We will need the following explicit constructions of multi-source extractors:

Theorem 3.2 ([[CG88](#), [Vaz85](#)]). *For every constant $\delta > 0$, and for all $n, k \in \mathbb{N}$ with $k \geq (1/2 + \delta)n$, there exists an explicit 2-source extractor $\text{Had} : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^m$ for entropy k with output length $m = \Omega(n)$ and error $\epsilon = 2^{-\Omega(n)}$.*

Theorem 3.3 ([[Li15c](#), [Coh16a](#)]). *For all $n, k \in \mathbb{N}$ with $k \geq \log^8 n$, there exists an explicit 3-source extractor $3\text{Ext} : (\{0, 1\}^n)^3 \rightarrow \{0, 1\}^m$ for entropy k with output length $m = 0.9k$ and error $\epsilon = 2^{-k^{\Omega(1)}}$.*

We will also need a weaker notion called a *condenser*, which only guarantees that its output is close to a high entropy source, instead of being close to uniform. In particular, we will use the following explicit construction:

Theorem 3.4 ([BACDTS19]). *There exists a constant $C \geq 1$ such that for every $n, k, m \in \mathbb{N}$ and $\epsilon > 0$ such that $n \geq k \geq (m \log(n/\epsilon))^C$, there exists an explicit function $2\text{Cond} : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^m$ such that for any two independent (n, k) -sources $\mathbf{X}_1, \mathbf{X}_2$, with probability $1 - \epsilon$ over $x_2 \sim \mathbf{X}_2$, the output $2\text{Cond}(\mathbf{X}_1, x_2)$ is $2^{-k/2}$ -close to an $(m, m - o(\log(1/\epsilon)))$ -source, \mathbf{Y} .*

3.2 Two-source non-malleable extractors

Next, we need a stronger notion of two-source extraction that arises in cryptography and was first defined in [CG14], known as a two-source non-malleable extractor.

Definition 3.5. *We call a function $2\text{nmExt} : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^m$ a $(2, t)$ -non-malleable extractor for entropy k , output length m , and error ϵ , if, given any two (n, k) -sources $\mathbf{X}_1, \mathbf{X}_2$, and t pairs of tampering functions $\{(f_i, g_i)\}_{i \in [t]}$, where each $f_i, g_i : \{0, 1\}^n \rightarrow \{0, 1\}^n$ have no fixed points,*

$$\begin{aligned} & |2\text{nmExt}(\mathbf{X}_1, \mathbf{X}_2) \circ 2\text{nmExt}(f_1(\mathbf{X}_1), g_1(\mathbf{X}_2)) \circ \cdots \circ 2\text{nmExt}(f_t(\mathbf{X}_1), g_t(\mathbf{X}_2)) \\ & - \mathbf{U}_m \circ 2\text{nmExt}(f_1(\mathbf{X}_1), g_1(\mathbf{X}_2)) \circ \cdots \circ 2\text{nmExt}(f_t(\mathbf{X}_1), g_t(\mathbf{X}_2))| \leq \epsilon. \end{aligned}$$

We will in fact need a more robust non-malleable extractor whose output $2\text{nmExt}(\mathbf{X}_1, \mathbf{X}_2)$ looks uniform, even if conditioned on tamperings of the form $2\text{nmExt}(g_i(\mathbf{X}_2), f_i(\mathbf{X}_1))$. We define this new object under the same name, and will only be referring to this robust variant throughout the paper.

Definition 3.6. *We call a function $2\text{nmExt} : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^m$ a $(2, t)$ -non-malleable extractor for entropy k , output length m , and error ϵ , if the following holds. Let $\mathbf{X}_1, \mathbf{X}_2$ be any two (n, k) -sources, let $\{(f_i, g_i)\}_{i \in [t]}$ be any t pairs of tampering functions where each $f_i, g_i : \{0, 1\}^n \rightarrow \{0, 1\}^n$ have no fixed points, and let $b \in \{0, 1\}^n$ be any bitstring. Then if we define \mathbf{Y}_i as $(f_i(\mathbf{X}_1), g_i(\mathbf{X}_2))$ if the i^{th} bit of b is 0, and we define \mathbf{Y}_i as $(g_i(\mathbf{X}_2), f_i(\mathbf{X}_1))$ otherwise, then:*

$$\begin{aligned} & |2\text{nmExt}(\mathbf{X}_1, \mathbf{X}_2) \circ 2\text{nmExt}(\mathbf{Y}_1) \circ \cdots \circ 2\text{nmExt}(\mathbf{Y}_t) \\ & - \mathbf{U}_m \circ 2\text{nmExt}(\mathbf{Y}_1) \circ \cdots \circ 2\text{nmExt}(\mathbf{Y}_t)| \leq \epsilon. \end{aligned}$$

We say that a $(2, t)$ -non-malleable extractor has tampering degree t .

We note that the above extractor is a special case of the more general (s, t) -non-malleable extractor which we define and prove its existence in [Appendix A](#). As it turns out, however, the existing constructions of $(2, t)$ -non-malleable extractors also have this more robust property, as the constructions of these objects use *alternating extraction*, which is symmetric in the way it deals with sources. Thus, we have:

Theorem 3.7 ([CGL16]). *There exists a constant $\gamma > 0$ such that for all $n, k \in \mathbb{N}$ with $k \geq n - n^\gamma$, and all $t \leq n^\gamma$, there exists an explicit seedless $(2, t)$ -non-malleable extractor $2\text{nmExt} : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^m$ with output length $m = n^{\Omega(1)}$ and error $\epsilon = 2^{-n^{\Omega(1)}}$.*

Throughout the paper, we will use the following shared parameter ranges that can be achieved while explicitly constructing two-source condensers and two-source non-malleable extractors.

Lemma 3.8 ([CGL16, BACDTS19]). *There exist universal constants $C_0, \gamma_0 > 0$ such that for all $n, k \in \mathbb{N}$ satisfying $k \geq (\log n)^{C_0}$, there exists:*

- An explicit function $2\text{Cond} : (\{0, 1\}^n)^2 \rightarrow \{0, 1\}^{m_1}$ such that for any two independent (n, k) -sources $\mathbf{X}_1, \mathbf{X}_2$, with probability $1 - \epsilon_1$ over $x_2 \sim \mathbf{X}_2$, the output $2\text{Cond}(\mathbf{X}_1, x_2)$ is ϵ_2 -close to an $(m_1, m_1 - o(\log(1/\epsilon_1)))$ -source, where $m_1 = k^{1/C_0}$, and $\epsilon_1 = \epsilon_2 = 2^{-m_1^{\gamma_0/2}}$.
- For any $m_2 \geq m_1, t \leq m_2^{\gamma_0}$, an explicit $(2, t)$ -non-malleable extractor $2\text{nmExt} : (\{0, 1\}^{m_2})^2 \rightarrow \{0, 1\}^m$ for entropy at least $m_2 - m_2^{\gamma_0}$ with output length $m \in [m_1^{\Omega(1)}, m_1^{\gamma_0/2}]$ and error $\epsilon_3 = 2^{-m_2^{\Omega(1)}}$.

3.3 Conditional min-entropy

Finally, we need the following two lemmas about conditional min-entropy.

Lemma 3.9 ([MW97]). *Let \mathbf{X}, \mathbf{Y} be random variables such that \mathbf{Y} takes at most ℓ values. Then:*

$$\Pr_{y \sim \mathbf{Y}} [H_\infty(\mathbf{X} \mid \mathbf{Y} = y) \geq H_\infty(\mathbf{X}) - \log \ell - \log(1/\epsilon)] \geq 1 - \epsilon.$$

Lemma 3.10 ([Li15a]). *Let \mathbf{X}, \mathbf{Y} be random variables such that \mathbf{X} is ϵ -close to a source with min-entropy at least k and \mathbf{Y} takes at most ℓ values. Then:*

$$\Pr_{y \sim \mathbf{Y}} [(\mathbf{X} \mid \mathbf{Y} = y) \text{ is } 2\epsilon^{1/2}\text{-close to a source with min-entropy at least } k - \log \ell - \log(1/\epsilon)] \geq 1 - 2\epsilon^{1/2}.$$

4 Extracting from 0-locality

In this section, we construct extractors for adversarial sources of locality 0. In particular, we prove [Theorems 1](#) and [2](#). As a warm-up, we start off with a construction that has worse parameters, but motivates our more complicated constructions.

4.1 Extractors from Steiner triple systems and cap set bounds

In this section, we prove the following result:

Theorem 4.1. *There is a constant $C \geq 1$ such that for all $N, K, n, k \in \mathbb{N}$ satisfying $k \geq \log^C n$ and $K \geq CN^{0.923}$, there exists an explicit extractor $\text{Ext} : (\{0, 1\}^n)^N \rightarrow \{0, 1\}^m$ for (N, K, n, k) -sources of locality 0, with output length $m = 0.9k$ and error $\epsilon = 2^{-k^{\Omega(1)}}$.*

In order to obtain this explicit construction, we will call the three-source extractor from [Theorem 3.3](#) over certain triples of sources, and take the XOR of the results. We will use the following object from combinatorics to select these triples:

Definition 4.2 (partial Steiner triple system). *A partial Steiner triple system over N points, or an $\text{STS}(N)$, is a 3-uniform hypergraph over N vertices such that no two hyperedges share more than one vertex.*

Remark 4.3. *While an $\text{STS}(N)$ is unordered, each hyperedge will need to (arbitrarily) order its vertices so that the inputs to each 3Ext call is well-defined. One way to do so is by simply identifying the vertices with $[N]$, and having each hyperedge list its vertices in increasing order.*

We define *activating* and *independent* sets with respect to Steiner triple systems as follows.

Definition 4.4 (STS-activating set). Let $H = (V, \mathcal{E})$ be an STS(N). An STS-activating set in H is a subset $S \subseteq V$ that contains at least one hyperedge $e \in \mathcal{E}$.

Definition 4.5 (STS-independent set). Let $H = (V, \mathcal{E})$ be an STS(N). An STS-independent set in H is a subset $S \subseteq V$ that is not an STS-activating set. The STS-independence number of H is the size of its max STS-independent set, and is denoted by α_{STS} .

Note that an STS-independent set is the same definition as an standard independent set in a 3-uniform hypergraph. We can now formally define our extractor over partial Steiner triple systems.

Definition 4.6 (STS-extractor). Let $H = (V, \mathcal{E})$ be an STS(N), and let $3\text{Ext} : (\{0, 1\}^n)^3 \rightarrow \{0, 1\}^m$ be the three-source extractor from [Theorem 3.3](#). We define the STS-extractor over H for (N, K, n, k) -sources, $\text{stsExt}_H : (\{0, 1\}^n)^N \rightarrow \{0, 1\}^m$, as:

$$\text{stsExt}_H(\mathbf{X}) := \bigoplus_{(v_1, v_2, v_3) \in \mathcal{E}(H)} 3\text{Ext}(\mathbf{X}_{v_1}, \mathbf{X}_{v_2}, \mathbf{X}_{v_3}).$$

For an illustration, we refer the reader to [Figure 1a](#). This extractor takes an STS(N) as advice, and we will prove a general lemma showing that the performance of our extractor will depend on the value of α_{STS} for STS(N). Then, we will explicitly construct STS(N)'s with small α_{STS} , which will immediately yield [Theorem 4.1](#).

Lemma 4.7. *There is a constant $C \geq 1$ such that for all $N, K, n, k \in \mathbb{N}$ satisfying $k \geq \log^C n$, the following holds. If $H = (V, \mathcal{E})$ is an STS(N) with STS-independence number $\alpha_{\text{STS}} < K$, then $\text{stsExt}_H : (\{0, 1\}^n)^N \rightarrow \{0, 1\}^m$ is an extractor for (N, K, n, k) -sources of locality 0, with output length $m = 0.9k$ and error $\epsilon = 2^{-k^{\Omega(1)}}$.*

Proof. Identify the vertices of H with $[N]$, and define $I := \{i \in [N] : \mathbf{X}_i \text{ is a good source}\}$. We know that $|I| \geq K$, and thus because H has independence number $\alpha_{\text{STS}} < K$, there exists some hyperedge $F := (v_1, v_2, v_3) \in \mathcal{E}$ such that $F \subseteq I$. Now, for each $(h, i, j) \in \mathcal{E}$, define the random variable $\mathbf{Y}_{(h, i, j)} := 3\text{Ext}(\mathbf{X}_h, \mathbf{X}_i, \mathbf{X}_j)$, and note that $\text{stsExt}_H(\mathbf{X}) = \bigoplus_{(h, i, j) \in \mathcal{E}} \mathbf{Y}_{(h, i, j)}$. Observe that because H is an STS(N), we can partition $\mathcal{E} \setminus (v_1, v_2, v_3)$ into $\mathcal{E}_1, \mathcal{E}_2, \mathcal{E}_3, \mathcal{E}_4$ such that the triples in \mathcal{E}_4 have an empty intersection with F , while the triples in \mathcal{E}_ℓ , for $\ell \in [3]$, intersect with F at exactly v_i . Thus, if we define $\mathbf{Z}_\ell := \bigoplus_{(h, i, j) \in \mathcal{E}_\ell} \mathbf{Y}_{(h, i, j)}$, for each $\ell \in [4]$, we have:

$$\text{stsExt}_H(\mathbf{X}) = \mathbf{Y}_{(v_1, v_2, v_3)} \oplus \mathbf{Z}_1 \oplus \mathbf{Z}_2 \oplus \mathbf{Z}_3 \oplus \mathbf{Z}_4.$$

Now, fix all random variables \mathbf{X}_i such that $i \notin \{v_1, v_2, v_3\}$. Notice that this fixes \mathbf{Z}_4 . Next, fix \mathbf{Z}_1 . By [Lemma 3.9](#), the min-entropy of \mathbf{X}_{v_1} is still at least $k - \log(2^m) - \log(1/\epsilon) = k - 0.9k - 0.09k = 0.01k$ with probability at least $1 - 2^{-0.09k}$. Furthermore, because \mathbf{Z}_1 is uncorrelated with $\mathbf{X}_{v_2}, \mathbf{X}_{v_3}$, no correlation is introduced between $\mathbf{X}_{v_1}, \mathbf{X}_{v_2}, \mathbf{X}_{v_3}$. Next, do the same for \mathbf{Z}_2 and \mathbf{Z}_3 . $\mathbf{X}_{v_1}, \mathbf{X}_{v_2}, \mathbf{X}_{v_3}$ remain independent, and with probability at least $1 - 3 \cdot 2^{-0.09k}$, they each have min-entropy at least $0.01k$. Thus, as long as $0.01k \geq \log^{12} n$, $\mathbf{Y}_{(v_1, v_2, v_3)}$ is $2^{-k^{\Omega(1)}}$ -close to \mathbf{U}_m and the other random variables are fixed with probability $1 - 3 \cdot 2^{-0.09k}$, and so $\text{stsExt}_H(\mathbf{X})$ is $(2^{-k^{\Omega(1)}} + 3 \cdot 2^{-0.09k}) = 2^{-k^{\Omega(1)}}$ -close to the uniform distribution on $m = 0.9k$ bits. \square

As discussed, in order for [Lemma 4.7](#) to yield an explicit extractor, we need to explicitly construct an $\text{STS}(N)$ with a small STS-independence number α_{STS} . It is known that such objects exist for $\alpha = \Theta(\sqrt{N \log N})$ [[RŠ94](#)], but such upper bounds on α_{STS} are probabilistic, and there does not appear to be known explicit constructions. We use a recent result on the cap set problem in order to explicitly construct an $\text{STS}(N)$ with no big STS-independent set.

Lemma 4.8. *There is a constant $C \geq 1$ such that for all $N \in \mathbb{N}$, there exists an explicit construction of an $\text{STS}(N)$ with STS-independence number $\alpha_{\text{STS}} \leq CN^{0.9228}$.*

Proof. For now, assume that N is a power of 3. Define a hypergraph $H = (V, \mathcal{E})$ where $V = \mathbb{F}_3^{\log_3 N}$, and \mathcal{E} are the lines in $\mathbb{F}_3^{\log_3 N}$. H is clearly an $\text{STS}(N)$, because no two lines intersect at more than one point. Furthermore, notice that an independent set S in H is exactly a set of points in $\mathbb{F}_3^{\log_3 N}$ such that no three are in a line. Such a set is called a *cap set*, and it is known that a cap set in $\mathbb{F}_3^{\log_3 N}$ has size at most $2.756^{\log_3 N} \leq N^{0.9228}$ [[CLP17, EG17](#)]. It is straightforward to extend this construction to when N is not a power of three, by writing N in base-3, doing the above construction for every component of its base-3 representation, and adding the max cap set sizes from each component. \square

Combining [Lemmas 4.7](#) and [4.8](#) immediately gives us [Theorem 4.1](#).

Remark 4.9. *In order to reduce the number of required good sources, the natural idea is to look for better explicit constructions of $\text{STS}(N)$'s with small α_{STS} . As mentioned, it appears that this area has not yet been explored. Furthermore, we know that upper bounding the size of cap sets can only get us so far: there exist cap sets in $\mathbb{F}_3^{\log_3 N}$ of size $2.2174^{\log_3 N}$ [[Ede04](#)], so an $\text{STS}(N)$ constructed in the above manner will have $\alpha_{\text{STS}} \geq N^{0.724}$. Thus, even optimal cap set bounds would not allow us to extract from fewer than $K = N^{0.724}$ good sources. Thus, if we want to reach $K = O(\sqrt{N \log N})$, the theoretical limit that can be achieved using $\text{STS}(N)$'s, we need significantly new ideas.*

In the next section, we present new ideas that, in some settings, allow us to bring down the requirement on the number of good sources, K , to a function that is *quasi-polylogarithmic* in N . Then, in the next section, we show how to remove any restriction on the setting and extract from just $K = \sqrt{N \cdot \mathcal{R}_N} = N^{0.5+o(1)}$. In fact, this requirement on K would become exactly $O(\sqrt{N \log N})$ if we had access to explicit optimal Ramsey graphs.

4.2 The wedge extractor

In this section, we prove our first main theorem:

Theorem 4.10 ([Theorem 1](#), restated). *There exist universal constants $C, \gamma > 0$ such that for all $N, K, n, k \in \mathbb{N}$ satisfying $k \geq \log^C n$ and $K \geq \mathcal{R}_N^2$, there exists an explicit extractor $\text{Ext} : (\{0, 1\}^n)^N \rightarrow \{0, 1\}^m$ for (N, K, n, k) -sources of locality 0, with output length $m = k^{\Omega(1)}$ and error $\epsilon = 2^{-k^{\Omega(1)}}$, provided $N \leq k^\gamma$.*

In order to obtain this explicit construction, we will construct a more robust three-source extractor, call it over certain triples of sources, and take the XOR of the results. We will use the following type of induced subgraph to select these triples:

Definition 4.11 (wedge). *Let $G = (V, E)$ be a graph over N vertices. A wedge is a subset $S \subseteq V$ of size three such that the induced subgraph $G[S]$ is the complete bipartite graph $K_{1,2}$. The terminals of S are the two vertices in $G[S]$ of degree 1.*

Remark 4.12. While a wedge is unordered, we will need each wedge to order its vertices so that the inputs to each call of our more robust three-source extractor are well-defined. We order the vertices in each wedge in G by identifying the vertices of G with $[N]$, and having the wedge list its terminals first, where the terminals are listed in increasing order.

We define *activating* and *independent* sets with respect to wedges as follows.

Definition 4.13 (wedge-activating set). Let $G = (V, E)$ be a graph over N vertices. A wedge-activator in G is a subset $S \subseteq V$ that contains at least one wedge $W \subseteq V$.

Definition 4.14 (wedge-independent set). Let $G = (V, E)$ be a graph over N vertices. A wedge-independent set in G is a subset $S \subseteq V$ that is not a wedge-activating set. The wedge-independence number of G is the size of its max wedge-independent set, and is denoted by α_W .

We can now formally define our extractor over the wedges of a graph.

Definition 4.15 (wedge-extractor). Let $G = (V, E)$ be a graph over N vertices and let \mathcal{W} be the collection of all wedges in G . Let $2\text{Cond} : (\{0, 1\}^n)^2 \rightarrow \{0, 1\}^{m_1}$, $2\text{nmExt} : (\{0, 1\}^{m_2})^2 \rightarrow \{0, 1\}^m$ be the two-source condenser and two-source non-malleable extractor from [Lemma 3.8](#), where $m_2 = m_1 + \log \binom{N}{3}$. Let $\tau : \mathcal{W} \rightarrow \{0, 1\}^{\log \binom{N}{3}}$ be a tagging function that assigns a unique label to each wedge. We define the wedge-extractor over G for (N, K, n, k) -sources, $\text{wExt}_G : (\{0, 1\}^n)^N \rightarrow \{0, 1\}^m$, as:

$$\text{wExt}_G(\mathbf{X}) := \bigoplus_{W := (h, i, j) \in \mathcal{W}} 2\text{nmExt}(\tau(W) \circ 2\text{Cond}(\mathbf{X}_h, \mathbf{X}_j), \tau(W) \circ 2\text{Cond}(\mathbf{X}_i, \mathbf{X}_j)).$$

For an illustration, we refer the reader to [Figure 1b](#). Like the STS-extractor, the wedge-extractor takes a graph G as advice, and we will prove a general lemma showing that the performance of our extractor will depend on the value of α_W for G . Then, we will explicitly construct graphs with small α_W , which will immediately yield [Theorem 4.10](#).

Lemma 4.16. There exist constants $C, \gamma > 0$ such that for all $N, K, n, k \in \mathbb{N}$ satisfying $k \geq \log^C n$ and $N \leq k^\gamma$, the following holds. If $G = (V, E)$ is a graph over N vertices with wedge-independence number $\alpha_W < K$, then $\text{wExt}_G : (\{0, 1\}^n)^N \rightarrow \{0, 1\}^m$ is an extractor for (N, K, n, k) -sources of locality 0, with output length $m = k^{\Omega(1)}$ and error $\epsilon = 2^{-k^{\Omega(1)}}$.

Proof. Identify the vertices of G with $[N]$, and define $I := \{i \in [N] : \mathbf{X}_i \text{ is a good source}\}$. We know that $|I| \geq K$, and thus because G has wedge-independence number $\alpha_W < K$, there exists some wedge $W := (v_1, v_2, v_3) \in \mathcal{W}$ such that $W \subseteq I$.

Next, recall that \mathcal{W} holds *wedges*, and wedges have the property that if one wedge A has terminal vertices that both appear in another wedge B , then they must also be B 's terminal vertices (otherwise, because wedges are induced subgraphs, this would imply one of the wedges has more than one non-edge, contradicting the definition of a wedge). Furthermore, we originally had the wedges order their vertices with terminals appearing first (with the lower number terminal appearing first), and we did not select multiple copies of the same wedge.

Thus, we can partition $\mathcal{W} \setminus (v_1, v_2, v_3)$ into $\mathcal{W}_1, \mathcal{W}_2, \mathcal{W}_3, \mathcal{W}_4$ such that: the triples in \mathcal{W}_1 each hold vertex v_1 but not v_2 ; the triples in \mathcal{W}_2 each hold vertex v_2 but not v_1 ; the triples in \mathcal{W}_3 each

hold vertex v_1 and v_2 as the first two elements of the triple; and the triples in \mathcal{W}_4 each do not hold v_1 nor v_2 . Note that because a wedge has three vertices, and we do not have multiple copies of the same wedge, $|\mathcal{W}_3| \leq N - 2 < N$. Next, for each $\ell \in \{1, 2, 4\}$, define the random variable

$$\mathbf{Z}_\ell := \bigoplus_{B:=(h,i,j) \in \mathcal{W}_\ell} 2\text{nmExt}(\tau(B) \circ 2\text{Cond}(\mathbf{X}_h, \mathbf{X}_j), \tau(B) \circ 2\text{Cond}(\mathbf{X}_i, \mathbf{X}_j)),$$

and for each $A := (v_1, v_2, j) \in (\mathcal{W}_3 \cup (v_1, v_2, v_3))$, define the random variables

$$\begin{aligned} \mathbf{Y}_A^{(1)} &:= \tau(A) \circ 2\text{Cond}(\mathbf{X}_{v_1}, \mathbf{X}_j), \\ \mathbf{Y}_A^{(2)} &:= \tau(A) \circ 2\text{Cond}(\mathbf{X}_{v_2}, \mathbf{X}_j). \end{aligned}$$

Note that we can now rewrite

$$\text{wExt}_G(\mathbf{X}) = \mathbf{Z}_1 \oplus \mathbf{Z}_2 \oplus \mathbf{Z}_4 \oplus 2\text{nmExt}(\mathbf{Y}_W^{(1)}, \mathbf{Y}_W^{(2)}) \oplus \bigoplus_{B \in \mathcal{W}_3} 2\text{nmExt}(\mathbf{Y}_B^{(1)}, \mathbf{Y}_B^{(2)}).$$

We are now ready to do some fixings. First, briefly note that since we are using the $2\text{Cond}, 2\text{nmExt}$ from [Lemma 3.8](#), we will be importing our parameters for these objects from that lemma. Using these parameters, we set $C = C_0, \gamma = \gamma_0/C_0$.

Now, fix all random variables \mathbf{X}_i such that $i \notin (v_1, v_2, v_3)$. Next, fix random variable \mathbf{X}_{v_3} . As a result, \mathbf{Z}_4 is fixed, and $\mathbf{Y}_W^{(1)}, \mathbf{Y}_W^{(2)}$ become independent. Furthermore, by the strength of 2Cond in the second source, $\mathbf{Y}_W^{(1)}, \mathbf{Y}_W^{(2)}$ both become ϵ_2 -close to an $(\log \binom{N}{3} + m_1, m_1 - o(\log(1/\epsilon_1)))$ -source with probability $1 - 2\epsilon_1$.

Next, fix \mathbf{Z}_1 . By [Lemma 3.10](#), $\mathbf{Y}_W^{(1)}$ is $2\epsilon_2^{1/2}$ -close to a source with min-entropy at least $k' := m_1 - o(\log(1/\epsilon_1)) - m - \log(1/\epsilon_2)$ with probability at least $1 - 2\epsilon_2^{1/2}$. Next, fix \mathbf{Z}_2 , and the same thing holds for $\mathbf{Y}_W^{(2)}$. Note that as per the parameter settings in [Lemma 3.8](#), we have:

$$k' := m_1 - o(\log(1/\epsilon_1)) - m - \log(1/\epsilon_2) \geq m_1 - o(m_1^{\gamma_0/2}) - m_1^{\gamma_0/2} - m_1^{\gamma_0/2} \geq m_1 - 3m_1^{\gamma_0/2},$$

and because we have set $N \leq k^\gamma = k^{\gamma_0/C_0} = m_1^{\gamma_0}$, this becomes

$$k' \geq m_1 - 3m_1^{\gamma_0/2} \geq m_1 + \log \binom{N}{3} - m_1^{\gamma_0} \geq m_2 - m_2^{\gamma_0}.$$

Finally, recall that $|\mathcal{W}_3| < N$ and that $N \leq m_1^{\gamma_0} \leq m_2^{\gamma_0}$. So to summarize our fixings thus far: $\mathbf{Z}_1, \mathbf{Z}_2, \mathbf{Z}_4$ are all fixed, $\mathbf{Y}_W^{(1)}, \mathbf{Y}_W^{(2)}$ are independent, and with probability at least $1 - (2\epsilon_1 + 4\epsilon_2^{1/2})$, they are each $2\epsilon_2^{1/2}$ -close to a source with min-entropy at least $k' \geq m_2 - m_2^{\gamma_0}$. Additionally, we have $|\mathcal{W}_3| < m_2^{\gamma_0}$. Note also that for each $B \in \mathcal{W}_3$, the following holds:

1. $\mathbf{Y}_W^{(1)}$ is correlated with at most one of $\mathbf{Y}_B^{(1)}$ and $\mathbf{Y}_B^{(2)}$, and the same is true for $\mathbf{Y}_W^{(2)}$.
2. $\mathbf{Y}_B^{(1)}$ is correlated with at most one of $\mathbf{Y}_W^{(1)}$ and $\mathbf{Y}_W^{(2)}$, and the same is true for $\mathbf{Y}_B^{(2)}$.
3. $\text{support}(\mathbf{Y}_W^{(h)}) \cap \text{support}(\mathbf{Y}_B^{(i)}) = \emptyset$, for all $h, i \in [2]$.

The first two items are true because B lists v_1, v_2 as the first two elements of its tuple, by definition of \mathcal{W}_3 . The last item holds because of our tagging function, τ . Thus, for each $B \in \mathcal{W}_3$, either: $\mathbf{Y}_B^{(1)}$ is a function of only $\mathbf{Y}_W^{(1)}$, and $\mathbf{Y}_B^{(2)}$ is a function of only $\mathbf{Y}_W^{(2)}$; or $\mathbf{Y}_B^{(1)}$ is a function of only $\mathbf{Y}_W^{(2)}$, and $\mathbf{Y}_B^{(2)}$ is a function of only $\mathbf{Y}_W^{(1)}$. Furthermore, these functions have no fixed points. **Definition 3.6** says that we are now in a good position to use the non-malleability of 2nmExt .

The non-malleability of 2nmExt tells us that if $\mathbf{Y}_W^{(1)}, \mathbf{Y}_W^{(2)}$ were both guaranteed to have entropy k' , then with probability $1 - \epsilon_3^{1/2}$, $2\text{nmExt}(\mathbf{Y}_W^{(1)}, \mathbf{Y}_W^{(2)})$ is at most $\epsilon_3^{1/2}$ -far from \mathbf{U}_m , even after fixing $2\text{nmExt}(\mathbf{Y}_B^{(1)}, \mathbf{Y}_B^{(2)})$ for every $B \in \mathcal{W}_3$. We have instead that with probability at least $1 - (2\epsilon_1 + 4\epsilon_2^{1/2})$, the product distribution $(\mathbf{Y}_W^{(1)}, \mathbf{Y}_W^{(2)})$ is $4\epsilon_2^{1/2}$ -close to the product distribution of two independent sources that have entropy k' . Thus with probability at least $1 - (2\epsilon_1 + 4\epsilon_2^{1/2} + \epsilon_3^{1/2})$, $2\text{nmExt}(\mathbf{Y}_W^{(1)}, \mathbf{Y}_W^{(2)})$ is at most $(4\epsilon_2^{1/2} + \epsilon_3^{1/2})$ -far from \mathbf{U}_m after these fixings.

Thus, everything except $2\text{nmExt}(\mathbf{Y}_W^{(1)}, \mathbf{Y}_W^{(2)})$ has been fixed, so with probability at least $1 - (2\epsilon_1 + 4\epsilon_2^{1/2} + \epsilon_3^{1/2})$ over these fixings, $\text{wExt}_G(\mathbf{X})$ is at most $(4\epsilon_2^{1/2} + \epsilon_3^{1/2})$ -far from \mathbf{U}_m . In other words, $\text{wExt}_G(\mathbf{X})$ is at most ϵ -far from \mathbf{U}_m , where $\epsilon = 2\epsilon_1 + 4\epsilon_2^{1/2} + \epsilon_3^{1/2} + 4\epsilon_2^{1/2} + \epsilon_3^{1/2}$. Recalling that $\epsilon_1 = \epsilon_2 = 2^{-m_1^{\gamma_0/2}}$, $\epsilon_3 = 2^{-m_2^{\Omega(1)}}$, $m_2 \geq m_1 = k^{1/C}$, and $m = m_1^{\Omega(1)}$, we see that

$$\epsilon = 2^{-k^{\Omega(1)}} \text{ and } m = k^{\Omega(1)},$$

completing the proof. \square

In order for **Lemma 4.16** to yield an explicit extractor, we need to explicitly construct a graph G over N vertices whose max wedge-independent set has small size, α_W . We show that we can explicitly construct such objects using classical Ramsey graphs.

Lemma 4.17. *For all $N \in \mathbb{N}$, there exists an explicit construction of a graph G with N vertices that has max wedge-independent set size $\alpha_W < \mathcal{R}_N^2$.*

Proof. Let $G = (V, E)$ be a graph over N vertices. We claim that if G has no clique nor independent set of size K , then it has no wedge-independent set of size K^2 . To see why this is true, we prove the contrapositive. Let $S \subseteq V$ be a wedge-independent set of size K^2 . Partition S into sets S_1, S_2, \dots, S_ℓ such that $\{G[S_i]\}_{i \in [\ell]}$ are the connected components of $G[S]$. Observe that each $G[S_i]$ must be a clique: if not, the shortest path between some two vertices in $G[S_i]$ uses more than one edge; but any two consecutive edges on a shortest path must create a wedge (or else it is not a shortest path), which contradicts S being a wedge-independent set. Now, note that because $|S| = K^2$, either some S_i has size $\geq K$, or $\ell \geq K$. In the former case, we have a clique of size K ; in the latter case, we can pick one vertex from each S_i to create an independent set of size K , which proves the claim. By definition of \mathcal{R}_N , we can explicitly construct a Ramsey graph G over N vertices with no clique nor independent set of size \mathcal{R}_N . Thus, we can explicitly construct a graph H with $\alpha_W < \mathcal{R}_N^2$, by simply taking $H = G$. \square

Combining **Lemmas 4.16** and **4.17** immediately gives us **Theorem 4.10**. Next, we construct an extractor which generalizes all of the constructions we've seen thus far. It will allow us to remove the restriction between N and k seen in **Theorem 4.10**.

4.3 The main extractor

In this section, we prove our second main theorem:

Theorem 4.18 (Theorem 2, restated). *There exists a universal constant $C > 0$ such that for all $N, K, n, k \in \mathbb{N}$ satisfying $k \geq \log^C n$ and $K \geq \sqrt{N \cdot \mathcal{R}_N}$, there exists an explicit extractor $\text{Ext} : (\{0, 1\}^n)^N \rightarrow \{0, 1\}^m$ for (N, K, n, k) -sources of locality 0, with output length $m = k^{\Omega(1)}$ and error $\epsilon = 2^{-k^{\Omega(1)}}$.*

In order to obtain this explicit construction, we will construct a more general version of the robust three-source extractor developed in the previous section (via a composition of a two-source condenser with a seedless non-malleable extractor). We will invoke this new function over certain sets of sources, and take the XOR of the results. We will use the following type of hypergraph to select sets of sources:

Definition 4.19 (fragile set system). *Given a graph $G = (V, E)$ over N vertices, a fragile set is a subset $S \subseteq V$ of size at least 3 such that $G[S]$ contains exactly one edge, denoted $\pi(S)$. A fragile set system is a tuple $H = (G, \mathcal{S})$ where $G = (V, E)$ is a graph, and \mathcal{S} is a collection of some (not necessarily all) fragile sets in G .*

Remark 4.20. *While a fragile set system is unordered, we will need each fragile set to order its vertices so that the inputs to each call of our new function are well-defined. We identify the vertices of the fragile set system with $[N]$, and order the vertices in each fragile set S as a triple of the form (u, v, S') , where u, v are the endpoints of the edge in S listed in increasing order, and $S' = S \setminus \{u, v\}$.*

We define *activating* and *independent* sets with respect to fragile set systems as follows.

Definition 4.21 (FSS-activating set). *Let $G = (V, E)$ be a graph over N vertices, and let $H = (G, \mathcal{S})$ be a fragile set system. An FSS-activating set in H is a subset $T \subseteq V$ such that for some fragile set $S \in \mathcal{S}$, it holds that $|T \cap S| = 3$ and $\pi(S) \subseteq T$.*

Definition 4.22 (FSS-independent set). *Let $G = (V, E)$ be a graph over N vertices, and let $H = (G, \mathcal{S})$ be a fragile set system. An FSS-independent set in H is a subset $T \subseteq V$ that is not an FSS-activating set. The FSS-independence number of H is the size of its max FSS-independent set, and is denoted by α_{FSS} .*

We can now formally define our extractor over fragile set systems.

Definition 4.23 (FSS-extractor). *Let $H = (G, \mathcal{S})$ be a fragile set system over N vertices. Let $2\text{Cond} : (\{0, 1\}^n)^2 \rightarrow \{0, 1\}^{m_1}$, $2\text{nmExt} : (\{0, 1\}^{m_2})^2 \rightarrow \{0, 1\}^m$ be the two-source condenser and non-malleable extractor from Lemma 3.8, where $m_2 = m_1 + \log(\deg(H))$. Let $\tau : \mathcal{S} \rightarrow \{0, 1\}^{\log(\deg(H))}$ be a tagging function that assigns a unique label to each fragile set that shares a common edge. We define the FSS-extractor over H for (N, K, n, k) -sources, $\text{fssExt}_H : (\{0, 1\}^n)^N \rightarrow \{0, 1\}^m$, as:*

$$\text{fssExt}_H(\mathbf{X}) := \bigoplus_{S=(u,v,S') \in \mathcal{S}} 2\text{nmExt}(\tau(S) \circ 2\text{Cond}(\mathbf{X}_u, \bigoplus_{i \in S'} \mathbf{X}_i), \tau(S) \circ 2\text{Cond}(\mathbf{X}_v, \bigoplus_{i \in S'} \mathbf{X}_i)).$$

For an illustration, we refer the reader to [Figure 2](#). Like the previous two extractors, the FSS-extractor takes an object (here, a fragile set system) H as advice, and we will prove a general lemma showing that the performance of our extractor will depend on the values $\alpha_{\text{FSS}}, \deg(H)$ of H . Then, we will explicitly construct graphs with small α_{FSS} and very small $\deg(H)$, which will immediately yield [Theorem 4.18](#).

Lemma 4.24 (Main extractor). *There exist constants $C, \gamma > 0$ such that for all $N, K, n, k \in \mathbb{N}$ satisfying $k \geq \log^C n$, the following holds. If $H = (G, \mathcal{S})$ is a fragile set system over N vertices with FSS-independence number $\alpha_{\text{FSS}} < K$, and degree $t = \deg(H)$, then $\text{fssExt}_H : (\{0, 1\}^n)^N \rightarrow \{0, 1\}^m$ is an extractor for (N, K, n, k) -sources of locality 0, with output length $m = k^{\Omega(1)}$ and error $\epsilon = 2^{-k^{\Omega(1)}}$, provided $t - 1 \leq k^\gamma$.*

Proof. Identify the vertices of G with $[N]$, and define $I := \{i \in [N] : \mathbf{X}_i \text{ is a good source.}\}$ We know that $|I| \geq K$, and thus because G has activation independence number $\alpha_{\text{FSS}} < K$, there exists some fragile set $S := (v_1, v_2, S') \in \mathcal{S}$ and some vertex $v_3 \in S'$ such that $v_1, v_2, v_3 \in I$.

Next, recall that \mathcal{S} holds *fragile sets*, and a fragile set is defined to contain exactly one edge. Furthermore, we have specified that when interpreting a fragile set as a triple of the form (u, v, T') , u, v should denote the endpoints of the edge in the fragile set, and it should be the case that $u < v$ under the identification of $V(G)$ by $[N]$. Thus, we know that if any other fragile set $T \neq S \in \mathcal{S}$ contains both vertices v_1, v_2 , then it must contain them in that order, and as the first two elements of the triple used to represent T .

Thus, we can partition $\mathcal{S} \setminus S$ into $\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3, \mathcal{S}_4$ such that: each fragile set in \mathcal{S}_1 holds vertex v_1 but not v_2 ; each fragile set in \mathcal{S}_2 holds vertex v_2 but not v_1 ; each fragile set in \mathcal{S}_3 holds vertex v_1 and v_2 , but as the first two elements of its triple representation; and each fragile set in \mathcal{S}_4 does not hold v_1 nor v_2 . Note that by our definition of degree of a fragile set system, we have $|\mathcal{S}_3| \leq t - 1$.

Next, for each $\ell \in \{1, 2, 4\}$, define the random variable

$$\mathbf{Z}_\ell := \bigoplus_{B := (h, i, J) \in \mathcal{S}_\ell} 2\text{nmExt}(\tau(B) \circ 2\text{Cond}(\mathbf{X}_h, \bigoplus_{j \in J} \mathbf{X}_j), \tau(B) \circ 2\text{Cond}(\mathbf{X}_i, \bigoplus_{j \in J} \mathbf{X}_j)),$$

and for each $A := (v_1, v_2, J) \in (\mathcal{S}_3 \cup S)$, define the random variables

$$\begin{aligned} \mathbf{Y}_A^{(1)} &:= \tau(A) \circ 2\text{Cond}(\mathbf{X}_{v_1}, \bigoplus_{j \in J} \mathbf{X}_j), \\ \mathbf{Y}_A^{(2)} &:= \tau(A) \circ 2\text{Cond}(\mathbf{X}_{v_2}, \bigoplus_{j \in J} \mathbf{X}_j). \end{aligned}$$

Note that we can now rewrite

$$\text{fssExt}_H(\mathbf{X}) = \mathbf{Z}_1 \oplus \mathbf{Z}_2 \oplus \mathbf{Z}_4 \oplus 2\text{nmExt}(\mathbf{Y}_S^{(1)}, \mathbf{Y}_S^{(2)}) \oplus \bigoplus_{B \in \mathcal{S}_3} 2\text{nmExt}(\mathbf{Y}_B^{(1)}, \mathbf{Y}_B^{(2)}).$$

The remainder of the proof is identical to the proof of [Lemma 4.16](#), with the exception that because $|\mathcal{S}_3| \leq t - 1$ and our tags have length $\log t$, we only need the restriction $t - 1 \leq k^\gamma$ (as opposed to $N \leq k^\gamma$) for the final conditioning step that uses the non-malleability of 2nmExt . \square

To yield [Theorem 4.18](#), we will now construct a fragile set system with not-too-large α_{FSS} and very small $\deg(H)$. It is worth noting that because we achieve $\deg(H) = 1$, our fssExt would have actually worked even if we replaced the non-malleable extractor calls with calls to the Hadamard extractor, Had , because the set \mathcal{S}_3 from our analysis will be empty.

Lemma 4.25. *For all sufficiently large $N \in \mathbb{N}$, there exists an explicit construction of a fragile set system $H = (G, \mathcal{S})$ over N vertices with activation independence number $\alpha_{\text{FSS}} < \sqrt{N \cdot \mathcal{R}_N}$ and degree $\deg(H) = 1$.*

Proof. In order to construct H , we first construct G , and then describe how to pick the fragile sets, \mathcal{S} . To construct G , place N vertices into $r := \sqrt{N/\mathcal{R}_N}$ disjoint sets (call them clouds) of size $\ell := \sqrt{N \cdot \mathcal{R}_N}$, labeled S_1, S_2, \dots, S_r . For every $1 \leq i < j \leq r$, draw edges between clouds S_i, S_j so that $G[S_i \cup S_j]$ is the best possible bipartite Ramsey graph that can be explicitly constructed.

To construct \mathcal{S} , iterate over every $1 \leq i < j \leq r$, and every pair of vertices $u \in S_i, v \in S_j$. Let $N_j^-(u) := \{y \in S_j : uy \notin E(G)\}$ denote the non-neighbors of u in set S_j . If there is an edge from u to v , and $N_j^-(u)$ is non-empty, then add to \mathcal{S} the set $T := \{u, v\} \cup N_j^-(u)$. This completes the construction of $H = (G, \mathcal{S})$.

First, observe that H is a fragile set system, since every set we add to \mathcal{S} contains a single edge and at least three vertices. Next, since our construction iterates over each edge once, and this is the only time a fragile set may be added containing this edge (since fragile sets contain exactly one edge), H has degree 1.

Finally, consider any subset $A \subseteq V$ such that there exist two clouds S_i, S_j with $i < j$ such that A shares at least $2\mathcal{R}_\ell$ vertices with S_i and at least $2\mathcal{R}_\ell$ vertices with S_j . Call these intersections B_i and B_j , respectively. By definition of \mathcal{R} (Definition 1.3), we know that $G' := G[B_i \cup B_j]$ contains no bipartite clique nor independent set of size \mathcal{R}_ℓ . Thus, there must be at least one vertex $v \in B_i$, and two vertices $x, y \in B_j$ such that v is adjacent to x , and not adjacent to y : otherwise, either half the vertices in B_i have degree 0 in G' , or half the vertices in B_i have degree $|B_j|$ in G' , which guarantees that we can find a bipartite clique or independent set in G' of size $|B_i|/2 = \mathcal{R}_\ell$.

Thus, we see that any subset $A \subseteq V$ that is guaranteed to share at least $2\mathcal{R}_\ell$ vertices with two different clouds each is guaranteed to be an activating set of H . We know this will happen if $|A| \geq \ell + 2\mathcal{R}_\ell \cdot r$. Given our setting of ℓ, r , we have $\sqrt{N \cdot \mathcal{R}_N} \geq \ell + 2\mathcal{R}_\ell \cdot r$ for sufficiently large N . Thus, for sufficiently large N , any set of size at least $\sqrt{N \cdot \mathcal{R}_N}$ is activating. \square

Combining Lemmas 4.24 and 4.25 immediately yields Theorem 4.18. Our second main theorem removes any restriction between N and k , and requires just $K = \sqrt{N \cdot \mathcal{R}_N}$ good sources. We also note the following two remarks:

Remark 4.26. *Given an explicit optimal Ramsey graph, our construction exactly matches, up to constant factors, the original requirement on K promised by non-explicit STS(N)'s.*

Remark 4.27. *It is worth noting the generality of the above extractor. Depending on the trade-off one can achieve between $\alpha_{\text{FSS}}(H)$ and $\deg(H)$ in explicit constructions of fragile set systems, one can successfully extract from a wide range of regimes. For example: if we can achieve very small α_{FSS} and large $\deg(H)$, then we can extract when K is very small, but the restriction $\deg(H) \leq k^\gamma$ will be more restrictive, since $\deg(H)$ will be a larger function of N . Indeed, this is the case of wExt_G , which is simply fssExt_H where H is the fragile set system made out of the complements of wedges in the complement of G . On the other hand, if one is willing to allow slightly larger α_{FSS} , it is possible to completely get rid of any restriction between N, k if one can make $\deg(H)$ a constant. Indeed, this is the case of stsExt_H , which is simply $\text{fssExt}_{H'}$ where H' is the fragile set system made out of putting a single edge in each hyperedge of H . As seen with Theorem 4.18, however, we can exercise the full generality of fragile set systems to find an explicit construction that yields much better parameters than the fragile set systems that arise from explicitly constructible STS(N)'s.*

In the next section, we show that, in fact, our FSS-extractor works even if we our (N, K, n, k) -sources have polynomial locality, d .

5 Extracting from polynomial locality

In this section, we will show that our main extractor, the FSS-extractor from [Definition 4.23](#), can extract from from polynomial locality. In fact, we show the stronger result that our less-general wedge-extractor ([Definition 4.15](#)) can do the same. In particular, we prove the following:

Theorem 5.1. *There exist universal constants $C, \gamma > 0$ such that for all $N, K, n, k, d \in \mathbb{N}$ satisfying $k \geq \log^C n$ and $d/K \leq (\gamma/N) \min\{k^\gamma, N^{1/24}\}$, there exists an explicit extractor $\text{Ext} : (\{0, 1\}^n)^N \rightarrow \{0, 1\}^m$ for (N, K, n, k) -sources of locality d , with output length $m = k^{\Omega(1)}$ and error $\epsilon = 2^{-k^{\Omega(1)}}$.*

We note that one parameter setting yields our third main theorem:

Theorem 5.2 ([Theorem 3](#), restated). *There exist universal constants $C, \gamma > 0$ such that for all $N, K, n, k, d \in \mathbb{N}$ satisfying $k \geq \log^C n$ and $K \geq N^{1-\gamma}$, and $d \leq K^\gamma$, there exists an explicit extractor $\text{Ext} : (\{0, 1\}^n)^N \rightarrow \{0, 1\}^m$ for (N, K, n, k) -sources of locality d , with output length $m = k^{\Omega(1)}$ and error $\epsilon = 2^{-k^{\Omega(1)}}$, provided $N \leq k^\gamma$.*

In order to prove [Theorem 5.1](#), we recall that our wedge-extractor can be instantiated with any graph G , and that its performance will be related to independence properties of G . We show that, for some new notion of independence, this remains true even in the 1-local case. Then, we show how to reduce from the d -local case to the 1-local case.

Previously, we were concerned with constructing a graph so that no matter how K good sources were placed on the vertices, they were guaranteed to cover a wedge. In the case of 1-locality, we must defend against a stronger adversary. In particular, we want to make sure that no matter how K good sources are placed on the vertices of our graph, some wedge is covered, and the two good sources on the terminals of the wedge do not each influence one bad source placed on a distinct endpoint of the same edge. We capture this notion with the following definition.

Definition 5.3 (cloud _{w} -wedge). *Let $G = (V, E)$ be a graph over N vertices, $\{h, i, j\} \subseteq V$ be a subset of size 3, and $A_h, A_i, A_j \subseteq V$ be three nonempty disjoint subsets (clouds) of size at most w such that A_h contains h , A_i contains i , and A_j contains j . We call $(\{h, i, j\}, \{A_h, A_i, A_j\})$ a cloud _{w} -wedge if $\{h, i, j\}$ is a wedge and there exist no edges crossing between the two clouds holding the terminals of $\{h, i, j\}$.*

We define *activating* and *independent* sets with respect to cloud-wedges as follows.

Definition 5.4 (cloud _{w} -wedge-activating set). *Let $G = (V, E)$ be a graph over N vertices. A cloud _{w} -wedge-activating set in G is a subset $S \subseteq V$ such that the following holds. For any family of $|S|$ nonempty disjoint clouds $\mathcal{B} = \{A_s \subseteq V\}_{s \in S}$ where A_s contains s and $|A_s| \leq w$ for all $s \in S$, there exist three vertices $h, i, j \in S$ such that $(\{h, i, j\}, \{A_h, A_i, A_j\})$ is a cloud _{w} -wedge.*

Definition 5.5 (cloud _{w} -wedge-independent set). *Let $G = (V, E)$ be a graph over N vertices. A cloud _{w} -wedge-independent set in G is a subset $S \subseteq V$ that is not a cloud _{w} -wedge-activating set. The cloud _{w} -wedge-independence number in G is the size of its max cloud _{w} -wedge-independent set, and is denoted by $\alpha_{C_w, W}$.*

For an illustration, we refer the reader to [Figure 3](#). We now prove a general lemma showing that our wedge-extractor can extract from 1-local adversarial sources, given a graph G with small $\alpha_{C_w\mathcal{W}}$ as advice. Then, we will show how to explicitly construct graphs with small $\alpha_{C_w\mathcal{W}}$, which will immediately yield an explicit extractor for 1-local adversarial sources. We then provide a lemma that shows d -local adversarial sources can be reduced to 1-local adversarial sources. Using this reduction, we show how to extract from polynomial locality, and ultimately prove [Theorem 5.1](#).

Lemma 5.6. *There exist constants $C, \gamma > 0$ such that for all $N, K, n, k \in \mathbb{N}$ satisfying $k \geq \log^C n$, the following holds. If \mathbf{X} is an (N, K, n, k) -source of locality 1 such that at most w sources depend on a single good source, and $G = (V, E)$ is a graph over N vertices with $\alpha_{C_w\mathcal{W}} < K$, then the wedge-extractor from [Definition 4.15](#), $\text{wExt}_G : (\{0, 1\}^n)^N \rightarrow \{0, 1\}^m$, is an extractor for \mathbf{X} , with output length $m = k^{\Omega(1)}$ and error $\epsilon = 2^{-k^{\Omega(1)}}$, provided $\max\{w^2, \log N\} \leq k^\gamma$.*

Proof. Identify the vertices of G with $[N]$, and define $I := \{i \in [N] : \mathbf{X}_i \text{ is a good source}\}$. We know that $|I| \geq K$. For each $i \in I$, define $D_i := \{j \in [N] : \mathbf{X}_i, \mathbf{X}_j \text{ are not independent}\}$ to be the sources depending on \mathbf{X}_i , including \mathbf{X}_i itself. Let \mathcal{W} be the collection of all wedges in G , and let $W = (v_1, v_2, v_3) \in \mathcal{W}$ be such that $W \subseteq I$ and $(\{v_1, v_2, v_3\}, D_{v_1}, D_{v_2}, D_{v_3})$ is a cloud- w -wedge in G . We know such a wedge exists, because $\alpha_{C_w\mathcal{W}} < K$. Furthermore, we know that $|D_{v_1}| \leq w$ and $|D_{v_2}| \leq w$ by the statement of our lemma, and we know that D_{v_1}, D_{v_2} are disjoint and have no crossing edges (i.e. edges with one endpoint in D_{v_1} and the other in D_{v_2}), because our source is 1-local and by definition of cloud-wedge.

Partition $\mathcal{W} \setminus (v_1, v_2, v_3)$ into $\mathcal{W}_1, \mathcal{W}_2, \mathcal{W}_3, \mathcal{W}_4$ such that: the triples in \mathcal{W}_1 have some vertex in D_{v_1} but no vertex in D_{v_2} ; the triples in \mathcal{W}_2 have some vertex in D_{v_2} but no vertex in D_{v_1} ; the triples in \mathcal{W}_3 have some vertex in D_{v_1} and some vertex in D_{v_2} ; and the triples in \mathcal{W}_4 have no vertex in D_{v_1} nor D_{v_2} .

Now, let $(u_1, u_2, u_3) \in \mathcal{W}_3$, and suppose for contradiction $u_3 \in D_{v_1}$. By definition of \mathcal{W}_3 , and because D_{v_1}, D_{v_2} are disjoint, we know that either u_1 or u_2 is in D_{v_2} . Furthermore, recall that our triple representation of a wedge lists the two non-adjacent vertices (terminals) first, and so u_1u_3 and u_2u_3 are both edges in G . But D_{v_1}, D_{v_2} have no crossing edges, so this is a contradiction. So $u_3 \notin D_{v_1}$, and identical reasoning shows that $u_3 \notin D_{v_2}$. Moreover, since we established that each wedge in \mathcal{W}_3 has one terminal in D_{v_1} and the other terminal in D_{v_2} , it must be the case that $|\mathcal{W}_3| \leq |D_{v_1}| \cdot |D_{v_2}| \leq w^2$: if not, then two wedges would share the same terminals, which would create a cycle of length 4 (contradicting that G is C_4 -free).

Next, just as in the proof to [Lemma 4.16](#), for each $\ell \in \{1, 2, 4\}$, define the random variable

$$\mathbf{Z}_\ell := \bigoplus_{B:=(h,i,j) \in \mathcal{W}_\ell} 2\text{nmExt}(\tau(B) \circ 2\text{Cond}(\mathbf{X}_h, \mathbf{X}_j), \tau(B) \circ 2\text{Cond}(\mathbf{X}_i, \mathbf{X}_j)),$$

and for each $A := (u_1, u_2, u_3) \in (\mathcal{W}_3 \cup (v_1, v_2, v_3))$, define the random variables

$$\begin{aligned} \mathbf{Y}_A^{(1)} &:= \tau(A) \circ 2\text{Cond}(\mathbf{X}_{u_1}, \mathbf{X}_{u_3}), \\ \mathbf{Y}_A^{(2)} &:= \tau(A) \circ 2\text{Cond}(\mathbf{X}_{u_2}, \mathbf{X}_{u_3}). \end{aligned}$$

Note that we can now rewrite

$$\text{wExt}_G(\mathbf{X}) = \mathbf{Z}_1 \oplus \mathbf{Z}_2 \oplus \mathbf{Z}_4 \oplus 2\text{nmExt}(\mathbf{Y}_W^{(1)}, \mathbf{Y}_W^{(2)}) \oplus \bigoplus_{B \in \mathcal{W}_3} 2\text{nmExt}(\mathbf{Y}_B^{(1)}, \mathbf{Y}_B^{(2)}).$$

By our selection of $\mathcal{W}_1, \mathcal{W}_2, \mathcal{W}_3, \mathcal{W}_4$, we know that if we fix all *good* sources except $\mathbf{X}_{v_1}, \mathbf{X}_{v_2}$, then \mathbf{Z}_4 is fixed, \mathbf{Z}_1 is correlated only with \mathbf{X}_{v_1} , and \mathbf{Z}_2 is correlated only with \mathbf{X}_{v_2} . Furthermore, for each $B = (u_1, u_2, u_3) \in \mathcal{W}_3$, we established that $u_3 \notin D_{v_1} \cup D_{v_2}$, and one terminal of B is in D_{v_1} while the other terminal is in D_{v_2} , and D_{v_1}, D_{v_2} are disjoint. Thus for each $B \in \mathcal{W}_3$, either: $\mathbf{Y}_B^{(1)}$ is a function of only $\mathbf{Y}_W^{(1)}$, and $\mathbf{Y}_B^{(2)}$ is a function of only $\mathbf{Y}_B^{(2)}$; or $\mathbf{Y}_B^{(1)}$ is a function of only $\mathbf{Y}_W^{(2)}$, and $\mathbf{Y}_B^{(2)}$ is a function of only $\mathbf{Y}_W^{(1)}$; and none of these functions have fixed points.

Since we know that $|\mathcal{W}_3| \leq w^2$, we know that as long as $w^2 \leq k^\gamma$ (and $\log \binom{N}{3} \leq k^\gamma/2$, though this is rarely the dominant constraint), we can perform the same conditioning steps as in the proof to [Lemma 4.16](#) to show that $\text{wExt}_G(\mathbf{X})$ is ϵ -close to \mathbf{U}_m , for $m = k^{\Omega(1)}, \epsilon = 2^{-k^{\Omega(1)}}$. \square

In order to use the above lemma, we need to find and explicitly construct graphs that have small cloud-wedge-independence numbers. We recall the definition of \mathbf{C}_4 -free graphs, and show that \mathbf{C}_4 -free graphs with no big (standard) independent set have exactly this property.

Definition 5.7 (\mathbf{C}_4 -free graphs). *A graph $G = (V, E)$ is \mathbf{C}_4 -free if it contains no cycle of length 4 as a subgraph (induced or not induced).*

Lemma 5.8. *Let $G = (V, E)$ be a \mathbf{C}_4 -free graph over N vertices with independence number α . For any $w \in \mathbb{N}$, the cloud_w -wedge-independence number of G is $\alpha_{\mathbf{C}_w\mathbf{W}} < 64w^4\alpha$.*

To prove [Lemma 5.8](#), we will need the following generalization of a wedge.

Definition 5.9 (star). *Let $G = (V, E)$ be a graph. A subset $S \subseteq V$ is an ℓ -star if $G[S]$ is the complete bipartite graph $\mathbf{K}_{1,\ell}$ with one vertex on the left and ℓ on the right. The vertices of degree 1 are known as the terminals of the star.*

The idea is to show that given any large enough subset of vertices in a \mathbf{C}_4 -free graph with small α , some large star (say, with ℓ terminals) must be covered. Then, if we show that the structure of \mathbf{C}_4 -free graphs forbids one from finding ℓ nonempty disjoint subsets such that each pair has a crossing edge, we will have found exactly a cloud-wedge. Formally, we must prove the following two structural lemmas.

Lemma 5.10 (Structural Lemma 1). *Let $G = (V, E)$ be a \mathbf{C}_4 -free graph over N vertices with independence number α . Then any subset $S \subseteq V$ of size K must contain an ℓ -star, where $\ell = \lceil \frac{K}{2\alpha} - 1 \rceil \geq K/(4\alpha)$.*

Proof. Consider the subgraph $H = G[S]$ induced by S . If $\Delta(H) \leq 2\ell - 1$, then we can greedily find an independent set of size $\alpha' = \lceil \frac{K}{2\ell} \rceil$ by iteratively selecting a vertex and discarding its neighbors. But since $\ell = \lceil \frac{K}{2\alpha} \rceil - 1$, we have $\alpha' > \alpha$, which is not possible. Thus, we know $\Delta(H) \geq 2\ell$. Consider a max degree vertex v and 2ℓ of its neighbors, denoted $N(v)$. The induced subgraph $H[N(v)]$ must have max degree at most 1, because otherwise we can find a cycle of length 4 in $H[\{v\} \cup N(v)]$. Thus $H[N(v)]$ must have an independent set of size at least ℓ , which forms the star $\mathbf{K}_{1,\ell}$ with v . \square

Lemma 5.11 (Structural Lemma 2). *Let $G = (V, E)$ be a \mathbf{C}_4 -free graph over N vertices, and let $D_1, D_2, \dots, D_\ell \subseteq V$ be ℓ nonempty disjoint subsets of size at most w . If $\ell \geq 16w^4$, then for some $i \neq j \in [\ell]$, D_i, D_j have no crossing edge.*

Proof. We prove the contrapositive. Let $D = \bigcup_{i \in [\ell]} D_i$, and let $H = G[D]$ be the subgraph of G , over at most $w\ell$ vertices, induced by D . Suppose that each pair of subsets has some crossing edge. Then for every $i \in [\ell]$, there exists some $v_i \in D_i$ of degree at least $(\ell - 1)/w$ in H . For each $i \in [\ell]$, let N_i denote the set of neighbors of v_i . By our selection of v_i , we know that each $|N_i| \geq (\ell - 1)/w$. Furthermore, we know that for each $i \neq j \in [\ell]$, $|N_i \cap N_j| \leq 1$, because otherwise v_i, v_j share at least two common neighbors, which produces a cycle of length 4. We can now use a standard technique from the theory of Steiner systems to upper bound ℓ . First, note that the number of vertex pairs in H is at most $\binom{w\ell}{2}$. Next, the number of vertex pairs *within* a subset N_i is at least $\binom{(\ell-1)/w}{2}$. Because of our intersection property, no vertex pair can be counted in more than one N_i . Thus, it must be the case that $\ell \cdot \binom{(\ell-1)/w}{2} \leq \binom{w\ell}{2}$, and in particular that $16w^4 > \ell$. \square

Equipped with these two lemmas, it is straightforward to prove [Lemma 5.8](#).

Proof of Lemma 5.8. Let S be an arbitrary subset of size $K \geq 64w^4\alpha$ in G , and let $\mathcal{B} = \{A_s \subseteq V\}_{s \in S}$ be an arbitrary collection of nonempty disjoint subsets such that for each $s \in S$, A_s contains s and has size at most w . By [Lemma 5.10](#), S contains an ℓ -star of size $\ell \geq 16w^4$. Let v_1, v_2, \dots, v_ℓ denote the terminals of this star, and let $u \in S$ denote the other vertex in the star. By [Lemma 5.11](#), we know that for some $i, j \in [\ell]$, A_{v_i}, A_{v_j} have no crossing edge. Thus $(\{u, v_i, v_j\}, \{A_u, A_{v_i}, A_{v_j}\})$ is a cloud_w -wedge, and so $\alpha_{\text{C}_w} < K$, as desired. \square

Thus, we want to explicitly construct C_4 -free graphs with small independence number. Fortunately, such explicit constructions exist: in particular, Alon showed in [\[Alo86\]](#) that well-known graphs [\[ER62\]](#) which can be explicitly constructed from finite projective planes have no C_4 and have independence number $\alpha \leq 2N^{3/4}$.

Now, note that a simple Markov type argument shows that any (N, K, n, k) -source of locality 1 must contain at least $K/2$ good sources, such that at most $w = 2N/K$ sources in \mathbf{X} depend on any one of these good sources. Thus, by fixing every other good source, we see that an (N, K, n, k) -source of locality 1 is a convex combination of $(N, K/2, n, k)$ -sources of locality 1 where at most $2N/K$ sources depend on a single good source. Thus, [Lemma 5.8](#) tells us that we can combine the above explicit C_4 -free graphs with [Lemma 5.6](#) to obtain the following theorem:

Theorem 5.12. *There exist constants $C, \gamma > 0$ such that for all $N, K, n, k \in \mathbb{N}$ satisfying $k \geq \log^C n$ and $K \geq CN^{19/20}$, there exists an explicit extractor $\text{Ext} : (\{0, 1\}^n)^N \rightarrow \{0, 1\}^m$ for (N, K, n, k) -sources of locality 1, with output length $m = k^{\Omega(1)}$ and error $\epsilon = 2^{-k^{\Omega(1)}}$, provided $N/K \leq k^\gamma$.*

In order to extend this result to higher locality, we show that we can reduce such adversarial sources to the 1-local setting, and therefore use the tools we have developed above to extract. Our reduction, below, is a simpler version of the reduction used by [\[Vio14\]](#) to reduce samplable sources to affine sources.

Lemma 5.13. *Let $\mathbf{X} = \mathbf{X}_1, \dots, \mathbf{X}_N$ be an (N, K, n, k) -source of locality d . Then \mathbf{X} is a convex combination of (N, K', n, k) -sources of locality 1, where $K' = K^2/(4Nd^2)$ and at most $w = 2Nd/K$ sources depend on a single good source (including itself).*

Proof. Identify \mathbf{X} with a bipartite graph $G = (V, E)$ over $K + N$ vertices as follows: partition V into L, R such that $|L| = K$ and $|R| = N$. Identify all N sources of \mathbf{X} with R , identify the good sources of \mathbf{X} with L , and draw an edge between $u \in L$ and $v \in R$ if $\mathbf{X}_u, \mathbf{X}_v$ are correlated. We will

find a subset $L' \subseteq L$ of size K' such that every vertex in R has at most one neighbor in L' , and every vertex in L' has at most $2Nd/K$ neighbors in R . Then, if we fix all sources in $\{\mathbf{X}_i\}_{i \in L \setminus L'}$, we are done: the locality is reduced to 1, K' good sources remain unfixed, and each unfixed good source has at most $2Nd/K$ bad sources depending on it.

First, note that because \mathbf{X} has locality d , every vertex in R has degree at most d , and thus $|E| \leq Nd$. Next, observe that there must be some subset $L^* \subseteq L$ of size at least $K/2$ where each $u \in L^*$ has degree at most $2Nd/K$; otherwise L contains at least $K/2$ vertices of degree $> 2Nd/K$, contradicting $|E| \leq Nd$. We will now show how to select L' from L^* .

For a vertex $v \in V$, we let $N(v)$ denote its neighbors (adjacent vertices), and for a set of vertices $S \subseteq V$, we let $N(S) = \bigcup_{v \in V} N(v)$. Now, we will greedily grow L' as follows: first, initialize $U \leftarrow L^*$. Then arbitrarily pick a vertex u from U and add it to L' , remove $N(N(u))$ from U , and repeat while U is not empty. Because each $v \in R$ has degree at most d , and each $v \in L^*$ has degree at most $2Nd/K$, the above process will produce a set L' of size at least $|L^*|/(2Nd^2/K) \geq K^2/(4Nd^2) = K'$. Furthermore, our selection process ensures that we never select two vertices from L that share a neighbor, so no vertex in R has more than one neighbor in L' . Lastly, because $L' \subseteq L^*$, each vertex in L' has at most $2Nd/K$ neighbors in R , so we are done. \square

It is now straightforward to use [Lemma 5.13](#) with [Theorem 5.12](#) to obtain extractors for d -local adversarial sources. However, if we avoid using [Theorem 5.12](#) as a black box, we can get slightly better parameters. In particular, by using the bound on w provided in [Lemma 5.13](#), using Alon's explicit C_4 -free graphs in [Lemma 5.6](#) immediately yields [Theorem 5.1](#).

Lastly, it is worth noting that new explicit constructions of the combinatorial objects introduced here would immediately imply better parameters for [Theorem 5.1](#). In particular, it would be interesting to explicitly construct C_4 -free graphs with smaller α , or, more generally, some other class of graphs with a structure that yields a stronger upper bound on α_{C_w} .

6 Extracting from many short sources

As discussed, the primary focus of our paper is negligible-error extraction from adversarial sources. In particular, given an (N, K, n, k) -source of locality d , we would like to extract $m = (Kk)^{\Omega(1)}$ bits with error $\epsilon = 2^{-(Kk)^{\Omega(1)}}$. In order to obtain such parameters m, ϵ that depend on *both* K, k , one might consider constructing extractors for the following two (slightly overlapping) regimes as separate tasks.

1. The regime $K \geq k^\gamma$, for an arbitrarily small constant $\gamma > 0$. In this regime, the adversarial source has most of its entropy distributed *across* many sources, instead of *within* a few sources.
2. The regime $k \geq K^\gamma$, for an arbitrarily small constant $\gamma > 0$. In this regime, the adversarial source has most of its entropy distributed *within* a few sources, instead of *across* many sources.

Roughly, the first regime corresponds to extracting from many small sources, while the latter regime corresponds to extracting from a few large sources. Notice that in the first regime we have $K = (Kk)^{\Omega(1)}$, and in the second regime we have $k = (Kk)^{\Omega(1)}$. Thus, if we want to construct explicit extractors that work for all (N, K, n, k) -sources, it makes sense to treat these two regimes separately. In particular, one might try constructing an extractor for the first regime that works with parameters $m = K^{\Omega(1)}, \epsilon = 2^{-K^{\Omega(1)}}$, and an extractor for the second regime that works with

parameters $m = k^{\Omega(1)}, \epsilon = 2^{-k^{\Omega(1)}}$. Together, these extractors can be used to output $m = (Kk)^{\Omega(1)}$ bits with error $\epsilon = 2^{-(Kk)^{\Omega(1)}}$ from any (N, K, n, k) -source.

Henceforth, when we discuss extracting from the first regime, we mean constructing extractors for adversarial sources that have output and error parameters m, ϵ that depend on K . When we discuss extracting from the second regime, we mean constructing extractors for adversarial sources that have output and error parameters m, ϵ that depend on k . It is worth noting that extractors constructed for either regime can work across all regimes, but their output and error are most impressive in the regime for which they are intended (i.e., because in such regimes the output and error can be written as $m = (Kk)^{\Omega(1)}, \epsilon = 2^{-(Kk)^{\Omega(1)}}$).

The main focus of our paper (outside this section) is to extract from the second regime $k \geq K^\gamma$, and thus produce extractors that have output and error parameters m, ϵ that depend on k . The purpose of the current section is to justify this focus, by showing a straightforward way to construct extractors for the first regime. In particular, the following is the main result of the section.

Theorem 6.1. *For all fixed $\gamma > 0$ and all $N, K, n, k, d \in \mathbb{N}$ satisfying $K/d \geq N^{2/3+\gamma} n^{1/3+\gamma}$, there exists an explicit extractor $\text{Ext} : (\{0, 1\}^n)^N \rightarrow \{0, 1\}^m$ for (N, K, n, k) -sources of locality d , with output length $m = K^{\Omega(1)}$ and error $\epsilon = 2^{-K^{\Omega(1)}}$.*

As discussed in the introduction, the work of Kamp et al. [KRVZ06] gives explicit low-error extractors for (N, K, n, k) -sources of locality 0 as long as $Kk = \omega(2^n \sqrt{nN})$. Theorem 6.1 greatly improves the dependence of n in this result, and furthermore works for polynomially high locality. To prove this result, we will show that constructing extractors for adversarial sources in the first regime simply reduces to constructing extractors for the following class of sources, which generalizes to a well-studied class of sources.

Definition 6.2. *A d -local non-oblivious bit-fixing (NOBF) source \mathbf{X} over $\{0, 1\}^n$ with min-entropy k has the following structure:*

1. *There exists a set $S \subseteq [n]$ of size k of good coordinates of \mathbf{X} , which are sampled uniformly and independently at random.*
2. *Each bit outside S is computed by a deterministic function of up to d bits inside S .*

We proceed by showing how to reduce d -local adversarial sources to d -local NOBF sources. Then, we show how d -local NOBF sources generalize to well-studied classes of sources, which will immediately give us Theorem 6.1.

Lemma 6.3. *Let $N, K, n, k, d \in \mathbb{N}$, and let $\mathbf{X} = \mathbf{X}_1, \dots, \mathbf{X}_N$ be an (N, K, n, k) -source of locality d . Then \mathbf{X} is a convex combination of d -local NOBF sources of length Nn and min-entropy K .*

Proof. By a standard convex combination argument, we may, without loss of generality, assume that each good source \mathbf{X}_i is a flat source (i.e., a uniform distribution over 2^k points in $\{0, 1\}^n$). Because $k \geq 1$, each good source \mathbf{X}_i can be written as a convex combination of flat sources of entropy 1 $\{\mathbf{Y}_{i,j}\}_{j \in \mathcal{J}}$, for some index set \mathcal{J} . Notice that each $\mathbf{Y}_{i,j}$ is simply a uniform distribution over two distinct strings $u, v \in \{0, 1\}^n$. Let $e \in [n]$ be the coordinate where u, v differ. Now $\mathbf{Y}_{i,j}$ can be seen as a 1-local NOBF source over $\{0, 1\}^n$ of entropy 1, where the single good bit is indexed by e (call this the “representative bit” of the good source). If we write every good source in this way, then in each component of the convex combination, any bit of a bad source that previously depended on

up to d good sources now depends on just the representative bits of those good sources (since every other bit in the good source is a deterministic function of the representative bit). Finally, we note that there are K representative bits, each of which are uniform, which completes the proof. \square

In a line of work initialized by Trevisan and Vadhan [TV00], Viola [Vio14] studied extraction from a class of sources that could be called *d-locally samplable sources*. A d -locally samplable source \mathbf{X} over $\{0, 1\}^n$ with min-entropy k has the following structure: for each coordinate $i \in [n]$, there exists a deterministic function $f_i : \{0, 1\}^k \rightarrow \{0, 1\}$ such that $\mathbf{X} = f_1(\mathbf{U}_k), \dots, f_n(\mathbf{U}_k)$, where each \mathbf{U}_k is the *same* copy of a random variable equal to the uniform distribution over $\{0, 1\}^k$. It is straightforward to show that a d -local NOBF source is a d -locally samplable source. Thus, by Lemma 6.3, any extractor for d -locally samplable sources over Nn bits that works at min-entropy K with output length $m = m(K)$ and error $\epsilon = \epsilon(K)$ immediately gives an extractor for (N, K, n, k) adversarial sources of locality d , with the same output and error parameters m, ϵ , *even if the min-entropy of each good sources is just $k = 1$* .

Thus, if one is interested in extracting from adversarial sources of the first regime, it makes sense to continue the current research program on constructing extractors for locally samplable sources (or, easier, d -local NOBF sources), instead of treating adversarial sources as a new class. In fact, by combining Lemma 5.13 (inspired by [Vio14]) with Lemma 6.3, we get the following lemma, which shows that extracting from adversarial sources in the first regime can be reduced to extracting from affine sources (with some loss in parameters).

Lemma 6.4. *Let $N, K, n, k, d \in \mathbb{N}$, and let $\mathbf{X} = \mathbf{X}_1, \dots, \mathbf{X}_N$ be an (N, K, n, k) -source of locality d . Then \mathbf{X} is a convex combination of 1-local NOBF sources of length Nn and min-entropy $K^2/(4Nd^2)$.*

A straightforward argument shows that a 1-local NOBF source is a special type of affine source [Vio14], and thus extractors for affine sources give extractors for adversarial sources in the first regime. We conclude by showing what sort of parameters are possible, given the best known low-error affine extractors (applied to 1-local NOBF sources).

1-local NOBF sources were introduced by [Vio14], under the name of *bit-block sources*. There, Viola says that a 1-local NOBF source \mathbf{X} has weight w if at most w bits in \mathbf{X} depend on the same good bit. He notes that a refinement of the best known low-error affine extractors gives the following extractors for 1-local NOBF sources:

Lemma 6.5 ([Rao09b, Vio14]). *There exists a universal constant $C > 0$ such that for all fixed $\gamma > 0$ and all $n, k \in \mathbb{N}$ such that $k \geq \log^C n$, there exists an explicit extractor $\text{Ext} : \{0, 1\}^n \rightarrow \{0, 1\}^m$ for 1-local NOBF sources of weight $w \leq k^{1-\gamma}$, with output length $m = k(1 - o(1))$ and error $\epsilon = 2^{-k^{\Omega(1)}}$.*

We note that Viola reduces locally samplable sources to 1-local NOBF sources, and thus provides Lemma 6.5 to construct extractors for locally samplable sources. As we have seen through our reductions, extractors for locally samplable sources and extractors for 1-local NOBF sources both provide extractors for adversarial sources in the first regime. However, it turns out that directly using Lemma 6.5 (instead of using Viola's extractors for locally samplable sources) will give us better parameters.³

³This is because adversarial sources look more like d -local NOBF sources instead of the more general d -locally samplable sources, as they offer the extra guarantee that there exist regions of high entropy. The reduction we provide from d -local adversarial sources to 1-local adversarial sources, Lemma 5.13, which mirrors Viola's reduction from locally samplable sources to 1-local NOBF sources, takes advantage of this.

In particular, we can apply [Lemma 6.5](#) to get extractors for adversarial sources as follows. First, we note that it is straightforward to modify the proof of [Lemma 6.3](#) so that the lemma statement additionally says: *furthermore, if at most w sources in \mathbf{X} depend on the same good source, then at most wn bits in each NOBF source of the convex combination depend on the same good bit.* Then, by combining this with [Lemma 5.13](#), we can obtain the following, more precise, statement of [Lemma 6.4](#):

Lemma 6.6. *Let $N, K, n, k, d \in \mathbb{N}$, and let $\mathbf{X} = \mathbf{X}_1, \dots, \mathbf{X}_N$ be an (N, K, n, k) -source of locality d . Then \mathbf{X} is a convex combination of 1-local NOBF sources of length Nn , min-entropy $K^2/(4Nd^2)$, and weight $2Ndn/K$.*

Combining [Lemma 6.6](#) with [Lemma 6.5](#), and removing redundant constraints, immediately gives us [Theorem 6.1](#).

Lastly, a few remarks are in order. First, we note that the requirement on K in [Theorem 6.1](#) can be slightly improved if extracting from (N, K, n, k) -sources of locality 0 or 1, since one can simply combine [Lemma 6.5](#) with [Lemma 6.3](#) instead of with [Lemma 6.4](#) or [Lemma 6.6](#). Second, we note the extractor in [Lemma 6.5](#) is an affine extractor, yet all that we need (just like in [[Vio14](#)]) is an extractor for 1-local NOBF sources, which have considerably more structure. This provides more motivation for the construction of low-error extractors for 1-local NOBF sources (a.k.a. bit-block sources). Third, we reiterate that improved extractors for locally samplable sources (perhaps using different techniques than reducing them to 1-local sources) would greatly improve the parameters in [Theorem 6.1](#).

7 Existential and impossibility results

Recall that given a 2-source extractor, we can extract from an $(N, 2, n, k)$ -source of locality 0 by simply applying the 2-source extractor on every pair our sources, and computing the XOR of the obtained outputs (see [Section 2](#)). A standard probabilistic argument shows that there exists a 2-source extractor for (n, k) -sources with output length m and error ϵ as long as $k \geq m + \log n + 2 \log(1/\epsilon) + 5$. Using these observations with [Lemma 3.9](#), it is straightforward to apply the same arguments used throughout this paper to obtain the following.

Theorem 7.1. *For all $N, n, k, m \in \mathbb{N}$ and $\epsilon > 0$ satisfying $k \geq 3m + \log n + 4 \log(1/\epsilon) + 12$, there exists an extractor $\text{Ext} : (\{0, 1\}^n)^N \rightarrow \{0, 1\}^m$ for $(N, 2, n, k)$ -sources of locality 0, with output length m and error ϵ .*

We now generalize the above existential result to (≥ 1) -locality.

Theorem 7.2. *For all $N, K, n, k, d, m, s \in \mathbb{N}$ and $\epsilon > 0$, there exists an extractor $\text{Ext} : (\{0, 1\}^n)^N \rightarrow \{0, 1\}^m$ for (N, K, n, k) -sources of locality d , with output length m and error ϵ , as long as $s > 1$ and the following hold:*

- $\binom{d}{s} < \frac{1}{N-K} \cdot \binom{K}{s}$,
- $k > g(n + \lceil \log N \rceil, s, \binom{N}{s}, m, \epsilon^2/4)$, where $g()$ is the function from [Theorem A.2](#).

Proof. Let $\mathbf{X} = \mathbf{X}_1 \circ \dots \circ \mathbf{X}_N$ be an (N, K, n, k) source of locality d , and define $\mathbf{X}' = \mathbf{X}'_1 \circ \dots \circ \mathbf{X}'_N$ to be such that $\mathbf{X}'_i = i \circ \mathbf{X}_i$, where we take i to be the binary encoding of i . Define $t = \binom{N}{s}$, and let

$\text{snmExt} : (\{0, 1\}^{n+\lceil \log N \rceil})^N \rightarrow \{0, 1\}^m$ be an optimal (non-explicit) generalized (s, t) -non-malleable extractor for entropy k , with output length m and error $\epsilon^2/4$. Indeed such a non-malleable extractor exists by [Theorem A.2](#). Now let $\mathcal{S} \subseteq [N]^s$ denote the set of all s -tuples of $[N]$ formed by taking a subset of $[N]$ of size s and writing it in increasing order; i.e., $|\mathcal{S}| = \binom{N}{s}$. Our extractor simply computes the following function:

$$\bigoplus_{(i_1, \dots, i_s) \in \mathcal{S}} \text{snmExt}(\mathbf{X}'_{i_1}, \dots, \mathbf{X}'_{i_s}).$$

Suppose now that there exists some $I = (i_1, \dots, i_s) \in \mathcal{S}$ such that $(\mathbf{X}_{i_1}, \dots, \mathbf{X}_{i_s})$ are all good sources and no bad source \mathbf{X}_j depends on all of them. Then after fixing all good sources outside I , we see that for all $J \neq I \in \mathcal{S}$, $(\mathbf{X}'_j)_{j \in J}$ is a tampered version of $(\mathbf{X}'_i)_{i \in I}$ with no fixed points, and such that each individual source in $(\mathbf{X}'_j)_{j \in J}$ depends on at most $s-1$ of $(\mathbf{X}'_i)_{i \in I}$. Thus by the non-malleability of snmExt and the setting of k , we see that with probability at least $1 - \sqrt{\epsilon^2/4}$ over fixing all $\text{snmExt}((\mathbf{X}'_j)_{j \in J})$, $J \neq I \in \mathcal{S}$, $\text{snmExt}((\mathbf{X}'_i)_{i \in I})$ is $\sqrt{\epsilon^2/4}$ -close to \mathbf{U}_m , the uniform distribution over m bits. In particular, this means that the output of our construction is $2\sqrt{\epsilon^2/4} = \epsilon$ -close to \mathbf{U}_m , as desired.

To complete the proof, we just need to argue that such an I exists. Consider the bipartite graph with K good sources on left and $N-K$ bad sources on right, such that a bad source is adjacent to every good source it depends on. Thus the right degree is bounded by d . Note that if $I \subseteq [N]$ indexes a set $\{\mathbf{X}_i\}_{i \in I}$ of s good sources, and \mathbf{X}_j is a bad source depending on all of $\{\mathbf{X}_i\}_{i \in I}$, then $(\{\mathbf{X}_i\}_{i \in I}, \mathbf{X}_j)$ creates an s -star. We let Δ denote the number of such s -stars in the graph (with the central vertex on the right). We will prove our result by providing upper and lower bounds on Δ .

Note that if no such I exists, then $\Delta \geq \binom{K}{s}$. However, because the right degree is bounded by d , we know that $\Delta \leq (N-K) \binom{d}{s}$, and thus $\binom{K}{s} \leq (N-K) \binom{d}{s}$, which contradicts our setting of $\binom{d}{s} < \frac{1}{N-K} \cdot \binom{K}{s}$. \square

This theorem immediately gives the following corollary.

Corollary 7.3 ([Theorem 4](#), restated). *For any constant $0 < \gamma < 1$ there exists a constant $\alpha > 0$ such that for all $N, K, n, k, d \in \mathbb{N}$ satisfying $k \geq (1 + \gamma) \log n$ and $K \geq N^\gamma$, and $d \leq K^{1-\gamma}$, there exists a (possibly non-explicit) extractor for (N, K, n, k) -sources of locality d with output length $m = k^{\Omega(1)}$ and error $\epsilon = 2^{-\Omega(k)}$, provided $N \leq k^\alpha$.*

Proof. In order to apply [Theorem 7.2](#), note that for any $0 < \gamma < 1$ with $K \geq N^\gamma$ and $d \leq K^{1-\gamma}$, there is a sufficiently large constant $s \in \mathbb{N}$ such that $\binom{d}{s} < \frac{1}{N-K} \cdot \binom{K}{s}$ holds. By resetting γ in the first step as necessary (to a smaller constant), the condition $k > g(n + \lceil \log N \rceil, s, \binom{N}{s}, m, \epsilon^2/4)$ holds by choice of α , the output length $m = k^{\Omega(1)}$, the error $\epsilon = 2^{-\Omega(k)}$, and the conditions $k \geq (1 + \gamma) \log n$ and $N \leq k^\alpha$. \square

We now prove two impossibility results on extraction from adversarial sources. First, we show that if just half of the sources are good, it is impossible to extract from adversarial sources of unbounded locality.

Theorem 7.4. *There does not exist an extractor $\text{Ext} : (\{0, 1\}^n)^N \rightarrow \{0, 1\}$ for $(N, N/2, n, n)$ -sources of locality $d = N/2$ with error $\epsilon < 1/2$.*

Proof. Fix an arbitrary $\text{Ext} : (\{0, 1\}^n)^N \rightarrow \{0, 1\}$. We consider two cases. If there is some $u \in (\{0, 1\}^n)^{N/2}$ such that for all $v \in (\{0, 1\}^n)^{N/2}$ we have $\text{Ext}(u, v) = 1$, then clearly $\text{Ext}(u, \mathbf{U}_{nN/2}) = 1$ with probability 1, and we have found an 0-local $(N, N/2, n, n)$ -source that makes the extractor constant. Otherwise, for each $u \in (\{0, 1\}^n)^{N/2}$, there is some $f(u) \in (\{0, 1\}^n)^{N/2}$ such that $\text{Ext}(u, f(u)) = 0$. In this case, $\text{Ext}(\mathbf{U}_{nN/2}, f(\mathbf{U}_{nN/2})) = 0$ with probability 1, and we have found an $(N/2)$ -local $(N, N/2, n, n)$ -source that makes the extractor constant. \square

Next, we show an impossibility result for when there is just one bad source, but it depends on all the good sources. The result rules out negligible-error extractors for adversarial sources in the extreme setting where the number of sources is exponentially larger than the length of each source. We first introduce a useful definition.

Definition 7.5. *Let Σ be some alphabet. For any function $f : \Sigma^n \rightarrow \{0, 1\}$, and any $i \in [n]$, define the influence of i on f , denoted by $I_i(f)$ to be the probability that on uniformly sampling the coordinates $[n] \setminus \{i\}$, the function f is still not determined.*

For any function $f : \Sigma^n \rightarrow \{0, 1\}$, we use $\mathbb{E}[f]$ to denote the average of f under the uniform distribution on the domain of f . We need the following generalization of the KKL theorem.

Theorem 7.6 ([BKK⁺92]). *Let $\Sigma = \{0, 1\}^m$. For any function $f : \Sigma^n \rightarrow \{0, 1\}$, there exists a coordinate $i \in [n]$ such that $I_i(f) \geq \Omega(p \log n/n)$, where $p = \min\{\mathbb{E}[f], 1 - \mathbb{E}[f]\}$.*

Theorem 7.7. *There does not exist an extractor $\text{Ext} : (\{0, 1\}^n)^N \rightarrow \{0, 1\}$ for $(N, N - 1, n, n)$ -sources of locality $d = N - 1$ with error $\epsilon \leq o((\log N)/(N2^n))$.*

Proof. The proof is straightforward using [Theorem 7.6](#). Let $\Sigma = \{0, 1\}^n$. Assume that there exists an extractor $f : \Sigma^N \rightarrow \{0, 1\}$ for the class of $(N, N - 1, n, n)$ -sources of locality $(N - 1)$ with error $\epsilon = o((\log N)/(N2^n))$. Without loss of generality assume $\mathbb{E}[f] \geq 1/2$. Using [Theorem 7.6](#), there exists $i \in [N]$ such that $I_i(f) = \Omega((\log N)/N)$. Now we can define a source $\mathbf{X} = \mathbf{X}_1, \dots, \mathbf{X}_N$ that is $(N, N - 1, n, n)$ -source of locality $N - 1$ in the following way: \mathbf{X}_i corresponds to the bad source and is set by looking at the values of $\{\mathbf{X}_j : j \in [N] \setminus \{i\}\}$ such that whenever f is not already determined by the good sources, \mathbf{X}_i is set so that $f(\mathbf{X}) = 1$. Using the lower bound on $I_i(f)$, it follows that $\mathbb{E}[f(\mathbf{X})] - 1/2 = \Omega((\log N)/(N2^n))$, contradicting our assumption. \square

8 Improved extractors for small-space and total entropy sources

Our results for adversarial sources directly imply improved extractors for sources that are sampled by small-space algorithms. This class of sources was first studied by Kamp et al. [\[KRVZ06\]](#), and fits into the line of work initiated by Trevisan and Vadhan [\[TV00\]](#) on constructing extractors for sources sampled by algorithms of bounded complexity.

Definition 8.1. *A source \mathbf{X} over $\{0, 1\}^n$ is called a space s source if it is sampled by a random walk on a width 2^s branching program of length n , where each edge of the branching program is labeled by a bit and an associated transition probability.*

Probabilistically, it is known that there are small space extractors for space s sources on $\{0, 1\}^n$ with min-entropy $k \geq O(s + \log s + \log(n/\epsilon))$ and error ϵ . The best known explicit extractor for the negligible error regime $\epsilon = 2^{-n^{\Omega(1)}}$ is from Kamp et al. [\[KRVZ06\]](#), who gave explicit extractors

for space s sources on $\{0, 1\}^n$ that require min-entropy $k \geq n^{1-\gamma}$ and space $s \leq \gamma \cdot (k/n)^3 \cdot n$, where γ is some tiny constant. Chattopadhyay and Li [CL16b] reduced the entropy requirement, but also significantly reduced the allowed space and increased the error to $\epsilon = n^{-\Omega(1)}$, which is no longer negligible.

Our contribution is a new extractor for the negligible error regime $\epsilon = 2^{-n^{\Omega(1)}}$. In particular, we construct an explicit extractor that can handle effectively the same space as the extractor from [KRVZ06], but significantly smaller entropy.

Theorem 8.2 (Theorem 5, restated). *For any fixed $\gamma > 0$ and all $n, k, s \in \mathbb{N}$ satisfying $k \geq n^{2/3+\gamma}$ and $s \leq (k/n)^{3+\gamma} \cdot n$, there exists an explicit extractor $\text{Ext} : \{0, 1\}^n \rightarrow \{0, 1\}^m$ for space s sources of min-entropy k , with output length $m = n^{\Omega(1)}$ and error $\epsilon = 2^{-n^{\Omega(1)}}$.*

Following [KRVZ06], we derive our results for small-space sources by first reducing to an intermediate model called total entropy sources that was first studied by Koenig and Maurer [KM05].

Definition 8.3. *A source X over $(\{0, 1\}^\ell)^r$ is called an (r, ℓ, k) -total entropy source if $\mathbf{X} = \mathbf{X}_1, \dots, \mathbf{X}_r$, where each \mathbf{X}_i is an independent random variable over $\{0, 1\}^\ell$, and $\sum_{i=1}^r H_\infty(\mathbf{X}_i) \geq k$.*

The best known extractor with negligible error for (r, ℓ, k) -total entropy sources (that doesn't restrict ℓ to be exponentially smaller than r) requires $k \geq (r\ell)^{1-\gamma}$, for a tiny constant γ [KRVZ06]. We show that our new constructions can extract from total entropy sources with significantly less entropy.

Theorem 8.4. *For any fixed $\gamma > 0$ and all $r, \ell, k \in \mathbb{N}$ satisfying $k \geq (r\ell)^{1-\alpha}$, where*

$$\alpha := \min \{ (1/3 - \gamma), (1/2 - \gamma) \log r / \log(r\ell) \},$$

there exists an explicit extractor $\text{Ext} : (\{0, 1\}^\ell)^r \rightarrow \{0, 1\}^m$ for (r, ℓ, k) -total entropy sources, with output length $m = (r\ell)^{\Omega(1)}$ and error $\epsilon = 2^{-(r\ell)^{\Omega(1)}}$.

Proof. Let $\mathbf{X} = \mathbf{X}_1, \dots, \mathbf{X}_r$ be an (r, ℓ, k) -total entropy source. We will consider two cases over r, ℓ , and show that in each case we may select certain $N, n \in \mathbb{N}$ such that \mathbf{X} can be viewed as an (N, n, k) -total entropy source. Given such a source, a standard Markov type argument says that if $k \geq N^{1/2+\gamma}n + n^\gamma N$, then \mathbf{X} is in fact a 0-local $(N, N^{1/2+\gamma}, n, n^\gamma)$ adversarial source. Thus if we selected N, n to ensure this entropy guarantee, and to ensure that $n^\gamma = (r\ell)^{\Omega(1)}$, then our extractor from Theorem 2 produces $m = (r\ell)^{\Omega(1)}$ bits from \mathbf{X} with error $\epsilon = 2^{-(r\ell)^{\Omega(1)}}$. We show how to select such N, n , below.

If $r \geq \ell^{(2-2\gamma)/(1-2\gamma)}$, we set $N = (r\ell)^{(2-2\gamma)/(3-4\gamma)}$, and $n = (r\ell)^{(1-2\gamma)/(3-4\gamma)}$. Notice that because $N \leq r$, we may bucket the sources $\mathbf{X}_1, \dots, \mathbf{X}_r$ into N consecutive buckets, each containing $r/N \geq 1$ independent sources. Thus, we may rewrite $\mathbf{X} = \mathbf{X}_1, \dots, \mathbf{X}_r$ as $\mathbf{X}_1, \dots, \mathbf{X}_N$, where each \mathbf{X}_i has length $r\ell/N = n$ and is independent of every other \mathbf{X}_j . And thus \mathbf{X} is also an (N, n, k) -total entropy source. Now, by our theorem statement, we know $k \geq (r\ell)^{2/3+\gamma}$ (by plugging in the first option for α). Thus, resetting γ to be a sufficiently small constant, we know that for sufficiently large r, ℓ (allowed by the asymptotic expression in the error), we have $k \geq N^{1/2+\gamma}n + n^\gamma N$. Furthermore, by the current setting of n , we clearly have $n^\gamma = (r\ell)^{\Omega(1)}$.

If $r < \ell^{(2-2\gamma)/(1-2\gamma)}$, we set $N = r$ and $n = \ell$, and thus \mathbf{X} is an (N, n, k) -total entropy source. By our theorem statement, we know $k \geq r^{1/2+\gamma}\ell$ (by plugging in the second option for α). Thus we have $k \geq N^{1/2+\gamma}n$ and $k \geq n^\gamma N$. Resetting γ to be a sufficiently small constant, we know that for sufficiently large r, ℓ , we have $k \geq N^{1/2+\gamma}n + n^\gamma N$. Furthermore, by the current setting of n and the upper bound on r imposed by this case, we have $n^\gamma = (r\ell)^{\Omega(1)}$, as desired. \square

We now recall a reduction from small-space sources to total entropy sources.

Lemma 8.5 ([KRVZ06]). *Let \mathbf{X} be a space s source on $\{0,1\}^n$ with min-entropy k . Then \mathbf{X} is $2^{-k/4}$ -close to a convex combination of $(r, \ell, k/2)$ -total entropy sources, where $r = k/(4s), \ell = 4sn/k$.*

It is now straightforward to combine [Lemma 8.5](#) with [Theorem 8.4](#) to prove [Theorem 8.2](#):

Proof of [Theorem 8.2](#). Set $\beta = \gamma/8$. By [Lemma 8.5](#) and [Theorem 8.4](#), we can extract $m = n^{\Omega(1)}$ bits from \mathbf{X} with error $\epsilon = 2^{-k/4} + 2^{-n^{\Omega(1)}} = 2^{-n^{\Omega(1)}}$ if $k/2 \geq n^{2/3+\beta}$ and $k/2 \geq nr^{\beta-1/2} = n(k/(4s))^{\beta-1/2}$. The former holds for sufficiently large n because we have $k \geq n^{2/3+\gamma}$. A straightforward calculation shows that the latter holds for sufficiently large n because we have $s \leq (k/n)^{3+\gamma}n$. \square

9 Future directions

In this work, we initiate a systemic study of *adversarial sources*, which generalize the well-studied setting of independent sources in extractor theory. We present explicit constructions for a wide range of parameters in this new setting, and give existential results that show there is still much room for improvement. For instance, it would be particularly interesting to extend our techniques to handle adversarial sources with the following parameters, in the negligible error regime: (1) 0-locality, and a sub-polynomial number of good sources, K , each with sub-polynomial entropy, k ; and (2) $K^{0.99}$ -locality, and an arbitrary polynomial number of good sources, K , each with polylogarithmic entropy, k . Explicit constructions for (1) would yield much improved extractors for small-space sources, and constructions for (2) would allow for extraction in a much more robust setting.

We introduce a new framework for extracting from multiple sources, based on new connections between extremal combinatorics and randomness extraction. In particular, all of our explicit constructions are built on extremal hypergraphs that exhibit a specific structure capable of controlling dependency between sources, and on *non-malleable extractors* which are capable of breaking these dependencies once they are nicely controlled. It would be interesting to see how much further these connections can be pushed, by constructing explicit hypergraphs that exhibit stronger extremal properties, or by constructing more powerful non-malleable extractors (which would allow the use of simpler hypergraphs). In particular, it's an interesting open problem to give explicit constructions of generalized s -source non-malleable extractors (as we define in [Definition A.1](#)).

References

- [Alo86] Noga Alon. Eigenvalues, geometric expanders, sorting in rounds, and ramsey theory. *Combinatorica*, 6(3):207–219, 1986.
- [AOR⁺19] Divesh Aggarwal, Maciej Obremski, Joao Ribeiro, Luisa Siniscalchi, and Ivan Visconti. How to extract useful randomness from unreliable sources. *Cryptology ePrint Archive: Report 2019/1156*, 2019.
- [BACD⁺18] Avraham Ben-Aroya, Eshan Chattopadhyay, Dean Doron, Xin Li, and Amnon Ta-Shma. A new approach for constructing low-error, two-source extractors. In

33rd Computational Complexity Conference (CCC 2018). Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.

- [BACDTS19] Avraham Ben-Aroya, Gil Cohen, Dean Doron, and Amnon Ta-Shma. Two-source condensers with low error and small entropy gap via entropy-resilient functions. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019.
- [BADTS17] Avraham Ben-Aroya, Dean Doron, and Amnon 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*, pages 1185–1194. ACM, 2017.
- [BDT18] Avraham Ben-Aroya, Dean Doron, and Amnon Ta-Shma. Near-optimal strong dispersers, erasure list-decodable codes and friends. *Electronic Colloquium on Computational Complexity (ECCC)*, 25:65, 2018.
- [BIW06] Boaz Barak, Russell Impagliazzo, and Avi Wigderson. Extracting randomness using few independent sources. *SIAM Journal on Computing*, 36(4):1095–1118, 2006.
- [BKK⁺92] Jean Bourgain, Jeff Kahn, Gil Kalai, Yitzhak Katznelson, and Nathan Linial. The influence of variables in product spaces. *Israel Journal of Mathematics*, 77(1-2):55–64, 1992.
- [BKS⁺05] Boaz Barak, Guy Kindler, Ronen Shaltiel, Benny Sudakov, and Avi Wigderson. Simulating independence: New constructions of condensers, ramsey graphs, dispersers, and extractors. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 1–10. ACM, 2005.
- [Bou05] J. Bourgain. More on the sum-product phenomenon in prime fields and its applications. *International Journal of Number Theory*, 01(01):1–32, 2005.
- [Bou07] Jean Bourgain. On the construction of affine extractors. *GAFSA Geometric And Functional Analysis*, 17(1):33–57, 2007.
- [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.
- [CG14] Mahdi Cheraghchi and Venkatesan Guruswami. Non-malleable coding against bit-wise and split-state tampering. In *Theory of Cryptography Conference*, pages 440–464. Springer, 2014.
- [CGH⁺85] Benny Chor, Oded Goldreich, Johan Hasted, Joel Freidmann, Steven Rudich, and Roman Smolensky. The bit extraction problem or t-resilient functions. In *26th Annual Symposium on Foundations of Computer Science (sfcs 1985)*, pages 396–407. IEEE, 1985.

- [CGL16] Eshan Chattopadhyay, Vipul Goyal, and Xin Li. Non-malleable extractors and codes, with their many tampered extensions. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 285–298. ACM, 2016.
- [CL16a] Eshan Chattopadhyay and Xin Li. Explicit non-malleable extractors, multi-source extractors, and almost optimal privacy amplification protocols. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 158–167. IEEE, 2016.
- [CL16b] Eshan Chattopadhyay and Xin Li. Extractors for sumset sources. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 299–311. ACM, 2016.
- [CLP17] Ernie Croot, Vsevolod F Lev, and Péter Pál Pach. Progression-free sets in are exponentially small. *Annals of Mathematics*, pages 331–337, 2017.
- [Coh16a] Gil Cohen. Local correlation breakers and applications to three-source extractors and mergers. *SIAM Journal on Computing*, 45(4):1297–1338, 2016.
- [Coh16b] Gil Cohen. Making the most of advice: New correlation breakers and their applications. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 188–196. IEEE, 2016.
- [Coh17] Gil Cohen. Towards optimal two-source extractors and ramsey graphs. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1157–1170. ACM, 2017.
- [CPS07] Ran Canetti, Rafael Pass, and Abhi Shelat. Cryptography from sunspots: How to use an imperfect reference string. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS’07)*, pages 249–259. IEEE, 2007.
- [CZ19] Eshan Chattopadhyay and David Zuckerman. Explicit two-source extractors and resilient functions. *Annals of Mathematics*, 189(3):653–705, 2019.
- [DGW09] Zeev Dvir, Ariel Gabizon, and Avi Wigderson. Extractors and rank extractors for polynomial sources. *Computational Complexity*, 18(1):1–58, 2009.
- [DKSS13] Zeev Dvir, Swastik Kopparty, Shubhangi Saraf, and Madhu Sudan. Extensions to the method of multiplicities, with applications to takeya sets and mergers. *SIAM Journal on Computing*, 42(6):2305–2328, 2013.
- [DOPS04] Yevgeniy Dodis, Shien Jin Ong, Manoj Prabhakaran, and Amit Sahai. On the (im) possibility of cryptography with imperfect randomness. In *45th Annual IEEE Symposium on Foundations of Computer Science*, pages 196–205. IEEE, 2004.
- [Dvi12] Zeev Dvir. Extractors for varieties. *Computational complexity*, 21(4):515–572, 2012.
- [Ede04] Yves Edel. Extensions of generalized product caps. *Designs, Codes and Cryptography*, 31(1):5–14, 2004.

- [EG17] Jordan S Ellenberg and Dion Gijswijt. On large subsets of with no three-term arithmetic progression. *Annals of Mathematics*, pages 339–343, 2017.
- [ER62] Paul Erdos and Alfréd Rényi. On a problem in the theory of graphs. *Publ. Math. Inst. Hungar. Acad. Sci.*, 7:215–235, 1962.
- [GGJS11] Sanjam Garg, Vipul Goyal, Abhishek Jain, and Amit Sahai. Bringing people of different beliefs together to do uc. In *Theory of Cryptography Conference*, pages 311–328. Springer, 2011.
- [GK08] Vipul Goyal and Jonathan Katz. Universally composable multi-party computation with an unreliable common reference string. In *Theory of Cryptography Conference*, pages 142–154. Springer, 2008.
- [GO14] Jens Groth and Rafail Ostrovsky. Cryptography in the multi-string model. *Journal of Cryptology*, 27(3):506–543, 2014.
- [GRS06] Ariel Gabizon, Ran Raz, and Ronen Shaltiel. Deterministic extractors for bit-fixing sources by obtaining an independent seed. *SIAM Journal on Computing*, 36(4):1072–1094, 2006.
- [GSZ19] Vipul Goyal, Akshayaram Srinivasan, and Chenzhi Zhu. Multi-source non-malleable extractors and applications. *manuscript*, 2019.
- [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.
- [KKL88] Jeff Kahn, Gil Kalai, and Nathan Linial. The influence of variables on boolean functions. In *[Proceedings 1988] 29th Annual Symposium on Foundations of Computer Science*, pages 68–80. IEEE, 1988.
- [KM05] Robert Koenig and Ueli Maurer. Generalized strong extractors and deterministic privacy amplification. In *IMA International Conference on Cryptography and Coding*, pages 322–339. Springer, 2005.
- [KRVZ06] Jesse Kamp, Anup Rao, Salil Vadhan, and David Zuckerman. Deterministic extractors for small-space sources. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 691–700. ACM, 2006.
- [KZ06] Jesse Kamp and David Zuckerman. Deterministic extractors for bit-fixing sources and exposure-resilient cryptography. *SIAM Journal on Computing*, 36(5):1231–1247, 2006.
- [Lew19] Mark Lewko. An explicit two-source extractor with min-entropy rate near $4/9$. *Mathematika*, 65(4):950–957, 2019.
- [Li11a] Xin Li. Improved constructions of three source extractors. In *2011 IEEE 26th Annual Conference on Computational Complexity*, pages 126–136. IEEE, 2011.

- [Li11b] Xin Li. A new approach to affine extractors and dispersers. In *Proceedings of the 26th Annual IEEE Conference on Computational Complexity*, pages 137–147, 2011.
- [Li13a] Xin Li. Extractors for a constant number of independent sources with polylogarithmic min-entropy. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, 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.
- [Li15a] Xin Li. Non-malleable condensers for arbitrary min-entropy, and almost optimal protocols for privacy amplification. In *Theory of Cryptography Conference*, pages 502–531. Springer, 2015.
- [Li15b] Xin Li. Three-source extractors for polylogarithmic min-entropy. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 863–882. IEEE, 2015.
- [Li15c] Xin Li. Three-source extractors for polylogarithmic min-entropy. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 863–882. IEEE, 2015.
- [Li16] Xin Li. Improved two-source extractors, and affine extractors for polylogarithmic entropy. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 168–177. IEEE, 2016.
- [Li17] Xin Li. Improved non-malleable extractors, non-malleable codes and independent source extractors. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1144–1156. ACM, 2017.
- [Li19] Xin Li. Non-malleable extractors and non-malleable codes: Partially optimal constructions. In *34th Computational Complexity Conference, CCC 2019, July 18-20, 2019, New Brunswick, NJ, USA.*, pages 28:1–28:49, 2019.
- [LPV09] Huijia Lin, Rafael Pass, and Muthuramakrishnan Venkitasubramaniam. A unified framework for concurrent security: universal composability from stand-alone non-malleability. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 179–188. ACM, 2009.
- [LRVW03] Chi-Jen Lu, Omer Reingold, Salil Vadhan, and Avi Wigderson. Extractors: Optimal up to constant factors. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 602–611. ACM, 2003.
- [Mek17] Raghu Meka. Explicit resilient functions matching ajtai-linial. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1132–1148. SIAM, 2017.
- [MW97] Ueli Maurer and Stefan Wolf. Privacy amplification secure against active adversaries. In *Annual International Cryptology Conference*, pages 307–321. Springer, 1997.

- [Rao09a] Anup Rao. Extractors for a constant number of polynomially small min-entropy independent sources. *SIAM Journal on Computing*, 39(1):168–194, 2009.
- [Rao09b] Anup Rao. Extractors for low-weight affine sources. In *2009 24th Annual IEEE Conference on Computational Complexity*, pages 95–101. IEEE, 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.
- [RŠ94] Vojtěch Rödl and Edita Šinajová. Note on independent sets in steiner systems. *Random Structures & Algorithms*, 5(1):183–190, 1994.
- [TV00] Luca Trevisan and Salil Vadhan. Extracting randomness from samplable distributions. In *Proceedings 41st Annual Symposium on Foundations of Computer Science*, pages 32–42. IEEE, 2000.
- [Vaz85] Umesh V Vazirani. Towards a strong communication complexity theory or generating quasi-random sequences from two communicating slightly-random sources. In *Proceedings of the seventeenth annual ACM symposium on Theory of computing*, pages 366–378. ACM, 1985.
- [Vio14] Emanuele Viola. Extractors for circuit sources. *SIAM Journal on Computing*, 43(2):655–672, 2014.
- [Yeh11] Amir Yehudayoff. Affine extractors over prime fields. *Combinatorica*, 31(2):245–256, 2011.

A Existential bounds on generalized s -source non-malleable extractors

We define a generalized s -source non-malleable extractor with tampering degree t as follows.

Definition A.1. *We call a function $\text{snmExt} : (\{0,1\}^n)^s \rightarrow \{0,1\}^m$ a generalized (s,t) -non-malleable extractor for entropy k , output length m , and error ϵ , if the following holds. Let $\mathbf{X}_1, \dots, \mathbf{X}_s$ be any s independent (n,k) -sources, and let $h_i, i \in [t] : \{0,1\}^{ns} \rightarrow \{0,1\}^{ns}$ be t tampering functions of the form $h_i = (f_i^1, \dots, f_i^s)$, where each $f_i^j : \{0,1\}^{ns} \rightarrow \{0,1\}^n$ is a tampering function that depends on at most $s - 1$ of the sources. Suppose that each h_i has no fixed point. Then:*

$$|\text{snmExt}(\mathbf{X}_1, \dots, \mathbf{X}_s) \circ \text{snmExt}(h_1(\mathbf{X}_1, \dots, \mathbf{X}_s)) \circ \dots \circ \text{snmExt}(h_t(\mathbf{X}_1, \dots, \mathbf{X}_s)) - \mathbf{U}_m \circ \text{snmExt}(h_1(\mathbf{X}_1, \dots, \mathbf{X}_s)) \circ \dots \circ \text{snmExt}(h_t(\mathbf{X}_1, \dots, \mathbf{X}_s))| \leq \epsilon.$$

We say that a generalized (s,t) -non-malleable extractor has tampering degree t .

Cheraghchi and Guruswami [CG14] proved existential bounds on 2-source non-malleable extractors with tampering degree 1. We extend this result to prove existential bounds for generalized s -source non-malleable extractors with tampering degree t .

Theorem A.2. For all $n, k, s, t, m \in \mathbb{N}$ and $\epsilon > 0$ satisfying $s > 1$ and $k > g(n, s, t, m, \epsilon)$, there exists a generalized (s, t) -non-malleable extractor $\text{snmExt} : (\{0, 1\}^n)^s \rightarrow \{0, 1\}^m$ for entropy k , output length m , and error ϵ , where

$$g(n, s, t, m, \epsilon) = \frac{m(t+1)}{s} + \log(n) + 2\log(1/\epsilon) + 2\log(t(t+1)) + \log(s) + 3.$$

To prove the above result, we record a result that follows from the proof of Theorem A.1 in [BACD⁺18].

Lemma A.3 ([BACD⁺18]). Let \mathbf{X} be a flat (n', k') -source, let $h_i, i \in [t] : \{0, 1\}^{n'} \rightarrow \{0, 1\}^{n'}$ be arbitrary tampering functions with no fixed points, and let $\mathcal{D} : (\{0, 1\}^m)^{t+1} \rightarrow \{0, 1\}$ be an arbitrary distinguisher function. Let $\text{snmExt} : \{0, 1\}^{n'} \rightarrow \{0, 1\}^m$ be a function sampled uniformly at random, and define the distributions $\mathbf{D}_1 := (\text{snmExt}(\mathbf{X}), \text{snmExt}(h_1(\mathbf{X})), \dots, \text{snmExt}(h_t(\mathbf{X})))$ and $\mathbf{D}_2 := (\mathbf{U}_m, \text{snmExt}(h_1(\mathbf{X})), \dots, \text{snmExt}(h_t(\mathbf{X})))$. Then with probability at least $1 - \beta(t, k', \epsilon)$ over sampling snmExt , we have $|\Pr[\mathcal{D}(\mathbf{D}_1) = 1] - \Pr[\mathcal{D}(\mathbf{D}_2) = 1]| \leq \epsilon$, where

$$\beta(t, k', \epsilon) = (t+1) \exp\left(\frac{-\epsilon^2 2^{k'}}{2(t+1)^2}\right).$$

Given the above lemma, Theorem A.2 is straightforward to derive via the probabilistic method:

Proof of Theorem A.2. We note the following bounds.

- The number of sources of the form $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_s)$, where each \mathbf{X}_i is a flat independent (n, k) -source, is at most $\binom{2^n}{2^k}^s \leq 2^{sn2^k}$.
- The number of distinguishers $\mathcal{D} : (\{0, 1\}^m)^{t+1} \rightarrow \{0, 1\}$ is at most $2^{2^{m(t+1)}}$.
- For any given source of the form $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_s)$, where each \mathbf{X}_i is a flat independent (n, k) -source, the number of distinct tampering functions $f : \{0, 1\}^{sn} \rightarrow \{0, 1\}^n$, restricted to the support of \mathbf{X} and depending on at most $s-1$ of the sources in \mathbf{X} , is at most $s \cdot (2^n)^{2^{k(s-1)}} = 2^{n2^{k(s-1)} + \log s}$. Thus the number of distinct t -tuples of adversaries (h_1, \dots, h_t) , with each adversary of the form $h_i = (f_i^1, \dots, f_i^s)$, where each f_i^j is a tampering function as above, is at most $2^{ts(n2^{k(s-1)} + \log s)}$.

Thus, by applying Lemma A.3 with $k' = sk$, the probabilistic method and a standard union bound argument tells us that the claimed generalized (s, t) -non-malleable extractor exists if

$$\beta(t, k', \epsilon) \cdot 2^{sn2^k} \cdot 2^{2^{(t+1)m}} \cdot 2^{ts(n2^{k(s-1)} + \log s)} < 1,$$

which holds for our selection of k . □