THEORY OF COMPUTING, Volume 21 (8), 2025, pp. 1–38
www.theoryofcomputing.org

Seed-Protecting Extractors

Gil Cohen[*] Dean Doron[†] Shahar Samocha[‡]

Received November 23, 2021; Revised December 18, 2022; Published October 28, 2025

Abstract. We introduce a new type of seeded extractors we dub *seed-protecting* extractors. Informally, a seeded extractor is seed protecting against a class C of functions, mappings seeds to seeds, if the seed Y remains close to uniform even after observing the output $\mathsf{Ext}(X, A(Y))$ for every choice of $A \in C$ (or, more generally, observing the outputs corresponding to several adversaries from C).

The results of this paper are structural. We establish what we believe to be surprising relations, in fact, *equivalences* between seed-protecting extractors and each of the well-studied strengthenings of seeded extractors: strong extractors, non-malleable extractors (albeit only against permutations), and two-source extractors, where each case is classified by a suitable class C.

Our work motivates the study of non-malleable extractors against permutations and puts forth a novel approach for their construction. Indeed, the existing machinery developed for constructing non-malleable extractors focuses on the *output* and so it is aimed towards breaking correlations. Instead, our work suggests developing techniques for protecting the *seed*.

[*] The research leading to these results has received funding from the Israel Science Foundation (grant number 1569/18) and from the Azrieli Faculty Fellowship

[†] The work was done while at Stanford University, supported by Supported by NSF award CCF-1763311.

[‡] Supported by Blavatnik Research Fund.

ACM Classification: G.2.1, G.3, F.0

AMS Classification: 68Q87, 68R05

Key words and phrases: pseudorandomness, extractors

DOI: 10.4086/toc.2025.v021a008

1 Introduction

Informally, a *seeded extractor* is a function that "purifies" defective randomness using few fresh random bits. A defective random source is modelled by a distribution X that has some lower bound on its *min-entropy*. A random variable X is said to have min-entropy k, denoted $H_\infty(X) \geq k$, if for every x, $\Pr[X = x] \leq 2^{-k}$. When X is supported over n-bit strings, we call X an (n, k)-source. A function $\mathsf{Ext}\colon \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ is a (k, ε)*-seeded extractor* [26] if for every (n, k)-source X it holds that $\mathsf{Ext}(X, Y)$ is ε-close, in statistical distance, to the uniform distribution over m-bit string. Here, Y is a random variable, independent of X, that is uniformly distributed over d-bit strings. We write this as $\mathsf{Ext}(X, Y) \approx_\varepsilon U_m$. Informally, using the "fresh" randomness in the, hopefully short, string Y, the function Ext extracts the randomness from X to a nearly perfect form, namely, to a distribution that is close to uniform. We refer to Y as the *seed* of the extractor.

The notion of seeded extractors can be strengthened in different ways. Three such strengthenings that emerged from the study of seeded extractors are strong seeded extractors, non-malleable extractors [18], and two-source extractors [9]. The latter is the oldest notion, in fact, two-source extractors predate the explicit definition of seeded extractors. Nevertheless, such extractors proved to be the most challenging to construct. In a span of about a decade, strong seeded extractors with nearly optimal parameters were constructed using sophisticated algebraic and combinatorial ideas (see, e. g., [31, 25, 21, 19, 30] as well as [29] and [32, Chapter 6]). Non-malleable extractors were introduced more recently, and despite their syntactic resemblance to strong seeded extractors (see Section 1.1 below for the formal definitions), their constructions required completely different techniques. Furthermore, it was the insight regarding the connection between non-malleable extractors and the seemingly unrelated two-source extractors that enabled the breakthrough work of Chattopadhyay and Zuckerman [8] who constructed two-source extractors for polylogarithmic min-entropy.

A brief and informal summary of our contribution

In this article we introduce a new, very natural, variant of seeded extractors we dub *seed-protecting extractors*. Informally, a seeded extractor is seed protecting against a class $\mathcal{C}$ of functions, mapping seeds to seeds, if the seed Y remains close to uniform even after observing the output $\mathsf{Ext}(X, A(Y))$ for every choice of $A \in \mathcal{C}$ (and, of course, a source X with sufficient min-entropy). See Definition 1.1 below for the formal definition. We establish what we believe to be surprising and insightful relations, in fact, *equivalences* between seed-protecting extractors and strong extractors, non-malleable extractors against permutations, and two-source extractors, where each case is classified by a suitable class $\mathcal{C}$.

This fresh point of view on non-malleable extractors suggests, in particular, a novel approach for constructing such extractors as, indeed, the focus shifts from breaking output correlations to protecting the seed. We first recall the definitions of strong and non-malleable extractors (Section 1.1). Then, in Section 1.2, we give the definition of seed-protecting extractors and present out results.

1.1 Strong seeded extractors and non-malleable extractors

1.1.1 Strong seeded extractors

A (k, ε)-seeded extractor Ext is called *strong* if the output distribution $\mathsf{Ext}(X, Y)$ is close to uniform even given the seed Y used for the extraction. This can be expressed by writing $(\mathsf{Ext}(X, Y), Y) \approx_\varepsilon (U_m, Y)$. Seeded extractors and, more so, their strong counterparts have found many applications. As mentioned, in a beautiful and deep line of work, efficiently computable strong seeded extractors were constructed for any min-entropy k, having seed-length $O(\log \frac{n}{\varepsilon})$ (see [21, 19, 30] and references therein). Furthermore, connections between strong seeded extractors and other objects of study such as list decodable codes, samplers, and expander graphs were found and enabled many applications.

1.1.2 Non-malleable extractors

A *non-malleable extractor* is a strong seeded extractor that has the following additional property. The output of the extractor remains close to uniform even after observing the output of the extractor on any altered seed. Formally, let $A\colon \{0,1\}^d \to \{0,1\}^d$ be an arbitrary function with no fixed points, that is, $A(y) \neq y$ for all y. The reader should think of A as an adversarially chosen way of altering the seed. A function $\mathsf{Ext}\colon \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ is a (k, ε)-*non-malleable extractor* if for every (n, k)-source X and A as above,

$$(\mathsf{Ext}(X, Y), \mathsf{Ext}(X, A(Y)), Y) \approx_\varepsilon (U_m, \mathsf{Ext}(X, A(Y)), Y)\,, \tag{1.1}$$

where again Y is uniform over $\{0,1\}^d$ and is independent of X.

Non-malleable extractors were introduced by Dodis and Wichs [18]. The original motivation for studying such extractors was for the classic problem of devising privacy amplification protocols against active adversaries. Indeed, strong seeded extractors yield a solution to the passive adversary variant. As we discuss later on, non-malleable extractors proved key for the construction of good two-source extractors. More precisely, one requires a certain generalization obtained by considering more than one adversarial function [15]. Let $t \geq 1$ be an integer. The function Ext above is called a (k, ε) t-non-malleable extractor if for every t-tuple of functions $A_1, \ldots, A_t\colon \{0,1\}^d \to \{0,1\}^d$ with no fixed points, it holds that

$$(\mathsf{Ext}(X, Y), \{\mathsf{Ext}(X, A_i(Y))\}_{i=1}^t, Y) \approx_\varepsilon (U_m, \{\mathsf{Ext}(X, A_i(Y))\}_{i=1}^t, Y)\,.$$

In a large body of work, non-malleable extractors were constructed (see [17, 15, 22, 10, 6, 12, 16, 7, 11, 23] and references therein). The state-of-the-art construction of (k, ε)-non-malleable extractors [24] has seed length $d = O(\log n) + O(\log \frac{1}{\varepsilon}) \cdot 2^{O(a \cdot (\log\log \frac{1}{\varepsilon})^{1/a})}$ for min-entropy as low as $k = O(\log\log n + a \log \frac{1}{\varepsilon})$ for every choice of $a \geq 2$. All of these constructions generalize to t-non-malleable extractors. Alternatively, a black-box reduction from t-non-malleable extractors to non-malleable extractors [12] can be invoked to give explicit t-non-malleable extractors with seed length $\mathrm{poly}(t) \cdot d$.

1.1.3 Non-malleable extractors against permutations

In this article we initiate the study of non-malleable extractors against permutations. These are functions that are ought to satisfy Equation (1.1) only for A a permutation with no fixed points. While existing applications of non-malleable extractors (for the construction of two-source extractors and for the design of privacy amplification protocols) consider a general adversarial function A with no fixed points, we believe that the new notion of non-malleable extractors against permutations is natural and interesting both in its own right as well for as a step towards constructing full-fledged non-malleable extractors. It is interesting to note that non-malleable *two-source* extractors against permutations were found useful in independent work by Goldreich and Wigderson [20], where they were studied in connection to robustly self-ordered graphs.

1.2 Strong, non-malleable, and seed-protecting extractors

As mentioned, in this paper we introduce a new type of randomness extractors which we call *seed-protecting extractors*. To give the formal definition, for an integer d, let $\mathcal{A}_d$ be the set of all functions from $\{0,1\}^d$ to $\{0,1\}^d$. When d is clear from context, we simply write $\mathcal{A}$.

Definition 1.1 (seed-protecting extractors)**.** Let $\mathsf{Ext}\colon \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ be a (k,ε)-seeded extractor. Let $C \subseteq \mathcal{A}_d$. We say that Ext is *seed protecting against* C if for every $A \in C$ and every (n,k)-source X it holds that

$$(Y, \mathsf{Ext}(X, A(Y))) \approx_{\varepsilon} (U_d, \mathsf{Ext}(X, A(Y))) \ , \tag{1.2}$$

where Y is uniformly distributed and is independent of X.

Both non-malleable extractors and seed-protecting extractors have a certain resilience property against tampering with the seed. While non-malleable extractors focus on the *output*, seed-protecting extractors are concerned about the *seed*. This shift of focus induces inherent differences. Indeed, while fixed points trivially rule out the possibility of non-malleability (hence, functions with fixed points are excluded by definition), fixed points turn out to be a non-issue for seed-protecting extractors. Indeed, consider the extreme case – the identity function. Clearly, non-malleability cannot be achieved against this function. However, note that to be seed protecting against the identity function precisely means to be a strong seeded extractor. In fact, the first observation we make in this preliminary discussion is that strong extractors are equivalent to seed-protecting extractors against a class of permutations which we denote by $\Pi \subseteq \mathcal{A}_d$.

Claim 1.2 (strong $\iff$ seed protecting against Π)**.** *Let* $\mathsf{Ext}\colon \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ *be a* (k,ε)*-seeded extractor. Then,*

1. *If* Ext *is a* (k,ε)*-seed-protecting extractor against* Π *then* Ext *is a* $(k,2\varepsilon)$*-strong seeded extractor.*

2. *If* Ext *is a* (k,ε)*-strong seeded extractor then* Ext *is a* $(k,2\varepsilon)$*-seed-protecting extractor against* Π*.*

Proof. For the first item, a seed-protecting extractor against Π can be seen to be a strong seeded extractor by taking A to be the identity function. Indeed, with this choice, since Ext is a seeded extractor the right hand side of Equation (1.2) is ε-close to U_{d+m}, and so $(Y, \mathsf{Ext}(X,Y)) \approx_{2\varepsilon} (Y, U_m)$. For the other direction, as Ext is (k,ε)-strong we have that $(Y, \mathsf{Ext}(X,Y)) \approx_{\varepsilon} (Y, U_m)$. Thus, for any $A \in \Pi_d$, it holds that

$$(A(Y), \mathsf{Ext}(X, A(Y))) \approx_{\varepsilon} (A(Y), U_m)\,.$$

By the data processing inequality, one can apply any function to the first component $A(Y)$ of both sides and maintain the ε-closeness. In particular, by applying A^{-1} we get

$$(Y, \mathsf{Ext}(X, A(Y))) \approx_{\varepsilon} (Y, U_m)\,. \tag{1.3}$$

Now, (Y, U_m) has the same distribution as U_{d+m}, and $\mathsf{Ext}(X, A(Y)) \approx_{\varepsilon} U_m$. This, together with Equation (1.3), implies that $(Y, \mathsf{Ext}(X, A(Y))) \approx_{2\varepsilon} (U_d, \mathsf{Ext}(X, A(Y)))$, completing the proof. □

Going back to non-malleable extractors, by the discussion above, it is not a priori clear whether non-malleability is in any way related to seed protection. Nonetheless, one of the results of this paper is an *equivalence* between the property of non-malleability and seed protection, at least when focusing on permutation adversaries. By saying that Ext is a non-malleable extractor against $C \subseteq \mathcal{A}$, we mean that Equation (1.1) holds for every $A \in C$ (but not necessarily for other functions) that, in addition, has no fixed points.

For stating our result, we generalize seed protection to several adversarial functions. This should be done with some care. Indeed, naively, one's first suggestion might require that for every two functions $A_1, A_2 \in C$, it holds that

$$(Y, \mathsf{Ext}(X, A_1(Y)), \mathsf{Ext}(X, A_2(Y))) \approx_{\varepsilon} (U_d, \mathsf{Ext}(X, A_1(Y)), \mathsf{Ext}(X, A_2(Y)))\ .$$

This definition, we observe, is moot. Indeed, consider two functions $A_1, A_2 \colon \{0,1\}^d \to \{0,1\}^d$ that according to the first bit of the seed, Y_1, decide whether to "behave" exactly the same or very differently. More concretely, sample two permutations π_0, π_1 on $\{0,1\}^d$ at random, and define $A_1(y) = \pi_0(y)$ and $A_2(y) = \pi_{y_1}(y)$. As π_0, π_1 were chosen at random, and thus are usually disagree, by observing $\mathsf{Ext}(X, A_1(Y))$ and $\mathsf{Ext}(X, A_2(Y))$ one can distinguish Y, in fact its first bit Y_1, from uniform by checking whether both outputs are equal.

As we prove, this "collusion," in which two or more adversarial functions attain the same value, is the only obstacle for seed protection against permutations. (Already here we stress that there are other obstacles when considering functions other than permutations, as we discuss in Section 1.3.) Given $A_1, \ldots, A_t \colon \{0,1\}^d \to \{0,1\}^d$, we say $A_1, \ldots, A_t$ are *non-colluding* if for every $y \in \{0,1\}^d$, all the evaluations $A_1(y), \ldots, A_t(y)$ are distinct. We denote by $\mathcal{X}^t \subseteq \mathcal{A}_d^t$ the set of t-tuples of functions that are non-colluding. With hindsight, we give the following generalization of seed-protecting extractors to several adversarial functions.

Definition 1.3. Let $\mathsf{Ext} \colon \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ be a (k,ε)-seeded extractor. Let $C \subseteq \mathcal{A}$. We say that Ext is *t-seed protecting against* C if for every (n,k)-source X and $(A_1, \ldots, A_t) \in \mathcal{X}^t \cap C^t$ it holds that

$$\left(Y, \{\mathsf{Ext}(X, A_i(Y))\}_{i=1}^t\right) \approx_{\varepsilon} \left(U_d, \{\mathsf{Ext}(X, A_i(Y))\}_{i=1}^t\right)\ .$$

We also express this by saying that Ext is seed protecting against C^t.

The following lemma gives the equivalence's easier direction, showing that non-malleable extractors against permutations (with no fixed points) are seed protecting against (non-colluding) permutations. We refer the reader to Lemma 4.4 for a more general statement.

Lemma 1.4 (non-malleable $\Longrightarrow$ seed protecting)**.** *Let $t \geq 1$ and assume* Ext *is a (k, ε)-non-malleable extractor against Π^t. Then,* Ext *is a $(k, 4t\varepsilon)$-seed-protecting extractor against Π^{t+1}.*

As a warm-up, in Section 2, we prove Lemma 1.4 for $t = 1$ as well as its converse which is indeed more surprising and difficult to prove (see Theorem 2.1).

Towards stating the other direction, for $\Delta \geq 0$, define $\mathcal{F}_\Delta$ to be the subset of $\mathcal{A}_d$ containing all functions $A\colon \{0,1\}^d \to \{0,1\}^d$ such that $H_\infty(A(U_d)) \geq d - \Delta$. For example, note that $\mathcal{F}_0 = \Pi$. We give a reduction from non-malleable extractors against $\mathcal{F}_\Delta$ to seed-protecting extractors against $\mathcal{F}_\Delta$.

Theorem 1.5 (seed protecting $\Longrightarrow$ non-malleable)**.** *Let $t \geq 1$ be an integer, and $\Delta \geq \max(1, \log t)$. Let* Ext: $\{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ *be a (k, ε)-seed-protecting extractor against $\mathcal{F}_\Delta^{t+1} \cap \mathcal{X}^{t+1}$. Assume further that $d = \Omega(\log t)$. Then,* Ext *is (k', ε')-non-malleable against $\mathcal{F}_\Delta^t$ with $k' = k + mt + \log \frac{1}{\varepsilon}$ and $\varepsilon' = O(\varepsilon^{1/3})$.*

We prove Theoren 1.5 in Section 3. We note that for $\Delta = 0$ and $t = 1$, seed protection against permutations is enough, and one does not need to devise an extractor against $\mathcal{F}_1$. Thus, for $\Delta = 0$ and $t = 1$ there is a strong equivalence, which we prove as a warm-up, in Theorem 2.1.

1.3 Two-source extractors as seed-protecting extractors

Discussing seed protection against permutations sufficed for characterizing both strong and non-malleable extractors against permutations. But, what about other adversarial functions? Is it the case that seed protection is achievable against any single function? (Of course, collusion is irrelevant in such a setting.) The quick answer is "no." Surprisingly, our next result is a characterization we obtain for two-source extractors as 1-seed-protecting extractors i. e., seed-protecting extractors with a single adversarial function) against a suitable family. In particular, known impossibility results on two-source extractors translate to impossibility results on 1-seed-protecting extractors.

Before recalling the formal definition of two-source extractors and describing this family, we believe it is instructive to first consider an extreme case and ask whether one can seed-protect against an adversarial function A that, unlike a permutation, is allowed to "focus" on seeds of its choice. The ultimate case is where A has range of size one. However, in such case, $A(y)$ gives no information about the seed y, and so seed protection trivially follows. What about a range of size two? We have the following easy claim that establishes the impossibility of seed protection against such functions. Let $\mathcal{T} \subseteq \mathcal{A}_d$ be the set of all functions $A\colon \{0,1\}^d \to \{0,1\}^d$ with range of size precisely two.

Claim 1.6. *Let* Ext: $\{0,1\}^n \times \{0,1\}^d \to \{0,1\}$ *be a (k, ε)-seeded extractor. Then, for $k \leq n - 1$ and $\varepsilon < \frac{1}{6}$,* Ext *is not (k, ε)-seed protecting against $\mathcal{T}$.*

Proof. Fix an arbitrary $w \in \{0,1\}^d$ and assume without loss of generality that $\Pr[\mathsf{Ext}(U_n, w) = 0] \geq \frac{1}{2}$. Define X to be the random variable that is uniformly distributed over all $x \in \{0,1\}^n$ such that $\mathsf{Ext}(x, w) = 0$. Since Ext is a (k, ε)-seeded extractor for min-entropy k and as $H_\infty(X) \geq n - 1 \geq k$, we have that $\mathsf{Ext}(X, Y) \approx_\varepsilon U_1$. Thus,

$$\Pr[\mathsf{Ext}(X, Y) = 1] \geq \frac{1}{2} - \varepsilon\,.$$

By an averaging argument, there exists $z \in \{0,1\}^d$ such that

$$\Pr[\mathsf{Ext}(X, z) = 1] \geq \frac{1}{2} - \varepsilon\,.$$

Note that $w \neq z$. Define the function $A : \{0,1\}^d \longrightarrow \{0,1\}^d$ by

$$A(y) = \begin{cases} w & y_1 = 0, \\ z & y_1 = 1, \end{cases}$$

and note that $A \in \mathcal{T}$. Denote $Z = \mathsf{Ext}(X, Y)$ and $Z' = \mathsf{Ext}(X, A(Y))$. We turn to show that

$$\gamma = \mathsf{SD}\left((Y, Z'), (U_d, Z')\right) \geq \frac{1}{4} - \frac{\varepsilon}{2}\,.$$

To see this, note that

$$\gamma \geq \mathsf{SD}\left((Y_1, Z'), (U_1, Z')\right) \geq \Pr[Z' = Y_1] - \Pr[Z' = U_1]\,. \tag{1.4}$$

We have that

$$\begin{aligned} \Pr[Z' = Y_1] &= \frac{1}{2}\Pr[Z' = Y_1 \mid Y_1 = 0] + \frac{1}{2}\Pr[Z' = Y_1 \mid Y_1 = 1] \\ &= \frac{1}{2}\left(\Pr[\mathsf{Ext}(X, w) = 0] + \Pr[\mathsf{Ext}(X, z) = 1]\right) \\ &\geq \frac{1}{2}\left(1 + \frac{1}{2} - \varepsilon\right) \\ &= \frac{3 - 2\varepsilon}{4}\,. \end{aligned}$$

On the other hand, $\Pr[Z' = U_1] = \frac{1}{2}$, and so by Equation (1.4),

$$\gamma \geq \frac{3 - 2\varepsilon}{4} - \frac{1}{2} \geq \frac{1}{4} - \frac{\varepsilon}{2}\,.$$

Thus, as we assume $\varepsilon < \frac{1}{6}$, we get that $\gamma > \frac{1}{6}$ which concludes the proof. □

We next recall the definition of a two-source extractor or, more generally, of an unbalanced two-source extractor, and then present our characterization of two-source extractors as seed-protecting extractors. A (k_1, k_2, ε)-*two-source extractor* is a function $\mathsf{Ext}\colon \{0,1\}^{n_1} \times \{0,1\}^{n_2} \to \{0,1\}^m$ such that for every (n_1, k_1)-source X and an independent (n_2, k_2)-source Y it holds that $\mathsf{Ext}(X, Y) \approx_\varepsilon U_m$. The existence of a two-source extractor for min-entropies $k_1 = \log n_2 + O(\log \frac{1}{\varepsilon})$ and $k_2 = \log n_1 + O(\log \frac{1}{\varepsilon})$ with $m = k_1 + k_2 - O(\log \frac{1}{\varepsilon})$ output bits was established in [9] but the problem of explicitly constructing a (k, k, ε)-two-source extractor with $n_1 = n_2 = n$ even for min-entropy as high as $k = 0.49n$ remained open for three decades [9, 5, 28].

Over the last few years, there has been remarkable progress on this problem. In particular, in a breakthrough paper Chattopadhyay and Zuckerman [8] obtained a $(k, k, \frac{1}{n})$-two-source extractor $\mathsf{Ext}\colon \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ for min-entropy as low as $k = \mathrm{poly}(\log n)$. Subsequent work [4, 13, 24] improved the entropy requirement even further to $k = \widetilde{O}(\log n)$. Constructing two-source extractors for min-entropy $O(\log n)$ (or, more ambitiously, $\log(n) + O(1)$) is highly motivated by the problem of constructing explicit Ramsey graphs [1, 2, 14, 13]. A second important open problem is constructing two-source extractors with low error. Current techniques do not yield explicit two-source extractors when $\varepsilon = 1/n^{\omega(1)}$.

Our main result here is proving an equivalence of two-source extractors and seed-protecting extractors for the class $\mathcal{F}_\Delta$.

Theorem 1.7 (two-source extractors as seed-protecting extractors)**.** *Let* $\mathsf{Ext}\colon \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$*.*

1. *If* Ext *is* (k_1, ε)*-seed-protecting extractor against* $\mathcal{F}_{d-k_2}$ *then* Ext *is a* $(k_1, k_2, 3\varepsilon)$*-two-source extractor.*
2. *If* Ext *is a* (k_1, k_2, ε)*-two-source extractor which is strong in the second source*[1] *then* Ext *is* $(k_1, 2\varepsilon)$*-seed protecting against* $\mathcal{F}_{d-k_2}$*.*

Theorem 1.7 is proven in Section 5 (see Lemma 5.1 and Lemma 5.3 for the proof of each direction. By invoking a known lower bound result on the amount of min-entropy required for two-source extractors, Theorem 1.7 in particular implies that seed protection cannot be achieved against adversaries that are allowed to have small range. Claim 1.6 gives a direct proof of that for the extreme case of range size two. We find it insightful that the natural notion of seed protection gives a characterization of the three most well-studied types of randomness extractors: strong, two-source, and non-malleable (albeit against adversaries with large range).

1.4 Open problems

For which other classes C do seed-protecting extractors against C exist? Can we get full-fledged non-malleability (i. e., against general adversaries with no fixed points) from seed-protecting extractors against these prospective classes C ? More generally, extending the connection

[1] Note that this strength requirement has the undesired effect of breaking the equivalence. However, one can always assume that a two-source extractor is strong in each of its sources provided one is willing to increase the error by a multiplicative factor of $2^{O(m)}$. See also a remark in the footnote of [28] just prior to Definition 1.3.

between seed protecting and non-malleable extractors is an intriguing open question left for future research.

For $C = \Pi$, i. e., for permutation adversaries, we proved that indeed seed protection implies non-malleability. An interesting open problem is whether non-malleable extractors for permutations (or even involutions) imply full-fledged non-malleability. A second interesting open problem is whether non-malleability against permutations (or involutions) implies that for each source X there exists a small set $B_X \subset \{0,1\}^d$ of "bad" seeds such that $(\mathsf{Ext}(X, y_1), \mathsf{Ext}(X, y_2)) \approx U$ for every two distinct $y_1, y_2 \in \{0,1\}^d \setminus B_X$. Note that this weaker notion of non-malleability already suffices for Chattopadhyay and Zuckerman's construction of two-source extractors [8].

Also, note that Theorem 1.5 incurs tm entropy loss. Recall that non-malleable extractors must satisfy $k \geq (t+1)m$, as the $t+1$ outputs must be independent for a proper choice of adversaries. On the other hand, intuitively, this requirement is unnecessary for seed-protecting extractors, and it seems that the only requirement should be $k \geq m$ (we ignore the additive error dependence for simplicity). Having said that, we do not know how to formalize the intuition above regarding the entropy loss of seed-protecting extractors. Indeed, in Section 6 we prove the existence of seed-protecting extractors via a probabilistic argument, and our proof technique requires $k \geq tm$. We leave the question of understanding the entropy loss of seed-protecting extractors to future research.

Lastly, it would be interesting to investigate the connection between seed-protecting extractors to other types of extractors. Concretely, can *non-malleable two-source* extractors be characterized by seed-protecting extractors?

2 Warm-up

As a warm-up, in this section we give a proof sketch for the equivalence between non-malleable extractors and seed-protecting extractors for permutations. For simplicity we focus on the case $t = 1$. The case $t > 1$ follows by similar ideas but is somewhat more involved.

Theorem 2.1 (Π non-malleability $\iff$ Π^2 seed protection)**.** *Let* $\mathsf{Ext}\colon \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$.

1. *If* Ext *is* (k, ε)*-seed protecting against* Π^2 *then* Ext *is* $(k', 14\varepsilon^{1/3})$*-non-malleable against* Π*, where* $k' = k + 2m + O(\log(1/\varepsilon))$.
2. *If* Ext *is* (k, ε)*-non-malleable against* Π *then* Ext *is* $(k, 3\varepsilon)$*-seed-protecting extractor against* Π^2.

Proof. We start with the first and more difficult item.

Proof sketch of the first item

Set $\delta = 14\varepsilon^{1/3}$. Assume towards a contradiction that Ext is not non-malleable against Π. Consider then a source $X \sim \{0,1\}^n$ and adversarial permutation $A \in \Pi$ with no fixed points for which

$$\mathsf{SD}\Big(\big(\mathsf{Ext}(X,Y), \mathsf{Ext}(X, A(Y)), Y\big), \big(U_m, \mathsf{Ext}(X, A(Y)), Y\big)\Big) > \delta\,,$$

where $Y \sim \{0,1\}^d$ is uniform and independent of X. This implies

$$\mathop{\mathbf{E}}_{y\sim Y}\left[\mathsf{SD}\big((\mathsf{Ext}(X,y),\mathsf{Ext}(X,A(y))),(U_m,\mathsf{Ext}(X,A(y)))\big)\right] > \delta\,. \tag{2.1}$$

For every $y, y_1 \in \{0,1\}^d$ we define the distributions

$$\begin{aligned}\mathcal{D}_{y,y_1} &= (\mathsf{Ext}(X,y),\mathsf{Ext}(X,y_1))\,,\\ \mathcal{I}_{y_1} &= (U_m,\mathsf{Ext}(X,y_1))\;.\end{aligned}$$

Define the function $T\colon \{0,1\}^d \times \{0,1\}^d \to [0,1]$ by $T(y,y_1) = \mathsf{SD}\left(\mathcal{D}_{y,y_1},\mathcal{I}_{y_1}\right)$. With this notation, Equation (2.1) can be written as $\mathop{\mathbf{E}}_{y\sim Y}[T(y,A(y))] > \delta$. By an averaging argument, there exists a set $H \subseteq \{0,1\}^d$ of size $|H| = \delta/2 \cdot 2^d$ such that $T(y,A(y)) > \delta/2$ for every $y \in H$.

Based on the fact that Ext is strong, we prove that there exists a subset $B_1 \subseteq \{0,1\}^d$ of density $2\varepsilon^{1/3}$ such that for every $y \notin B_1$, $\mathbf{E}_{y_1\sim U_d}[T(y,y_1)] \le \varepsilon^{2/3}$. We remark that this is where one pays $2m + O(\log(1/\varepsilon))$ in the entropy loss. We choose to skip the proof of this fact (see Claim 3.4).

A main part of the proof is extending $A|_H$, the restriction of A to H, to a new permutation $\widehat{A}$ over $\{0,1\}^d$ such that for almost every y outside of H it holds that $T(y,\widehat{A}(y))$ is small, in particular, bounded by $\varepsilon^{1/3}$. This is done via a greedy algorithm. We arrange the elements of $\{0,1\}^d \setminus (H \cup B_1)$ in some order $y_1,\ldots,y_\ell$. By an averaging argument, for every $y \notin B_1$, there are at most $\varepsilon^{1/3}$ fraction of seeds y_1 for which $T(y,y_1) \ge \varepsilon^{1/3}$. Denote this set by $B(y)$. We proceed iteratively, starting from $i = 1$, and choose an element $z_i \notin B(y_i)$, also different from y_i, that has not been assigned already as an element of the range of (the partially defined) $\widehat{A}$, and set $\widehat{A}(y_i) = z_i$. This can be done for most elements y_i. When i gets very close to ℓ we may have to assign the remaining elements of the range in any way, but which will still guarantee that $\widehat{A}$ is a permutation. At any rate, for simplicity, let us assume that for all elements $y_1,\ldots,y_\ell$ we have that $\widehat{A}(y_i) \notin B(y_i) \cup \{y_i\}$.

Define the random variables $\mathcal{D}_{\widehat{A}} = \mathcal{D}_{Y,\widehat{A}(Y)}$ and $\mathcal{I}_{\widehat{A}} = \mathcal{I}_{\widehat{A}(Y)}$, where Y is uniformly distributed over $\{0,1\}^d$. Using the fact that Ext is seed protecting, we prove the following.

Claim 2.2. *There exists a set $B_2 \subseteq \{0,1\}^d$ of density at most $\sqrt{\varepsilon}$ such that for every $y \in \{0,1\}^d \setminus B_2$, $\mathsf{SD}(\mathcal{D}_{\widehat{A}},\mathcal{D}_{y,\widehat{A}(y)}) \le \sqrt{\varepsilon}$.*

Proof. Denote

$$\widehat{Z}(y) = \mathsf{Ext}(X,\widehat{A}(y))\,.$$

With this notation, we have that

$$\begin{aligned}\mathcal{D}_{\widehat{A}} &= (\mathsf{Ext}(X,Y),\widehat{Z}(Y)),\\ \mathcal{D}_{y,\widehat{A}(y)} &= (\mathsf{Ext}(X,y),\widehat{Z}(y))\,.\end{aligned}$$

By construction, $\widehat{A}$ is a permutation with no fixed points (more importantly, it does not collude with the identity function). Observe that as Ext is seed protecting against $\Pi^2 \cap \mathcal{X}^2$ with error ε,

$$\mathop{\mathbf{E}}_{y \sim U_d}[\mathsf{SD}(\mathcal{D}_{y,\widehat{A}(y)}, \mathcal{D}_{\widehat{A}})] \le \varepsilon.$$

By Markov's inequality, the set $B_2 \subseteq \{0,1\}^d$ of all the y satisfying $\mathsf{SD}((\mathcal{D}_{y,\widehat{A}(y)}, \mathcal{D}_{\widehat{A}})) \ge \sqrt{\varepsilon}$ has density at most $\sqrt{\varepsilon}$, as stated.

□

Using that Ext is a strong seeded extractor, one can prove that

Claim 2.3. *There exists a set $B_3 \subseteq \{0,1\}^d$ of density at most $\sqrt{\varepsilon}$ such that for every $y \in \{0,1\}^d \setminus B_3$, $\mathsf{SD}(\mathcal{I}_{\widehat{A}}, \mathcal{I}_{\widehat{A}(y)}) \le \sqrt{\varepsilon}$.*

Write $B = B_1 \cup B_2 \cup B_3$. Recall that $|H| = \delta/2 \cdot 2^d = 7\varepsilon^{1/3} \cdot 2^d$. By the above claims we can bound $|B| < 7\varepsilon^{1/3} \cdot 2^d$, and so there exists $y_h \in H \setminus B$. Take $y_\ell \in \{0,1\}^d \setminus (H \cup B)$. By Claim 2.2, since $y_\ell, y_h \notin B_2$, $\mathsf{SD}(\mathcal{D}_{y_h,\widehat{A}(y_h)}, \mathcal{D}_{\widehat{A}}) \le \sqrt{\varepsilon}$ and $\mathsf{SD}(\mathcal{D}_{y_\ell,\widehat{A}(y_\ell)}, \mathcal{D}_{\widehat{A}}) \le \sqrt{\varepsilon}$. Hence,

$$\mathsf{SD}\left(\mathcal{D}_{y_\ell,\widehat{A}(y_\ell)}, \mathcal{D}_{y_h,\widehat{A}(y_h)}\right) \le 2\sqrt{\varepsilon}\,. \tag{2.2}$$

By Claim 2.3 and since $y_\ell, y_h \notin B_3$, $\mathsf{SD}(\mathcal{I}_{\widehat{A}(y_h)}, \mathcal{I}_{\widehat{A}}) \le \sqrt{\varepsilon}$ and $\mathsf{SD}(\mathcal{I}_{\widehat{A}(y_\ell)}, \mathcal{I}_{\widehat{A}}) \le \sqrt{\varepsilon}$. Thus,

$$\mathsf{SD}\left(\mathcal{I}_{\widehat{A}(y_\ell)}, \mathcal{I}_{\widehat{A}(y_h)}\right) \le 2\sqrt{\varepsilon}\,. \tag{2.3}$$

Now, since $y_\ell \in \{0,1\}^d \setminus (H \cup B_1)$,

$$\mathsf{SD}\left(\mathcal{D}_{y_\ell,\widehat{A}(y_\ell)}, \mathcal{I}_{\widehat{A}(y_\ell)}\right) \le \varepsilon^{1/3}\,. \tag{2.4}$$

By Equations (2.2), (2.3), and (2.4), and the triangle inequality, $\mathsf{SD}(\mathcal{D}_{y_h,\widehat{A}(y_h)}, \mathcal{I}_{\widehat{A}(y_h)}) \le 5\varepsilon^{1/3}$. However, $y_h \in H$ and so

$$\mathsf{SD}\left(\mathcal{D}_{y_h,\widehat{A}(y_h)}, \mathcal{I}_{\widehat{A}(y_h)}\right) = T\left(y_h, \widehat{A}(y_h)\right) = T(y_h, A(y_h)) > \frac{\delta}{2}\,,$$

contradicting our choice of δ. This proves the first item.

Proof sketch of the second item

Moving on to the second item, let $A_1, A_2 \in \Pi$ be non-colluding. Let X be an (n,k)-source, and let $Y \sim \{0,1\}^d$ be a uniform random variable, independent of X. Denote

$$Z(y) = (\mathsf{Ext}(X, A_1(y)), \mathsf{Ext}(X, A_2(y))).$$

As A_1^{-1} is a permutation, $A_1^{-1}(Y)$ distributes the same as Y does. Thus,

$$\mathsf{SD}((U_d, Z(Y)), (Y, Z(Y))) = \mathsf{SD}((U_d, Z(A_1^{-1}(Y))), (A_1^{-1}(Y), Z(A_1^{-1}(Y)))).$$

By the data-processing inequality,[2] we can apply A_1 on the prefix of both random variables without increasing the statistical distance; The above equation then becomes

$$\mathsf{SD}((U_d, Z(Y)), (Y, Z(Y))) = \mathsf{SD}((U_d, Z(A_1^{-1}(Y))), (Y, Z(A_1^{-1}(Y)))) . \tag{2.5}$$

Define $P = A_2 \circ A_1^{-1}$. Then,

$$Z(A_1^{-1}(Y)) = (\mathsf{Ext}(X, Y), \mathsf{Ext}(X, P(Y))) .$$

Note that P is a permutation. Furthermore, observe that as $(A_1, A_2) \in \mathcal{X}^2$, P has no fixed points. Indeed, assuming towards a contradiction that $P(y) = y$ for some $y \in \{0,1\}^d$, we get that $A_1^{-1}(y) = A_2^{-1}(y)$, which is impossible. By the non-malleability of Ext, we know that

$$\mathsf{SD}((Y, Z(A_1^{-1}(Y))), (Y, U_m, \mathsf{Ext}(X, P(Y)))) \le \varepsilon ,$$

Thus, by the triangle inequality, we get

$$\begin{aligned} \mathsf{SD}((U_d, Z(A_1^{-1}(Y))), (Y, Z(A_1^{-1}(Y)))) \le\ & \mathsf{SD}((U_d, Z(A_1^{-1}(Y))), (U_d, U_m, \mathsf{Ext}(X, P(Y))))+ \\ & \mathsf{SD}((U_d, U_m, \mathsf{Ext}(X, P(Y))), (Y, U_m, \mathsf{Ext}(X, P(Y))))+ \\ & \mathsf{SD}((Y, U_m, \mathsf{Ext}(X, P(Y))), (Y, Z(A_1^{-1}(Y)))) \\ \le\ & 3\varepsilon . \end{aligned} \tag{2.6}$$

The proof then follows by Equation (2.5). □

3 Seed-protecting extractors for high entropy adversaries are non-malleable

In this section we prove Theorem 1.5, showing that a seed-protecting extractor with adversarial entropy parameter Δ is non-malleable against roughly the same adversarial function class.

Definition 3.1 (non-colluding functions)**.** We say that a tuple $(A_1, \ldots, A_t) \in \mathcal{A}_d^t$ *do not collude*, if for every $y \in \{0,1\}^d$, $A_1(y), \ldots, A_t(y)$ are pairwise distinct. When d is clear from context, we denote by $\mathcal{X}^t$ the set of t-tuples of non-colluding functions from $\{0,1\}^d$ to $\{0,1\}^d$.

We prove the following slight restatement of Theorem 1.5. Henceforth, we denote by $\mathcal{N} \subseteq \mathcal{A}_d$ the class of functions with no fixed points. Recall that $\mathcal{F}_\Delta$ is the set of functions $A \in \mathcal{A}_d$ for which $H_\infty(A(U_d)) \ge d - \Delta$.

[2] By this we refer to the fact that for any random variables $X, Y \sim \Omega_1$, and every $f : \Omega_1 \to \Omega_2$, possibly randomized, $\mathsf{SD}(f(X), f(Y)) \le \mathsf{SD}(X, Y)$. Equality is attained when f is injective.

Theorem 3.2. *Let $t \geq 1$ be an integer, and $\Delta \geq 0$. Let $\mathsf{Ext}\colon \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ be a (k,ε)-seeded extractor with $d \geq 2\log t + \log\frac{1}{\varepsilon} + 2$ and $\varepsilon \leq 10^{-4}$. Assume that Ext is seed protecting against $\mathcal{F}^{t+1}_{\max(\Delta,\log t,1)}$. Then, Ext is non-malleable against $(\mathcal{F}_\Delta \cap \mathcal{N})^t \cap \mathcal{X}^t$ for min-entropy $k+mt+\log\frac{1}{\varepsilon}$ and error guarantee $14\varepsilon^{1/3}$.*

Note that the non-malleability guaranteed by Theorem 3.2 is with respect to $(\mathcal{F}_\Delta \cap \mathcal{N})^t \cap \mathcal{X}^t$. That is, non-colluding is still required. In Theorem 3.1 we show how to get rid of this requirement in a "black-box" fashion.

Proof of Theorem 3.2. Assume towards a contradiction that Ext is not non-malleable against $(\mathcal{F}_\Delta \cap \mathcal{N})^t$ for min-entropy $k' = k + mt + \log(1/\varepsilon)$. Then, there exists an (n,k') source X and functions $A_1,\dots,A_t \in \mathcal{F}_\Delta \cap \mathcal{N}$ with $(A_1,\dots,A_t) \in \mathcal{X}^t$, such that

$$\begin{aligned}\mathsf{SD}\Big(&\big(\mathsf{Ext}(X,Y),\mathsf{Ext}(X,A_1(Y)),\dots,\mathsf{Ext}(X,A_t(Y)),Y\big),\\ &\big(U_m,\mathsf{Ext}(X,A_1(Y)),\dots,\mathsf{Ext}(X,A_t(Y)),Y\big)\Big) > \delta\,,\end{aligned}$$

for $\delta = 14\varepsilon^{1/3}$, where $Y \sim \{0,1\}^d$ is uniform and independent of X. That is,

$$\begin{aligned}\mathop{\mathbf{E}}_{y\sim Y}\Big[&\mathsf{SD}\big((\mathsf{Ext}(X,y),\mathsf{Ext}(X,A_1(y)),\dots,\mathsf{Ext}(X,A_t(y))),\\ &(U_m,\mathsf{Ext}(X,A_1(y)),\dots,\mathsf{Ext}(X,A_t(y)))\big)\Big] > \delta\,. \qquad (3.1)\end{aligned}$$

For $y, y_1, \dots, y_t \in \{0,1\}^d$ define the distributions

$$\begin{aligned}\mathcal{D}_{y,y_1,\dots,y_t} &= (\mathsf{Ext}(X,y),\mathsf{Ext}(X,y_1),\dots,\mathsf{Ext}(X,y_t))\,,\\ \mathcal{I}_{y_1,\dots,y_t} &= (U_m,\mathsf{Ext}(X,y_1),\dots,\mathsf{Ext}(X,y_t))\,.\end{aligned}$$

Define the function $T\colon \left(\{0,1\}^d\right)^{t+1} \to [0,1]$ as follows. For $y, y_1, \dots, y_t \in \{0,1\}^d$,

$$T(y,y_1,\dots,y_t) = \mathsf{SD}\left(\mathcal{D}_{y,y_1,\dots,y_t},\mathcal{I}_{y_1,\dots,y_t}\right)\,.$$

With this notation, Equation (3.1) can be written as $\mathop{\mathbf{E}}_{y\sim Y}[T(y,A_1(y),\dots,A_t(y))] > \delta$. By an averaging argument, there exists a set $H \subseteq \{0,1\}^d$ of size $|H| = (\delta/2)\cdot 2^d$ such that for every $y \in H$,

$$T(y,A_1(y),\dots,A_t(y)) > \frac{\delta}{2}\,.$$

Let $Y_1,\dots,Y_t$ be independent random variables, that are jointly independent of X, and each Y_i is uniformly distributed over $\{0,1\}^d$. Denote

$$Z = \mathsf{Ext}(X,Y_1) \circ \cdots \circ \mathsf{Ext}(X,Y_t)\,.$$

For $z = (z_1,\dots,z_t) \in (\{0,1\}^m)^t$ define the random variable $X_z = X \mid \{Z = z\}$.

Claim 3.3. *There exists a set $B' \subseteq (\{0,1\}^m)^t$ such that:*

1. $\Pr[Z \in B'] \leq \varepsilon$*, and,*
2. *For every* $z \notin B'$, $H_\infty(X_z) \geq H_\infty(X) - tm - \log(1/\varepsilon) \geq k$.

Proof. Fix $z \in (\{0,1\}^m)^t$ and observe that

$$H_\infty(X_z) \geq H_\infty(X) - \log\left(\frac{1}{\Pr[Z = z]}\right) . \tag{3.2}$$

Define $B' = \left\{z : \Pr[Z = z] \leq 2^{-mt}\varepsilon\right\}$. By Equation (3.2), for every $z \notin B'$, item (2) holds. As $B' \subseteq \{0,1\}^{mt}$,

$$\Pr[Z \in B'] = \sum_{z \in B'} \Pr[Z = z] \leq |B'| \cdot 2^{-mt}\varepsilon \leq \varepsilon ,$$

and so item (1) follows as well. □

Claim 3.4. *There exists a subset $B_1 \subseteq \{0,1\}^d$ of size $|B_1| \leq 2\varepsilon^{1/3} \cdot 2^d$ such that for every $y \notin B_1$,*

$$\mathop{\mathbf{E}}_{(y_1,\ldots,y_t)\sim U_{td}} [T(y, y_1, \ldots, y_t)] \leq \varepsilon^{2/3} .$$

Proof. Recall the definition of X_z and B' from Lemma 3.3. Fix $z \notin B'$. Let Y be uniform over $\{0,1\}^d$ and independent of $(X, Y_1, \ldots, Y_t)$. Note that conditioned on the event $Z = z$, the random variables X_z, Y are independent, and, furthermore, Y is uniformly distributed over $\{0,1\}^d$. In addition, as $z \notin B'$, $H_\infty(X_z) \geq k$. Thus, as Ext is a strong (k, ε) seeded extractor,

$$(\mathsf{Ext}(X_z, Y), Y) \approx_\varepsilon (U_m, Y) .$$

Let Z' be the distribution obtained by sampling $z \sim Z$ conditioned on $z \notin B'$. Then,

$$\begin{aligned} T(Y, Y_1, \ldots, Y_t) &= \mathsf{SD}\big((\mathsf{Ext}(X, Y), Z, Y) , (U_m, Z, Y) \big) \\ &= \mathop{\mathbf{E}}_{z\sim Z} [\mathsf{SD}\left((\mathsf{Ext}(X_z, Y), Y) , (U_m, Y)\right)] \\ &\leq \Pr[Z \in B'] + \mathop{\mathbf{E}}_{z'\sim Z'} [\mathsf{SD}\left((\mathsf{Ext}(X_{z'}, Y), Y) , (U_m, Y)\right)] \leq 2\varepsilon . \end{aligned}$$

By Markov's inequality, there exists a subset $B_1 \subseteq \{0,1\}^d$ of size $|B_1| \leq 2\varepsilon^{1/3} \cdot 2^d$ such that for every $y \notin B_1$,

$$T(y, Y_1, \ldots, Y_t) \leq \varepsilon^{2/3} .$$

Thus, as $Y_1, \ldots, Y_t$ are independent, we get that for every $y \notin B_1$,

$$\mathop{\mathbf{E}}_{(y_1,\ldots,y_t)\sim U_{td}} [T(y, y_1, \ldots, y_t)] \leq \varepsilon^{2/3} ,$$

concluding the proof of the claim. □

Proposition 3.5. *There exist* $(\widehat{A}_1, \dots, \widehat{A}_t) \in \mathcal{F}^t_{\max(\Delta, \log t, 1)} \cap \mathcal{X}^t$ *and* $B_2 \subseteq \{0,1\}^d$ *of size* $|B_2| \leq 3\varepsilon^{1/3} \cdot 2^d + t$ *such that*

1. *For every* $y \in H$, $\widehat{A}_j(y) = A_j(y)$*, and,*
2. *For every* $y \in \{0,1\}^d \setminus (H \cup B_2)$ *it holds that* $T\left(y, \widehat{A}_1(y), \dots, \widehat{A}_t(y)\right) \leq \varepsilon^{1/3}$.

Proof. Denote $\beta = 2^{-d} \cdot |B_1 \setminus H|$. By Claim 3.4, $\beta \leq 2\varepsilon^{1/3}$. Moreover, by Markov's inequality, for every $y \notin B_1$, there exists $B(y) \subseteq (\{0,1\}^d)^t$ of size at most $|B(y)| \leq \varepsilon^{1/3} \cdot 2^{dt}$ such that for every $(y_1, \dots, y_t) \notin B(y)$,

$$T(y, y_1, \dots, y_t) \leq \varepsilon^{1/3} .$$

Let $L = \{0,1\}^d \setminus (H \cup B_1)$. Fix an (arbitrary) ordering of the elements in L and denote them by $y^1, y^2, \dots, y^{|L|}$. Let c be the least integer larger than $\max(\varepsilon^{1/3} \cdot 2^d, t)$, and set $\ell = |L| - c$. Denote

$$B_2 = (B_1 \setminus H) \cup \{y^{\ell+1}, \dots, y^{|L|}\} ,$$

observing that indeed

$$|B_2| \leq |B_1| + |L| - \ell = |B_1| + c \leq 3\varepsilon^{1/3} \cdot 2^d + t .$$

We define a family of functions

$$\left\{\widehat{A}_{j,m} : 1 \leq j \leq t, \, 0 \leq m \leq \ell\right\} ,$$

where $\widehat{A}_{j,0} \colon H \to \{0,1\}^d$ and for $m \geq 1$,

$$\widehat{A}_{j,m} \colon H \cup \{y^1, \dots, y^m\} \to \{0,1\}^d .$$

The above functions are constructed via the following algorithm which proceeds iteratively on m.

The construction algorithm for the $\widehat{A}_{j,m}$ functions

1. For every $j \in [t]$, set $\widehat{A}_{j,0} = A_j|_H$.
2. For $m = 1, \dots, \ell$, we will show in Claim 3.6 below that there exists $(y^m_1, \dots, y^m_t) \in (\{0,1\}^d)^t \setminus B(y^m)$ such that $y^m, y^m_1, \dots, y^m_t$ are pairwise distinct; furthermore, for every $j \in [t]$, $\left|\widehat{A}^{-1}_{j,m-1}(y^m_j)\right| < t$. Under these assumptions, for every $j \in [t]$ and $y \in H \cup \{y^1, \dots, y^m\}$, set

$$\widehat{A}_{j,m}(y) = \begin{cases} y^m_j & y = y^m; \\ \widehat{A}_{j,m-1}(y) & \text{otherwise.} \end{cases}$$

Claim 3.6. *The underlying assumption of step (2) in the algorithm above holds for every* $m \in [\ell]$.

Proof. We begin by setting notation. For every $m \in [\ell]$ and $j \in [t]$, we define the set $G_j^m \subseteq \{0,1\}^d \setminus \{y^m\}$ of elements whose preimage is of size strictly less than t with respect to the function $\widehat{A}_{j,m-1}$. Formally,

$$G_j^m = \left\{ y \in \{0,1\}^d \setminus \{y^m\} : \left| \widehat{A}_{j,m-1}^{-1}(y) \right| < t \right\} .$$

With this notation, the hypothesis underlying step (2) holds at iteration m if there exists an element in $(G_1^m \times \cdots \times G_t^m) \setminus B(y^m)$ with pairwise distinct entries. To establish this, we start by bounding the size of G_j^m from below for every fixed $j \in [t]$.

For an element y not to be contained in G_j^m there must be at least t elements whose image under $\widehat{A}_{j,m-1}$ is y. At the beginning of the m^{th} iteration, $m-1$ elements have been assigned an image at step (2) and additional $|H| = \delta/2 \cdot 2^d$ elements were assigned at Equation (1). Hence,

$$\begin{aligned} \left| G_j^m \right| &\geq |\{0,1\}^d \setminus \{y^m\}| - \frac{|H| + m - 1}{t} \\ &\geq 2^d - 1 - \frac{|H| + \ell - 1}{t}. \end{aligned} \tag{3.3}$$

As $\ell = |L| - c$ and since $L \cap H = \emptyset$ we have that

$$|H| + \ell = |H| + |L| - c \leq |H \cup L| - c \leq 2^d - c \,,$$

and so

$$\left| G_j^m \right| \geq \left(1 - \frac{1}{t}\right) 2^d + \frac{c+1}{t} - 1.$$

Thus, as $y^m \notin B_1$,

$$\left| \left(G_1^m \times \cdots \times G_t^m \right) \setminus B(y^m) \right| \geq \left(\left(1 - \frac{1}{t}\right) 2^d + \frac{c+1}{t} - 1 \right)^t - \varepsilon^{1/3} \cdot 2^d \,. \tag{3.4}$$

Let $\mathrm{NE} \subseteq (\{0,1\}^d)^t$ be the largest set of vectors $v \in (\{0,1\}^d)^t$ such that $v_i \neq v_j$ for every pair of distinct $i, j \in [t]$. With this notation, to prove that the assumption underlying step (2) holds at step m, one must show that

$$\left(\left(G_1^m \times \cdots \times G_t^m \right) \setminus B(y^m) \right) \cap \mathrm{NE} \neq \emptyset \,. \tag{3.5}$$

To this end, note that $\left| (\{0,1\}^d)^t \setminus \mathrm{NE} \right| \leq \binom{t}{2} 2^{(t-1)d}$. Thus, by Equation (3.4), it suffices to show that

$$\left(\left(1 - \frac{1}{t}\right) 2^d + \frac{c+1}{t} - 1 \right)^t > \varepsilon^{1/3} \cdot 2^d + \binom{t}{2} 2^{(t-1)d} \,. \tag{3.6}$$

For $t = 1$, Equation (3.6) is equivalent to $c > \varepsilon^{1/3} \cdot 2^d$ which readily follows by the definition of c. Consider then $t \geq 2$. By definition, $c \geq t$ and so $(c+1)/t - 1 > 0$. Note further that $\left(1 - \frac{1}{t}\right)^t \geq \frac{1}{4}$. Therefore, to satisfy Equation (3.6), it suffices to establish that

$$\frac{1}{4} \cdot 2^{dt} \geq \varepsilon^{1/3} \cdot 2^d + \binom{t}{2} 2^{(t-1)d} \,.$$

It is straightforward to verify that the above equation follows per our assumption $d \geq 2\log t + 2$. □

We turn to extend the functions $\widehat{A}_{1,\ell}, \ldots, \widehat{A}_{t,\ell}$ to the domain $\{0,1\}^d$. To this end, let $D = \{0,1\}^d \setminus (H \cup \{y^1, \ldots, y^\ell\})$ and denote its elements by $y^{\ell+1}, \ldots, y^{\ell+e}$. For $m = 1, \ldots, e$ we define the function

$$\widehat{A}_{j,\ell+m} \colon D \cup \{y^{\ell+1}, \ldots, y^{\ell+m}\} \to \{0,1\}^d$$

using the following iterative algorithm.

The algorithm for extending $\widehat{A}_{j,\ell}$ to $\widehat{A}_j$

1. For $m = 1, \ldots, e$, we will show in Claim 3.7 below that there exist pairwise distinct $y_1^{\ell+m}, \ldots, y_t^{\ell+m} \in \{0,1\}^d \setminus \{y^{\ell+m}\}$ such that for every $j \in [t]$, $\left|\widehat{A}_{j,\ell+m-1}^{-1}(y_j^{\ell+m})\right| \leq 1$. Under these assumptions, for every $j \in [t]$ and $y \in H \cup \{y^1, \ldots, y^{\ell+m}\}$, set

$$\widehat{A}_{j,\ell+m}(y) = \begin{cases} y_j^{\ell+m} & y = y^{\ell+m}; \\ \widehat{A}_{j,\ell+m-1}(y) & \text{otherwise.} \end{cases}$$

2. For $j \in [t]$, set $\widehat{A}_j = \widehat{A}_{j,e}$.

Claim 3.7. *The underlying assumption in step (1) of the algorithm above holds.*

Proof. Fix $m \in [e]$. For every $j \in [t]$, the number of elements z for which $|\widehat{A}_{j,\ell+m-1}^{-1}(z)| \geq 2$ is bounded above by $\frac{1}{2} \cdot 2^d$. Thus, when setting $\widehat{A}_{j,\ell+m}(y^{\ell+m})$ one has at least $2^d/2$ choices for an image with respect to this restriction. Recall that we also need to guarantee that

$$y^{\ell+m}, \widehat{A}_{1,\ell+m}(y^{\ell+m}), \ldots, \widehat{A}_{t,\ell+m}(y^{\ell+m})$$

are pairwise distinct. This can be achieved as $\frac{1}{2} \cdot 2^d + t + 1 < 2^d$ per our assumption $d \geq 2\log t + 2$. □

Analyzing the $\widehat{A}$ functions. First, note that the domain of each $\widehat{A}_1, \ldots, \widehat{A}_t$ is $\{0,1\}^d$. Recall that each of $A_1, \ldots, A_t$ has no fixed points by assumption and, furthermore, are non-colluding. Thus, $\widehat{A}_1, \ldots, \widehat{A}_t$ have no fixed points and are non-colluding when restricted to H. Moreover, by construction, the functions $\widehat{A}_1, \ldots, \widehat{A}_t$ are defined to have no fixed points and to be non-colluding outside of H as well. Thus, $(\mathrm{id}, \widehat{A}_1, \ldots, \widehat{A}_t) \in \mathcal{X}^{t+1}$, where $\mathrm{id} \colon \{0,1\}^d \to \{0,1\}^d$ is the identity function.

We turn to show that $\widehat{A}_j \in \mathcal{F}_{\max(\Delta, \log t, 1)}$ for every $j \in [t]$. Recall that $A_j \in \mathcal{F}_\Delta$. Both algorithms above assure that for any $j \in [t]$, and $y \notin \mathrm{Im}(\widehat{A}_j|_H)$ it holds that, $\left|\widehat{A}_j^{-1}(y)\right| \leq \max(t, 2)$. Observe that $f \in \mathcal{F}_{\log t}$ if and only if for every $y \in \mathrm{Im}(f)$, $|f^{-1}(y)| \leq t$, and so it holds that

$$\widehat{A}_j \in \mathcal{F}_\Delta \cup \mathcal{F}_{\log \max(t,2)} = \mathcal{F}_{\max(\Delta, \log t, 1)}\,.$$

This concludes the proof of Proposition 3.5. □

We turn back to the proof of Theorem 3.2. Define the random variables

$$\mathcal{D}_{\widehat{A}} = \mathcal{D}_{Y,\widehat{A}(Y)},$$
$$\mathcal{I}_{\widehat{A}} = \mathcal{I}_{\widehat{A}(Y)},$$

where Y is uniformly distributed over $\{0,1\}^d$ and $\widehat{A}(Y) = (\widehat{A}_1(Y), \dots, \widehat{A}_t(Y))$.

Claim 3.8. *There exists a set $B_3 \subseteq \{0,1\}^d$ of size $|B_3| \le \sqrt{\varepsilon} \cdot 2^d$ such that for every $y \in \{0,1\}^d \setminus B_3$,*

$$\mathsf{SD}\left(\mathcal{D}_{\widehat{A}}, \mathcal{D}_{y,\widehat{A}(y)}\right) \le \sqrt{\varepsilon}.$$

Proof. Denote

$$\widehat{Z}(y) = \left(\mathsf{Ext}(X, \widehat{A}_1(y)), \dots, \mathsf{Ext}(X, \widehat{A}_1(y))\right).$$

With this notation,

$$\mathcal{D}_{\widehat{A}} = \left(\mathsf{Ext}(X,Y), \widehat{Z}(Y)\right),$$
$$\mathcal{D}_{y,\widehat{A}(y)} = \left(\mathsf{Ext}(X,y), \widehat{Z}(y)\right).$$

By Proposition 3.5, $\widehat{A}_1, \dots, \widehat{A}_t \in \mathcal{F}_{\max(\Delta,\log t,1)}$; moreover, note that $\mathrm{id} \in \mathcal{F}_{\max(\Delta,\log t,1)}$. Observe that as Ext is seed protecting against $\mathcal{F}^{t+1}_{\max(\Delta,\log t,1)}$, and since $(\mathrm{id}, \widehat{A}_1, \dots, \widehat{A}_t) \in \mathcal{X}^{t+1}$. it holds that

$$\mathop{\mathbf{E}}_{y\sim U_d}\left[\mathsf{SD}\big((\mathsf{Ext}(X,y), \widehat{Z}(y)), (\mathsf{Ext}(X,Y), \widehat{Z}(Y))\big)\right] \le \varepsilon. \tag{3.7}$$

Let $B_3 \subseteq \{0,1\}^d$ be the set of all elements y satisfying

$$\mathsf{SD}\big((\mathsf{Ext}(X,y), \widehat{Z}(y)), (\mathsf{Ext}(X,Y), \widehat{Z}(Y))\big) \ge \sqrt{\varepsilon}.$$

By Markov's inequality, it follows that $|B_3| \le \sqrt{\varepsilon} \cdot 2^d$, as stated.

□

Claim 3.9. *There exists a set $B_4 \subseteq \{0,1\}^d$ of size $|B_4| \le \sqrt{\varepsilon} \cdot 2^d$ such that for every $y \in \{0,1\}^d \setminus B_4$,*

$$\mathsf{SD}\left(\mathcal{I}_{\widehat{A}}, \mathcal{I}_{\widehat{A}(y)}\right) \le \sqrt{\varepsilon}.$$

Proof. Note that for every $y \in \{0,1\}^d$ it holds that

$$\mathsf{SD}\left(\mathcal{I}_{\widehat{A}}, \mathcal{I}_{\widehat{A}(y)}\right) = \mathsf{SD}\big(\widehat{Z}(Y), \widehat{Z}(y)\big).$$

By Equation (3.7),

$$\mathop{\mathbf{E}}_{y\sim U_d}\left[\mathsf{SD}\left(\mathcal{I}_{\widehat{A}},\mathcal{I}_{\widehat{A}(y)}\right)\right]\leq\varepsilon\,.$$

Let $B_4\subseteq\{0,1\}^d$ be the set of all elements y satisfying

$$\mathsf{SD}\big(\mathcal{I}_{\widehat{A}},\mathcal{I}_{\widehat{A}(y)}\big)\geq\sqrt{\varepsilon}\,.$$

By Markov's inequality, it follows that $|B_4|\leq\sqrt{\varepsilon}\cdot 2^d$, as stated. □

We are now ready to complete the proof. Write $B=B_1\cup B_2\cup B_3\cup B_4$. Recall that

$$|H|=\frac{\delta}{2}\cdot 2^d=7\varepsilon^{1/3}\cdot 2^d\,.$$

By the above claims and using our hypothesis on d,

$$|B|\leq(5\varepsilon^{1/3}+2\sqrt{\varepsilon})\cdot 2^d+t<7\varepsilon^{1/3}\cdot 2^d\,,$$

and so there exists $y_h\in H\setminus B$. On the other hand,

$$|H\cup B|\leq\left(\frac{\delta}{2}+5\varepsilon^{1/3}+2\sqrt{\varepsilon}\right)2^d+t<2^d\,,$$

where the last inequality follows as $\varepsilon\leq 10^{-4}$ and, again, using our hypothesis on d. Hence, there exists $y_\ell\in\{0,1\}^d\setminus(H\cup B)$.

By Claim 3.8, since $y_\ell,y_h\notin B_3$,

$$\mathsf{SD}\left(\mathcal{D}_{y_h,\widehat{A}(y_h)},\mathcal{D}_{\widehat{A}}\right)\leq\sqrt{\varepsilon},$$
$$\mathsf{SD}\left(\mathcal{D}_{y_\ell,\widehat{A}(y_\ell)},\mathcal{D}_{\widehat{A}}\right)\leq\sqrt{\varepsilon}\,,$$

and so

$$\mathsf{SD}\left(\mathcal{D}_{y_\ell,\widehat{A}(y_\ell)},\mathcal{D}_{y_h,\widehat{A}(y_h)}\right)\leq 2\sqrt{\varepsilon}\,. \tag{3.8}$$

By Claim 3.9 and since $y_\ell,y_h\notin B_4$,

$$\mathsf{SD}\left(\mathcal{I}_{\widehat{A}(y_h)},\mathcal{I}_{\widehat{A}}\right)\leq\sqrt{\varepsilon},$$
$$\mathsf{SD}\left(\mathcal{I}_{\widehat{A}(y_\ell)},\mathcal{I}_{\widehat{A}}\right)\leq\sqrt{\varepsilon}\,.$$

Thus,

$$\mathsf{SD}\left(\mathcal{I}_{\widehat{A}(y_\ell)},\mathcal{I}_{\widehat{A}(y_h)}\right)\leq 2\sqrt{\varepsilon}\,. \tag{3.9}$$

Now, since $y_\ell\in\{0,1\}^d\setminus(H\cup B_2)$, by item (2) of Proposition 3.5,

$$\mathsf{SD}\left(\mathcal{D}_{y_\ell,\widehat{A}(y_\ell)},\mathcal{I}_{\widehat{A}(y_\ell)}\right)\leq\varepsilon^{1/3}\,. \tag{3.10}$$

By Equation (3.8), Equation (3.9), Equation (3.10), and the triangle inequality,

$$\mathsf{SD}\left(\mathcal{D}_{y_h,\widehat{A}(y_h)},\mathcal{I}_{\widehat{A}(y_h)}\right)\leq 5\varepsilon^{1/3}\,.$$

However, $y_h\in H$ and so

$$\begin{aligned}\mathsf{SD}\left(\mathcal{D}_{y_h,\widehat{A}(y_h)},\mathcal{I}_{\widehat{A}(y_h)}\right)&=T\left(y_h,\widehat{A}(y_h)\right)\\&=T(y_h,A(y_h))>\frac{\delta}{2}\,,\end{aligned}$$

contradicting our choice of δ. □

Observe that for $\Delta=0$ and $t=1$, Theorem 2.1 shows that seed protection against permutations is enough, and one does not need to devise an extractor against $\mathcal{F}_1$. For a general $t>1$ however, we suspect it is not true, and one needs to to handle $\mathcal{F}_\Delta$ for $\Delta>0$ to get non-malleability against permutations. We note however, that if one is willing to tolerate "smoothness" (in the sense of Definition 4.3 and Definition 4.1 below), or a slight error degradation, we can get non-malleability against permutations from seed-protecting extractors against permutations, as long as t is small enough. We omit the details.

3.1 Colluding does not harm non-malleability

Theorem 3.2 established non-malleability against $(\mathcal{F}_\Delta\cap\mathcal{N})^t\cap\mathcal{X}^t$.[3] However, in sharp contrast to seed-protecting extractors, colluding cannot help adversaries in breaking non-malleable extractors. Intuitively this should be clear, as redundant information should not help the adversary in distinguishing $\mathsf{Ext}(X,Y)$ from uniform. Here we make it formal.

Lemma 3.10. *Let* $\mathsf{Ext}\colon\{0,1\}^n\times\{0,1\}^d\to\{0,1\}^m$ *be a* (k,ε)*-non-malleable extractor against* $\mathcal{N}^t\cap\mathcal{X}^t$ *(i. e., non-colluding functions with no fixed points), so that* $d\geq\log(t+2)$. *Then,* Ext *is* (k,ε)*-non-malleable against* $\mathcal{N}^t$ *(i. e., a* t*-non-malleable extractor).*

Proof. Let $A_1,\dots,A_t\in\mathcal{A}_d\cap\mathcal{N}$ be any adversarial functions with no fixed points. Fix an (n,k)-source X and let $Y\sim\{0,1\}^d$ be a uniform random variable, independent of X.

We define the tuple of non-colluding functions $A'_1,\dots,A'_t$ as follows. Let $\mathbf{B}\subseteq\{0,1\}^d$ be the set of y-s in which a colluding occurs. Namely, for each $y\in\mathbf{B}$ there exist distinct $i,j\in[t]$ for which $A_i(y)=A_j(y)$. Note that it is possible that different set of functions collude separately, say $A_i(y)=A_j(y)=z$ and $A_{i'}(y)=A_{j'}(y)=z'$ for $z\neq z'$. Given $y\in\{0,1\}^d$, let $B(y)\subseteq[t]$ be the set of "redundant" adversaries for y. Formally,

$$B(y)=\left\{i\in[t]:\text{there exists }j<i\text{ such that }A_j(y)=A_i(y)\right\}\,.$$

[3] We recall that $\mathcal{F}_\Delta$ is the set of functions $A\in\mathcal{A}_d$ for which $H_\infty(A(U_d))\geq d-\Delta$, $\mathcal{N}$ is the set of functions with no fixed points, and $\mathcal{X}^t$ is the set of t-tuples of non-colluding functions.

Note that if $y \notin \mathbf{B}$, $B(y)$ is empty. Also, given $y \in \mathbf{B}$, we denote by $I(y) = \{A_1(y), \ldots, A_t(y)\}$, and take $E_1(y), \ldots, E_t(y)$ to be the first t elements in $\{0,1\}^d \setminus (I(y) \cup \{y\})$ in some fixed order. As $d \geq \log(t+2)$, we can indeed do so. For every $i \in [t]$, we define

$$A'_i(y) = \begin{cases} A_i(y) & y \notin \mathbf{B}, \\ A_i(y) & y \in \mathbf{B} \wedge i \notin B(y), \\ E_i(y) & \text{otherwise.} \end{cases}$$

It then follows that $\left(A'_1, \ldots, A'_t\right) \in \mathcal{X}^t$.

Denote $Z(y) = (\mathsf{Ext}(X, A_1(y)), \ldots, \mathsf{Ext}(X, A_t(y)))$, and likewise,

$$Z'(y) = (\mathsf{Ext}(X, A'_1(y)), \ldots, \mathsf{Ext}(X, A'_t(y)))\,.$$

We further define $Z_{\mathsf{reduce}}(y)$ to be joint distribution of the $\mathsf{Ext}(X, A_i(y))$ for i-s which are not in $B(y)$. Namely,

$$Z_{\mathsf{reduce}}(y) = \bigcirc_{i \in [t] \setminus B(y)} \mathsf{Ext}(X, A_i(y))$$

where $\circ$ denotes concatenation. We record the following two easy claims.

Claim 3.11. *For every $y \in \{0,1\}^d$, it holds that*

$$\begin{aligned} &\mathsf{SD}((\mathsf{Ext}(X, y), Z(y)), (U_m, Z(y))) = \\ &\mathsf{SD}((\mathsf{Ext}(X, y), Z_{\mathsf{reduce}}(y)), (U_m, Z_{\mathsf{reduce}}(y)))\,. \end{aligned}$$

Proof. The claim follows from the following observation: For every three random variables A, B, and C, it holds that $\mathsf{SD}((A, B, C, C), (U, B, C, C)) = \mathsf{SD}((A, B, C), (U, B, C))$, where U is uniform over the support of A and independent of all other random variables. □

Claim 3.12. *For every $y \in \{0,1\}^d$ it holds that*

$$\begin{aligned} &\mathsf{SD}((\mathsf{Ext}(X, y), Z_{\mathsf{reduce}}(y)), (U_m, Z_{\mathsf{reduce}}(y))) \leq \\ &\mathsf{SD}((\mathsf{Ext}(X, y), Z'(y)), (U_m, Z'(y)))\,. \end{aligned}$$

Proof. The claim readily follows from the data processing inequality, observing that $Z'(y) = (Z_{\mathsf{reduce}(y)}, A)$ for some random variable A. □

We can now finish the proof. As Ext is non-malleable against non-colluding functions, we know that

$$\mathsf{SD}((\mathsf{Ext}(X, Y), Z'(Y), Y), (U_m, Z'(Y), Y)) \leq \varepsilon\,.$$

But

$$\mathsf{SD}((\mathsf{Ext}(X, Y), Z'(Y), Y), (U_m, Z'(Y), Y)) = \mathop{\mathbf{E}}_{y \sim U_d}\left[\mathsf{SD}((\mathsf{Ext}(X, y), Z'(y)), (U_m, Z'(y)))\right],$$

so combining the above with Claim 3.11 and Claim 3.12, it follows that

$$\begin{aligned}\mathsf{SD}((\mathsf{Ext}(X,Y),Z(Y),Y),(U_m,Z(Y),Y)) &= \mathop{\mathbf{E}}_{y\sim U_d}[\mathsf{SD}((\mathsf{Ext}(X,y),Z(y)),(U_m,Z(y)))]\\ &= \mathop{\mathbf{E}}_{y\sim U_d}[\mathsf{SD}((\mathsf{Ext}(X,y),Z_{\mathsf{reduce}}(y)),(U_m,Z_{\mathsf{reduce}}(y)))]\\ &\le \mathop{\mathbf{E}}_{y\sim U_d}[\mathsf{SD}((\mathsf{Ext}(X,y),Z'(y)),(U_m,Z'(y)))]\\ &\le \varepsilon\,,\end{aligned}$$

as desired. □

The above lemma can be adapted to cases when the non-malleability is against more restricted family of functions. In particular, we will need the following lemma.

Lemma 3.13. *Let* $\mathsf{Ext}\colon \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ *be a* (k,ε)*-non-malleable extractor against* $(\mathcal{F}_\Delta \cap \mathcal{N})^t \cap \mathcal{X}^t$ *for* $\Delta \ge 1$*, so that* $d \ge \log t + 3$*. Then,* Ext *is* (k,ε)*-non-malleable against* $(\mathcal{F}_\Delta \cap \mathcal{N})^t$*.*

Proof. Inspecting the proof of Lemma 3.10, we just need to make sure that $A'_1,\dots,A'_t$ stay inside the family $\mathcal{F}_\Delta$. For every $i \in [t]$, let $\mathbf{C}_i \subseteq \{0,1\}^d$ be the set of seeds we re-wired. Namely,

$$\mathbf{C}_i = \left\{y \in \{0,1\}^d : y \in \mathbf{B} \land i \in B(y)\right\}\,.$$

Recall that if $y \in \mathbf{C}_i$ then $A'_i(y) = E_i(y)$. We will modify the definition of E_i to guarantee that $A'_i \in \mathcal{F}_\Delta$, while still satisfying $(A'_1,\dots,A'_t) \in \mathcal{X}^t \cap \mathcal{N}^t$.

Let $\mathbf{G}_1(y) = \{0,1\}^d \setminus (I(y) \cup \{y\})$ be the set of "safe" seeds. In Lemma 3.10 we simply set $E_i(y)$ to be the i-th element of $\mathbf{G}_1(y)$. Now, we need to be just a bit more careful. Let

$$\mathbf{G}_2 = \left\{z \in \{0,1\}^d : \left|A_i^{-1}(z)\right| < 2^{\Delta}\right\}\,.$$

By an averaging argument, $|\mathbf{G}_2| \ge 2^{d-\Delta}$. Thus,

$$|\mathbf{G}_1(y) \cap \mathbf{G}_2| \ge 2^d - t - 1 - 2^{d-\Delta} \ge t\,,$$

and we can set $E_i(y)$ to be the i-th element of $\mathbf{G}_1(y) \cap \mathbf{G}_2$. Both properties now hold. □

Combining Therorem 3.2 and Lemma 3.13, we get our main result.

Corollary 3.14. *Let* $t \ge 1$ *be an integer, and* $\Delta \ge 1$*. Let* $\mathsf{Ext}\colon \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ *be a* (k,ε)*-seeded extractor with* $d \ge 2\log t + \log\frac{1}{\varepsilon} + 2$ *and* $\varepsilon \le 10^{-4}$*. Assume that* Ext *is seed protecting against* $\mathcal{F}^{t+1}_{\max(\Delta,\log t)} \cap \mathcal{X}^{t+1}$*. Then,* Ext *is non-malleable against* $(\mathcal{F}_\Delta \cap \mathcal{N})^t$ *for min-entropy* $k + mt + \log\frac{1}{\varepsilon}$ *and error guarantee* $14\varepsilon^{1/3}$*.*

4 Seed-protecting extractors from non-malleable extractors

In this section, we show that non-malleability against permutations implies seed protection for the class of permutations. We also prove a similar claim for a more restricted kind of permutations – t-cliques, which we define below. For the sake of generality, we consider the "smooth" variants of these classes.

Definition 4.1 (smooth permutations)**.** Given $A \in \mathcal{A}_d$ and $\tau \in [0,1]$, we say that $A \in \Pi_\tau$ if there exists a set $G \subseteq \{0,1\}^d$ with $|G| \geq (1-\tau)\cdot 2^d$ such that $A|_G$ is injective.

To describe our next family of structured adversaries, we introduce the following notation. We say that a function $A \in \mathcal{A}$ is τ-close to an involution if for all but τ-fraction of $y \in \{0,1\}^d$ it holds that $A(A(y)) = y$ and $A(y) \neq y$. That is, the directed graph induced by the function A has, but for one component of density τ, an involution structure (i. e., a perfect matching). More generally, for $t \geq 1$, we formalize what it means for an adversarial function to be τ-close to $(t+1)$-cliques, or clusters. Fix functions $A_1, \ldots, A_t \in \mathcal{A}$. For $y \in \{0,1\}^d$, we define the neighborhood of y by $\Gamma(y) = \{y, A_1(y), \ldots, A_t(y)\}$. We say that $(A_1, \ldots, A_t) \in \mathcal{M}^t_\tau$ if there exists $G \subseteq \{0,1\}^d$ with $|G| \geq (1-\tau)\cdot 2^d$ such that for every $y \in G$ it holds that $\Gamma(\Gamma(y)) = \Gamma(y)$, and $|\Gamma(y)| = t+1$. So, intuitively, but for a density-τ component, the vertices are partitioned to cliques, or clusters, of size $t+1$. In fact, we consider a more structured variant, which we now formally define.

Definition 4.2 (t-cliques)**.** Given $A_1, \ldots, A_t \in \mathcal{A}_d$ and $p \geq t$, we say $(A_1, \ldots, A_t) \in \mathcal{M}^t[p]$ if there exists a partition $\{0,1\}^d = C_1 \uplus \cdots \uplus C_\ell$, each $t \leq |C_i| \leq p$, such that the following holds. For every $i \in [\ell]$ let $c_i = |C_i|$ and denote $C_i = \left\{y^0, \ldots, y^{c_i - 1}\right\}$, and $A_1\left(y^j\right) = z^j$. Then,

1. A_1 restricted to C_i is a permutation. That is, $\left\{y^0, \ldots, y^{c_i-1}\right\} = \left\{z^0, \ldots, z^{c_i-1}\right\}$.
2. For any integers $j \in [c_i]$ and $2 \leq r \leq t$, $A_r\left(y^j\right) = z^{j+r-1 \bmod c_i}$.

Note, in particular, that $\mathcal{M}^t[p] \subset \Pi^t$. For brevity, we denote $\mathcal{M}^t[t] = \mathcal{M}^t$.

Definition 4.3 (smooth t-cliques)**.** For positive integers $p \geq t$, and $\tau \in [0,1]$, we define the set $\mathcal{M}^t_\tau[p] \subseteq \mathcal{A}^t_d$ as follows. A tuple $\boldsymbol{A} = (A_1, \ldots, A_t) \in \mathcal{M}^t_\tau[p]$ if there exists $G \subseteq \{0,1\}^d$ with $|G| \geq (1-\tau)\cdot 2^d$ such that

$$(A_1|_G, \ldots, A_t|_G) \in \mathcal{M}^t[p]$$

where by the latter we mean formally that there exists a partition $G = C_1 \uplus \cdots \uplus C_\ell$, each $t \leq |C_i| \leq p$, such that the conditions of Definition 4.2 are met. For brevity, we denote $\mathcal{M}^t_\tau[t] = \mathcal{M}^t_\tau$.

We begin with permutations, and recall that $\mathcal{X}^t$ denotes the set of t-tuples of non-colluding functions.

Lemma 4.4. *Let* $\mathsf{Ext}\colon \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ *be a* (k,ε)*-non-malleable extractor against* $(\Pi \cap \mathcal{N})^t$ *(i. e., permutations with no fixed points), and fix any* $\tau \geq 0$*. Then,* Ext *is* (k,ε')*-seed protecting against* $\Pi^{t+1}_\tau \cap \mathcal{X}^{t+1}$ *for* $\varepsilon' = 2(t+1)(\tau+\varepsilon)$*.*

Proof. Let $A_1, \dots, A_{t+1} \in \Pi_\tau$ be non-colluding. Let X be an (n,k)-source, and let $Y \sim \{0,1\}^d$ be a uniform random variable, independent of X. For each $i \in [t+1]$, let $G_i \subseteq \{0,1\}^d$ be such that $A_i|_{G_i}$ is injective. Writing $G = G_1 \cap \cdots \cap G_{t+1}$, we have that $|G| \geq (1-(t+1)\tau) \cdot 2^d$. For each $i \in [t+1]$, denote by $\widetilde{A}_i \in \Pi$ the permutation that is obtained by keeping the function's value on G and completing it to a permutation on $\{0,1\}^d$ in such a way that $\widetilde{A}_1, \dots, \widetilde{A}_{t+1}$ do not collude. Observe that it is possible to do so as long as $\left|\{0,1\}^d \setminus G\right| \geq t+1$, which certainly holds.[4] In what follows, we denote

$$\begin{aligned} Z_1(Y) &= (\mathsf{Ext}(X, A_1(Y)), \dots, \mathsf{Ext}(X, A_{t+1}(Y))), \\ \widetilde{Z}_1(Y) &= (\mathsf{Ext}(X, \widetilde{A}_1(Y)), \dots, \mathsf{Ext}(X, \widetilde{A}_{t+1}(Y))) \,. \end{aligned}$$

First, note that

$$\mathsf{SD}((Y, Z_1(Y)), (Y, \widetilde{Z}_1(Y))) \leq \Pr[Y \notin G] \leq (t+1)\tau \,. \tag{4.1}$$

Next, as $\widetilde{A}_1^{-1}$ is a permutation, $\widetilde{A}_1^{-1}(Y)$ distributes the same as Y does. Thus,

$$\mathsf{SD}((U_d, \widetilde{Z}_1(Y)), (Y, \widetilde{Z}_1(Y))) = \mathsf{SD}((U_d, \widetilde{Z}_1(\widetilde{A}_1^{-1}(Y))), (\widetilde{A}_1^{-1}(Y), \widetilde{Z}_1(\widetilde{A}_1^{-1}(Y)))) \,. \tag{4.2}$$

By the data-processing inequality, we can apply $\widetilde{A}_1$ on the prefix of both random variables without increasing the statistical distance; Equation (4.2) becomes

$$\mathsf{SD}((U_d, \widetilde{Z}_1(Y)), (Y, \widetilde{Z}_1(Y))) = \mathsf{SD}((U_d, \widetilde{Z}_1(\widetilde{A}_1^{-1}(Y))), (Y, \widetilde{Z}_1(\widetilde{A}_1^{-1}(Y)))) \,. \tag{4.3}$$

For $i \in [t]$, define $P_i = \widetilde{A}_{i+1} \circ \widetilde{A}_1^{-1}$. Then,

$$\begin{aligned} \widetilde{Z}_1(\widetilde{A}_1^{-1}(Y)) &= \left(\mathsf{Ext}(X, Y), \mathsf{Ext}\left(X, \widetilde{A}_2(\widetilde{A}_1^{-1}(Y))\right), \dots, \mathsf{Ext}\left(X, \widetilde{A}_{t+1}(\widetilde{A}_1^{-1}(Y))\right)\right) \\ &= \left(\mathsf{Ext}\left(X, Y\right), \mathsf{Ext}\left(X, P_1(Y)\right), \dots, \mathsf{Ext}\left(X, P_t(Y)\right)\right) \,. \end{aligned}$$

Note that for every $i \in [t]$, P_i is a permutation. Furthermore, observe that as $(\widetilde{A}_1, \dots, \widetilde{A}_{t+1}) \in \mathcal{X}^{t+1}$, P_i has no fixed points. Indeed, assuming towards a contradiction that P_i satisfies $P_i(y) = y$ for some $y \in \{0,1\}^d$, we get that $\widetilde{A}_1^{-1}(y) = \widetilde{A}_{i+1}^{-1}(y)$, which is impossible since $\widetilde{A}_1(z) \neq \widetilde{A}_{i+1}(z)$ for any $z \in \{0,1\}^d$. By the non-malleability of Ext, we know that

$$\mathsf{SD}((Y, \widetilde{Z}_1(\widetilde{A}_1^{-1}(Y))), (Y, U_m, Z_2(Y))) \leq \varepsilon \,,$$

for $Z_2(Y) = (\mathsf{Ext}(X, P_1(Y)), \dots, \mathsf{Ext}(X, P_t(Y)))$, and the same is true without conditioning on Y. Thus, by the triangle inequality, we get

$$\begin{aligned} \mathsf{SD}((U_d, \widetilde{Z}_1(\widetilde{A}_1^{-1}(Y))), (Y, \widetilde{Z}_1(\widetilde{A}_1^{-1}(Y)))) \leq{} & \mathsf{SD}((U_d, \widetilde{Z}_1(\widetilde{A}_1^{-1}(Y))), (U_d, U_m, Z_2(Y))) + \\ & \mathsf{SD}((U_d, U_m, Z_2(Y)), (Y, U_m, Z_2(Y))) + \\ & \mathsf{SD}((Y, U_m, Z_2(Y)), (Y, \widetilde{Z}_1(\widetilde{A}_1^{-1}(Y)))) \\ \leq{} & 2\varepsilon + \mathsf{SD}((U_d, U_m, Z_2(Y)), (Y, U_m, Z_2(Y))) \,. \end{aligned} \tag{4.4}$$

[4] Otherwise, $\tau < 2^{-d}$ which implies $\tau = 0$ and we can simply take $\widetilde{A}_i = A_i$ for each $i \in [t+1]$.

We continue in the same manner. Observing that $P_1^{-1}(Y)$ distributes the same as Y does, and using the data-processing inequality, we have

$$\mathsf{SD}((U_{d+m}, Z_2(Y)), (Y, U_m, Z_2(Y))) = \mathsf{SD}((U_{d+m}, Z_2(P_1^{-1}(Y))), (Y, U_m, Z_2(P_1^{-1}(Y)))) . \tag{4.5}$$

For $i \in [t-1]$ define the permutation $Q_i = P_{i+1} \circ P_1^{-1}$. Similarly to the above argument,

$$\begin{aligned} Z_2(P_1^{-1}(Y)) &= \left(\mathsf{Ext}(X, Y), \mathsf{Ext}(X, P_2(P_1^{-1}(Y))), \dots, \mathsf{Ext}(X, P_t(P_1^{-1}(Y)))\right) \\ &= (\mathsf{Ext}(X, Y), \mathsf{Ext}(X, Q_1(Y)), \dots, \mathsf{Ext}(X, Q_{t-1}(Y))) \ . \end{aligned}$$

This time, the fact that every Q_i has no fixed points follows from the fact that $\widetilde{A}_2(z) \neq \widetilde{A}_{i+2}(z)$ for any $z \in \{0,1\}^d$. Using the non-malleability, together with Equations (4.4), Equations (4.3), and Equations (4.5), we get

$$\begin{aligned} \mathsf{SD}((U_d, \widetilde{Z}_1(Y)), (Y, \widetilde{Z}_1(Y))) \le 2\varepsilon + {} & \mathsf{SD}((U_d, U_m, Z_2(Y)), (U_d, U_{2m}, Z_3(Y))) + \\ & \mathsf{SD}((U_d, U_{2m}, Z_3(Y)), (Y, U_{2m}, Z_3(Y))) + \\ & \mathsf{SD}((Y, U_{2m}, Z_3(Y)), (Y, U_m, Z_2(Y))) \\ \le {} & 4\varepsilon + \mathsf{SD}((U_d, U_{2m}, Z_3(Y)), (Y, U_{2m}, Z_3(Y))) , \end{aligned}$$

for $Z_3(Y) = (\mathsf{Ext}(X, Q_1(Y)), \dots, \mathsf{Ext}(X, Q_{t-1}(Y)))$. We continue this process inductively, and eventually obtain

$$\mathsf{SD}((U_d, \widetilde{Z}_1(Y)), (Y, \widetilde{Z}_1(Y))) \le 2(t+1)\varepsilon + \mathsf{SD}((U_d, U_{(t+1)m}), (Y, U_{(t+1)m})) = 2(t+1)\varepsilon . \tag{4.6}$$

Combining Equations (4.6) and Equation (4.1), we get

$$\begin{aligned} \mathsf{SD}((U_d, Z_1(Y)), (Y, Z_1(Y))) \le {} & \mathsf{SD}((Y, Z_1(Y)), (Y, \widetilde{Z}_1(Y))) + \\ & \mathsf{SD}((Y, \widetilde{Z}_1(Y)), (U_d, \widetilde{Z}_1(Y))) + \\ & \mathsf{SD}((U_d, \widetilde{Z}_1(Y)), (U_d, Z_1(Y))) \\ \le {} & (t+1)(2\varepsilon + \tau) + \mathsf{SD}(\widetilde{Z}_1(Y), Z_1(Y)) . \end{aligned}$$

To bound the last term of the above inequality, note that

$$\mathsf{SD}(\widetilde{Z}_1(Y), Z_1(Y)) \le \mathsf{SD}((Y, \widetilde{Z}_1(Y)), (Y, Z_1(Y))) \le (t+1)\tau ,$$

where the last inequality follows by Equation (4.1). This concludes the proof. □

Next, we prove a similar lemma for t-cliques.

Lemma 4.5. *Let* $\mathsf{Ext}\colon \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ *be a* (k, ε)*-non-malleable extractor against* $\mathcal{M}^t \cap \mathcal{N}^t$*, and fix any* $\tau \ge 0$*. Then,* Ext *is a* (k, ε')*-seed-protecting extractor against* $\mathcal{M}_\tau^{t+1} \cap \mathcal{X}^{t+1}$ *for* $\varepsilon' = 2(t+1)\varepsilon + 2\tau$*.*

Proof. Given non-colluding $\boldsymbol{A} = (A_1, \dots, A_{t+1}) \in \mathcal{M}_\tau^{t+1}$, an (n,k)-source X and a uniform $Y \sim \{0,1\}^d$ independent of X, the proof proceeds similarly to Lemma 4.4. Let $G \subseteq \{0,1\}^d$ be the set defined in Definition 4.3 with respect to $(A_1, \dots, A_{t+1}) \in \mathcal{M}_\tau^{t+1}$, and recall that $|G| \le \tau \cdot 2^d$. We define

$$\widetilde{\boldsymbol{A}} = \left(\widetilde{A}_1, \dots, \widetilde{A}_{t+1}\right) \in \mathcal{M}^{t+1}$$

by setting $\widetilde{A}_i$ to agree with A_i on G for every i. Completing $\widetilde{A}_i|_G$ to functions on $\{0,1\}^d$ while maintaining the $\mathcal{M}^{t+1}$ property can be done by arbitrarily selecting $t+1$ inputs that were not assigned yet, iteratively, and assigning them to form a clique.[5] Note that by the definition of $\mathcal{M}^{t+1}$, the functions in $\widetilde{\boldsymbol{A}}$ do not collude.

Define, inductively, the following set of functions. For $i \in [t+1]$, $P_i^{(0)} = \widetilde{A}_i$. For every $j \in [t]$ and $i \in [t+1-j]$, we define

$$P_i^{(j)} = P_{i+1}^{(j-1)} \circ \left(P_1^{(j-1)}\right)^{-1} . \tag{4.7}$$

Next, define, exactly as in Lemma 4.4,

$$\widetilde{Z}_1(Y) = \left(\mathsf{Ext}\left(X, P_1^{(0)}(Y)\right), \dots, \mathsf{Ext}\left(X, P_{t+1}^{(0)}(Y)\right)\right) .$$

Moreover, for every $j \in [t+1]$, we define

$$Z_{j+1}(Y) = \left(\mathsf{Ext}\left(X, P_1^{(j)}(Y)\right), \dots, \mathsf{Ext}\left(X, P_{t-j+1}^{(j)}(Y)\right)\right) .$$

The crux of the proof is establishing the following inequality for every $j \ge 2$.

$$\begin{aligned}\mathsf{SD}((U_d, \widetilde{Z}_1(Y)), (Y, \widetilde{Z}_1(Y))) \le 2(j-1)\varepsilon + \\ \mathsf{SD}((U_d, U_{(j-1)m}, Z_j(Y)), (Y, U_{(j-1)m}, Z_j(Y))) .\end{aligned} \tag{4.8}$$

Following the same reasoning as in Lemma 4.4, Equation (4.8) holds if the following conditions are met:

1. For every $j \ge 0$, $\left(P_1^{(j)}\right)^{-1}(Y)$ distributes the same as Y does.
2. For every $j \ge 1$ and $i \le t+1-j$, $P_i^{(j)}$ has no fixed points.
3. For every $j \ge 0$ it holds that $(P_1^{(j)}, \dots, P_{t-j+1}^{(j)}) \in \mathcal{M}^{t-j+1}[t+1]$. Note that a non-malleable extractor against $\mathcal{M}^{t+1}$ is also non-malleable against $\mathcal{M}^{t'}[t+1]$ for any $t' \le t+1$.

[5] Formally, choose $y_1, \dots, y_{t+1}$ that were not assigned by $\widetilde{A}_1$ yet, namely, with no preimage in $\widetilde{A}_1|_G$. Assign $y_1 \to y_1, \dots, y_{t+1} \to y_{t+1}$ in $\widetilde{A}_1$, $y_1 \to y_2, y_2 \to y_3, \dots, y_{t+1} \to y_1$ in $\widetilde{A}_2$, and so on. Note that when $t+1$ does not divide 2^d, we can make some cliques larger. We do not address this issue formally and it does not affect the statement.

Item (1) readily holds, since each $\widetilde{A}_i$ is a permutation, and permutations are closed under inversion and composition. To see that item (2) holds, fix some $j \in [t]$ and $i \in [t+1-j]$ and consider $P_i^{(j)}$. Assume towards a contradiction that $P_i^{(j)}(y) = y$ for some $y \in \{0,1\}^d$. Thus,

$$P_{i+1}^{(j-1)}\left(\left(P_1^{(j-1)}\right)^{-1}(y)\right) = y\,,$$

so

$$\left(P_1^{(j-1)}\right)^{-1}(y) = \left(P_{i+1}^{(j-1)}\right)^{-1}(y)\,,$$

which means there exists $z \in \{0,1\}^d$ such that $P_{i+1}^{(j-1)}(z) = P_1^{(j-1)}(z)$. But due to item (3) which we now prove, $\left(P_1^{(j-1)}, \ldots, P_{t-j+2}^{(j-1)}\right) \in \mathcal{M}^{t-j+2}$, so in particular they do not collude, so we have a contradiction. Indeed, what is left is to prove item (3).

Claim 4.6. *For every $j \geq 0$ it holds that $(P_1^{(j)}, \ldots, P_{t-j+1}^{(j)}) \in \mathcal{M}^{t-j+1}[t+1]$.*

Proof. Starting from $(t+1)$-cliques in the 0-th level, in general, the j-th level will also form $(t+1)$-cliques. We will prove this by induction on j. For $j = 0$, it follows by our construction. Assuming that it holds for some $j \geq 0$, we inspect the $(j+1)$-th level.

We characterize the cliques in the j-th level as follows. We can partition the domain (and codomain) to $\{0,1\}^d = C_1 \uplus \cdots \uplus C_\ell$. By definition, a full characterization of $P_1^{(j)}, \ldots, P_{t-j+1}^{(j)}$ can be given by a permutation $\phi_i \colon C_i \to C_i$, for each $i \in [\ell]$. Indeed, for any $y \in C_i$ for some $i \in [\ell]$, $P_1^{(j)}(y) = \phi_i(y)$, and for $r \geq 2$, $P_r^{(j)} = \phi_i^r$. Namely, ϕ_i generates the permutations $P_1^{(j)}, \ldots, P_{t-j+1}^{(j)}$ restricted to C_i. We claim that this structure is preserved for the next level, with the same clique structure. Indeed, for each $i \in [\ell]$ observe that the permutation ϕ_i is also a generator for $P_1^{(j+1)}, \ldots, P_{t-j}^{(j+1)}$ restricted to C_i as, by Equation (4.7) applied with $i = 1$,

$$P_1^{(j+1)} = P_2^{(j)} \circ \left(P_1^{(j)}\right)^{-1} = \phi_i^2 \circ \phi_i^{-1} = \phi_i\,.$$

Recalling that $P_r^{(j+1)} = P_{r+1}^{(j)} \circ \left(P_1^{(j)}\right)^{-1}$, we see that indeed computing $P_r^{(j+1)}(y)$ amounts to finding the i for which $y \in C_i$ and computing $\phi_i^r(y)$. Thus, the $(j+1)$-th level belong to $\mathcal{M}^{t-j}[t+1]$. □

□

5 1-seed protecting and two-source extractors

In this section we prove Theorem 1.7. We prove each direction separately in Lemma 5.1 and Lemma 5.3 below. Recall that $\mathcal{F}_\Delta$ is the set of functions $A \in \mathcal{A}_d$ for which $H_\infty(A(U_d)) \geq d - \Delta$.

Lemma 5.1. *Let $\mathsf{Ext} \colon \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ be a (k_1, ε) 1-seed-protecting extractor against $\mathcal{F}_{d-k_2}$ for $k_2 \leq d-1$. Then, Ext is a $(k_1, k_2, 3\varepsilon)$ two-source extractor.*

Proof. Assume towards a contradiction that Ext is not such a two-source extractor, and let X_1,X_2 with $H_\infty(X_1) \ge k_1$ and $H_\infty(X_2) \ge k_2$ be such that

$$\mathsf{SD}\big(\mathsf{Ext}(X_1, X_2), U_m\big) > 3\varepsilon\,,$$

and we assume without loss of generality that both X_1 and X_2 are flat.[6] Let Y be the uniform distribution over $\{0,1\}^d$, independent of (X_1, X_2). Recall that Ext is a (k_1, ε) extractor, so

$$\mathsf{SD}(\mathsf{Ext}(X_1, Y), U_m) \le \varepsilon.$$

Order the y-s according to $\mathsf{SD}(\mathsf{Ext}(X_1, y), U_m)$, and let $G_1 \subseteq \{0,1\}^d$ be the 2^{d-1} bottom ones (i. e., for which $\mathsf{SD}(\mathsf{Ext}(X_1, y), U_m)$ is smaller). As G_1 can be indexed using $d-1$ bits, there exists an injection $A_0\colon \{0,1\}^{d-1} \to \{0,1\}^d$ that maps to G_1 uniformly, and satisfies

$$\mathsf{SD}(\mathsf{Ext}(X_1, A_0(Y')), U_m) \le \varepsilon\,, \tag{5.1}$$

where Y' is the uniform distribution over $\{0,1\}^{d-1}$, independent of all other variables.

Denote $G_2 = \mathrm{Supp}(X_2)$ and recall that $|G_2| = 2^{k_2}$. Again, since $k_2 \le d-1$, we can define an injection $A_1\colon \{0,1\}^{d-1} \to \{0,1\}^d$ so that $A_1(U_{d-1}) = U_{G_2}$, and then

$$\mathsf{SD}(\mathsf{Ext}(X_1, A_1(Y')), U_m) > 3\varepsilon\,. \tag{5.2}$$

Next, define $A \in \mathcal{A}_d$ such that

$$A(y) = A_{y_1}(y_{[2:n]})\,.$$

Claim 5.2. *It holds that $A \in \mathcal{F}_{d-k_2}$.*

Proof. Let $Z = A(U_d)$. Clearly, $\mathrm{Supp}(Z) \subseteq G_1 \cup G_2$. If $z \in G_1$,

$$\begin{aligned}\Pr_{y\sim U_d}[A(y) = z] &= \frac{1}{2}\cdot \Pr_{y\sim U_d}[A_0(y_{[2,n]}) = z] + \frac{1}{2}\cdot \Pr_{y\sim U_d}[A_1(y_{[2,n]}) = z]\\ &= \frac{1}{2}\cdot\frac{1}{|G_1|} + \frac{1}{2}\cdot\frac{\mathbf{1}_{z\in G_2}}{|G_2|} \le 2^{-k_2}\,.\end{aligned}$$

The same bound applies for $z \in G_2$ in a similar manner. □

We will now show that using A, an adversary can learn the first bit of the seed, in contradiction to the fact that Ext is seed protecting. Define

$$R_{Y_1} \triangleq \mathsf{Ext}(X_1, A_{Y_1}(Y_{[2:d]}))\,,$$

and note that by Equation (5.1), Equation (5.2) and the triangle inequality, it holds that $\mathsf{SD}(R_0, R_1) > 2\varepsilon$. Then,

$$\begin{aligned}\mathsf{SD}((Y, \mathsf{Ext}(X_1, A(Y))), (U_d, \mathsf{Ext}(X_1, A(Y)))) &\ge \mathsf{SD}((Y_1, \mathsf{Ext}(X_1, A(Y))), (U_1, \mathsf{Ext}(X_1, A(Y))))\\ &= \mathsf{SD}((Y_1, R_{Y_1}), (U_1, R_{Y_1}))\,.\end{aligned}$$

[6] A k-source is *flat* if it is uniformly distributed over a set of size 2^k. It is well-known that one can assume X_1 and X_2 are flat, since any k-source is a convex combination of flat k-sources [9].

Finally, observe that

$$\begin{aligned}\mathsf{SD}((Y_1, R_{Y_1}), (U_1, R_{Y_1}) &= \frac{1}{2}\sum_r \left|\Pr[Y_1 = 1 \wedge R_1 = r] - \frac{1}{2}\Pr[R_{Y_1} = r]\right| \\ &\quad + \frac{1}{2}\sum_r \left|\Pr[Y_1 = 0 \wedge R_0 = r] - \frac{1}{2}\Pr[R_{Y_1} = r]\right| \\ &= \frac{1}{4}\sum_r |\Pr[R_0 = r] - \Pr[R_1 = r]| \\ &= \frac{1}{2}\cdot \mathsf{SD}(R_0, R_1) \\ &> \varepsilon\,,\end{aligned}$$

which is a contradiction to that Ext is 1-seed protecting per our hypothesis.

□

The other direction also holds.

Lemma 5.3. *Let* Ext: $\{0,1\}^n \times \{0,1\}^d \longrightarrow \{0,1\}^m$ *be a* (k_1, k_2, ε) *two-source extractor which is strong in the second source. Then,* Ext *is a* $(k_1, 2\varepsilon)$ *1-seed protecting against* $\mathcal{F}_{d-k_2}$.

Proof. Assume towards a contradiction that Ext is not 1-seed protecting, so there exists an (n, k_1)-source X_1 and $A \in \mathcal{F}_{d-k_2}$ such that

$$\mathsf{SD}((\mathsf{Ext}(X_1, A(Y)), Y), (\mathsf{Ext}(X_1, A(Y)), Y')) > 2\varepsilon\,,$$

where $Y, Y' \sim \{0,1\}^d$ are uniform and independent random variables, also independent of X_1. By the triangle inequality, it follows that at least one of the following holds:

1. $\mathsf{SD}((\mathsf{Ext}((X_1, A(Y)), Y'), (U_m, Y)) > \varepsilon$,
2. $\mathsf{SD}((\mathsf{Ext}(X_1, A(Y)), Y), (U_m, Y)) > \varepsilon$.

First assume the first inequality holds. As Y' is independent of X_1 and Y,

$$\begin{aligned}\mathsf{SD}((\mathsf{Ext}((X_1, A(Y)), Y'), (U_m, Y)) &= \mathsf{SD}(\mathsf{Ext}(X_1, A(Y)) \times Y', U_m \times Y) \\ &= \mathsf{SD}(\mathsf{Ext}(X_1, A(Y)) \times Y', U_m \times Y') \\ &= \mathsf{SD}(\mathsf{Ext}(X_1, A(Y)), U_m)\,.\end{aligned}$$

As $A \in \mathcal{F}_{d-k_2}$, $H_\infty(A(Y)) \geq k_2$ and by the fact that Ext is a two-source extractor, it follows that $\mathsf{SD}(\mathsf{Ext}(X_1, A(Y)), U_m) \leq \varepsilon$, contradicting item (1).

Next, assume that the second inequality holds. We have that

$$\begin{aligned}\mathsf{SD}((\mathsf{Ext}(X_1, A(Y)), A(Y)), (U_m, A(Y))) &= \mathop{\mathbf{E}}_{z\sim A(Y)}[\mathsf{SD}(\mathsf{Ext}(X_1, z), U_m)] \\ &= \mathop{\mathbf{E}}_{y\sim Y}[\mathsf{SD}(\mathsf{Ext}(X_1, A(y)), U_m)] \\ &= \mathsf{SD}((\mathsf{Ext}(X_1, A(Y)), Y), (U_m, Y)) > \varepsilon\,.\end{aligned}$$

But since Ext is a strong (k_1, k_2, ε) two-source extractor,

$$\mathsf{SD}((\mathsf{Ext}(X_1, A(Y)), A(Y)), (U_m, A(Y))) \le \varepsilon\,,$$

in contradiction. □

5.1 Lower bounds for 1-seed-protecting extractors

In light of the above result, lower bounds for (unbalanced) two-source extractors imply lower bounds for 1-seed-protecting extractors. We use the following standard lower bound.[7]

Theorem 5.4. *Let* $\mathsf{Ext}\colon \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ *be a* (k_1, k_2, ε) *two-source extractor which is strong in the second source. Then,* $m \le k_1 - 2\log(1/\varepsilon) + O(1)$*,* $k_2 \ge \log(n - k_1) + 2\log(1/\varepsilon) - O(1)$ *and* $k_1 \ge \log(d - k_2) + 2\log(1/\varepsilon) - O(1)$*.*

We can thus conclude:

Corollary 5.5. *Let* $\mathsf{Ext}\colon \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ *be a* (k, ε) *1-seed protecting against* $\mathcal{F}_\Delta$ *for some* $\Delta > 0$*. Then,* $m \le k - 2\log(1/\varepsilon) + O(1)$*,* $d \ge \log(n - k) + \Delta + 2\log(1/\varepsilon) - O(1)$ *and* $k \ge \log \Delta + 2\log(1/\varepsilon) - O(1)$*.*

The fact that no 1-seed-protecting extractors against $\mathcal{F}_\Delta$ exist for Δ approaching d can be established in a more straightforward way. Indeed, in Claim 1.6 we showed why $\Delta = d - 1$ is unattainable.

6 Non-explicit t-seed-protecting extractors

In this section we prove the existence of a seed-protecting extractor against non-colluding, entropy-preserving, adversaries via a probabilistic argument.

Theorem 6.1. *Let* $n, k, m, d, t \in \mathbb{N}$*,* $\Delta \ge 0$ *and* $\varepsilon > 0$ *be such that*

$$k \ge tm + 2\Delta + 2\log\frac{1}{\varepsilon} + O(\log d + \log t)\,.$$

Then, there exists a (k, ε)*-seed-protecting extractor* $\mathsf{Ext}\colon \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ *against* $\mathcal{F}_\Delta^t \cap \mathcal{X}^t$ *with*

$$d = \log(n - k) + 2\Delta + 2\log\frac{1}{\varepsilon} + O(\log d + \log t)\,.$$

Remark 6.2. Theorem 6.1 tells us we cannot take Δ to be larger than $d/2$, meaning we are at the *weak seeds* regime. In contrast, for $t = 1$, we know from the equivalence to two-source extractors (see Lemma 5.3) that we can take Δ to be much larger, roughly $d - \log n$. An interesting open problem is whether there is a real barrier going from $t = 1$ to larger t-s or whether it is a mere artifact of our proof.

[7] The lower bound for $m = 1$ and any nontrivial ε follows from bounds on strong dispersers and their connection to Ramsey graphs [27, 3]. The entropy loss and the $2\log(1/\varepsilon)$ factor in d and k follows from lower bounds on strong seeded extractors [27].

Proof. Choose $\mathsf{Ext}\colon \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ uniformly at random. Fix an (n,k) source X, and we assume without loss of generality that it is flat. Also, fix non-colluding functions $A_1,\dots,A_t \in \mathcal{F}_\Delta$. Let $Y \sim \{0,1\}^d$ be uniform and independent of X. Write $N = 2^n$, $K = 2^k$, $D = 2^d$, and

$$Z(x,y) = (\mathsf{Ext}(x,A_1(y)),\dots,\mathsf{Ext}(x,A_t(y)))\ ,$$

observing that for every fixed x and y, Z is a random variable whose randomness comes from Ext. We want to bound the probability that

$$\mathsf{SD}((Y,Z(X,Y)),(U_d,Z(X,Y))) > \varepsilon\,.$$

Fix $T\colon \{0,1\}^{d+mt} \to \{0,1\}$ and for each $y \in [D]$, denote by $T_y\colon \{0,1\}^{mt} \to \{0,1\}$ the corresponding restriction of T. Then, we want to bound the probability over Ext that

$$\mathop{\mathbf{E}}_{w\sim[D]}\left[\mathop{\mathbf{E}}_{X}[T_w(Z(X,w))]\right] - \mathop{\mathbf{E}}_{w\sim[D]}\left[\mathop{\mathbf{E}}_{X,Y}[T_w(Z(X,Y))]\right] > \varepsilon\,.$$

Write the expression on the left hand side as

$$\mathop{\mathbf{E}}_{x\sim X,w\sim[D]}\left[T_w(Z(x,w)) - \mathop{\mathbf{E}}_{Y}[T_w(Z(x,Y))]\right]\ ,$$

and define

$$Q(x,w) = T_w(Z(x,w)) - \mathop{\mathbf{E}}_{Y}[T_w(Z(x,Y))]\ .$$

First, we argue,

Claim 6.3. *For any $x \in \{0,1\}^n$ and $w \in [D]$ it holds that $\mathbf{E}[Q(x,w)] = 0$.*

Proof. By our assumption on $A_1,\dots,A_t$, the values $A_1(w),\dots,A_t(w)$ are distinct, so $Z(x,w)$ is uniform over $\{0,1\}^{mt}$, and thus $\mathbf{E}[T_w(Z(x,w))] = \mu(T_w)$. The claim now follows from the linearity of expectation. □

Write

$$Q(x,w) = (T_w(Z(x,w)) - \mu(T_w)) - \left(\mathop{\mathbf{E}}_{Y}[T_w(Z(x,Y))] - \mu(T_w)\right) \triangleq Q_1(x,w) - Q_2(x,w)\,,$$

each $Q_i(x,w)$ being a random variable with expectation zero. We handle each term separately.

Handling Q_1. Define the random variable

$$\mathbf{Q}_1 = \mathop{\mathbf{E}}_{x\sim X,w\sim[D]}[Q_1(x,w)]\,.$$

Unfortunately, the random variables $Q_1(x,w)$, for $x \in X$ and $w \in [D]$, are not independent. In particular, it may be the case that $Q_1(x,w_1)$ and $Q_1(x,w_2)$ query the same input to Ext. The next observation will help us overcome this issue.

Claim 6.4. *Assume there exists a subset $V \subseteq [D]$ such that $\{A_i(w)\}_{i \in [t], w \in V}$ are all distinct, and enumerate $\{Q_1(x,w)\}_{x \in X, w \in V} = Q_1, \ldots, Q_{s=K|V|}$ arbitrarily. Then, for every $i \in [s]$,*

$$\mathbf{E}[Q_i \mid Q_1, \ldots, Q_{i-1}] = 0\,.$$

Proof. Fix $i \in [s]$. First note that for $x_a \neq x_b$, $Q_a = Q_1(x_a, w_a)$ and $Q_b = Q_1(x_b, w_b)$ are independent. Assume that $Q_i = Q_1(x,w)$ and let $\{i_1, \ldots, i_\ell\} \subseteq [i-1]$ be the indices that correspond to the same x, i. e., each $Q_{i_j} = Q_1(x, w_j)$ for some $w_j \neq w$. Thus,

$$\mathbf{E}[Q_i \mid Q_1, \ldots, Q_{i-1}] = \mathbf{E}[Q_i \mid Q_{i_1}, \ldots, Q_{i_\ell}]\,.$$

Next, fix all values of Ext queried by the Q_{i_j}-s. Keeping the notation $Q_{i_j} = Q_1(x, w_j)$, this means we fix every $\mathsf{Ext}(x, A_r(w_j))$ for $j \in [\ell]$ and $r \in [t]$. These fixings do not affect Q_i, by our assumption on V. Thus, under these fixings, $\mathbf{E}[Q_i] = 0$, as desired. □

We now argue that we can partition $[D]$ to a bounded number of such V-s.

Lemma 6.5. *There exists a partition $[D] = V_1 \cup \cdots \cup V_L$ for $L = O(dt^2 2^\Delta)$ such that for every $i \in [L]$, $\left\{A_j(w)\right\}_{j \in [t], w \in V_i}$ are all distinct.*

Proof. Let $G' = (W = [D], U = [D], E_0)$ be the bipartite graph in which each $w \in W$ is connected to $A_1(w), \ldots, A_t(w)$. Thus, G is left-regular with degree t and its right-degree is bounded by $t \cdot 2^\Delta$. Let $G = (V = [D], E)$ be the two-step walk graph of G'. Namely, $(x,y) \in E$ if and only if there exists a path $x \sim z \sim y$ in G', where $x, y \in W$ and $z \in U$. Note that the maximal degree in G is at most $t^2 \cdot 2^\Delta$. We will repeatedly use the following standard claim, which can be shown by a simple greedy algorithm.

Claim 6.6. *Let G be an undirected graph over n vertices with maximal degree δ. Then, the size of the largest independent set in G is at least $n/(\delta+1)$.*

The crucial observation is that an independent set in G corresponds to a valid partition. To see this, take any $w_1, w_2 \in V$ such that $(w_1, w_2) \notin E$. By definition, there are no $r_1, r_2 \in [t]$ such that $A_{r_1}(w_1) = A_{r_2}(w_2)$. In light of this observation, we can greedily define $V_1, \ldots, V_L$ in the following manner.

1. Set $G_0 \leftarrow G$, $i \leftarrow 0$ and $\delta = t^2 2^\Delta$.
2. As long as G_i has more than 2δ vertices,
   - Let V_{i+1} be the largest independent set in G_i.
   - Remove V_{i+1} and all its adjacent edges and denote the resulting graph by G_{i+1}. Set $i \leftarrow i+1$.
3. The graph G_i has $b \leq 2\delta$ vertices. Put each of these vertices in a separate set.
4. The resulting partition is $V_1, \ldots, V_i, \ldots, V_{L=i+b}$.

Claim 6.6 guarantees that at each iteration, V_{i+1} contains at least $1 - 1/(\delta+1)$ fraction of the remaining vertices. Let j be the smallest integer for which

$$\left(1 - \frac{1}{\delta+1}\right)^j \cdot 2^d \le 2\delta\,.$$

One can verify that $j = O(\delta d)$, so overall $L \le j + 2\delta = O\left(dt^2 2^{\Delta}\right)$, as desired. □

In light of the above lemma and Claim 6.4, we can define, for each $i \in [L]$,

$$\mathbf{S}_i = \sum_{x \in X, w \in V_i} Q_1(x,w)\,,$$

so $\mathbf{Q}_1 = \frac{1}{KD}\sum_{i\in[L]} \mathbf{S}_i$. Note that every sequence in $\mathbf{S}_i$ is a martingale, and also, that $|Q_1(x,w)| \le 1$ with probability 1. Thus, using Azuma's inequality,

$$\begin{aligned}\Pr\left[|\mathbf{Q}_1| > \frac{\varepsilon}{2}\right] &= \Pr\left[\left|\sum_{i\in[L]} \mathbf{S}_i\right| > \frac{\varepsilon}{2}KD\right] \le \sum_{i\in[L]} \Pr\left[|\mathbf{S}_i| > \frac{\varepsilon}{2L}KD\right] \\ &\le \sum_{i\in[L]} 2\exp\left(-\frac{\left(\frac{\varepsilon}{2L}KD\right)^2}{2K|V_i|}\right) \le 2L \cdot \mathrm{e}^{-\frac{KD}{8L^2}\varepsilon^2}\,. \end{aligned} \tag{6.1}$$

Handling Q_2. Similarly, we define $\mathbf{Q}_2 = \mathbf{E}_{x\sim X, w\sim[D]}[Q_2(x,w)]$, but we write it as

$$\mathbf{Q}_2 = \mathop{\mathbf{E}}_{x\sim X, y\sim Y}\left[\mathop{\mathbf{E}}_{w\sim[D]}[T_w(Z(x,y))] - \mu(T)\right],$$

i. e., we switched the order of w and y. Now, define

$$Q_2'(x,y) = \mathop{\mathbf{E}}_{w\sim[D]}[T_w(Z(x,y))] - \mu(T)\,,$$

so $\mathbf{Q}_2 = \mathbf{E}_{x\sim X, y\sim[D]}[Q_2'(x,y)]$. We follow the same reasoning as before: For arbitrary $y_1 \ne y_2$, $Q_2'(x,y_1)$ may depend on $Q_2'(x,y_2)$, but with the same partitioning we can overcome the dependencies. Also, for any $x \in \{0,1\}^n$ and $y \in [D]$ it holds that

$$\mathbf{E}\left[\mathop{\mathbf{E}}_{w\sim[D]}[T_w(Z(x,y))]\right] = \mu(T)\,,$$

so overall,

$$\Pr\left[|\mathbf{Q}_2| > \frac{\varepsilon}{2}\right] \le 2L \cdot \mathrm{e}^{-\frac{KD}{8L^2}\varepsilon^2} \tag{6.2}$$

as well.

Putting it all together. Combining Equation (6.1) and Equation (6.2), we get

$$\Pr\left[\mathop{\mathbf{E}}_{x\sim X, w\sim[D]}[Q(x,w)] > \varepsilon\right] \le 4L\cdot \mathrm{e}^{-\frac{KD}{8L^2}\varepsilon^2} \le 2^{-c_1\frac{KD}{d^2t^42^{2\Delta}}\varepsilon^2+\log d+2\log t+c_2}$$

for some universal constants $c_1, c_2 > 0$. To complete our analysis, we require Ext to work for any $X, A_1, \dots, A_t$ and T. By the union bound, the probability for a random Ext to fail, denote it by p, is at given by

$$\begin{aligned} p &\le \binom{N}{K} D^{tD} 2^{DM^t} 2^{-c_1\frac{KD}{d^2t^42^{2\Delta}}\varepsilon^2+\log d+2\log t+c_2} \\ &\le 2^{K\log\left(\frac{Ne}{K}\right)+tDd+DM^t-c_1\frac{KD}{d^2t^42^{2\Delta}}\varepsilon^2+\log d+2\log t+c_2} \\ &\le 2^{K(n-k+2)+tDd+DM^t+\log d+2\log t+c_2-c_1\frac{KD}{d^2t^42^{2\Delta}}\varepsilon^2}\,. \end{aligned}$$

To prove that $p < 1$ (in fact, we will show that $p \ll 1$) it is sufficient to argue that:

1. $K(n-k+2) \le \frac{c_1}{4}\frac{KD}{d^2t^42^{2\Delta}}\varepsilon^2$, and,
2. $D(td + M^t) + \log d + 2\log t + c_2 \le \frac{c_1}{4}\frac{KD}{d^2t^42^{2\Delta}}\varepsilon^2$, or, $D(4td + M^t) \le \frac{c_1}{4}\frac{KD}{d^2t^42^{2\Delta}}\varepsilon^2$.

Item (1) is true whenever

$$D \ge \frac{4}{c_1}\cdot\frac{(n-k+2)d^2t^42^{2\Delta}}{\varepsilon^2}\,.$$

Item (2) it true whenever

$$K \ge \frac{4}{c_1}\cdot\frac{(4td+M^t)d^2t^42^{2\Delta}}{\varepsilon^2}\,.$$

The bounds on d and k follow from the above two inequalities. □

References

[1] Boaz Barak, Guy Kindler, Ronen Shaltiel, Benny Sudakov, and Avi Wigderson: Simulating independence: New constructions of condensers, Ramsey graphs, dispersers, and extractors. *J. ACM*, 57(4):20:1–52, 2010. [doi:10.1145/1734213.1734214] 8

[2] Boaz Barak, Anup Rao, Ronen Shaltiel, and Avi Wigderson: 2-source dispersers for $n^{o(1)}$ entropy, and Ramsey graphs beating the Frankl–Wilson construction. *Ann. Math.*, 176(3):1483–1543, 2012. [doi:10.4007/annals.2012.176.3.3] 8

[3] Avraham Ben-Aroya, Dean Doron, and Amnon Ta-Shma: Near-optimal erasure list-decodable codes. In *Proc. 35th Comput. Complexity Conf. (CCC'20)*, volume 169, pp. 1:1–27. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2020. [doi:10.4230/LIPIcs.CCC.2020.1] 30

[4] AVRAHAM BEN-AROYA, DEAN DORON, AND AMNON TA-SHMA: An efficient reduction from two-source to nonmalleable extractors: Achieving near-logarithmic min-entropy. *SIAM J. Comput.*, 51(2):31–49, 2022. [doi:10.1137/17M1133245] 8

[5] JEAN BOURGAIN: More on the sum-product phenomenon in prime fields and its applications. *Int. J. Number Theory*, 1(1):1–32, 2005. [doi:10.1142/S1793042105000108] 8

[6] ESHAN CHATTOPADHYAY, VIPUL GOYAL, AND XIN LI: Non-malleable extractors and codes, with their many tampered extensions. In *Proc. 48th STOC*, pp. 285–298. ACM Press, 2016. [doi:10.1145/2897518.2897547] 3

[7] ESHAN CHATTOPADHYAY AND XIN LI: Explicit non-malleable extractors, multi-source extractors and almost optimal privacy amplification protocols. In *Proc. 57th FOCS*, pp. 158–167. IEEE Comp. Soc., 2016. [doi:10.1109/FOCS.2016.25] 3

[8] ESHAN CHATTOPADHYAY AND DAVID ZUCKERMAN: Explicit two-source extractors and resilient functions. *Ann. Math.*, 189(3):653–705, 2019. [doi:10.4007/annals.2019.189.3.1] 2, 8, 9

[9] BENNY CHOR AND ODED GOLDREICH: Unbiased bits from sources of weak randomness and probabilistic communication complexity. *SIAM J. Comput.*, 17(2):230–261, 1988. [doi:10.1137/0217015] 2, 8, 28

[10] GIL COHEN: Local correlation breakers and applications to three-source extractors and mergers. *SIAM J. Comput.*, 45(4):1297–1338, 2016. [doi:10.1137/15M1029837] 3

[11] GIL COHEN: Making the most of advice: new correlation breakers and their applications. In *Proc. 57th FOCS*, pp. 188–196. IEEE Comp. Soc., 2016. [doi:10.1109/FOCS.2016.28] 3

[12] GIL COHEN: Non-malleable extractors—New tools and improved constructions. In *Proc. 31st Comput. Complexity Conf. (CCC'16)*, pp. 8:1–29. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2016. [doi:10.4230/LIPIcs.CCC.2016.8] 3

[13] GIL COHEN: Towards optimal two-source extractors and Ramsey graphs. In *Proc. 49th STOC*, pp. 1157–1170. ACM Press, 2017. [doi:10.1145/3055399.3055429] 8

[14] GIL COHEN: Two-source dispersers for polylogarithmic entropy and improved Ramsey graphs. *SIAM J. Comput.*, 50(3):30–67, 2021. Preliminary version in STOC'16. [doi:10.1137/16M1096219] 8

[15] GIL COHEN, RAN RAZ, AND GIL SEGEV: Nonmalleable extractors with short seeds and applications to privacy amplification. *SIAM J. Comput.*, 43(2):450–476, 2014. Preliminary version in CCC'12. [doi:10.1137/130908634] 3

[16] GIL COHEN AND LEONARD J. SCHULMAN: Extractors for near logarithmic min-entropy. In *Proc. 57th FOCS*, pp. 178–187. IEEE Comp. Soc., 2016. [doi:10.1109/FOCS.2016.27] 3

[17] YEVGENIY DODIS, XIN LI, TREVOR D. WOOLEY, AND DAVID ZUCKERMAN: Privacy amplification and non-malleable extractors via character sums. In *Proc. 52nd FOCS*, pp. 668–677. IEEE Comp. Soc., 2011. [doi:10.1109/FOCS.2011.67] 3

[18] YEVGENIY DODIS AND DANIEL WICHS: Non-malleable extractors and symmetric key cryptography from weak secrets. In *Proc. 41st STOC*, pp. 601–610. ACM Press, 2009. [doi:10.1145/1536414.1536496] 2, 3

[19] ZEEV DVIR, SWASTIK KOPPARTY, SHUBHANGI SARAF, AND MADHU SUDAN: Extensions to the method of multiplicities, with applications to Kakeya sets and mergers. *SIAM J. Comput.*, 42(6):2305–2328, 2013. [doi:10.1137/100783704] 2, 3

[20] ODED GOLDREICH AND AVI WIGDERSON: Robustly self-ordered graphs: Constructions and applications to property testing. *TheoretiCS*, 1(8788), 2022. Preliminary version in CCC'21. [doi:10.46298/theoretics.22.1] 4

[21] VENKATESAN GURUSWAMI, CHRISTOPHER UMANS, AND SALIL VADHAN: Unbalanced expanders and randomness extractors from Parvaresh–Vardy codes. *J. ACM*, 56(4):20:1–34, 2009. [doi:10.1145/1538902.1538904] 2, 3

[22] XIN LI: Three-source extractors for polylogarithmic min-entropy. In *Proc. 56th FOCS*, pp. 863–882. IEEE Comp. Soc., 2015. [doi:10.1109/FOCS.2015.58] 3

[23] XIN LI: Improved non-malleable extractors, non-malleable codes and independent source extractors. In *Proc. 49th STOC*, pp. 1144–1156. ACM Press, 2017. [doi:10.1145/3055399.3055486] 3

[24] XIN LI: Non-malleable extractors and non-malleable codes: partially optimal constructions. In *Proc. 34th Comput. Complexity Conf. (CCC'19)*, pp. 28:1–49. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2019. [doi:10.4230/LIPIcs.CCC.2019.28] 3, 8

[25] CHI-JEN LU, OMER REINGOLD, SALIL VADHAN, AND AVI WIGDERSON: Extractors: Optimal up to constant factors. In *Proc. 35th STOC*, pp. 602–611. ACM Press, 2003. [doi:10.1145/780542.780630] 2

[26] NOAM NISAN AND DAVID ZUCKERMAN: Randomness is linear in space. *J. Comput. System Sci.*, 52(1):43–52, 1996. [doi:10.1006/jcss.1996.0004] 2

[27] JAIKUMAR RADHAKRISHNAN AND AMNON TA-SHMA: Bounds for dispersers, extractors, and depth-two superconcentrators. *SIAM J. Discr. Math.*, 13(1):2–24, 2000. [doi:10.1137/S0895480197329508] 30

[28] RAN RAZ: Extractors with weak random seeds. In *Proc. 37th STOC*, pp. 11–20. ACM Press, 2005. [doi:10.1145/1060590.1060593] 8

[29] RONEN SHALTIEL: An introduction to randomness extractors. In *Proc. 38th Internat. Colloq. on Automata, Languages, and Programming (ICALP'11)*, pp. 21–41. Springer, 2011. [doi:10.1007/978-3-642-22012-8_2] 2

[30] AMNON TA-SHMA AND CHRISTOPHER UMANS: Better condensers and new extractors from Parvaresh–Vardy codes. In *Proc. 27th IEEE Conf. on Comput. Complexity (CCC'12)*, pp. 309–315. IEEE Comp. Soc., 2012. [doi:10.1109/CCC.2012.25] 2, 3

[31] LUCA TREVISAN: Extractors and pseudorandom generators. *J. ACM*, 48(4):860–879, 2001. [doi:10.1145/502090.502099] 2

[32] SALIL P. VADHAN: Pseudorandomness. *Found. Trends Theor. Comp. Sci.*, 7(1–3):1–336, 2012. [doi:10.1561/0400000010] 2

AUTHORS

Gil Cohen
Assistant professor
Department of Computer Science
Tel Aviv University
Tel Aviv, Israel
gil@tauex.tau.ac.il
https://www.gilcohen.org

Dean Doron
Assistant professor
Department of Computer Science
Ben Gurion University of the Negev
Beer Sheva, Israel
deand@bgu.ac.il
https://www.cs.bgu.ac.il/~deand

Shahar Samocha
Software engineer
StarkWare Industries Ltd
shahar.samocha@gmail.com

ABOUT THE AUTHORS

GIL COHEN has been a faculty member at the school of computer science in Tel Aviv University since 2018. He completed his Ph. D. under the guidance of Ran Raz at the Weizmann Institute of Science in 2015. He then proceeded to postdoctoral positions at Caltech and Princeton University, hosted by Leonard Schulman and Thomas Vidick at Caltech, and by Mark Braverman at Princeton. His research focuses on pseudorandomness, derandomization, and explicit constructions.

DEAN DORON is an assistant professor at the Computer Science Department at Ben Gurion University of the Negev. He received his B. Sc. from the Technion, and spent his graduate studies at Tel Aviv University under the supervision of Amnon Ta-Shma. After receiving his Ph. D. in 2018, he spent a year as a postdoc at UT Austin hosted by Dana Moshkovitz and David Zuckerman, and two years as a Motwani postdoc at Stanford. He is interested in computational complexity theory, randomness in computing, coding theory, and also in beautiful hikes and good food.

SHAHAR SAMOCHA is a software engineer and researcher in cryptography and blockchain. He received his M. Sc. degree from Tel Aviv University under the supervision of Gil Cohen, where his research mainly focused on interactive error-correcting codes.