2012 IEEE 53rd Annual Symposium on Foundations of Computer Science

Non-Malleable Extractors, Two-Source Extractors and Privacy Amplification

Xin Li
University of Washington
lixints@cs.washington.edu

Abstract— In [1], Dodis and Wichs introduced the notion of a non-malleable extractor. A non-malleable extractor is a much stronger version of a seeded extractor. Dodis and Wichs showed that such an object can be used to give optimal privacy amplification protocols with an active adversary.

Previously, there are only two known constructions of non-malleable extractors [2], [3]. Both constructions only work for (n, k)-sources with $k > n/2$. Interestingly, both constructions are also two-source extractors.

In this paper, we present a strong connection between non-malleable extractors and two-source extractors. The first part of the connection shows that non-malleable extractors can be used to construct two-source extractors. This partially explains why previous constructions of non-malleable extractors only work for entropy rate $> 1/2$, and why explicit non-malleable extractors for small min-entropy may be hard to get.

The second part of the connection shows that certain two-source extractors can be used to construct non-malleable extractors. Using this connection, we obtain the first construction of non-malleable extractors for $k < n/2$.

Finally, despite the lack of explicit non-malleable extractors for arbitrarily linear entropy, we give the first 2-round privacy amplification protocol with asymptotically optimal entropy loss and communication complexity for (n, k) sources with $k = \alpha n$ for any constant $\alpha > 0$. This dramatically improves previous results and answers an open problem in [2].

1. Introduction

The broad area of *randomness extraction* studies the problem of converting a weakly random source into a distribution that is close to the uniform distribution in statistical distance. Over the past decades extensive research has been conducted in this area. Among which, a long line of research ([4], [5], [6], [7], [8] to name a few) studies the so called "seeded extractors", as defined by Nisan and Zuckerman [9]. Seeded extractors have a variety of applications in computer science. We refer the reader to [10], [11], [12] for a survey on this subject. Nowadays we have nearly optimal constructions of seeded extractors [6], [7], [8].

Another line of research focuses on the problem of extracting random bits from several independent sources [13], [14], [15], [16], [17], [18], [19], [20]. In this case, however, the best known construction is far from optimal. Specifically, the probabilistic method shows that there exists an extractor for two independent sources on n bits with each having roughly $\log n$ bits of entropy, while the best two-source extractor to date can only achieve entropy slightly below $n/2$ [17]. The best known extractor for small entropy k requires $O(\log n/\log k)$ independent sources [18], [19]. Moreover, it seems hard to improve these results. Especially in the two-source case, after decades of efforts the entropy requirement only drops from anything above $n/2$ [13] to slightly below $n/2$ [17].

Recently, a new kind of seeded extractors, called *non-malleable extractors* were introduced in [1] to give protocols for the problem of privacy amplification with an active adversary. We now give the definition of a non-malleable extractor below. As a comparison, we also give the definition of a strong seeded extractor.

Notation. We let $[s]$ denote the set $\{1, 2, \ldots, s\}$. For ℓ a positive integer, U_ℓ denotes the uniform distribution on $\{0,1\}^\ell$, and for S a set, U_S denotes the uniform distribution on S. When used as a component in a vector, each U_ℓ or U_S is assumed independent of the other components. We say $W \approx_\varepsilon Z$ if the random variables W and Z have distributions which are ε-close in variation distance.

Definition 1.1. The *min-entropy* of a random variable X is

$$H_\infty(X) = \min_{x \in \mathsf{supp}(X)} \log_2(1/\Pr[X = x]).$$

For $X \in \{0,1\}^n$, we call X an $(n, H_\infty(X))$-source, and we say X has *entropy rate* $H_\infty(X)/n$.

Definition 1.2. A function $\mathsf{Ext} : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ is a *strong* (k, ε)*-extractor* if for every (n, k) source X and independent uniform Y on $\{0,1\}^d$,

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

Definition 1.3. A function $\mathsf{nmExt} : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ is a (k, ε)-non-malleable extractor if, for any (n, k) source X and any function $\mathcal{A} : \{0,1\}^d \to \{0,1\}^d$ such that $\mathcal{A}(y) \neq y$ for all y, the following holds. When

Partially supported by NSF Grants CCF-0634811, CCF-0916160, THECB ARP Grant 003658-0113-2007, and a Simons postdoctoral fellowship.

0272-5428/12 $26.00 © 2012 IEEE
DOI 10.1109/FOCS.2012.26

Y is independent and uniform on $\{0,1\}^d$,

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

As we can see from the definitions, a non-malleable extractor is a stronger version of the strong extractor, in the sense that it requires the output to be close to uniform even conditioned on both the seed Y and the output $\mathsf{nmExt}(X,\mathcal{A}(Y))$ on a different but arbitrarily correlated seed $\mathcal{A}(Y)$.

The motivation to study a non-malleable extractor, the privacy amplification problem, is a fundamental problem in symmetric cryptography that has been studied by many researchers. Bennett, Brassard, and Robert introduced this problem in [21]. The basic setting is that, two parties (Alice and Bob) share an n-bit secret key X, which is weakly random. This could happen because the secret comes from a password or biometric data, which are themselves weakly random, or because an adversary Eve managed to learn some partial information about an originally uniform secret, for example via side channel attacks. We measure the entropy of X by the min-entropy defined above. The goal is to have Alice and Bob communicate over a public channel so that they can convert X into a nearly uniform secret key. Generally, we also assume that Alice and Bob have local private uniform random bits. The problem is the presence of the adversary Eve, who can see every message transmitted in the channel and may or may not change the messages. We assume that Eve has unlimited computational power.

The case where Eve is *passive*, i.e., cannot change the messages, can be solved simply by using the above mentioned strong seeded extractors. The case where Eve is *active* (i.e., can change the messages in arbitrary ways), on the other hand, is much more difficult. Historically, Maurer and Wolf [22] gave the first non-trivial protocol in this case. Their protocol takes one round and works when the entropy rate of the weakly-random secret X is bigger than $2/3$. Dodis, Katz, Reyzin, and Smith [23] later improved this result to give protocols that work for entropy rate bigger than $1/2$. One drawback in both cases is that the final secret key R is much shorter than the min-entropy of X. Later, Dodis and Wichs [1] showed that no one-round protocol exists for entropy rate less than $1/2$. The first protocol that breaks the $1/2$ entropy rate barrier is due to Renner and Wolf [24], where they gave a protocol that works for essentially any entropy rate. However their protocol takes $O(s)$ rounds and only achieves entropy loss $O(s^2)$, where s in the security parameter of the protocol. Kanukurthi and Reyzin [25] simplified their protocol, but the parameters remain essentially the same.

In [1], Dodis and Wichs showed that explicit non-malleable extractors can be used to give privacy amplification protocols that take an optimal 2 rounds and achieve optimal entropy loss $O(s)$. They showed that non-malleable extractors exist when $k > 2m + 3\log(1/\varepsilon) + \log d + 9$ and $d > \log(n-k+1) + 2\log(1/\varepsilon) + 7$. However, they only constructed weaker forms of non-malleable extractors and they gave a protocol that takes 2 rounds but that still has entropy loss $O(s^2)$. Chandran, Kanukurthi, Ostrovsky and Reyzin [26] improved the entropy loss to $O(s)$ but the number of rounds becomes $O(s)$ as well.

Dodis, Li, Wooley and Zuckerman [2] constructed the first explicit non-malleable extractor. Their construction works for entropy $k > n/2$, but they use a large seed length $d = n$ and the efficiency when outputting more than $\log n$ bits relies on an unproven assumption. Cohen, Raz, and Segev [3] later gave an alternative construction that also works for $k > n/2$, but uses a short seed length and does not rely on any unproven assumption. The construction in [3] also allows multiple adversarial functions $\{\mathcal{A}_i\}$. By using the non-malleable extractors, these two papers thus gave 2-round privacy amplification protocols that achieve optimal entropy loss $O(s)$. However, since both constructions of non-malleable extractors are only shown to work for entropy $k > n/2$,[1] the protocols also only work for $k > n/2$. For any constant $\delta > 0$, [2] also gave a protocol for $k = \delta n$ than runs in $\mathrm{poly}(1/\delta)$ rounds and achieves optimal entropy loss $O(s)$. Recently, Li [27] introduced the notion of a non-malleable condenser, which is a relaxation of a non-malleable extractor. He showed that non-malleable condensers for (n,k) sources also give privacy amplification protocols that take an optimal 2 rounds and achieve optimal entropy loss $O(s)$. However, the non-malleable condensers constructed in [27] also only work for $k > n/2$. Thus the natural open question is whether we can construct non-malleable extractors or condensers for smaller min-entropy, and whether there are 2-round privacy amplification protocols with optimal entropy loss for smaller min-entropy.

One interesting aspect of the two known constructions of non-malleable extractors is that they are also both two-source extractors. Indeed, the construction in [2] is in fact one of the two-source extractors introduced in [13], which requires the sources to have min-entropy $> n/2$, and the construction in [3] is in fact the two-

[1] We note that the 1-bit case construction in [2] is a special case of the construction in [3]. Also, it is possible that the construction in [2] can work for entropy $k \le n/2$ (but until now nobody can prove it), but the construction in [3] in general cannot work for $k \le n/2$.

source extractor in [16], which requires at least one of the sources to have min-entropy $> n/2$. Coincidently, when used as non-malleable extractors, both constructions also require the weak source to have min-entropy $> n/2$. These facts suggest possible connections between these two kinds of extractors. However, before this work, no such connection is known.

1.1. Our results

In this paper, we present a strong connection between non-malleable extractors and two-source extractors. First, we show that non-malleable extractors can be used to construct two-source extractors. If the non-malleable extractor works for small min-entropy and has a short seed length (w.r.t. $\log(1/\epsilon)$ where ϵ is the error of the extractor), then the resulted two-source extractor beats the best known construction of two-source extractors.

Theorem 1.4. *Assume that for any $\epsilon > 0$, we have explicit constructions of (k, ϵ)-non-malleable extractors with seed length $d = 2\log(1/\epsilon)+o(n)$ and output length m. Then there exists a constant $\delta > 0$ and an explicit construction of two source extractors that take as input an $(n, (1/2 - \delta)n)$ source and an independent (n, k) source, and output m bits with error $2^{-\Omega(n)}$.*

Note that if k is small, say $k = n/3$ then this already beats the best known two-source extractors, but better results can be achieved if we have explicit constructions of generalized non-malleable extractors. We have the following definition (which already appears in [3]).

Definition 1.5. A function $\mathsf{nmExt} : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ is a (r, k, ε)-non-malleable extractor if, for any (n, k) source X and any r function $\mathcal{A}_i : \{0,1\}^d \to \{0,1\}^d, i = 1, \cdots, r$ such that $\mathcal{A}_i(y) \neq y$ for all i and y, the following holds. When Y is independent and uniform on $\{0,1\}^d$,

$$\begin{aligned}&(\mathsf{nmExt}(X, Y), \{\mathsf{nmExt}(X, \mathcal{A}_i(Y))\}, Y)\\ \approx_\varepsilon &(U_m, \{\mathsf{nmExt}(X, \mathcal{A}_i(Y))\}, Y).\end{aligned}$$

Here r is the number of adversarial seeds. Note that traditional non-malleable extractors are just $(1, k, \epsilon)$-non-malleable extractors according to our definition. In the full version [28] we show that for any constant r, (r, k, ε)-non-malleable extractors exist with seed length $d > \frac{3}{2}\log(n - k) + 3\log(1/\epsilon) + O(1)$. Now we have the following theorem.

Theorem 1.6. *For any constant $b > 2$ and any constant $0 < \delta < 1$, there exists a constant $C = C(\delta) = \mathrm{poly}(1/\delta)$ such that the following holds. Assume that for any $\epsilon > 0$ there exists an explicit construction of (C, k, ϵ)-non-malleable extractors with seed length $d = b\log(1/\epsilon) + o(n)$ and output length m. Then there exists an explicit construction of two source extractors that take as input an $(n, \delta n)$ source and an independent (n, k) source, and output m bits with error $2^{-\Omega(n)}$.*

Note that if we have a (C, k, ϵ)-non-malleable extractor for $k = \delta n$ and some constant $C = C(\delta) = \mathrm{poly}(1/\delta)$ then this will give us a two-source extractor for $(n, \delta n)$ sources. If δ is small this will be a big breakthrough for two-source extractors. This also implies that, given current techniques, the (r, k, ϵ)-non-malleable extractor in [3] is probably the best that we can achieve.

Next, we show that in the opposite direction, certain two-source extractors can be used to construct non-malleable extractors. The two-source extractors we will use are those that are constructed based on the inner product function. More specifically, we will consider two-source extractors of the form $\mathsf{TExt} = \mathsf{IP}(f(X), Y)$, where IP is the inner product function over $\mathbb{F}_2$ and $f(X)$ stands for some function (encoding) of the source X. We have the following theorem.

Theorem 1.7. *Given two integers r, ℓ such that $\ell > r$. Assume that we have a two-source extractor $\mathsf{TExt} = \mathsf{IP}(f(X), W)$ such that when given an (n, k)-source X and an independent $(n_2, n_2/(r+1)-\ell)$-source W, TExt outputs 1 bit with error ϵ. Then there exists an explicit construction of (r, k, ϵ')-non-malleable extractors that output 1 bit with error $\epsilon' = O(r2^{r-\ell} + 2^{\frac{3r}{2}}\epsilon)$.*

Using this theorem, and by combining known two-source extractors, we obtain new and improved constructions of non-malleable extractors. We give the first explicit constructions of non-malleable extractors that work for min-entropy $k < n/2$. One of them is unconditional and works for $k = (1/2 - \delta)n$ for some universal constant $\delta > 0$. The other is conditional but can potentially work for $k = \delta n$ for any constant $\delta > 0$. Specifically, we have the following theorems.

Theorem 1.8. *There exists a constant $0 < \delta < 1$ and an explicit (k, ϵ)-non-malleable extractor $\mathsf{nmExt} : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^m$ with $k = (1/2 - \delta)n$, $m = \Omega(n)$ and $\epsilon = 2^{-\Omega(n)}$.*

Our conditional result needs to use an affine extractor. Roughly speaking, here by an $[n, m, \rho, \epsilon]$ affine extractor we mean a deterministic function $f : \{0,1\}^n \to \{0,1\}^m$ such that whenever X is the uniform distribution over some affine subspace over $\mathbb{F}_2^n$ with dimension ρn, the output is within ϵ to the uniform distribution

in the ℓ^∞ norm. We let λ denote the entropy loss rate, i.e., $\lambda = 1 - \frac{m}{\rho n}$. We note that it is straightforward to show by the probabilistic method that such extractors exist for any constant $\rho, \lambda > 0$. However the state of art constructions only achieve λ bigger than $1/2$.

[29] also introduced the Approximate Duality conjecture (ADC), which basically says that if two independent sources X, Y with linear entropy are such that $\mathsf{IP}(X, Y)$ is not close to uniform, then there exist two subsources $X' \subset X, Y' \subset Y$ with small deficiency such that $\mathsf{IP}(X', Y')$ is constant. In [29] it is shown that ADC is implied by the well-known Polynomial Freiman-Ruzsa Conjecture in additive combinatorics. For a formal definition, see [29] or the full version.

Theorem 1.9. *Given a constant integer r, assume the ADC conjecture and we have an explicit $[n, m, \frac{r+1}{r+2}, 2^{-m}]$ affine extractor with $m = (1 - \lambda)\frac{r+1}{r+2}n$. Then there exists a semi-explicit (r, k, ϵ)-non-malleable extractor with $k = \frac{(r+2)\lambda}{1+(r+1)\lambda}n$, seed length $d = \frac{r+2}{r+1+(r+1)^2\lambda}n - 1$ and $\epsilon = 2^{-\Omega(n)}$.*

Remark 1.10. Here we use the "*semi-explicit*" to mean that the construction may run in time 2^n (note that an exhaustive search takes time 2^{2^n}). If we have affine extractors such that $\lambda \to 0$, then we can essentially achieve $k = \alpha n$ for any constant $\alpha > 0$.

Finally, we give a new privacy amplification protocol for min-entropy $k = \delta n$ for any constant $\delta > 0$. Although we don't have explicit non-malleable extractors or condensers for such small k, our protocol simultaneously achieves optimal round complexity (2 rounds), asymptotically optimal entropy loss and asymptotically optimal communication complexity. This is the first optimal privacy amplification protocol for arbitrarily linear min-entropy. We have the following theorem.

Theorem 1.11. *For any constant $0 < \delta < 1$ there exists a constant $0 < \beta < 1$ such that as long as $s \leq \beta n$, there is an efficient 2-round privacy amplification protocol for any $(n, \delta n)$ weak secret X with security parameter s, entropy loss $O(s+\log n)$ and communication complexity $O(s + \log n)$.*

Thus, in the case where $k = \delta n$, our result dramatically improves all previous results. Especially, it improves the round complexity in [2] from $\mathrm{poly}(1/\delta)$ to 2, and thus answers an open problem in [2].

2. Preliminaries

We often use capital letters for random variables and corresponding small letters for their instantiations. Let $|S|$ denote the cardinality of the set S. All logarithms are to the base 2.

Definition 2.1 (statistical distance)**.** Let W and Z be two distributions on a set S. Their *statistical distance* (variation distance) is

$$\Delta(W, Z) \stackrel{def}{=} \max_{T \subseteq S}(|W(T) - Z(T)|) = \frac{1}{2}\sum_{s \in S}|W(s) - Z(s)|.$$

We say W is ε-close to Z, denoted $W \approx_\varepsilon Z$, if $\Delta(W, Z) \leq \varepsilon$. For a distribution D on S and a function $h : S \to T$, let $h(D)$ be the distribution on T induced by choosing x according to D and outputting $h(x)$.

Definition 2.2. A function $\mathsf{TExt} : \{0,1\}^{n_1} \times \{0,1\}^{n_2} \to \{0,1\}^m$ is a *strong two source extractor* for min-entropy k_1, k_2 and error ϵ if for every independent (n_1, k_1) source X and (n_2, k_2) source Y, $|(\mathsf{TExt}(X, Y), X) - (U_m, X)| < \epsilon$ and $|(\mathsf{TExt}(X, Y), Y) - (U_m, Y)| < \epsilon$.

Definition 2.3. An elementary somewhere-k-source is a vector of sources $(X_1, \cdots, X_t)$, such that some X_i is a k-source. A somewhere k-source is a convex combination of elementary somewhere-k-sources.

Definition 2.4. A function $C : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ is a $(k \to l, \epsilon)$-condenser if for every k-source X, $C(X, U_d)$ is ϵ-close to some l-source. When convenient, we call C a rate-$(k/n \to l/m, \epsilon)$-condenser.

Definition 2.5. A function $C : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ is a $(k \to l, \epsilon)$-somewhere-condenser if for every k-source X, the vector $(C(X, y)_{y \in \{0,1\}^d})$ is ϵ-close to a somewhere-l-source. When convenient, we call C a rate-$(k/n \to l/m, \epsilon)$-somewhere-condenser.

Theorem 2.6 ([30])**.** *There exists a constant $\alpha > 0$ such that for any constant $0 < \delta < 0.9$, there is an efficient family of rate-$(\delta \to (1+\alpha)\delta, \epsilon = 2^{-\Omega(n)})$-somewhere condensers $\mathsf{Scond} : \{0,1\}^n \to (\{0,1\}^m)^2$ where $m = \Omega(n)$.*

Theorem 2.7 ([15], [30])**.** *For any constant $\beta, \delta > 0$, there is an efficient family of rate-$(\delta \to 1 - \beta, \epsilon = 2^{-\Omega(n)})$-somewhere condensers $\mathsf{Cond} : \{0,1\}^n \to (\{0,1\}^m)^D$ where $D = O(1)$ and $m = \Omega(n)$.*

Definition 2.8. The *average conditional min-entropy* is defined as

$$\widetilde{H}_\infty(X|W) = -\log\left(\mathrm{E}_{w \leftarrow W}\left[\max_x \Pr[X = x|W = w]\right]\right)$$

2.1. BCH codes

In this paper we will only focus on BCH codes over $\mathbb{F}_2$. Given two parameters $m, t \in \mathbb{N}$, a BCH code is a

linear code with block length $n = 2^m - 1$, message length roughly $n - mt$ and distance $d \geq 2t + 1$. Specifically, we have the following theorem.

Theorem 2.9. *For all integers m and t there exists an explicit $[n, n-mt, 2t+1]$-BCH code[2], with $n = 2^m - 1$.*

Since a BCH code is a linear code, we can take its parity check matrix. Note that this is a $mt \times n$ matrix. Let α be a primitive element in $\mathbb{F}^*_{2^m}$, the i'th column of the parity check matrix is of the form $(\alpha^i, (\alpha^i)^3, (\alpha^i)^5, \cdots, (\alpha^i)^{2t-1})$, for $i = 0, 1, \cdots, n-1$. Since α is a generator in $\mathbb{F}^*_{2^m}$, equivalently, for $y \in \mathbb{F}^*_{2^m}$ we can think of the y'th column to be $(y, y^3, \cdots, y^{2t-1})$.

3. The connection between non-malleable extractors and two-source extractors

In this section we give an overview of the ideas that we use to show the connection between non-malleable extractors and two-source extractors, as well as our improved constructions of non-malleable extractors. The detailed constructions and analysis appear in the full version [28].

3.1. Non-malleable extractors to two-source extractors

Given a (k, ϵ) non-malleable extractor nmExt with seed length $d = 2\log(1/\epsilon) + o(n)$, here is how we can get a two-source extractor. Assume that we have an (n, k) source X and an independent $(n, (1/2-\delta)n)$ source Y for some constant $\delta > 0$. Our first step is to use the 1-bit condenser in [30] to convert Y into two sources $\bar{Y}_1, \bar{Y}_2$ such that each of them has $l = \Omega(n)$ bits and one of them has min-entropy at least $(1/2 + \delta)l$. Note that for an appropriately chosen δ this is indeed possible. Without loss of generality assume that $\bar{Y}_1$ has min-entropy at least $(1/2 + \delta)l$.

Our key observation here is that $\bar{Y}_2$ can now be viewed as a function of $\bar{Y}_1$. More precisely, we show that the source Y is a convex combination of sources $\{Y^i\}$ such that for each Y^i, the corresponding $\bar{Y}_1^i$ also has min-entropy at least $(1/2 + \delta)l$, and $\bar{Y}_2^i$ is a deterministic function of $\bar{Y}_1^i$. Now this looks like the setting of a non-malleable extractor, where we have one seed and another correlated seed. However, there is a small problem: $\bar{Y}_1^i$ and $\bar{Y}_2^i$ may be equal sometimes. To solve this, we let $Y_1 = \bar{Y}_1 \circ 0$ and $Y_2 = \bar{Y}_2 \circ 1$. In this way we guarantee that Y_1^i and Y_2^i are different, and Y_2^i is still a function of Y_1^i. Finally, this only increases the length of the seed by 1.

Now we are all set, and we can take the two-source extractor to be $\mathsf{TExt}(X, Y) = \mathsf{nmExt}(X, Y_1) \oplus \mathsf{nmExt}(X, Y_2)$. Note that the seed Y_1 here is not uniform. However, a simple argument shows that a non-malleable extractor with seed length d and error ϵ remains a non-malleable extractor even if the seed only has min-entropy k', with error increased to $2^{d-k'}\epsilon$. In our case, with seed length $d = l + 1 = \Omega(n) = 2\log(1/\epsilon) + o(n)$ and $k' = (1/2 + \delta)l$, the error is $\epsilon' = 2^{d-k'}\epsilon \approx 2^{(1/2-\delta)l}2^{-l/2} = 2^{-\Omega(n)}$. By the non-malleability of nmExt, we get that $\mathsf{TExt}(X, Y)$ is $2^{-\Omega(n)}$-close to uniform.

Similarly, if we have (r, k, ϵ)-non-malleable extractors for larger r, then we can afford to have more correlated seeds Y_i, or equivalently, more sources in the output of the condenser. Thus we can deal with smaller entropy in Y. For example, for any constant $\delta > 0$, the condensers in [15], [16], [30] allow us to convert an $(n, \delta n)$ source into a constant D number of sources such that each of them has $l = \Omega(n)$ bits and one of them has min-entropy at least $0.9l$. If we have $(D-1, k, \epsilon)$-non-malleable extractors with suitable parameters, then we can get two-source extractors or an $(n, \delta n)$ source and an (n, k) source.

3.2. Two-source extractors to non-malleable extractors

As stated before, we focus on two-source extractors of the form $\mathsf{IP}(f(X), Y)$, where IP is the inner product function. First consider the simplest function $\mathsf{IP}(X, Y)$. Note that it is a good two-source extractor. For two independent sources on n bits, it works as long as the sum of the entropies of the two sources is greater than n. However, at first this function does not seem to be a good candidate for a non-malleable extractor. To see this, let X be a source that is obtained by concatenating the bit 0 with U_{n-1}, and let Y be an independent uniform seed over $\{0,1\}^n$. Now for any $y \in \{0,1\}^n$, let $\mathcal{A}(y)$ be y with the first bit flipped. Thus we see that for all x in the support of X, one has $\langle x, y\rangle = \langle x, \mathcal{A}(y)\rangle$. Therefore, the inner product function is not a non-malleable extractor even for weak sources with min-entropy $k = n - 1$.

In the above example, we have that for all x in the support of X, $\mathsf{IP}(x, y) = \mathsf{IP}(x, \mathcal{A}(y))$. Or equivalently, $\mathsf{IP}(x, y + \mathcal{A}(y)) = 0$. How does this happen? Looking closely at this example, our key observation is that this is because the range of Y is *too large*. Indeed, in this example the range of Y is the entire $\{0,1\}^n$, thus for any y the adversary can choose a

[2]In fact, the message length may not be exactly $n - mt$, but for simplicity we will assume that it is exactly $n - mt$. The small error does not affect our analysis. Also, for small t the message length is exactly $n - mt$.

different $\mathcal{A}(y)$ such that $y + \mathcal{A}(y) = 10 \cdots 0$ so that $\forall x \in \mathsf{Supp}(X), \mathsf{IP}(x, y + \mathcal{A}(y)) = 0$.

This observation suggests that we should choose the range of Y to be a subset $S \subset \{0,1\}^n$, so that for some y's, the adversary will be unable to choose the appropriate $\mathcal{A}(y)$ from S. Equivalently, we take a shorter seed length l, choose a uniform $y \in \{0,1\}^l$ and map y to an element in $\{0,1\}^n$. This is essentially an encoding. Now let us see what properties we need the encoding to have.

We start with a construction for min-entropy $k > n/2$. Assume that we have an (n,k) source X with $k = (1/2 + \delta)n$ for some constant $\delta > 0$. We take an independent and uniform $y \in \{0,1\}^l$ and encode y to $\bar{y} \in \{0,1\}^n$. For any function $\mathcal{A}$, let $\bar{y}'$ be the encoding of $\mathcal{A}(y)$. We will use an injective encoding, so that $\forall y, \bar{y}' \neq \bar{y}$. The output of the non-malleable extractor is then $\mathsf{IP}(X, \bar{Y})$.

To show that $\mathsf{IP}(X, \bar{Y})$ is a non-malleable extractor, it suffices to show that $\mathsf{IP}(X, \bar{Y})$ is close to uniform, and that $\mathsf{IP}(X, \bar{Y}) \oplus \mathsf{IP}(X, \bar{Y}')$ is close to uniform. The first part is easy. If X has min-entropy $k > n/2$, then we can take Y to be the uniform distribution over some $l \geq n/2$ bits. Since the encoding is injective, $\bar{Y}$ will have min-entropy $l \geq n/2$. Thus $\mathsf{IP}(X, \bar{Y})$ is close to uniform. For the second part, note that $\mathsf{IP}(X, \bar{Y}) \oplus \mathsf{IP}(X, \bar{Y}') = \mathsf{IP}(X, \bar{Y} + \bar{Y}')$. Thus now we need $\bar{Y} + \bar{Y}'$ to have large min-entropy, or at least large support size.

The ideal case would be that $\bar{Y} + \bar{Y}'$ also has support size $|S| = 2^l$. This can be achieved if the encoding has the following property: for every two different y_1, y_2, we have that $\bar{y}_1 + \bar{y}'_1 \neq \bar{y}_2 + \bar{y}'_2$, or equivalently, $\bar{y}_1 + \bar{y}'_1 + \bar{y}_2 + \bar{y}'_2 \neq 0$. Indeed, if this is true then $\bar{Y} + \bar{Y}'$ also has min-entropy $l \geq n/2$, and thus $\mathsf{IP}(X, \bar{Y}) \oplus \mathsf{IP}(X, \bar{Y}')$ is close to uniform. Looking carefully at this property, we see that it can be ensured (at least almost ensured, as we will explain shortly) if we have another property: the elements in S (when viewed as vectors in $\mathbb{F}_2^n$) are 4-wise linearly independent. Indeed, assume this is the case and for some y_1, y_2, we have $\bar{y}_1 + \bar{y}'_1 + \bar{y}_2 + \bar{y}'_2 = 0$, then the only possible situation is that $\bar{y}'_1 = \bar{y}_2$ and $\bar{y}'_2 = \bar{y}_1$. Thus there cannot be three different y_1, y_2, y_3 such that $\bar{y}_1 + \bar{y}'_1 = \bar{y}_2 + \bar{y}'_2 = \bar{y}_3 + \bar{y}'_3$. Thus the min-entropy of $\bar{Y} + \bar{Y}'$ is at least $l - 1$.

So now the question reduces to explicitly finding a large subset $S \subset \{0,1\}^n$ such that the elements in S are 4-wise linearly independent. Note that in particular this implies that the sum of any two different pairs of elements in S cannot be the same. Thus we have $\binom{|S|}{2} \leq 2^n$. Therefore $|S|$ can be at most roughly $2^{n/2}$. On the other hand, in order to work for any min-entropy $k > n/2$, we will need $l \geq n/2$ and thus $|S| = 2^l \geq 2^{n/2}$. These are very tight upper and lower bounds. Luckily, we have explicit constructions that meet these bounds. We will think of the elements in S as columns in a parity check matrix of some binary linear code. Thus we basically need a code with block length $2^{n/2}$ and message length $2^{n/2} - n$. The 4-wise linearly independent property basically is equivalent to saying that the code has distance at least 5. This is precisely the $[2^{n/2}, 2^{n/2} - n, 5]$-BCH code. Note that although the parity check matrix has $2^{n/2}$ columns, each column is (a, a^3) for a different element $a \in \mathbb{F}^*_{2^{n/2}}$. Thus the encoding from y to $\bar{y}$ can be computed efficiently.

Thinking about the above encoding for a moment, one realizes that the same encoding can be used in any two-source extractor of the form $\mathsf{IP}(f(X), Y)$. Specifically, assume that $\mathsf{IP}(f(X), Y)$ is a two-source extractor for an (n,k) source X and an independent $(n, n/2 - 1)$ source Y. Then by the same argument above, if we choose the seed Y to be the uniform distribution over $\{0,1\}^{n/2}$ and encode Y to $\bar{Y}$ as before, we have that both $\bar{Y}$ and $\bar{Y} + \bar{Y}'$ have min-entropy at least $n/2 - 1$. Thus both $\mathsf{IP}(f(X), \bar{Y})$ and $\mathsf{IP}(f(X), \bar{Y}) \oplus \mathsf{IP}(f(X), \bar{Y}')$ are close to uniform. Therefore we get a non-malleable extractor for min-entropy k.

By using a BCH code with larger distance, we can extend the above argument to give a similar connection between two-source extractors and (r, k, ϵ)-non-malleable extractors.

3.3. Non-malleable extractors for min-entropy $k < n/2$

We give the first construction of non-malleable extractors for min-entropy $k < n/2$ by observing that the encoding of sources in [17] gives a function f such that $\mathsf{IP}(f(X), Y)$ is a two-source extractor for an $(n, (1/2 - \delta)n)$ source X and an independent (n, k') source Y with $k' \approx n/2$.

Specifically, let X be a distribution over some vector space $\mathbb{F}_q^n$ and let cX be the distribution obtained by sampling $x_1, x_2, \cdots, x_c$ from c independent copies of X and computing $\sum x_i$. One can show that in order to prove $\mathsf{IP}(X, Y)$ is close to uniform, it suffices to prove that $\mathsf{IP}(cX, Y)$ is close to uniform with a smaller error, for some integer $c > 1$. In [17], Bourgain showed that for a weak source X with min-entropy rate $1/2 - \delta$ for some constant $\delta > 0$, one can encode X to $\mathsf{Enc}(X)$ such that $3\mathsf{Enc}(X)$ is close to having min-entropy rate $1/2 + \delta$. Thus $\mathsf{IP}(\mathsf{Enc}(X), Y)$ is a two-source extractor that meets our needs.

3.4. Non-malleable extractors for linear min-entropy

In [29], Ben-Sasson and Zewi showed that affine extractors with large output size can be used to construct two source extractors for min-entropy rate $< 1/2$. Their "preimage construction" can potentially achieve any constant min-entropy rate. We observe that their encoding gives a function f such that $\mathsf{IP}(f(X), Y)$ is a two-source extractor for two independent sources with min-entropy rate δ for any constant $\delta > 0$. Specifically, they showed that the affine extractor gives an injective mapping $F : \{0,1\}^n \to \{0,1\}^{n'}$ such that for any weak source X with min-entropy δn, $F(\mathsf{Supp}(X))$ is not contained in any affine subspace of dimension say $(1-\delta/2)n'$. Thus when Y is a $(n', \delta n')$ source, we have that $\mathsf{IP}(F(X), Y)$ is non-constant. Next, similar as in [29], the ADC conjecture implies that in fact $\mathsf{IP}(F(X), Y)$ is close to uniform. Thus $\mathsf{IP}(F(X), Y)$ is a two-source extractor that meets our needs.

4. An optimal privacy amplification protocol for $k = \delta n$

Following [25] and [2], we define a privacy amplification protocol (P_A, P_B). The protocol is executed by two parties Alice and Bob, who share a secret $X \in \{0,1\}^n$. An active, computationally unbounded adversary Eve might have some partial information E about X satisfying $\widetilde{H}_\infty(X|E) \geqslant k$.

We assume that Eve has full control of the communication channel between the two parties. Alice and Bob are assumed to have fresh, private and independent random bits Y and W, respectively. In the protocol we use $\perp$ as a special symbol to indicate rejection. At the end of the protocol, Alice outputs a key $R_A \in \{0,1\}^m \cup \{\perp\}$. Similarly, Bob outputs a key $R_B \in \{0,1\}^m \cup \{\perp\}$. We let E' denote the final view of Eve, which includes E and the communication transcripts of the protocol.

Definition 4.1. An interactive protocol (P_A, P_B) is a (k, m, ϵ)-*privacy amplification protocol* if it satisfies the following properties whenever $\widetilde{H}_\infty(X|E) \geq k$:

1) Correctness. If Eve is passive, then $\Pr[R_A = R_B \wedge\ R_A \neq \perp \wedge R_B \neq \perp] = 1$.
2) Robustness. If Eve is active, $\Pr[R_A \neq R_B \wedge\ R_A \neq \perp \wedge R_B \neq \perp] \leqslant \epsilon$.
3) Extraction. Given a string $r \in \{0,1\}^m \cup \{\perp\}$, let $\mathsf{purify}(r)$ be $\perp$ if $r = \perp$, and otherwise replace $r \neq \perp$ by a fresh m-bit random string U_m: $\mathsf{purify}(r) \leftarrow U_m$. We require that

$$\Delta((R_A, E'), (\mathsf{purify}(R_A), E')) \leq \epsilon \quad \text{and}$$
$$\Delta((R_B, E'), (\mathsf{purify}(R_B), E')) \leq \epsilon$$

Namely, whenever a party does not reject, its key looks like a fresh random string to Eve.

Here $k - m$ is called the *entropy loss* and $\log(1/\epsilon)$ is called the *security parameter* of the protocol.

4.1. Prerequisites from previous work

Definition 4.2. A function family $\{\mathsf{MAC}_R : \{0,1\}^d \to \{0,1\}^v\}$ is a ϵ-secure one-time MAC for messages of length d with tags of length v if given a uniform R over $\{0,1\}^\ell$, for any $w \in \{0,1\}^d$ and any function (adversary) $A : \{0,1\}^v \to \{0,1\}^d \times \{0,1\}^v$,

$$\Pr_R[\mathsf{MAC}_R(W') = T' \wedge\ W' \neq w \mid (W', T') = A(\mathsf{MAC}_R(w))] \leq \epsilon,$$

Theorem 4.3 ([25])**.** *For any message length d and tag length v, there exists an efficient family of $(\lceil \frac{d}{v} \rceil 2^{-v})$-secure* MAC*s with key length $\ell = 2v$. In particular, this* MAC *is ε-secure when $v = \log d + \log(1/\epsilon)$.*
More generally, this MAC *also enjoys the following security guarantee, even if Eve has partial information E about its key R. Let (R, E) be any joint distribution. Then, for all attackers A_1 and A_2,*

$$\Pr_{(R,E)}[\mathsf{MAC}_R(W') = T' \wedge W' \neq W \mid W = A_1(E), (W', T') = A_2(\mathsf{MAC}_R(W), E)] \leq \left\lceil \frac{d}{v} \right\rceil 2^{v - \widetilde{H}_\infty(R|E)}.$$

Remark 4.4. Note that this MAC works as long as the key R has average conditional min-entropy rate $> 1/2$.

Theorem 4.5 ([6])**.** *For every constant $\alpha > 0$, all $n, k \in \mathbb{N}$ and any $\epsilon > 0$, there is an explicit strong (k, ϵ)-extractor $\mathsf{Ext} : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ with $d = O(\log n + \log(1/\epsilon))$ and $m \geq (1-\alpha)k$.*

Theorem 4.6 ([16])**.** *For any n_1, n_2, k_1, k_2, m and any $0 < \delta < 1/2$ with*

- $n_1 \geq 6 \log n_1 + 2 \log n_2$
- $k_1 \geq (0.5 + \delta) n_1 + 3 \log n_1 + \log n_2$
- $k_2 \geq 5 \log(n_1 - k_1)$
- $m \leq \delta \min[n_1/8, k_2/40] - 1$

There is a polynomial time computable strong 2-source extractor $\mathsf{Raz} : \{0,1\}^{n_1} \times \{0,1\}^{n_2} \to \{0,1\}^m$ for min-entropy k_1, k_2 with error $2^{-1.5m}$.

Theorem 4.7. *[2], [3], [27] For every constant $\delta > 0$, there exists a constant $\beta > 0$ such that for every $n, k \in \mathbb{N}$ with $k \geq (1/2 + \delta)n$ and $\epsilon > 2^{-\beta n}$ there exists an explicit (k, ϵ) non-malleable extractor with seed length $d = O(\log n + \log \epsilon^{-1})$ and output length $m = \Omega(n)$.*

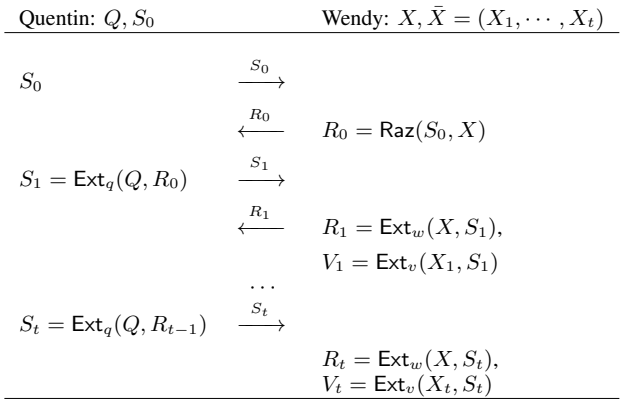

Figure 1. Alternating Extraction.

4.2. The privacy amplification protocol

We first define an alternating extraction protocol.

Alternating Extraction. Assume that we have two parties. Quentin has a source Q and a source S_0 with entropy rate $> 1/2$. Wendy has a source X and a source $\bar{X} = (X_1 \circ \cdots \circ X_t)$. Suppose that (Q, S_0) is kept secret from Wendy and $(X, \bar{X})$ is kept secret from Quentin. Let s, d be two parameters for the protocol. Let Ext_q, Ext_w, Ext_v be seeded extractors as in Theorem 4.5. Let Raz be the two-source extractor in Theorem 4.6. The *alternating extraction protocol* is an interactive process between Quentin and Wendy that runs in $t+1$ steps.

In the 0'th step, Quentin sends S_0 to Wendy, Wendy computes $R_0 = \mathsf{Raz}(S_0, X)$ and replies R_0 to Quentin, Quentin then computes $S_1 = \mathsf{Ext}_q(Q, R_0)$. In this step R_0, S_1 each outputs d bits. In the first step, Quentin sends S_1 to Wendy, Wendy computes $V_1 = \mathsf{Ext}_v(X_1, S_1)$ and $R_1 = \mathsf{Ext}_w(X, S_1)$. She sends R_1 to Quentin and Quentin computes $S_2 = \mathsf{Ext}_q(Q, R_1)$. In this step V_1 outputs $2^{t-1}s$ bits, and R_1, S_2 each outputs d bits. In each subsequent step i, Quentin sends S_i to Wendy, Wendy computes $V_i = \mathsf{Ext}_v(X_i, S_i)$ and $R_i = \mathsf{Ext}_w(X, S_i)$. She replies R_i to Quentin and Quentin computes $S_{i+1} = \mathsf{Ext}_q(Q, R_i)$. In step i, V_i outputs $2^{t-i}s$ bits, and R_i, S_{i+1} each outputs d bits. Thus, the process produces the following sequence:

$$S_0, R_0 = \mathsf{Raz}(S_0, X), S_1 = \mathsf{Ext}_q(Q, R_0),$$
$$V_1 = \mathsf{Ext}_v(X_1, S_1), R_1 = \mathsf{Ext}_w(X, S_1), \cdots,$$
$$S_t = \mathsf{Ext}_q(Q, R_{t-1}), V_t = \mathsf{Ext}_v(X_t, S_t), R_t = \mathsf{Ext}_w(X, S_t).$$

Look-Ahead Extractor. Let $Y = (Q, S_0)$ be a seed, the look-ahead extractor is defined as

$$\mathsf{laExt}((X, \bar{X}), Y) \stackrel{def}{=} V_1, \cdots, V_t.$$

4.2.1. The protocol: We assume that the error ϵ we seek satisfies $2^{-\Omega(\delta n)} < \epsilon < 1/n$. Let s be a parameter with $s = \log(C'/\epsilon) + O(1)$, so that $O(C')/2^s < \epsilon$, for a sufficiently large $O(C')$ constant related to the number of "bad" events. Let $d = O(\log n + s)$ be the seed length of a seeded extractor as in Theorem 4.5, with error 2^{-2s}. We will need the following building blocks:

- Let $\mathsf{Cond} : \{0,1\}^n \to (\{0,1\}^{n'})^C$ be a rate-$(\delta \to 0.9, 2^{-s})$-somewhere-condenser as in Theorem 2.7.
- Let $\mathsf{nmExt} : \{0,1\}^{n'} \times \{0,1\}^{d'} \to \{0,1\}^{m'}$ be a $(0.8n', 2^{-s})$-non-malleable extractor as in Theorem 4.7 with output length $m' = 6 \cdot 2^C s$.
- Let laExt be the look-ahead extractor defined above, with parameters $(2s, d)$.
- Let lrMAC be the MAC as in Theorem 4.3 for d-bit messages, with tag length $2^C(3s)$.

Using the above building blocks, the protocol is given in Figure 2. To emphasize the presence of Eve, we will use 'prime' to denote all the protocol values seen or generated by Bob; e.g., Bob picks W', but Alice sees potentially different W, etc.

The high-level idea of the protocol is as follows. We first use the condenser in [15], [16], [30] to convert the shared (n, k) source X into a somewhere rate-0.9 source $(X_1, \cdots, X_C)$ with $C = \mathrm{poly}(1/\delta)$ rows. In the first round, Alice samples a fresh random string Y_1 from her private random bits and sends it to Bob, where Bob receives a possibly modified version Y_1'. In the second round, Bob samples a fresh random string W' from his private random bits and sends it to Alice, where Alice receives a possibly modified version W. We want a protocol such that if Eve does not change Y_1, then with high probability Bob can authenticate W' to Alice and they can both output $\mathsf{Ext}(X, W')$ as the final outputs, by using a strong seeded extractor Ext. If Eve does change Y_1, then with high probability Alice should be able to detect this and reject.

The first goal is relatively easy to achieve. At the end of the first round, Alice and Bob compute $Z = \mathsf{Ext}(X, Y_1)$ and $Z' = \mathsf{Ext}(X, Y_1')$ respectively, using a strong extractor Ext. If Eve does not change Y_1 then $Z = Z'$ and is private and uniform. Thus in the second round Bob can authenticate W' to Alice by also sending a tag T' produced by a standard MAC with Z as the key. We now focus on the second goal. If the extractor Ext in computing Z and Z' is non-malleable then this can be done by using the protocol in [1]. However, we do not have explicit non-malleable extractors for $k = \delta n$.

Instead, we will have Alice and Bob each produce another variable V and V' respectively. We will ensure that, if Eve changes Y_1 to a different Y_1', then even given

Alice: X	Eve: E	Bob: X
$(X_1, \dots X_C) = \mathsf{Cond}(X)$.		$(X_1, \dots X_C) = \mathsf{Cond}(X)$.
Sample random $Y = (Y_1, Y_2, Y_3)$ $\lvert Y_1\rvert = d, \lvert Y_3\rvert = 30d + 3s$, $\lvert Y_2\rvert = 4Cd + 31d + 4s$		
	$(Y_1, Y_2, Y_3) \to (Y_1', Y_2', Y_3')$	
		Sample random W' with d bits. $Z' = \mathsf{Ext}(X; Y_1')$ with $2^C(6s)$ bits. $\bar{X}' = (\bar{X}_1', \dots, \bar{X}_C')$, where $\bar{X}_i' = \mathsf{nmExt}(X_i, Y_1')$. $V' = (V_1', \dots, V_C') = \mathsf{laExt}((X, \bar{X}'), (Y_2', Y_3'))$ $T' = \mathsf{lrMAC}_{Z'}(W')$. Set final $R_B = \mathsf{Ext}(X; W')$.
	$(W, T, \bar{V}) \leftarrow (W', T', V')$	
$Z = \mathsf{Ext}(X; Y_1)$ with $2^C(6s)$ bits. $\bar{X} = (\bar{X}_1, \dots, \bar{X}_C)$, where $\bar{X}_i = \mathsf{nmExt}(X_i, Y_1)$. $V = (V_1, \dots, V_C) = \mathsf{laExt}((X, \bar{X}), (Y_2, Y_3))$ **If** $T \neq \mathsf{lrMAC}_Z(W)$ or $V \neq \bar{V}$ *reject.* Set final $R_A = \mathsf{Ext}(X; W)$.		

Figure 2. 2-round Privacy Amplification Protocol for $\widetilde{H}_\infty(X|E) \geq \delta n$.

T' and V', with high probability Eve cannot come up with the correct V for Alice. If this is true then in the second round we can have Bob also send V' to Alice, where Alice receives a possibly modified version $\bar{V}$. Alice then checks both the tag T and whether $V = \bar{V}$. If either of them fails, Alice rejects. This will give us a privacy amplification protocol.

The first problem with the above strategy is that now V' may leak information about Z', thus now the MAC key may not be uniform. This is easy to solve since the MAC in Theorem 4.3 works as long as the key has entropy rate $> 1/2$. Thus by limiting the size of V' to be at most half the size of Z', we can ensure that if Eve does not change Y_1, Bob can still authenticate W' to Alice. We now explain how we produce V, V'.

We actually have Alice produce C variables $V = (V_1, \cdots, V_C)$. Similarly, Bob produces $V' = (V_1', \cdots, V_C')$. For this, we first choose a non-malleable extractor and have Alice and Bob each apply the extractor to the somewhere rate-0.9 source $(X_1, \cdots, X_C)$, using Y_1 and Y_1' as the seeds respectively. Let the outputs be $\bar{X} = (\bar{X}_1, \cdots, \bar{X}_C)$ and $\bar{X}' = (\bar{X}_1', \cdots, \bar{X}_C')$. Note that one of the X_i's, say X_g is a rate 0.9-source. Thus we can use the non-malleable extractors in [2], [3], [27]. Now we fix Y_1, Y_1', and we have that $\bar{X}_g$ is uniform and independent of $\bar{X}_g'$. Thus we can fix $\bar{X}_g'$ and $\bar{X}_g$ is still uniform. Next, we fix Z'. Since now Z' is a deterministic function of X, as long as the size of Z' is smaller than the size of $\bar{X}_g$, conditioned on this fixing $\bar{X}_g$ still has a lot of entropy left. Now to produce V, V', in the first round we also have Alice sample two other random strings (Y_2, Y_3) and send them to Bob, where Bob receives (Y_2', Y_3'). Note that after we fix (Y_1, Y_1'), (Y_2', Y_3') is a deterministic function of (Y_2, Y_3). We will now have Alice apply the look-ahead extractor to $(X, \bar{X})$, using (Y_2, Y_3) as the seed, and output $V = (V_1, \cdots, V_C)$. Similarly, Bob applies the look-ahead extractor to $(X, \bar{X}')$, using (Y_2', Y_3') as the seed, and output $V' = (V_1', \cdots, V_C')$. Note that we can indeed ensure that the size of $\bar{X}_g$ is bigger than Z', while the size of Z' is bigger than the size of $(V_1', \cdots, V_C')$ just by limiting the size of each V_i'.

Using properties of the alternating extraction protocol, we can show that V_g is close to uniform conditioned on $(V_1', \cdots, V_g')$ and Z'. Now, we can limit the size of $(V_{g+1}', \cdots, V_C')$ to be smaller than the size of V_g. Thus V_g still has a lot of entropy even conditioned on $V' = (V_1', \cdots, V_C')$. This will ensure that with high probability Eve cannot come up with the correct V_g. Since we do not know which one of $\{\bar{X}_i\}$ is $\bar{X}_g$, we will choose $(V_1, \cdots, V_C)$ such that the size of V_C is say $2s$, and for any i the size of V_i is twice the size of V_{i+1}. In this way, no matter what g is, the size of $(V_{g+1}', \cdots, V_C')$ is the size of V_g minus $2s$. Thus V_g still has $2s$ entropy left conditioned on V'.

This gives our whole protocol. Note that the entropy loss and communication complexity is $O(2^C s) = 2^{\mathrm{poly}(1/\delta)} s = O(s)$ for any constant $\delta > 0$.

Theorem 4.8. *The above protocol is a 2-round $(\delta n, \epsilon)$ privacy amplification protocol with entropy loss and communication complexity $2^{\mathrm{poly}(1/\delta)} \log(1/\epsilon)$.*

The detailed proof appears in the full version [28].

5. Open Problems

Our result shows a connection between non-malleable extractors and two-source extractors, and suggests that it may be hard to construct non-malleable extractors for small entropy with short seed length. However, it is still quite possible that we can get explicit non-malleable extractors for small entropy with large seed length. Moreover, the weaker notion of non-malleable condensers introduced in [27] is a hopeful alternative.

In our optimal privacy amplification protocol for $k = \delta n$, the entropy loss is $2^{\mathrm{poly}(1/\delta)}s$, which has a large hidden constant for small δ. As a comparison, the protocol in [2] runs in $\mathrm{poly}(1/\delta)$ rounds but only has entropy loss $\mathrm{poly}(1/\delta)s$. Thus for practical purposes it is interesting to see if we can reduce the hidden constant. In particular, it remains an interesting open problem to construct non-malleable extractors or non-malleable condensers for arbitrarily linear min-entropy.

Acknowledgments

We are grateful to David Zuckerman for many valuable discussions, especially for his suggestion to use the BCH code.

References

[1] Y. Dodis and D. Wichs, "Non-malleable extractors and symmetric key cryptography from weak secrets," in *STOC*, 2009, pp. 601–610.

[2] Y. Dodis, X. Li, T. D. Wooley, and D. Zuckerman, "Privacy amplification and non-malleable extractors via character sums," in *FOCS*, 2011.

[3] G. Cohen, R. Raz, and G. Segev, "Non-malleable extractors with short seeds and applications to privacy amplification," in *CCC*, 2012.

[4] L. Trevisan, "Extractors and pseudorandom generators," *Journal of the ACM*, pp. 860–879, 2001.

[5] R. Raz, O. Reingold, and S. Vadhan, "Extracting all the randomness and reducing the error in trevisan's extractors," *JCSS*, vol. 65, no. 1, pp. 97–128, 2002.

[6] V. Guruswami, C. Umans, and S. Vadhan, "Unbalanced expanders and randomness extractors from Parvaresh-Vardy codes," *Journal of the ACM*, vol. 56, no. 4, 2009.

[7] Z. Dvir and A. Wigderson, "Kakeya sets, new mergers and old extractors," in *FOCS*, 2008.

[8] Z. Dvir, S. Kopparty, S. Saraf, and M. Sudan, "Extensions to the method of multiplicities, with applications to kakeya sets and mergers," in *FOCS*, 2009.

[9] N. Nisan and D. Zuckerman, "Randomness is linear in space," *JCSS*, vol. 52, no. 1, pp. 43–52, 1996.

[10] L. Fortnow and R. Shaltiel, "Recent developments in explicit constructions of extractors," 2002.

[11] S. Vadhan, "Randomness extractors and their many guises: Invited tutorial," in *FOCS*, 2002.

[12] R. Shaltiel, "An introduction to randomness extractors," in *ICALP*, 2011.

[13] B. Chor and O. Goldreich, "Unbiased bits from sources of weak randomness and probabilistic communication complexity," *SIAM Journal on Computing*, vol. 17, no. 2, pp. 230–261, 1988.

[14] B. Barak, R. Impagliazzo, and A. Wigderson, "Extracting randomness using few independent sources," in *FOCS*, 2004, pp. 384–393.

[15] B. Barak, G. Kindler, R. Shaltiel, B. Sudakov, and A. Wigderson, "Simulating independence: New constructions of condensers, Ramsey graphs, dispersers, and extractors," in *STOC*, 2005, pp. 1–10.

[16] R. Raz, "Extractors with weak random seeds," in *STOC*, 2005, pp. 11–20.

[17] J. Bourgain, "More on the sum-product phenomenon in prime fields and its applications," *International Journal of Number Theory*, vol. 1, pp. 1–32, 2005.

[18] A. Rao, "Extractors for a constant number of polynomially small min-entropy independent sources," in *STOC*, 2006.

[19] B. Barak, A. Rao, R. Shaltiel, and A. Wigderson, "2 source dispersers for $n^{o(1)}$ entropy and Ramsey graphs beating the Frankl-Wilson construction," in *STOC*, 2006.

[20] X. Li, "Improved constructions of three source extractors," in *CCC*, 2011.

[21] C. Bennett, G. Brassard, and J.-M. Robert, "Privacy amplification by public discussion," *SIAM Journal on Computing*, vol. 17, pp. 210–229, 1988.

[22] U. M. Maurer and S. Wolf, "Privacy amplification secure against active adversaries," in *CRYPTO '97*, 1997.

[23] Y. Dodis, J. Katz, L. Reyzin, and A. Smith, "Robust fuzzy extractors and authenticated key agreement from close secrets," in *CRYPTO*, 2006, pp. 232–250.

[24] R. Renner and S. Wolf, "Unconditional authenticity and privacy from an arbitrarily weak secret," in *CRYPTO*, 2003, pp. 78–95.

[25] B. Kanukurthi and L. Reyzin, "Key agreement from close secrets over unsecured channels," in *EUROCRYPT*, 2009, pp. 206–223.

[26] N. Chandran, B. Kanukurthi, R. Ostrovsky, and L. Reyzin, "Privacy amplification with asymptotically optimal entropy loss," in *STOC*, 2010, pp. 785–794.

[27] X. Li, "Design extractors, non-malleable condensers and privacy amplification," in *STOC*, 2012.

[28] ——, "Non-malleable extractors, two-source extractors and privacy amplification," ECCC, Tech. Rep. 11-166.

[29] E. Ben-Sasson and N. Zewi, "From affine to two-source extractors via approximate duality," in *STOC*, 2011.

[30] D. Zuckerman, "Linear degree extractors and the inapproximability of max clique and chromatic number," in *Theory of Computing*, 2007, pp. 103–128.