Non-malleable Condensers for Arbitrary Min-entropy, and Almost Optimal Protocols for Privacy Amplification

Xin Li*

Department of Computer Science
Johns Hopkins University
Baltimore, MD 21218, USA
lixints@cs.jhu.edu

Abstract. Recently, the problem of privacy amplification with an active adversary has received a lot of attention. Given a shared n-bit weak random source X with min-entropy k and a security parameter s, the main goal is to construct an explicit 2-round privacy amplification protocol that achieves entropy loss $O(s)$. Dodis and Wichs [1] showed that optimal protocols can be achieved by constructing explicit *non-malleable extractors*. However, the best known explicit non-malleable extractor only achieves $k = 0.49n$ [2] and evidence in [2] suggests that constructing explicit non-malleable extractors for smaller min-entropy may be hard. In an alternative approach, Li [3] introduced the notion of a non-malleable condenser and showed that explicit non-malleable condensers also give optimal privacy amplification protocols.

In this paper, we give the first construction of non-malleable condensers for arbitrary min-entropy. Using our construction, we obtain a 2-round privacy amplification protocol with optimal entropy loss for security parameter up to $s = \Omega(\sqrt{k})$. This is the first protocol that simultaneously achieves optimal round complexity and optimal entropy loss for arbitrary min-entropy k. We also generalize this result to obtain a protocol that runs in $O(s/\sqrt{k})$ rounds with optimal entropy loss, for security parameter up to $s = \Omega(k)$. This significantly improves the protocol in [4]. Finally, we give a better non-malleable condenser for linear min-entropy, and in this case obtain a 2-round protocol with optimal entropy loss for security parameter up to $s = \Omega(k)$, which improves the entropy loss and communication complexity of the protocol in [2].

Keywords: privacy amplification, non-malleable, extractor, condenser.

1 Introduction

Modern cryptographic applications rely heavily on the use of randomness. Indeed, true randomness are provably necessary and key ingredients in even basic

* Most work was done while the author was a Simons postdoctoral fellow at University of Washington.

Y. Dodis and J.B. Nielsen (Eds.): TCC 2015, Part I, LNCS 9014, pp. 502–531, 2015.

tasks such as bit commitment and encryption. However, most of these applications require uniform random bits, yet real world random sources are rarely uniformly distributed. In addition, even initially uniform secret keys could be damaged by side channel attacks of an adversary. Naturally, the random sources we can use become imperfect, and it is therefore important to study how to run cryptographic applications using imperfect randomness. In [5], Dodis et. al showed that even slightly imperfect random sources cannot be used directly in many important cryptographic applications, thus we have to find a way to convert the imperfect random sources into nearly uniform random bits first.

In this general context, Bennett, Brassard, and Robert [6] introduced the basic cryptographic question of *privacy amplification*. The setting is as follows. Consider the simple model where two parties (Alice and Bob) share an n-bit secret key X, which is weakly random. They also have access to local (non-shared) uniform private random bits and share a public channel which is monitored by an adversary Eve. The goal now is for Alice and Bob to communicate over the channel to transform X into a nearly uniform secret key, so that Eve has negligible information about it. To measure the randomness in X, we use the standard min-entropy.

Definition 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 $H_\infty(X)/n$.

This problem arises naturally in several situations when two parties want to communicate with each other secretly (e.g., one-time pad). We note that shared randomness is an important resource and is often harder to obtain than local randomness. More importantly the quality of shared randomness generally may be much weaker than local randomness, thus it makes sense in the privacy amplification problem to assume that the parties have local uniform random bits and try to boost the quality of the shared weak random source.

Following [6], we assume the adversary Eve has unlimited computational power. If Eve is passive (i.e., can only see the messages but cannot change them), then this problem can be solved by using a well-studied combinatorial object called "strong 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.

Definition 2 (statistical distance). *Let W and Z be two distributions on a set S. Their* statistical distance *(variation distance) is*

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

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

Definition 3. *A function* $\mathsf{Ext} : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ *is a* strong (k,ε)-extractor *if for every source X with min-entropy k and independent Y which is uniform on $\{0,1\}^d$,*

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

Suppose we have a strong extractor Ext, we can then have Alice sample a fresh random string Y from her local random bits and send it to Bob. They then both compute $R = \mathsf{Ext}(X,Y)$. Since Eve only sees Y, the property of the strong extractor guarantees that the output is close to uniform even given this information.

However, if Eve is active (i.e., can arbitrarily change, delete and reorder messages), then the problem becomes much harder and the above simple solution fails. In this case, while one can show the task is still possible, the main goal is to try to use as few rounds as possible, and achieve a secret nearly uniform random string R that has length as close to $H_\infty(X)$ as possible. There has been a lot of effort in trying to achieve optimal parameters [7,8,1,9,10,4,11,12,3,2]. More specifically, [7] gave the first non-trivial protocol which takes one-round and works when the entropy rate of X is bigger than 2/3. [8] later improved this to work for entropy rate bigger than 1/2, yet both these results suffer from the drawback that the final secret key R is significantly shorter than the min-entropy of X. [1] showed that it is impossible to construct one-round protocol for if the entropy rate of X is less than 1/2. Moreover, one can show that the final output R has to be at least $O(s)$ shorter than $H_\infty(X)$, where s is the security parameter of the protocol (A protocol has security parameter s if Eve cannot predict with advantage more than 2^{-s} over random. When Eve is active, we also require that Eve cannot make Alice and Bob output different secrets and not abort with probability more than 2^{-s}.). This difference is call the *entropy loss* of the protocol. Thus in general the optimal protocol should take 2 rounds and have entropy loss $O(s)$.

The first protocol which works for entropy rate below 1/2 appeared in [9], which was simplified by [10] and shown to run in $O(s)$ rounds and achieve entropy loss $O(s^2)$. [1] improved the number of rounds to 2 but the entropy loss remains $O(s^2)$. [4] improved the entropy loss to $O(s)$ but the number of rounds increases to $O(s)$. The natural open question is therefore whether there is an explicit 2-round protocol with entropy loss $O(s)$. In the special case where X has entropy rate bigger than 1/2, [11,12,3] gave 2-round protocols with entropy loss $O(s)$. For any constant $0 < \delta < 1$, [11] also gave a protocol for the case where X has entropy rate δ, which runs in $\mathsf{poly}(1/\delta)$ rounds with entropy loss $\mathsf{poly}(1/\delta)s = O(s)$. Recently, [2] gave an improved protocol for the case of entropy rate δ, which runs in 2 rounds and achieves optimal entropy loss $2^{\mathsf{poly}(1/\delta)}s = O(s)$, although the hidden constant can be quite large.

In [1], Dodis and Wichs introduced the notion of a "non-malleable extractor" and showed that such an object can be used to construct 2-round privacy amplification protocols with optimal entropy loss.

Definition 4. [1] *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 source* X *with* $H_\infty(X) \geq k$ *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* Y *is chosen uniformly from* $\{0,1\}^d$ *and independent of* X*,*

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

Dodis and Wichs 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. The first explicit construction of non-malleable extractors appeared in [11], which works for entropy $k > n/2$. Later, various improvements appeared in [12,3,13]. However, the entropy requirement remains $k > n/2$. Recently, Li [2] gave the first explicit non-malleable extractor that breaks this barrier, which works for $k = (1/2 - \delta)n$ for some constant $\delta > 0$. [2] also showed a connection between non-malleable extractors and two-source extractors, which suggests that constructing explicit non-malleable extractors for smaller entropy may be hard.

Given the above background, an alternative approach seems promising. This is the notion of a non-malleable condenser introduced in [3]. While a non-malleable extractor requires the output to be close to uniform, a non-malleable condenser only requires the output to have enough min-entropy.

Definition 5. *[2] A* (k, k', ϵ) *non-malleable condenser is a function* $\mathsf{nmCond} : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ *such that given any* (n,k)*-source* X*, an independent uniform seed* $Y \in \{0,1\}^d$*, and any (deterministic) function* $\mathcal{A} : \{0,1\}^d \to \{0,1\}^d$ *such that* $\forall y, \mathcal{A}(y) \neq y$*, we have that with probability* $1 - \epsilon$ *over the fixing of* $Y = y$*,*

$$\Pr_{z' \leftarrow \mathsf{nmCond}(X,\mathcal{A}(y))}[\mathsf{nmCond}(X,y)|_{\mathsf{nmCond}(X,\mathcal{A}(y))=z'} \text{ is } \epsilon\text{−close to an } (m,k') \text{ source}] \geq 1-\epsilon.$$

As can be seen from the definition, a non-malleable condenser is a strict relaxation of a non-malleable extractor and thus it may be easier to construct. In [3], Li showed that non-malleable condensers can also be used to construct 2-round privacy amplification protocols with optimal entropy loss. Thus to give optimal privacy amplification protocols for smaller min-entropy, one can hope to first construct explicit non-malleable condensers for smaller min-entropy.

1.1 Our Results

In this paper, we indeed succeed in the above approach. We construct explicit non-malleable condensers for essentially any min-entropy. Our first theorem is as follows.

[1] Following [11], we define worst case non-malleable extractors, which is slightly different from the original definition of average case non-malleable extractors in [1]. However, the two definitions are essentially equivalent up to a small change of parameters.

Theorem 1. *There exists a constant $C > 0$ such that for any $n, k \in \mathsf{N}$ and $s > 0$ with $k \geq C(\log n + s)^2$, there is an explicit $(k, s, 2^{-s})$-non-malleable condenser with seed length*
$d = O(\log n + s)^2$ and output length $m = O(\log n + s)^2$.

Combining this theorem with the protocol in [3], we immediately obtain a 2-round privacy amplification protocol with optimal entropy loss for any security parameter up to $\Omega(\sqrt{k})$. This is the first explicit protocol that simultaneously achieves optimal parameters in both round complexity and entropy loss, for arbitrary min-entropy.

Theorem 2. *There exists a constant C such that for any $\epsilon > 0$ with $k \geq C(\log n + \log(1/\epsilon))^2$, there exists an explicit 2-round privacy amplification protocol for (n, k) sources with security parameter $\log(1/\epsilon)$, entropy loss $O(\log n + \log(1/\epsilon))$ and communication complexity $O(\log n + \log(1/\epsilon))^2$.*

We note that except the protocol in [4], all previous results that work for arbitrary min-entropy k only achieve security parameter up to $s = \Omega(\sqrt{k})$ like our protocol and all of them have entropy loss $\Omega(s^2)$. In this paper, we finally manage to reduce the entropy loss to $O(s)$. Thus, for this range of security parameter, ignoring the communication complexity, we essentially obtain optimal privacy amplification protocols.

For the special case where $k = \delta n$ for some constant $0 < \delta < 1$, we can do better. Here we have the following theorem.

Theorem 3. *For any constant $0 < \delta < 1$ and $k = \delta n$ there exists a constant $C = 2^{\mathsf{poly}(1/\delta)}$ such that given any $0 < s \leq k/C$, there is an explicit $(k, s, 2^{-s})$-non-malleable condenser with seed length $d = \mathsf{poly}(1/\delta)(\log n + s)$ and output length $m = 2^{\mathsf{poly}(1/\delta)}(\log n + s)$.*

Combined with the protocol in [3], this theorem yields:

Theorem 4. *There exists an absolute constant $C_0 > 1$ such that for any constant $0 < \delta < 1$ and $k = \delta n$ there exists a constant $C_1 = 2^{\mathsf{poly}(1/\delta)}$ such that given any $\epsilon > 0$ with $C_1 \log(1/\epsilon) \leq k$, there exists an explicit 2-round privacy amplification protocol for (n, k) sources with security parameter $\log(1/\epsilon)$, entropy loss $C_0(\log n + \log(1/\epsilon))$ and communication complexity $\mathsf{poly}(1/\delta)(\log n + \log(1/\epsilon))$.*

Note that for security parameter s, the 2-round protocol for $k = \delta n$ in [2] has entropy loss $2^{\mathsf{poly}(1/\delta)}s$ and communication complexity $2^{\mathsf{poly}(1/\delta)}s$. Here, we improve the entropy loss to $C_0 s$ for an absolute constant $C_0 > 1$ and the communication complexity to $\mathsf{poly}(1/\delta)s$.

Finally, one can ask what if for arbitrary min-entropy k, we want to achieve security parameter bigger than $\sqrt{k}$, as in [4]. Using our techniques combined with some techniques from [4], we obtain the following theorem.

Theorem 5. *There exists a constant $C > 1$ such that for any $n, k \in \mathsf{N}$ with $k \geq \log^4 n$ and any $\epsilon > 0$ with $k \geq C(\log(1/\epsilon))$ there exists an explicit $O((\log n + \log(1/\epsilon))/\sqrt{k})$ round privacy amplification protocol for (n, k) sources with security parameter $\log(1/\epsilon)$, entropy loss $O(\log n + \log(1/\epsilon))$ and communication complexity $O((\log n + \log(1/\epsilon))\sqrt{k})$.*

Thus, we can essentially achieve security parameter up to $s = \Omega(k)$ with optimal entropy loss, at the price of increasing the number of rounds to $O(s/\sqrt{k})$. Note that the protocol in [4], though also achieving optimal entropy loss, runs in $\Omega(s)$ rounds. Thus our protocol improves their round complexity by a $\sqrt{k}$ factor. For large k this is a huge improvement, especially in practice.

Table 1 summarizes our results compared to some previous results, assuming the security parameter is s.

Table 1. Summary of Results on Privacy Amplification with an Active Adversary

Construction	Entropy of X	Security parameter	Rounds	Entropy loss
Optimal non-explicit	$k > \log n$	$s \leq \Omega(k)$	2	$\Theta(s + \log n)$
[7]	$k > 2n/3$	$s = \Theta(k)$	1	$(n - k)$
[8]	$k > n/2$	$s = \Theta(k)$	1	$(n - k)$
[9,10]	$k \geq \mathsf{polylog}(n)$	$s \leq \Omega(\sqrt{k})$	$\Theta(s + \log n)$	$\Theta((s + \log n)^2)$
[1]	$k \geq \mathsf{polylog}(n)$	$s \leq \Omega(\sqrt{k})$	2	$\Theta((s + \log n)^2)$
[4]	$k \geq \mathsf{polylog}(n)$	$s \leq \Omega(k)$	$\Theta(s + \log n)$	$\Theta(s + \log n)$
[11]	$k \geq \delta n$	$s \leq k/\mathsf{poly}(1/\delta)$	$\mathsf{poly}(1/\delta)$	$\mathsf{poly}(1/\delta)(s + \log n)$
[2]	$k \geq \delta n$	$s \leq k/2^{\mathsf{poly}(1/\delta)}$	2	$2^{\mathsf{poly}(1/\delta)}(s + \log n)$
This work	$k \geq \mathsf{polylog}(n)$	$s \leq \Omega(\sqrt{k})$	2	$\Theta(s + \log n)$
This work	$k \geq \mathsf{polylog}(n)$	$s \leq \Omega(k)$	$\Theta((s + \log n)/\sqrt{k})$	$\Theta(s + \log n)$
This work	$k \geq \delta n$	$s \leq k/2^{\mathsf{poly}(1/\delta)}$	2	$\Theta(s + \log n)$

Subsequent Work. After the first version of this paper appeared online, Aggarwal et. al [14] made several improvements to our protocols to make them satisfy further security properties, such as *post-application robustness* and *source privacy*, at the cost of one or two extra rounds. In addition, they also applied techniques in our paper to the case of local computability and Bounded Retrieval Model [15,16].

2 Overview of the Constructions and Techniques

Here we give an informal overview of our constructions and the technique used. To give a clear description, we shall be imprecise sometimes.

2.1 Non-malleable Condenser for Arbitrary Min-entropy

For an (n, k) source X, our non-malleable condenser uses a uniform seed $Y = (Y_1, Y_2)$, where Y_2 has a bigger size than Y_1, say $|Y_1| = d$ and $|Y_2| = 10d$.

Consider now any function $\mathcal{A}(Y) = Y' = (Y_1', Y_2')$. In the following we will use letters with prime to denote variables produced with Y'. Since $Y' \neq Y$, we have two cases: $Y_1 = Y_1'$ or $Y_1 \neq Y_1'$. The output of our non-malleable condenser will be $Z = \mathsf{nmCond}(X, Y) = (V_1, V_2)$. Intuitively, V_1 handles the case where $Y_1 = Y_1'$ and V_2 handles the case where $Y_1 \neq Y_1'$. We now describe the two cases separately.

If $Y_1 = Y_1'$, then we take a strong extractor Ext and compute $W = \mathsf{Ext}(X, Y_1)$. Note that $W' = \mathsf{Ext}(X, Y_1') = W$ since $Y_1 = Y_1'$. Note that $Y' \neq Y$, thus we must have $Y_2' \neq Y_2$. We now fix Y_1 (and Y_1'). Note that conditioned on this fixing, $W = W'$ is still (close to) uniform since Ext is a strong extractor, and now Y_2' is a deterministic function of Y_2. At this point, we can take any non-malleable extractor nmExt from [11,12,3] and compute $V_1 = \mathsf{nmExt}(W, Y_2)$. Since W is uniform, by the property of the non-malleable extractor we have that V_1 is (close to) uniform even conditioned on the fixing of V_1' and (Y_2, Y_2'). Now let the size of V_1 be bigger than the size of V_2, say $|V_1| \geq |V_2| + s$. Thus the further conditioning on the fixing of V_2' will still leave V_1 with entropy roughly s. This takes care of our first case.

If $Y_1 \neq Y_1'$, then we first fix (Y_1, Y_1'). Note that fixing Y_1' may cause Y_2 to lose entropy. However, since $|Y_2| = 10|Y_1|$, conditioned on this fixing Y_2 still has entropy rate roughly $9/10$, and now Y_2' is a deterministic function of Y_2. We further fix $W' = \mathsf{Ext}(X, Y_1')$, which is now a deterministic function of X. As long as the entropy of X is larger than the size of W, conditioned on this fixing X still has a lot of entropy. Note that after these fixings X and Y_2 are still independent. Now, we use X and Y_2 to perform an alternating extraction protocol. Specifically, take the first $3d$ bits of Y_2 to be S_0, we compute the following random variables: $R_0 = \mathsf{Raz}(S_0, X), S_1 = \mathsf{Ext}(Y_2, R_0), R_1 = \mathsf{Ext}(X, S_1), S_2 = \mathsf{Ext}(Y_2, R_1), R_2 = \mathsf{Ext}(X, S_2), \cdots, S_t = \mathsf{Ext}(Y_2, R_{t-1}), R_t = \mathsf{Ext}(X, S_t)$. Here Raz is the strong two source extractor in [17], which works as long as the first source has entropy rate $> 1/2$, and Ext is a strong extractor. We take $t = 4d$ and let each R_i output s bits. Note that in the first step S_0 roughly has entropy rate $2/3$, thus we need to use the two-source extractor Raz. In all subsequent steps S_i, R_i are (close to) uniform, thus it suffices to use a strong extractor. The alternating extraction protocol is shown in Figure 1.

In the above alternating extraction protocol, as long as the size of each (S_i, R_i) is relatively small, one can show that for any j, R_j is (close to) uniform conditioned on $\{R_i, R_i', i < j\}$ and (Y_2, Y_2') (recall $\{R_j'\}$ are the random variables produced by Y_2' instead of Y_2). The intuitive reason is that in each step X still has enough entropy conditioned on all previous random variables produced, and we use a strong extractor which guarantees that the output is uniform even conditioned on the seed. Next, we borrow some ideas from [1]. Specifically, there they showed an efficient map f from a string with d bits to a subset of $[4d]$, such that for any $\mu \in \{0,1\}^d$, $f(\mu)$ has $2d$ elements. Moreover, for any $\mu \neq \mu'$, there exists a $j \in [4d]$ such that $|f(\mu)^{\geq j}| > |f(\mu')^{\geq j}|$, where $f(\mu)^{\geq j}$ denotes the subset of $f(\mu)$ which contains all the elements $\geq j$. Now, let $R = (R_1, \cdots, R_t)$ be the t random variables R_i produced in the above alternating extraction protocol. As

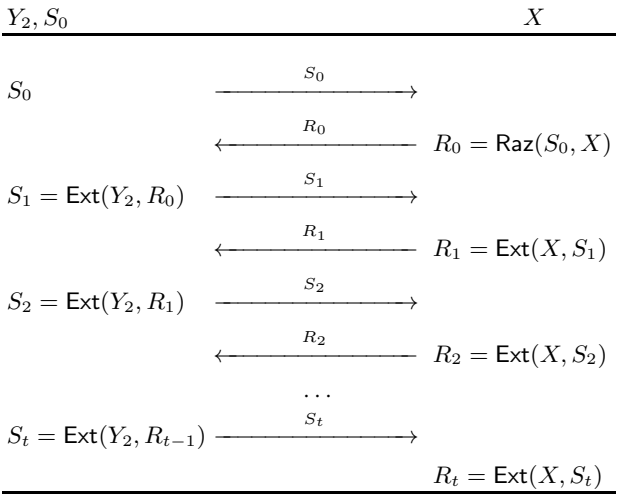

Fig. 1. Alternating Extraction

in [1], we define a "look-ahead" MAC (message authentication code) laMAC that uses R as the key. For any $\mu \in \{0,1\}^d$, we define $\mathsf{laMAC}_R(\mu) = \{R_i\}_{i \in f(\mu)}$. Now our V_2 is computed as $V_2 = \mathsf{laMAC}_R(Y_1)$.

Note that since we have fixed (Y_1, Y_1'), we can now view them as two different strings in $\{0,1\}^d$. Thus, there exists a $j \in [4d]$ such that $|f(Y_1)^{\geq j}| > |f(Y_1')^{\geq j}|$. We will now show that V_2 has entropy at least s conditioned on V_2'. To show this, let $\bar{R}$ be the concatenation of those R_is in $f(Y_1)^{\geq j}$ and $\bar{R}'$ be the concatenation of those R_i's in $f(Y_1')^{\geq j}$, then the size of $\bar{R}$ is bigger than the size of $\bar{R}'$ by at least s. Moreover, $\bar{R}$ is (close to) uniform conditioned on the fixing of $\{R_i', i < j\}$ and (Y_2, Y_2'). Thus $\bar{R}$ roughly has entropy s even conditioned on the fixing of $(\bar{R}', \{R_i', i < j\})$ and (Y_2, Y_2'), which also determines V_2'. Since $\bar{R}$ is part of V_2, we have that V_2 has entropy at least s conditioned on V_2'. Since we have fixed W' before, $V_1' = \mathsf{nmExt}(W', Y_2')$ is also fixed. Thus we have that $Z = (V_1, V_2)$ has entropy roughly s even conditioned on the fixing of $Z' = (V_1', V_2')$ and (Y_2, Y_2'). This takes care of our second case.

Thus, we obtain a non-malleable condenser for any min-entropy. However, since in the alternating extraction protocol each R_i outputs s bits, and we need $d = \Omega(s)$ to achieve error 2^{-s}, the entropy of X has to be larger than $4ds = \Omega(s^2)$. Thus we can only achieve s up to $\Omega(\sqrt{k})$.

2.2 Privacy Amplification Protocol

Combined with the techniques in [2], our non-malleable condenser immediately gives a 2-round privacy amplification protocol with optimal entropy loss for any min-entropy, with security parameter s up to $\Omega(\sqrt{k})$. To better illustrate the key idea, we also give a slightly simpler 2-round protocol with optimal entropy loss,

without using the non-malleable condenser. Assuming the security parameter we want to achieve is s, we now describe the protocol below.

In the first round, Alice samples 3 random strings (Y_1, Y_2, Y_3) from her local random bits and sends them to Bob, where Bob receives (Y_1', Y_2', Y_3'). We let $|Y_1| = d, |Y_2| = 10d, |Y_3| = 50d$. Take a strong extractor Ext, now Alice and Bob each computes $R_1 = \mathsf{Ext}(X, Y_1)$ and $R_1' = \mathsf{Ext}(X, Y_1')$ respectively. Let R_1, R_1' each output $4s$ bits. Next, Alice and Bob each uses (X, Y_2) and (X, Y_2') to perform the alternating extraction protocol we described above, where they compute $R_2 = (R_{21}, \cdots, R_{2t})$ and $R_2' = (R_{21}', \cdots, R_{2t}')$ respectively, with $t = 4d$. Finally, using R_2 and R_2' as the key, they compute $Z = \mathsf{laMAC}_{R_2}(Y_1)$ and $Z' = \mathsf{laMAC}_{R_2'}(Y_1')$ respectively as described before.

In the second round, Bob samples a random string W' from his local random bits and sends it to Alice, where Alice receives W. Together with W', Bob also sends two tags (T_1', T_2'), where Alice receives (T_1, T_2). For T_1', Bob takes the two-source extractor Raz and computes $T_1' = \mathsf{Raz}(Y_3', Z')$. Let T_1' output s bits. For T_2', Bob takes a standard message authentication code (MAC) and computes $T_2' = \mathsf{MAC}_{R_1'}(W')$, where R_1' is used as the key to authenticate the message W'. Bob then computes $R_B = \mathsf{Ext}(X, W')$ as the final output. When receiving (W, T_1, T_2), Alice will check whether $T_1 = \mathsf{Raz}(Y_3, Z)$ and $T_2 = \mathsf{MAC}_{R_1}(W)$. If either test fails, Alice rejects and aborts. Otherwise Alice computes $R_A = \mathsf{Ext}(X, W)$ as the final output. The protocol is shown in Figure 2.

As before, the analysis can be divided into two cases: $Y_1 = Y_1'$ and $Y_1 \neq Y_1'$. In the first case, we have $R_1 = R_1'$ and is (close to) uniform and private. Thus R_1 can be used in the MAC to authenticate W' to Alice. The MAC works by the property that if Eve changes W' to a different W, then with high probability Even cannot produce the correct tag $T_2 = \mathsf{MAC}_{R_1}(W)$ even given T_2'. This works except that here Eve also has additional information from T_1'. However, although T_1' may give some information about the MAC key R_1, note that R_1 has size $4s$ and T_1' has size s. Thus even conditioned on T_1', R_1 has entropy roughly $3s$. We note that the MAC works as long as the entropy rate of R_1 is bigger than $1/2$. Thus in this case Bob can indeed authenticate W' to Alice and they will agree on a uniform and private final output.

In the second case, again we can first fix (Y_1, Y_1') and R_1'. As before we have that after this fixing, Y_2 still has entropy rate roughly $9/10$, X still has a lot of entropy, and X is independent of (Y_2, Y_3). Now we can view (Y_1, Y_1') as two different strings and by the same analysis before, Z roughly has entropy s conditioned on the fixing of Z' and (Y_2, Y_2'). Note that after this fixing Y_3 still has entropy rate $> 1/2$, and Y_3' is a deterministic function of Y_3. Since Raz is a strong two-source extractor, we have that $\mathsf{Raz}(Y_3, Z)$ is (close to) uniform even given (Y_3', Z', R_1', W'), which also determines (T_1', T_2'). Thus, in this case Alice will reject with probability $1 - 2^{-s}$, since the probability that Eve guesses $\mathsf{Raz}(Y_3, Z)$ correctly is at most 2^{-s}.

We note that our protocol shares some similarities with the 2-round protocol in [1], as they both use the alternating extraction protocol and the "look-ahead" MAC. However, there is one important difference. The protocol in [1] uses the

Alice: X	Eve: E	Bob: X
Sample random $Y = (Y_1, Y_2, Y_3)$.		
Compute $R_2 = (R_{21}, \cdots, R_{2t})$		
by alternating extraction of (X, Y_2).		
$Z = \mathsf{laMAC}_{R_2}(Y_1)$.		
$R_1 = \mathsf{Ext}(X, Y_1)$ and output $4s$ bits.		
	(Y_1, Y_2, Y_3) $\longrightarrow$	
	(Y_1', Y_2', Y_3')	
		Sample random W' with d bits.
		Compute $R_2' = (R_{21}', \cdots, R_{2t}')$
		by alternating extraction of (X, Y_2').
		$Z' = \mathsf{laMAC}_{R_2'}(Y_1')$.
		$R_1' = \mathsf{Ext}(X, Y_1')$ and output $4s$ bits.
		$T_1' = \mathsf{Raz}(Y_3', Z')$ with s bits,
		$T_2' = \mathsf{MAC}_{R_1'}(W')$.
		Set final $R_B = \mathsf{Ext}(X, W')$.
	(W', T_1', T_2') $\longleftarrow$	
	(W, T_1, T_2)	
If $T_1 \neq \mathsf{Raz}(Y_3, Z)$ or		
$T_2 \neq \mathsf{MAC}_{R_1}(W)$ *reject.*		
Set final $R_A = \mathsf{Ext}(X, W)$.		

Fig. 2. 2-round Privacy Amplification Protocol

look-ahead MAC to authenticate the string W' that Bob sends to Alice in the second round. The look-ahead MAC has size $\Omega(s^2)$ and is revealed in the second round, which causes an entropy loss of $\Omega(s^2)$. Our protocol, on the other hand, uses the look-ahead MAC to authenticate the string Y_1 that Alice sends to Bob in the first round. Although in the protocol we do compute some variables that have size $\Omega(s^2)$ (namely (Z, Z')), they are computed locally by Alice and Bob, and are *never* revealed in the protocol to Eve. Instead, what is revealed to Eve is $T_1' = \mathsf{Raz}(Y_3', Z')$, which only has size $O(s)$. In other words, in the case where $Y_1 \neq Y_1'$, since we know that Z has entropy s conditioned on Z', we can apply another extractor Raz to Z and Z' respectively, such that the resulting variable T_1' only has size $O(s)$ and $\mathsf{Raz}(Y_3, Z)$ is (close to) uniform conditioned on T_1'. This is enough for the purpose of authentication, while bringing the entropy loss down to $O(s)$.

One might think that the same trick can also be applied to the protocol in [1] directly. However, this is not the case. The reason is that conditioned on (Y, Y'), all the random variables in our protocol that are used to authenticate W' are (R_1, T_1, R_1', T_1'), which are deterministic functions of X and have size $O(s)$. Thus in the case where Bob successfully authenticates W' to Alice, we can fix them and conditioned on the fixing, X and W are still independent so we can

apply a strong extractor to obtain the final output $\mathsf{Ext}(X, W)$. This results in a protocol with optimal entropy loss. In the protocol in [1], conditioned on (Y, Y'), the random variables that are used to authenticate W' include the output of the look-ahead extractor, which has size $\Omega(s^2)$. Thus conditioning on this random variable will cause X to lose entropy $\Omega(s^2)$. On the other hand, we cannot simply apply another extractor to this MAC to reduce the output size; since then the output will be a function of W and X, and thus conditioned on the fixing of it, W and X will no longer be independent.

We now describe our protocol for security parameter $s > \sqrt{k}$. The very high level strategy is as follows. At the beginning of the protocol, Alice samples a random string Y from her local random bits with $d_1 = O(s)$ bits and sends it to Bob, where Bob receives Y'. They each compute $R = \mathsf{Ext}(X, Y)$ and $R' = \mathsf{Ext}(X, Y')$ respectively, by using a strong extractor Ext. At the end of the protocol, Bob samples a random string W' from his local random bits with d_1 bits and sends it to Alice, together with a tag $T = \mathsf{MAC}_{R'}(W')$. Alice receives (W, T). Bob will compute $R_B = \mathsf{Ext}(X, W')$ as his final output and Alice will check if $T = \mathsf{MAC}_R(W)$. If the test fails then Alice rejects. Otherwise she will compute $R_A = \mathsf{Ext}(X, W)$ as her final output. In the case where $Y = Y'$, again we will have that $R = R'$ and is uniform and private. Thus in this case Bob can authenticate W' to Alice by using a MAC and R' as the key. We will now modify the protocol to ensure that if $Y \neq Y'$, then with probability $1 - 2^{-s}$ either Alice or Bob will reject.

If $s < \sqrt{k}$ then we can use our 2-round protocol described above. However, we want to achieve $s > \sqrt{k}$ and X does not have enough entropy for the 2-round protocol. On the other hand, we note that we can still use the 2-round protocol to authenticate a substring of Y with $s' = \Theta(\sqrt{k})$ bits to Bob, such that if Eve changes this string, then with probability $1 - 2^{-s'}$ Alice will reject. The key observation now is that after running this 2-round protocol, conditioned on the transcript revealed to Eve, X only loses $O(s')$ entropy. Thus X still has entropy $k - O(\sqrt{k})$ in Eve's view. Therefore, we can run the 2-round protocol again, using fresh random strings sampled from Alice and Bob's local random bits. This will authenticate another substring of Y with $s' = \Theta(\sqrt{k})$ bits to Bob. As long as X has enough entropy, we can keep doing this and it will take us $O(s/\sqrt{k})$ rounds to authenticate the entire Y to Bob, while the entropy loss is $O(s')O(s/\sqrt{k}) = O(s)$. Thus as long as $k \geq Cs$ for a sufficiently large constant C, the above approach will work.

However, the simple idea described above is not enough. The reason is that to change Y, Eve only needs to change one substring, and she can succeed with probability $2^{-s'} >> 2^{-s}$. To fix this, we modify the protocol to ensure that, if Eve changes Y to $Y' \neq Y$, then she has to change $\Omega(s/\sqrt{k})$ substrings, i.e., a constant fraction of the substrings. This is where we borrow some ideas from [4]. Specifically, instead of having Alice just authenticate substrings of Y to Bob, we will use an asymptotically good code for edit errors and have Alice authenticate substrings of the encoding of Y to Bob. More specifically, let $M = \mathsf{Edit}(Y)$ be the encoding of Y, which has size $O(d_1)$. At the beginning of the protocol, Alice will send Y to Bob, where Bob receives Y'. Next, our protocol will run

in $L = O(s/\sqrt{k})$ phases, with each phase consisting of two rounds. In phase i, Alice will send the i'th substring M_i of M to Bob, where M_i has $d_2 = \Theta(\sqrt{k})$ bits. In the first round of phase i, Alice samples two random strings (Y_{i2}, Y_{i3}) from her local random bits and sends them to Bob, together with M_i. Bob receives (M'_i, Y'_{i2}, Y'_{i3}). We will let $|Y_{i3}| \geq 10|Y_{i2}|$. As in the previous 2-round protocol, Alice will use X and Y_{i2} to perform an alternating extraction protocol, where she computes $R_i = (R_{i1}, \cdots, R_{it})$ with $t = 4d_2$ and $Z_i = \mathsf{laMAC}_{R_i}(M_i)$, where laMAC is the look-ahead MAC described before. Correspondingly, Bob will compute R'_i and $Z'_i = \mathsf{laMAC}_{R'_i}(M'_i)$, using X and Y'_{i2}. In the second round, Bob will send $T'_i = \mathsf{Raz}(Y'_{i3}, Z'_i)$ to Alice, where Alice receives T_i. Alice will now check if $T_i = \mathsf{Raz}(Y_{i3}, Z_i)$ and she rejects if the test fails. By the same analysis of the 2-round protocol, if Eve changes the substring M_i to $M'_i \neq M_i$, then with probability $1 - 2^{-\Omega(\sqrt{k})}$ Alice will reject.

One problem of the above approach is that Eve can first delay messages from Alice, send fake messages to Bob to get responses that contain additional information, and then resume execution with Alice. To avoid this problem, we need to synchronize between Alice and Bob. To achieve this, in the second round of phase i, we will also have Bob sample a fresh random string W'_i from his local random bits and send it as a challenge to Alice, together with T'_i. Alice will receive (W_i, T_i). Now if Alice does not reject, then she will also compute a response $V_i = \mathsf{Ext}(X, W_i)$ and send it back to Bob in the first round of phase $i + 1$. Bob will receive V'_i and then check if $V'_i = \mathsf{Ext}(X, W'_i)$. If the test fails then he rejects. Otherwise he proceeds as before. At the end of the protocol, Bob will first check if the received codeword $M' = M'_1 \circ \cdots \circ M'_L$ is indeed equal to $\mathsf{Edit}(Y')$. If the test fails he rejects. Otherwise he proceeds as before. This gives our whole protocol. The formal protocol appears in Section 6, Figure 4.

For the analysis, by the property of the code, if Eve wants to change $M = \mathsf{Edit}(Y)$ to $M' = \mathsf{Edit}(Y')$ with $Y' \neq Y$, then she has to make $\Omega(d_1)$ edit operations (insertion, deletion or altering). Since changing one substring costs at most $\sqrt{k}$ edit operations, Eve has to change at least $\Omega(s/\sqrt{k})$ substrings. As in [4], we then show that as long as X has an extra entropy of $O(s)$, for a constant fraction of these changes, conditioned on the event that Eve has successfully made all previous changes, the probability that Eve can make this change successfully is at most $2^{-\Omega(\sqrt{k})}$. Thus the overall probability that Eve can change M to M' without causing either Alice or Bob to reject is at most $(2^{-\Omega(\sqrt{k})})^{\Omega(s/\sqrt{k})} = 2^{-\Omega(s)}$. The round complexity is $O(s/\sqrt{k})$ and the communication complexity is $O(s\sqrt{k})$ since in each phase, the communication complexity is $O(k)$.

2.3 Non-malleable Condenser for Linear Min-entropy

Our non-malleable condenser for linear min-entropy is similar to the construction for arbitrary min-entropy, except we use a different alternating extraction protocol, namely that in [2]. Specifically, we will again use a seed $Y = (Y_1, Y_2)$, where $|Y_1| = d$ and $|Y_2| \geq 10d$. The output will also be $Z = (V_1, V_2)$. For any

function $\mathcal{A}(Y) = Y' = (Y_1', Y_2')$, we use V_1 to take care of the case where $Y_1 = Y_1'$ and use V_2 to take care of the case where $Y_1 \neq Y_1'$.

If $Y_1 = Y_1'$, then again we take a strong extractor Ext and compute $W = \mathsf{Ext}(X, Y_1)$ and $V_1 = \mathsf{nmExt}(W, Y_2)$. By the same argument before, as long as $|V_1| \geq |V_2| + s$, we have that V_1 roughly has min-entropy s conditioned on (V_1', V_2'). This takes care of our first case.

If $Y_1 \neq Y_1'$, then again we first fix (Y_1, Y_1') and W'. Conditioned on this fixing Y_2 still has entropy rate roughly $9/10$, and now Y_2' is a deterministic function of Y_2. Moreover X still has a lot of entropy (say δn for some constant $\delta > 0$) and is independent of Y_2. Now we use the alternating extraction protocol in [2]. More specifically, since X has min-entropy $k = \delta n$ we can apply a somewhere condenser in [18], [17], [19] to X and obtain $\bar{X} = (X_1, \cdots, X_C)$ with $C = \mathsf{poly}(1/\delta)$ such that at least one X_i has entropy rate 0.9. In [2], Li showed that as long as $k \geq 2^{\mathsf{poly}(1/\delta)}s$, one can use $X, \bar{X}, Y_2$ to perform an alternating extraction protocol and then use the output and Y_1 to obtain V_2 with size $2^{\mathsf{poly}(1/\delta)}s$, such that whenever $Y_1 \neq Y_1'$, V_2 roughly has entropy s conditioned on the fixing of V_2' and (Y_2, Y_2'). Since we have fixed (Y_1, Y_1') and W' before, this means that Z roughly has entropy s conditioned on the fixing of Z' and (Y, Y').

Combined with the protocol in [3], we thus reduce the entropy loss of the protocol in [2] to $O(s)$ for an absolute constant $O(\cdot)$ and the communication complexity to $\mathsf{poly}(1/\delta)s$.

Organization. in Section 3 we give the formal definition of the privacy amplification problem. We then give some preliminaries in Section 4 and define alternating extraction in Section 5. We give our non-malleable condenser for arbitrary min-entropy and the 2-round protocol in Section 6. The general multi-round protocol and non-malleable condenser for linear min-entropy are deferred to the full version. We conclude with some open problems in Section 7.

3 Privacy Amplification with an Active Adversary

In this section we formally define the privacy amplification problem. First we define average conditional min-entropy.

Definition 6. *The* average conditional min-entropy *is defined as*

$$\begin{aligned}\widetilde{H}_\infty(X|W) &= -\log\left(\mathrm{E}_{w\leftarrow W}\left[\max_x \Pr[X = x|W = w]\right]\right)\\ &= -\log\left(\mathrm{E}_{w\leftarrow W}\left[2^{-H_\infty(X|W=w)}\right]\right).\end{aligned}$$

We will follow [11] and 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 $\widetilde{H}_\infty(X|E) \geqslant k$. Since Eve is unbounded, we can assume without loss of generality that she is deterministic.

We assume that Eve has full control of the communication channel between the two parties. This means that Eve can arbitrarily insert, delete, reorder or modify messages sent by Alice and Bob to each other. In particular, Eve's strategy P_E defines two correlated executions (P_A, P_E) and (P_E, P_B) between Alice and Eve, and Eve and Bob, called "left execution" and "right execution", respectively. Alice and Bob are assumed to have fresh, private and independent random bits Y and W, respectively. Y and W are not known to Eve. In the protocol we use $\perp$ as a special symbol to indicate rejection. At the end of the left execution $(P_A(X,Y), P_E(E))$, Alice outputs a key $R_A \in \{0,1\}^m \cup \{\perp\}$. Similarly, Bob outputs a key $R_B \in \{0,1\}^m \cup \{\perp\}$ at the end of the right execution $(P_E(E), P_B(X,W))$. We let E' denote the final view of Eve, which includes E and the communication transcripts of both executions $(P_A(X,Y), P_E(E))$ and $(P_E(E), P_B(X,W)$. We can now define the security of (P_A, P_B).

Definition 7. *An interactive protocol (P_A, P_B), executed by Alice and Bob on a communication channel fully controlled by an active adversary Eve, 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. We start by defining the notion of* pre-application *robustness, which states that even if Eve is active,* $\Pr[R_A \neq R_B \wedge R_A \neq \perp \wedge R_B \neq \perp] \leqslant \epsilon$. *The stronger notion of* post-application *robustness is defined similarly, except Eve is additionally given the key R_A the moment she completed the left execution (P_A, P_E), and the key R_B the moment she completed the right execution (P_E, P_B). For example, if Eve completed the left execution before the right execution, she may try to use R_A to force Bob to output a different key $R_B \notin \{R_A, \perp\}$, and vice versa.*
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:* $\mathsf{purify}(r) \leftarrow U_m$. *Letting E' denote Eve's view of the protocol, we require that*

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

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

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

Remark 1. Our protocol, as well as many others in [1], [9], [10], [4], [11], [12], [3], [2] only achieve *pre-application* robustness. Recently, Aggarwal et. al [14] gave a general transformation that can convert any privacy amplification protocol with pre-application robustness into another privacy amplification protocol with *post-application* robustness at the cost of one extra round. Thus, using their transformation, our protocol can be turned into a 3-round post-application robust privacy amplification protocol with optimal entropy loss, for security parameter up to $s = \Omega(\sqrt{k})$ (as Aggarwal et. al did in [14]); or a $O(s/\sqrt{k})$ round

post-application robust privacy amplification protocol with optimal entropy loss, for security parameter up to $s = \Omega(k)$.

4 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.

4.1 Somewhere Random Sources, Extractors and Condensers

Definition 8 (Somewhere Random sources). *A source $X = (X_1, \cdots, X_t)$ is $(t \times r)$* somewhere-random *(SR-source for short) if each X_i takes values in $\{0,1\}^r$ and there is an i such that X_i is uniformly distributed.*

Definition 9. *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 10. *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 11. *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.*

Definition 12. *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

$$|(\mathsf{TExt}(X,Y), Y) - (U_m, Y)| < \epsilon,$$

where U_m is the uniform distribution on m bits independent of (X, Y).

4.2 Average Conditional Min-entropy

Dodis and Wichs originally defined non-malleable extractors with respect to average conditional min-entropy. However, this notion is essentially equivalent to the standard (worst-case) min-entropy, up to a small loss in parameters.

Lemma 1 ([20]). *For any $s > 0$, $\Pr_{w \leftarrow W}[H_\infty(X|W = w) \geq \widetilde{H}_\infty(X|W) - s] \geq 1 - 2^{-s}$.*

Lemma 2 ([20]). *If a random variable B has at most 2^ℓ possible values, then $\widetilde{H}_\infty(A|B) \geq H_\infty(A) - \ell$.*

To clarify which notion of min-entropy and non-malleable extractor we mean, we use the term *worst-case non-malleable extractor* when we refer to our Definition 4, which is with respect to traditional (worst-case) min-entropy, and *average-case non-malleable extractor* to refer to he original definition of Dodis and Wichs, which is with respect to average conditional min-entropy.

Corollary 1. *A (k, ε)-average-case non-malleable extractor is a (k, ε)-worst-case non-malleable extractor. For any $s > 0$, a (k, ε)-worst-case non-malleable extractor is a $(k + s, \varepsilon + 2^{-s})$-average-case non-malleable extractor.*

Throughout the rest of our paper, when we say non-malleable extractor, we refer to the worst-case non-malleable extractor of Definition 4.

4.3 Prerequisites from Previous Work

One-time message authentication codes (MACs) use a shared random key to authenticate a message in the information-theoretic setting.

Definition 13. *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 for any $w \in \{0,1\}^d$ and any function (adversary) $A : \{0,1\}^v \to \{0,1\}^d \times \{0,1\}^v$,*

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

where R is the uniform distribution over the key space $\{0,1\}^\ell$.

Theorem 6 ([10]). *For any message length d and tag length v, there exists an efficient family of $(\lceil \frac{d}{v} \rceil 2^{-v})$-secure MACs 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,

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

(In the special case when $R \equiv U_{2v}$ and independent of E, we get the original bound.)

Remark 2. Note that the above theorem indicates that the MAC works even if the key R has average conditional min-entropy rate $> 1/2$.

Sometimes it is convenient to talk about average case seeded extractors, where the source X has average conditional min-entropy $\widetilde{H}_\infty(X|Z) \geq k$ and the output of the extractor should be uniform given Z as well. The following lemma is proved in [20].

Lemma 3 ([20]). *For any $\delta > 0$, if Ext is a (k, ϵ) extractor then it is also a $(k + \log(1/\delta), \epsilon + \delta)$ average case extractor.*

Theorem 7 ([18,17,19]). *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 $m = \Omega(n)$.*

For a strong seeded extractor with optimal parameters, we use the following extractor constructed in [21].

Theorem 8 ([21]). *For every constant $\alpha > 0$, and all positive integers n, k and any $\epsilon > 0$, there is an explicit construction of a 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 $m \geq (1-\alpha)k$. It is also a strong (k, ϵ) average case extractor with $m \geq (1-\alpha)k - O(\log n + \log(1/\epsilon))$.*

We need the following construction of strong two-source extractors in [17].

Theorem 9 ([17]). *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 10 ([11,12,3]). *For every constant $\delta > 0$, there exists a constant $\beta > 0$ such that for every $n, k \in \mathsf{N}$ 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 $m = \Omega(n)$.*

The following theorem is proved in [3].

Theorem 11 ([3]). *There exists a constant $C > 1$ such that the following holds. For any $n, k \in \mathsf{N}$ and $\epsilon > 0$, assume that there is an explicit (k, k', ϵ)-non-malleable condenser with seed length d such that $k' \geq C(\log n + \log(1/\epsilon))$. Then there exists an explicit 2-round privacy amplification protocol for (n, k) sources with entropy loss $O(\log n + \log(1/\epsilon))$ and communication complexity $O(d + \log n + \log(1/\epsilon))$.*

The following standard lemma about conditional min-entropy is implicit in [22] and explicit in [7].

Lemma 4 ([7]). *Let X and Y be random variables and let $\mathcal{Y}$ denote the range of Y. Then for all $\epsilon > 0$, one has*

$$\Pr_Y \left[H_\infty(X|Y=y) \geq H_\infty(X) - \log|\mathcal{Y}| - \log\left(\frac{1}{\epsilon}\right) \right] \geq 1 - \epsilon.$$

We also need the following lemma.

Lemma 5. *Let (X,Y) be a joint distribution such that X has range $\mathcal{X}$ and Y has range $\mathcal{Y}$. Assume that there is another random variable X' with the same range as X such that $|X - X'| = \epsilon$. Then there exists a joint distribution (X',Y) such that $|(X,Y) - (X',Y)| = \epsilon$.*

Proof. First let (X'',Y) be the same probability distribution as (X,Y). For any $x \in \mathcal{X}$, let $p''_x = \Pr[X'' = x]$ and $p'_x = \Pr[X' = x]$. For any $y \in \mathcal{Y}$, let $p_y = \Pr[Y = y]$. Let $p''_{xy} = \Pr[X'' = x, Y = y]$. Let $W = \{x \in \mathcal{X} : p''_x > p'_x\}$ and $V = \{x \in \mathcal{X} : p''_x < p'_x\}$. Thus we have that $\sum_{x \in W} |p''_x - p'_x| = \sum_{x \in V} |p''_x - p'_x| = \epsilon$.

We now gradually change the probability distribution X'' into X', while keeping the distribution Y the same, as follows. While W is not empty or V is not empty, do the following.

1. Pick $x \in W \cup V$ such that $|p''_x - p'_x| = min\{|p''_x - p'_x|, x \in W \cup V\}$.
2. If $x \in W$, we decrease $\Pr[X'' = x]$ to p'_x. Let $\tau = p''_x - p'_x$. To ensure this is still a probability distribution, we also pick any $\bar{x} \in V$ and increase $\Pr[X'' = \bar{x}]$ to $\Pr[X'' = \bar{x}] + \tau$. To do this, we pick the elements $y \in \mathcal{Y}$ one by one in an arbitrary order and while $\tau > 0$, do the following. Let $\tau' = min(p''_{xy}, \tau)$, $\Pr[X'' = x, Y = y] = \Pr[X'' = x, Y = y] - \tau'$, $\Pr[X'' = \bar{x}, Y = y] = \Pr[X'' = \bar{x}, Y = y] + \tau'$ and $\tau = \tau - \tau'$. We then update the sets $\{p''_x\}$ and $\{p''_{xy}\}$ accordingly. Note that since $p''_x = \tau + p'_x \geq \tau$, this process will indeed end when $\tau = 0$ and now $\Pr[X'' = x] = p'_x$. Note that after this change we still have that $p''_{\bar{x}} \leq p'_{\bar{x}}$. Also, for any $y \in \mathcal{Y}$ the probability $\Pr[Y = y]$ remains unchanged. Finally, remove x from W and if $p''_{\bar{x}} = p'_{\bar{x}}$, remove $\bar{x}$ from V.
3. If $x \in V$, we increase $\Pr[X'' = x]$ to p'_x. Let $\tau = p'_x - p''_x$. To ensure that X'' is still a probability distribution, we also pick any $\bar{x} \in W$ and decrease $\Pr[X'' = \bar{x}]$ to $\Pr[X'' = \bar{x}] - \tau$. To do this, we pick the elements $y \in \mathcal{Y}$ one by one in an arbitrary order and while $\tau > 0$, do the following. Let $\tau' = min(p''_{\bar{x}y}, \tau)$, $\Pr[X'' = x, Y = y] = \Pr[X'' = x, Y = y] + \tau'$, $\Pr[X'' = \bar{x}, Y = y] = \Pr[X'' = \bar{x}, Y = y] - \tau'$ and $\tau = \tau - \tau'$. We then update the sets $\{p''_x\}$ and $\{p''_{xy}\}$ accordingly. Note that since $p''_{\bar{x}} \geq \tau + p'_{\bar{x}}$, this process will indeed end when $\tau = 0$ and we still have $p''_{\bar{x}} \geq p_{\bar{x}}$. Also, for any $y \in \mathcal{Y}$ the probability $\Pr[Y = y]$ remains unchanged. Finally, remove x from V and if $p''_{\bar{x}} = p_{\bar{x}}$, remove $\bar{x}$ from W.

Note that in each iteration, at least one element will be removed from $W \cup V$. Thus the iteration will end after finite steps. When it ends, we have that $\forall x, \Pr[x'' = x] = p'_x$, thus $X'' = X'$. Since in each step the probability $\Pr[Y = y]$ remains unchanged, the distribution Y remains the same. Finally, it is clear from the algorithm that $|(X'', Y) - (X, Y)| = \epsilon$.

Next we have the following lemma.

Lemma 6. *Let X and Y be random variables and let $\mathcal{Y}$ denote the range of Y. Assume that X is ϵ-close to having min-entropy k. Then for any $\epsilon' > 0$*

$$\Pr_Y\left[(X|Y=y) \text{ is } \epsilon'\text{-close to a source with min-entropy } k - \log|\mathcal{Y}| - \log\left(\frac{1}{\epsilon'}\right)\right] \geq 1 - \epsilon' - \frac{\epsilon}{\epsilon'}.$$

Proof. Let $\mathcal{X}$ denote the range of X. Assume that X' is a distribution on $\mathcal{X}$ with min-entropy k such that $|X - X'| \leq \epsilon$. Then by lemma 5, there exists a joint distribution (X', Y) such that

$$|(X, Y) - (X', Y)| \leq \epsilon.$$

Now for any $y \in \mathcal{Y}$, let $\Delta_y = \sum_{x\in\mathcal{X}} |\Pr[X = x, Y = y] - \Pr[X' = x, Y = y]|$. Then we have

$$\sum_{y\in\mathcal{Y}} \Delta_y \leq \epsilon.$$

For any $y \in \mathcal{Y}$, the statistical distance between $X|Y = y$ and $X'|Y = y$ is

$$\begin{aligned}\delta_y &= \sum_{x\in\mathcal{X}} |\Pr[X = x|Y = y] - \Pr[X' = x|Y = y]| \\ &= (\sum_{x\in\mathcal{X}} |\Pr[X = x, Y = y] - \Pr[X' = x, Y = y]|)/(\Pr[Y = y]) = \Delta_y/\Pr[Y = y].\end{aligned}$$

Thus if $\delta_y \geq \epsilon'$ then $\Delta_y \geq \epsilon' \Pr[Y = y]$. Let $B_Y = \{y : \delta_y \geq \epsilon'\}$ then we have

$$\epsilon' \Pr[y \in B_Y] = \sum_{y\in B_Y} \epsilon' \Pr[Y = y] \leq \sum_{y\in B_Y} \Delta_y \leq \sum_{y\in\mathcal{Y}} \Delta_y \leq \epsilon.$$

Thus $\Pr[y \in B_Y] \leq \frac{\epsilon}{\epsilon'}$. Note that when $y \notin B_y$ we have $|X|Y = y - X'|Y = y| < \epsilon'$. Thus by Lemma 4 we have the statement of the lemma.

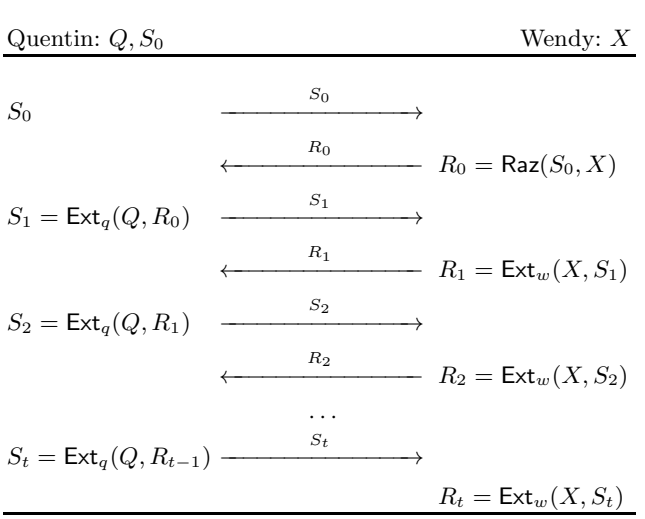

Fig. 3. Alternating Extraction

5 Alternating Extraction Protocol and Look Ahead Extractor

Recall that, an important ingredient in our construction is the following alternating extraction protocol modified from that in [1].

Alternating Extraction. Assume that we have two parties, Quentin and Wendy. Quentin has a source Q, Wendy has a source X. Also assume that Quentin has a weak source S_0 with entropy rate $> 1/2$ (which may be correlated with Q). Suppose that (Q, S_0) is kept secret from Wendy and X is kept secret from Quentin. Let Ext_q, Ext_w be strong seeded extractors with optimal parameters, such as that in Theorem 8. Let Raz be the strong two-source extractor in Theorem 9. Let d be an integer parameter for the protocol. For some integer parameter $t > 0$, 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 $R_1 = \mathsf{Ext}_w(X, S_1)$. She sends R_1 to Quentin and Quentin computes $S_2 = \mathsf{Ext}_q(Q, R_1)$. In this step R_1, S_2 each outputs d bits. In each subsequent step i, Quentin sends S_i to Wendy, Wendy computes $R_i = \mathsf{Ext}_w(X, S_i)$. She replies R_i to Quentin and Quentin computes $S_{i+1} = \mathsf{Ext}_q(Q, R_i)$. In step i, R_i, S_{i+1} each outputs d bits. Therefore, this process produces the following sequence:

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

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

$$\mathsf{laExt}(X, Y) = \mathsf{laExt}(X, (Q, S_0)) =: R_1, \cdots, R_t.$$

Note that the look-ahead extractor can be computed by each party (Alice or Bob) alone in our final protocol. We now have the following lemma.

Lemma 7. *In the alternating extraction protocol, assume that X has n bits and Q has at most n bits. Let $\epsilon > 0$ be a parameter and $d = O(\log n + \log(1/\epsilon)) > \log(1/\epsilon)$ be the number of random bits needed in Theorem 8 to achieve error ϵ. Assume that X has min-entropy at least $12d^2$, Q has min-entropy at least $11d^2$ and S_0 is a $(40d, 38d)$ source. Let Ext_w and Ext_q be strong extractors in Theorem 8 that use d bits to extract d bits. Let $t = 4d$.*

Let (Q', S_0') be another distribution on the same support of (Q, S_0) such that (Q, S_0, Q', S_0') is independent of X. Now run the alternating extraction protocol with X and (Q', S_0') where in each step we obtain S_i', R_i'. For any $i, 0 \leq i \leq t-1$, let $\overline{S_i} = (S_0, \cdots, S_i)$, $\overline{S_i'} = (S_0', \cdots, S_i')$, $\overline{R_i} = (R_0, \cdots, R_i)$ and $\overline{R_i'} = (R_0', \cdots, R_i')$. Then for any $i, 0 \leq i \leq t-1$, we have

$$\begin{aligned}&(R_i, \overline{S_{i-1}}, \overline{S_{i-1}'}, \overline{R_{i-1}}, \overline{R_{i-1}'}, S_i, S_i', Q, Q')\\ \approx_{(2i+2)\epsilon} &(U_d, \overline{S_{i-1}}, \overline{S_{i-1}'}, \overline{R_{i-1}}, \overline{R_{i-1}'}, S_i, S_i', Q, Q').\end{aligned}$$

Proof. We first prove the following claim.

Claim. In step 0, we have

$$(R_0, S_0, S_0', Q, Q') \approx_\epsilon (U_d, S_0, S_0', Q, Q')$$

and

$$(S_1, R_0, S_0, R_0', S_0') \approx_{3\epsilon} (U_d, R_0, S_0, R_0', S_0').$$

Moreover, conditioned on (S_0, S_0'), (R_0, R_0') are both deterministic functions of X; conditioned on (R_0, S_0, R_0', S_0'), (S_1, S_1') are deterministic functions of (Q, Q').

Proof (Proof of the claim.). Note that S_0 is a $(40d, 38d)$ source. Thus by Theorem 9 we have that

$$(R_0, S_0) \approx_\epsilon (U_d, S_0).$$

Since conditioned on S_0, R_0 is a deterministic function of X, which is independent of (Q, Q'), we also have that

$$(R_0, S_0, S_0', Q, Q') \approx_\epsilon (U_d, S_0, S_0', Q, Q').$$

Now we fix (S_0, S_0') and (R_0, R_0') are both deterministic functions of X. Since the size of (S_0, S_0') is at most $80d$, by Lemma 4 we have that with probability $1 - \epsilon$ over these fixings, Q is a source with entropy $10d^2$. Since R_0, R_0' are both

deterministic functions of X, they are independent of Q. Therefore by Theorem 8 we have

$$(S_1, R_0, R'_0) \approx_\epsilon (U_d, R_0, R'_0).$$

Thus altogether we have that

$$(S_1, R_0, S_0, R'_0, S'_0) \approx_{3\epsilon} (U_d, R_0, S_0, R'_0, S'_0)$$

Moreover, conditioned on (R_0, S_0, R'_0, S'_0), (S_1, S'_1) are deterministic functions of (Q, Q').

Now we fix (R_0, S_0, R'_0, S'_0). Note that after this fixing, S_1, S'_1 are are deterministic functions of (Q, Q'). Note that with probability $1 - \epsilon$ over this fixing, Q has min-entropy at least $10d^2$.

We now prove the lemma. In fact, we prove the following stronger claim.

Claim. For any i, we have that

$$\begin{aligned}&(R_i, \overline{S_{i-1}}, \overline{S'_{i-1}}, \overline{R_{i-1}}, \overline{R'_{i-1}}, S_i, S'_i, Q, Q')\\ \approx_{(2i+2)\epsilon} &(U_d, \overline{S_{i-1}}, \overline{S'_{i-1}}, \overline{R_{i-1}}, \overline{R'_{i-1}}, S_i, S'_i, Q, Q')\end{aligned}$$

and

$$(S_{i+1}, \overline{S_i}, \overline{S'_i}, \overline{R_i}, \overline{R'_i}) \approx_{(2i+3)\epsilon} (U_d, \overline{S_i}, \overline{S'_i}, \overline{R_i}, \overline{R'_i}).$$

Moreover, conditioned on $(\overline{S_{i-1}}, \overline{S'_{i-1}}, \overline{R_{i-1}}, \overline{R'_{i-1}}, S_i, S'_i)$, (R_i, R'_i) are both deterministic functions of X; conditioned on $(\overline{S_i}, \overline{S'_i}, \overline{R_i}, \overline{R'_i})$, (S_{i+1}, S'_{i+1}) are deterministic functions of (Q, Q').

We prove the claim by induction on i. When $i = 0$, the statements are already proved in Claim 5. Now we assume that the statements hold for $i = j$ and we prove them for $i = j + 1$.

We first fix $(\overline{S_j}, \overline{S'_j}, \overline{R_j}, \overline{R'_j})$. Since now (S_{j+1}, S'_{j+1}) are deterministic functions of (Q, Q'), they are independent of X. Moreover S_{j+1} is $(2j+3)\epsilon$-close to uniform. Note that the average conditional min-entropy of X is at least $12d^2 - 2d \cdot 4d \geq 4d^2$. Therefore by Theorem 8 we have that

$$(R_{j+1}, \overline{S_j}, \overline{S'_j}, \overline{R_j}, \overline{R'_j}, S_{j+1}, S'_{j+1}) \approx_{(2j+4)\epsilon} (U_d, \overline{S_j}, \overline{S'_j}, \overline{R_j}, \overline{R'_j}, S_{j+1}, S'_{j+1}).$$

Since (S_{j+1}, S'_{j+1}) are deterministic functions of (Q, Q'), we also have

$$(R_{j+1}, \overline{S_j}, \overline{S'_j}, \overline{R_j}, \overline{R'_j}, S_{j+1}, S'_{j+1}, Q, Q') \approx_{(2j+4)\epsilon} (U_d, \overline{S_j}, \overline{S'_j}, \overline{R_j}, \overline{R'_j}, S_{j+1}, S'_{j+1}, Q, Q').$$

Moreover, conditioned on $(\overline{S_j}, \overline{S'_j}, \overline{R_j}, \overline{R'_j}, S_{j+1}, S'_{j+1})$, (R_{j+1}, R'_{j+1}) are both deterministic functions of X.

Next, since conditioned on $(\overline{S_j}, \overline{S'_j}, \overline{R_j}, \overline{R'_j}, S_{j+1}, S'_{j+1})$, (R_{j+1}, R'_{j+1}) are both deterministic functions of X, they are independent of (Q, Q'). Moreover R_{j+1} is

$(2j+4)\epsilon$-close to uniform. Note that the average conditional min-entropy of Q is at least $10d^2 - 8d^2 = 2d^2$. Therefore by Theorem 8 we have that

$$\begin{aligned}&(S_{j+2}, \overline{S_j}, \overline{S'_j}, \overline{R_j}, \overline{R'_j}, S_{j+1}, S'_{j+1}, R_{j+1}, R'_{j+1})\\ \approx_{(2j+5)\epsilon} &(U_d, \overline{S_j}, \overline{S'_j}, \overline{R_j}, \overline{R'_j}, S_{j+1}, S'_{j+1}, R_{j+1}, R'_{j+1}).\end{aligned}$$

Namely,

$$(S_{j+2}, \overline{S_{j+1}}, \overline{S'_{j+1}}, \overline{R_{j+1}}, \overline{R'_{j+1}}) \approx_{(2(j+1)+3)\epsilon} (U_d, \overline{S_{j+1}}, \overline{S'_{j+1}}, \overline{R_{j+1}}, \overline{R'_{j+1}}).$$

Moreover, conditioned on $(\overline{S_{j+1}}, \overline{S'_{j+1}}, \overline{R_{j+1}}, \overline{R'_{j+1}})$, (S_{j+2}, S'_{j+2}) are deterministic functions of (Q, Q').

6 Non-malleable Condensers for Arbitrary Min-entropy

In this section we give our construction of non-malleable condensers for arbitrary min-entropy.

First, we need the following definitions and constructions from [1].

Definition 14. *[1] Given $S_1, S_2 \subseteq \{1, \cdots, t\}$, we say that the ordered pair (S_1, S_2) is top-heavy if there is some integer j such that $|S_1^{\geq j}| > |S_2^{\geq j}|$, where $S^{\geq j} =: \{s \in S | s \geq j\}$. Note that it is possible that (S_1, S_2) and (S_2, S_1) are both top-heavy. For a collection Ψ of sets $S_i \subseteq \{1, \cdots, t\}$, we say that Ψ is pairwise top-heavy if every ordered pair (S_i, S_j) of sets $S_i, S_j \in \Psi$ with $i \neq j$, is top-heavy.*

Now, for any m-bit message $\mu = (b_1, \cdots, b_m)$, consider the following mapping of μ to a subset $S \subseteq \{1, \cdots, 4m\}$:

$$f(\mu) = f(b_1, \cdots, b_m) = \{4i - 3 + b_i, 4i - b_i | i = 1, \cdots, m\}$$

i.e., each bit b_i decides if to include $\{4i-3, 4i\}$ (if $b_i = 0$) or $\{4i-2, 4i-1\}$ (if $b_i = 1$) in S.

We now have the following lemma.

Lemma 8. *[1] The above construction gives a pairwise top-heavy collection Ψ of 2^m sets $S \subseteq \{1, \cdots, t\}$ where $t = 4m$. Furthermore, the function f is an efficient mapping of $\mu \in \{0,1\}^m$ to S_μ.*

Now we have the following construction.

Let $r \in (\{0,1\}^d)^t$ be the output of the look-ahead extractor defined above, i.e., $r = (r_1, \cdots, r_t) = \mathsf{laExt}(X, (Q, S_0))$. Let $\Psi = \{S_1, \cdots, S_{2^m}\}$ be the pairwise top-heavy collection of sets constructed above. For any message $\mu \in \{0,1\}^m$, define the function $\mathsf{laMAC}_r(\mu) =: [r_i | i \in S_\mu]$, indexed by r.

Now we can describe our construction of the non-malleable condenser.

Algorithm 12 ($\mathsf{nmCond}(x, y)$)

Input: ℓ–an integer parameter. x — a sample from an (n, k)-source with $k \geq 60d^2$. y–an independent random seed with $y = (y_1, y_2)$ such that y_1 has size $d = O(\log n + \ell) > 5\ell$ and y_2 has size $12d^2$.
Output: z — an m bit string.

Sub-Routines and Parameters:
Let nmExt be the non-malleable extractor from Theorem 10, with error $2^{-4\ell}$.
Let Ext be the strong extractor with optimal parameters from Theorem 8, with error $2^{-5\ell}$.
Let laExt be the look-ahead extractor defined above, using Ext as Ext_q and Ext_s. laExt is set up to extract from x using seed (q, s_0) such that $q = y_2$ and s_0 is the string that contains the first $40d$ bits of y_2, and output a string $r \in (\{0,1\}^d)^t$ with $t = 4d$.
Let $\mathsf{laMAC}_r(\mu)$ be the function defined above.

1. Compute $w = \mathsf{Ext}(x, y_1)$ with output size $20d^2$ and $r = \mathsf{laExt}(x, (q, s_0))$.
2. Output $z = (\mathsf{nmExt}(w, y_2), \mathsf{laMAC}_r(y_1))$ such that $\mathsf{nmExt}(w, y_2)$ has size $8d^2$.

We can now prove the following theorem.

Theorem 13. *There exists a constant $C > 0$ such that given any $s > 0$, as long as $k \geq C(\log n + s)^2$, the above construction is a $(k, s, 2^{-s})$-non-malleable condenser with seed length $O(\log n + s)^2$ and output length $O(\log n + s)^2$.*

Proof. Let $\mathcal{A}$ be any (deterministic) function such that $\forall y \in \mathsf{Supp}(Y), \mathcal{A}(y) \neq y$. We will show that for most y, with high probability over the fixing of $\mathsf{nmCond}(X, \mathcal{A}(y))$, $\mathsf{nmCond}(X, y)$ is still close to having min-entropy at least ℓ. Let $Y' = \mathcal{A}(Y)$. Thus $Y' \neq Y$. In the following analysis we will use letters with prime to denote the corresponding random variables produced with Y' instead of Y. Let $V_1 = \mathsf{nmExt}(W, Y_2)$ and $V_2 = \mathsf{laMAC}_R(Y_1)$. Thus $Z = (V_1, V_2)$. We have the following two cases.

Case 1: $Y_1 = Y_1'$. In this case, since $Y' \neq Y$, we must have that $Y_2 \neq Y_2'$. Now by Theorem 8 we have that

$$(W, Y_1) \approx_{2^{-5\ell}} (U, Y_1).$$

Therefore, we can now fix Y_1 (and thus Y_1'), and with probability $1 - 2^{-\ell}$ over this fixing, W is $2^{-4\ell}$-close to uniform. Moreover, after this fixing W is a deterministic function of X and thus is independent of Y_2. Note also that after this fixing, Y_2' is a deterministic function of Y_2. Thus by Theorem 10 we have that

$$(V_1, V_1', Y_2, Y_2') \approx_{O(2^{-4\ell})} (U_{8d^2}, V_1', Y_2, Y_2').$$

Therefore, we can now further fix Y_2 (and thus Y_2') and with probability at least $1 - O(2^{-\ell})$ over this fixing, (V_1, V_1') is $2^{-3\ell}$-close to (U_{8d^2}, V_1'). Thus we can further fix V_1', and with probability at least $1 - 2^{-\ell}$ over this fixing, V_1 is $2^{-2\ell}$-close to uniform. Now note that V_1 has size $8d^2$ and V_2' has size $2d^2$. Thus by Lemma 6, we can further fix V_2', and with probability at least $1 - 2 \cdot 2^{-\ell}$ over this fixing, V_1 is 2^{ℓ}-close to having min-entropy at least $8d^2 - 2d^2 - \ell \geq 5d^2$.

Thus in this case we have shown that, with probability $1 - O(2^{-\ell})$ over the fixing of Y, with probability $1 - O(2^{-\ell})$ over the fixing of Z', Z is $2^{-\ell}$-close to having min-entropy at least $5d^2 > 5\ell^2$.

Case 2: $Y_1 \neq Y_1'$. In this case, we first fix Y_1 and Y_1'. Note that after this fixing, W and W' are now deterministic functions of X. We now further fix W and W' and after this fixing, X and Y_2 are still independent. Since the total size of (W, W') is $40d^2$, by Lemma 4 we have that with probability $1 - 2^{-2\ell}$ over this fixing, X still has min-entropy at least $60d^2 - 40d^2 - 2\ell > 12d^2$. Note also that after this fixing, Y_2' is a deterministic function of Y_2. However, since Y_1' may be a function of Y_2, fixing Y_1' may cause Y_2 to lose entropy. Note that Y_1' only has size d, thus by Lemma 4, with probability $1 - 2 \cdot 2^{-2\ell}$ over the fixing of (Y_1, Y_1'), we have that Y_2 has min-entropy at least $12d^2 - d - 2\ell > 11d^2$ and S_0 has min-entropy at least $40d - d - 2\ell > 38d$.

Now assume that X has min-entropy at least $12d^2$, Y_2 has min-entropy at least $11d^2$ and S_0 has min-entropy at least $38d$. This happens with probability at least $1 - O(2^{-\ell})$. For any $i, 0 \leq i \leq t-1$, let $\overline{S_i} = (S_0, \cdots, S_i)$, $\overline{S_i'} = (S_0', \cdots, S_i')$, $\overline{R_i} = (R_0, \cdots, R_i)$ and $\overline{R_i'} = (R_0', \cdots, R_i')$. Now by Lemma 7 (note that $Y_2 = (Q, S_0)$) we have that for any $i, 0 \leq i \leq t-1$,

$$\begin{aligned} &(R_i, \overline{S_{i-1}}, \overline{S'_{i-1}}, \overline{R_{i-1}}, \overline{R'_{i-1}}, S_i, S_i', Y_2) \\ \approx_{(2i+2)2^{-5\ell}} &(U_d, \overline{S_{i-1}}, \overline{S'_{i-1}}, \overline{R_{i-1}}, \overline{R'_{i-1}}, S_i, S_i', Y_2). \end{aligned}$$

Therefore, we have that for any i,

$$(R_i, \overline{R_{i-1}}, \overline{R'_{i-1}}, Y_2) \approx_{(2i+2)2^{-5\ell}} (U_d, \overline{R_{i-1}}, \overline{R'_{i-1}}, Y_2).$$

Thus, for any i, with probability $1 - 2^{-1.25\ell}$ over the fixing of Y_2, we have

$$(R_i, \overline{R_{i-1}}, \overline{R'_{i-1}}) \approx_{(2i+2)2^{-3.75\ell}} (U_d, \overline{R_{i-1}}, \overline{R'_{i-1}}).$$

By the union bound, we have that with probability $1 - t2^{-1.25\ell}$ over the fixing of Y_2, for any i,

$$(R_i, \overline{R_{i-1}}, \overline{R'_{i-1}}) \approx_{(2i+2)2^{-3.75\ell}} (U_d, \overline{R_{i-1}}, \overline{R'_{i-1}}).$$

Consider a typical fixing of Y_2. Now note that $V_2 = \mathsf{laMAC}_R(Y_1)$ and $V_2' = \mathsf{laMAC}_{R'}(Y_1')$. Let the two sets in Lemma 8 that correspond to Y_1 and Y_1' be H and H'. Since $Y_1 \neq Y_1'$, by definition there exists $j \in [4d]$ such that $|H^{\geq j}| > |H'^{\geq j}|$. Let $l = |H^{\geq j}|$. Thus $l \leq t$ and $|H'^{\geq j}| \leq l-1$. Let R_H be the concatenation of $\{R_i, i \in H^{\geq j}\}$ and $R'_{H'}$ be the concatenation of $\{R_i', i \in H'^{\geq j}\}$.

By the above equation and the hybrid argument we have that

$$(R_H, \overline{R_{j-1}}, \overline{R'_{j-1}}) \approx_{3t^2 \cdot 2^{-3.75\ell}} (U_{ld}, \overline{R_{j-1}}, \overline{R'_{j-1}}).$$

Thus now we can first fix $\overline{R'_{j-1}}$, and with probability $1 - 2^{-1.25\ell}$ over this fixing, we have

$$R_H \approx_{3t^2 \cdot 2^{-2.5\ell}} U_{ld}.$$

We now fix $R'_{H'}$. Since $|H'^{\geq j}| \leq l - 1$, the size of $R'_{H'}$ is at most $(l-1)d$. Thus by Lemma 6 we have that with probability at least $1 - (3t^2 + 1) \cdot 2^{-1.25\ell}$ over this fixing, R_H is $2^{-1.25\ell}$-close to having min-entropy $d - 1.25\ell > \ell$. Note that after we fix $\overline{R'_{j-1}}$ and $R'_{H'}$, we have also fixed V'_2. Since W' and Y'_2 are already fixed, V'_1 is also fixed. Thus Z' is fixed. Therefore altogether we have that with probability $1 - 2 \cdot 2^{-2\ell} - t2^{-1.25\ell} = 1 - O(2^{-\ell})$ over the fixings of Y, with probability $1 - 2^{-1.25\ell} - (3t^2 + 1) \cdot 2^{-1.25\ell} = 1 - O(2^{-\ell})$ over the fixings of Z', Z is $2^{-1.25\ell}$-close to having min-entropy ℓ.

Combining **Case 1** and **Case 2**, and notice that the fraction of "bad seeds" that an adversary can achieve is at most the sum of the fraction of bad seeds in both cases. Thus by choosing an appropriate $\ell = O(s)$ we have that the construction is a $(k, s, 2^{-s})$-non-malleable condenser with seed length $O(\log n + s)^2$.

Combining Theorem 11 and Theorem 13, we immediately get a 2-round privacy amplification protocol with optimal entropy loss for any (n, k) source.

Theorem 14. *There exists a constant C such that for any $\epsilon > 0$ with $k \geq C(\log n + \log(1/\epsilon))^2$, there exists an explicit 2-round privacy amplification protocol for (n, k) sources with security parameter $\log(1/\epsilon)$, entropy loss $O(\log n + \log(1/\epsilon))$ and communication complexity $O(\log n + \log(1/\epsilon))^2$.*

In fact, we have a slightly simpler protocol that uses the look-ahead extractor and MAC somewhat more directly, while achieving the same performance.

We assume that the shared weak random source has min-entropy k, and the error ϵ we seek satisfies $\epsilon < 1/n$ and $k > C(\log n + \log(1/\epsilon))^2$ for some constant $C > 1$. For convenience, in the description below we introduce an "auxiliary" security parameter s. Eventually, we will set $s = \log(C'/\epsilon) + O(1) = \log(1/\epsilon) + O(1)$, so that $C'/2^s < \epsilon$, for a sufficiently large constant C' related to the number of "bad" events we need to account for. We need the following building blocks:

- Let Ext be a $(k, 2^{-5s})$-extractor with optimal entropy loss and seed length $d = O(\log n + s) > 202s$, from Theorem 8. Assume that $k \geq 15d^2$.
- Let Raz be the two source extractor from Theorem 9.
- Let MAC be the ("leakage-resilient") MAC, as in Theorem 6, with tag length $v = 2s$ and key length $\ell = 2v = 4s$.
- Let laExt be the look-ahead extractor defined above, using Ext as Ext_q and Ext_s. laExt is set up to extract from x using seed (q, s_0) such that $q = y_2$ and s_0 is the string that contains the first $40d$ bits of y_2, and output a string $r \in (\{0,1\}^d)^t$ with $t = 4d$.

- Let $\mathsf{laMAC}_r(\mu)$ be the function defined above.
- In the protocol Alice will sample three random strings Y_1, Y_2, Y_3, with size d, $12d^2$ and $50d^2$ respectively.

Using the above building blocks, the protocol is given in Figure 4. To emphasize the adversary Eve, we use letters with 'prime' to denote all the variables seen or generated by Bob; e.g., Bob picks W', but Alice may see a different W.

Alice: X	Eve: E	Bob: X
Sample random $Y = (Y_1, Y_2, Y_3)$.		
Compute $R_2 = \mathsf{laExt}(X, Y_2)$.		
$Z = \mathsf{laMAC}_{R_2}(Y_1)$.		
$R_1 = \mathsf{Ext}(X, Y_1)$ and output $4s$ bits.		
	(Y_1, Y_2, Y_3) $\longrightarrow$	
	(Y_1', Y_2', Y_3')	
		Sample random W' with d bits.
		Compute $R_2' = \mathsf{laExt}(X, Y_2')$.
		$Z' = \mathsf{laMAC}_{R_2'}(Y_1')$.
		$R_1' = \mathsf{Ext}(X, Y_1')$ and output $4s$ bits.
		$T_1' = \mathsf{Raz}(Y_3', Z')$ with s bits,
		$T_2' = \mathsf{MAC}_{R_1'}(W')$.
		Set final $R_B = \mathsf{Ext}(X, W')$.
	(W', T_1', T_2') $\longleftarrow$	
	(W, T_1, T_2)	
If $T_1 \neq \mathsf{Raz}(Y_3, Z)$ or		
$T_2 \neq \mathsf{MAC}_{R_1}(W)$ *reject.*		
Set final $R_A = \mathsf{Ext}(X, W)$.		

Fig. 4. 2-round Privacy Amplification Protocol

Theorem 15. *Assume that $k > C(\log n + \log(1/\epsilon))^2$ for some constant $C > 1$. The above protocol is a privacy amplification protocol with security parameter $\log(1/\epsilon)$, entropy loss $O(\log(1/\epsilon))$ and communication complexity $O(\log(1/\epsilon)^2)$.*

Proof. The proof can be divided into two cases: whether the adversary changes Y_1 or not.

Case 1: The adversary does not change Y_1. In this case, note that $R_1 = R_1'$ and is 2^{-5s}-close to uniform in Eve's view (even conditioned on Y_1, Y_2, Y_3). Thus the property of the MAC guarantees that Bob can authenticate W' to Alice. However, one thing to note here is that Eve has some additional information, namely T_1' which can leak information about the MAC key. On the other hand, the size of T_1' is s, thus by Lemma 2 the average conditional min-entropy $H_\infty(R_1|T_1')$ is at least $3s$. Therefore by Theorem 6 the probability that Eve can change W' to a different W without causing Alice to reject is at most

$$\left\lceil \frac{d_1}{2s} \right\rceil 2^{2s-\widetilde{H}_\infty(R_1|T_1')} + 2^{-5s} \leq O(2^{2s-3s}) + 2^{-5s} \leq O(2^{-s}).$$

When $W = W'$, by Theorem 8 $R_A = R_B$ and is 2^{-5s}-close to uniform in Eve's view.

Case 2: The adversary does change Y_1. Thus we have $Y_1 \neq Y_1'$. Here the proof is similar to the proof of the non-malleable condenser. We first fix Y_1 and Y_1'. Note that after this fixing, R_1 and R_1' are now deterministic functions of X. We now further fix R_1 and R_1' and after this fixing, X and (Y_2, Y_3) are still independent. Since the total size of (R_1, R_1') is $8s$, by Lemma 4 we have that with probability $1 - 2^{-2s}$ over this fixing, X still has min-entropy at least $15d^2 - 8s - 2s > 12d^2$. Note also that after this fixing, Y_2' is a deterministic function of (Y_2, Y_3). However, since Y_1' may be a function of Y_2, fixing Y_1' may cause Y_2 to lose entropy. Note that Y_1' only has size d, thus by Lemma 4, with probability $1-2\cdot 2^{-2s}$ over the fixing of (Y_1, Y_1'), we have that Y_2 has min-entropy at least $12d^2 - d - 2s > 11d^2$ and S_0 has min-entropy at least $40d - d - 2s > 38d$.

Now assume that X has min-entropy at least $12d^2$, Y_2 has min-entropy at least $11d^2$ and S_0 has min-entropy at least $38d$. This happens with probability at least $1 - O(2^{-s})$. For any $i, 0 \leq i \leq t-1$, let $\overline{S_i} = (S_0, \cdots, S_i)$, $\overline{S_i'} = (S_0', \cdots, S_i')$, $\overline{R_i} = (R_0, \cdots, R_i)$ and $\overline{R_i'} = (R_0', \cdots, R_i')$. Again by Lemma 7 we have that for any i,

$$\begin{aligned}&(R_i, \overline{S_{i-1}}, \overline{S_{i-1}'}, \overline{R_{i-1}}, \overline{R_{i-1}'}, S_i, S_i', Y_2, Y_2')\\ \approx_{(2i+2)2^{-5s}} &(U_d, \overline{S_{i-1}}, \overline{S_{i-1}'}, \overline{R_{i-1}}, \overline{R_{i-1}'}, S_i, S_i', Y_2, Y_2').\end{aligned}$$

Thus for any i, we have

$$(R_i, \overline{R_{i-1}}, \overline{R_{i-1}'}, Y_2, Y_2') \approx_{(2i+2)2^{-5s}} (U_d, \overline{R_{i-1}}, \overline{R_{i-1}'}, Y_2, Y_2').$$

Now by the same analysis as in the proof of the non-malleable condenser (and recall that $Y_1 \neq Y_1'$), we have that with probability $1 - t2^{-1.25\ell}$ over the fixing of (Y_2, Y_2'), with probability at least $1 - (3t^2+1)\cdot 2^{-1.25s}$ over the fixing of Z', Z is $2^{-1.25s}$-close to having min-entropy $d - 1.25s > 200s$.

Note that we have now fixed (Y_1, Y_1', Y_2, Y_2') and (R_1, R_1', Z'). After all these fixings, Z is a deterministic function of X and is $2^{-1.25s}$-close to having min-entropy $200s$. Thus Z is independent of Y_3 (note that Z' is also a deterministic function of X, thus fixing Z' does not influence the independence of Z and Y_3). Note that after these fixings, Y_3' is a deterministic function of Y_3, and since the size of (Y_1', Y_2') is $d + 12d^2 < 13d^2$, by Lemma 4 Y_3 is 2^{-s}-close to having min-entropy $50d^2 - 13d^2 - s > 36d^2$. Thus by Theorem 9 we have

$$(\mathsf{Raz}(Y_3, Z), Y_3, Y_3') \approx_{O(2^{-s})} (U_s, Y_3, Y_3').$$

Since we already fixed (Y_1, Y_1', Y_2, Y_2') and (R_1, R_1', Z'), and W' is independent of all random variables above, this also implies that

$$(\mathsf{Raz}(Y_3, Z), R'_1, Z', Y, Y', W') \approx_{O(2^{-s})} (U_s, R'_1, Z', Y, Y', W').$$

Note that $T'_1 = \mathsf{Raz}(Y'_3, Z')$ and $T'_2 = \mathsf{MAC}_{R'_1}(W')$. Thus we have

$$(\mathsf{Raz}(Y_3, Z), T'_1, T'_2, Y, Y', W') \approx_{O(2^{-s})} (U_s, T'_1, T'_2, Y, Y', W').$$

Therefore, the probability that the adversary can guess the correct T_1 is at most $2^{-s} + O(2^{-s}) = O(2^{-s})$. For an appropriately chosen $s = \log(1/\epsilon) + O(1)$ this is at most ϵ. Note that conditioned on the fixing of Y, the random variables that are used to authenticate W' are (R_1, T_1), which are deterministic functions of X and have size $O(s)$, thus the entropy loss of the protocol is $O(\log(1/\epsilon))$. The communication complexity can be easily verified to be $O(\log(1/\epsilon)^2)$.

7 Conclusions and Open Problems

In this paper we construct explicit non-malleable condensers for arbitrary min-entropy, and use them to give an explicit 2-round privacy amplification protocol with optimal entropy loss for arbitrary min-entropy k, with security parameter up to $s = \Omega(\sqrt{k})$. This is the first explicit protocol that simultaneously achieves optimal parameters in both round complexity and entropy loss, for arbitrary min-entropy.

We then generalize this result to give a privacy amplification protocol that runs in $O(s/\sqrt{k})$ rounds and achieves optimal entropy loss for arbitrary min-entropy k, with security parameter up to $s = \Omega(k)$. This significantly improves the protocol in [4]. In the special case where $k = \delta n$ for some constant $\delta > 0$, we give better non-malleable condensers and a 2-round privacy amplification protocol with optimal entropy loss for security parameter up to $s = \Omega(k)$, which improves the entropy loss and communication complexity of the 2-round protocol in [2].

Some open problems include constructing better non-malleable extractors or non-malleable condensers, and to construct optimal privacy amplification protocols for security parameter bigger than $\sqrt{k}$. Another interesting problem is to find other applications of non-malleable extractors or non-malleable condensers.

References

1. Dodis, Y., Wichs, D.: Non-malleable extractors and symmetric key cryptography from weak secrets. In: Proceedings of the 41st Annual ACM Symposium on Theory of Computing, pp. 601–610 (2009)
2. Li, X.: Non-malleable extractors, two-source extractors and privacy amplification. In: Proceedings of the 53nd Annual IEEE Symposium on Foundations of Computer Science (2012)
3. Li, X.: Design extractors, non-malleable condensers and privacy amplification. In: Proceedings of the 44th Annual ACM Symposium on Theory of Computing (2012)
4. Chandran, N., Kanukurthi, B., Ostrovsky, R., Reyzin, L.: Privacy amplification with asymptotically optimal entropy loss. In: Proceedings of the 42nd Annual ACM Symposium on Theory of Computing, pp. 785–794 (2010)

5. Dodis, Y., Ong, S.J., Prabhakaran, M., Sahai, A.: On the (im)possibility of cryptography with imperfect randomness. In: FOCS 2004, pp. 196–205 (2004)
6. Bennett, C., Brassard, G., Robert, J.M.: Privacy amplification by public discussion. SIAM Journal on Computing 17, 210–229 (1988)
7. Maurer, U.M., Wolf, S.: Privacy amplification secure against active adversaries. In: Kaliski Jr., B.S. (ed.) CRYPTO 1997. LNCS, vol. 1294, pp. 307–321. Springer, Heidelberg (1997)
8. Dodis, Y., Katz, J., Reyzin, L., Smith, A.: Robust fuzzy extractors and authenticated key agreement from close secrets. In: Dwork, C. (ed.) CRYPTO 2006. LNCS, vol. 4117, pp. 232–250. Springer, Heidelberg (2006)
9. Renner, R.S., Wolf, S.: Unconditional authenticity and privacy from an arbitrarily weak secret. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 78–95. Springer, Heidelberg (2003)
10. Kanukurthi, B., Reyzin, L.: Key agreement from close secrets over unsecured channels. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 206–223. Springer, Heidelberg (2009)
11. Dodis, Y., Li, X., Wooley, T.D., Zuckerman, D.: Privacy amplification and non-malleable extractors via character sums. In: Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science (2011)
12. Cohen, G., Raz, R., Segev, G.: Non-malleable extractors with short seeds and applications to privacy amplification. In: Proceedings of the 27th Annual IEEE Conference on Computational Complexity (2012)
13. Dodis, Y., Yu, Y.: Overcoming weak expectations. Manuscript (September 2012)
14. Aggarwal, D., Dodis, Y., Jafargholi, Z., Miles, E., Reyzin, L.: Amplifying privacy in privacy amplification. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part II. LNCS, vol. 8617, pp. 183–198. Springer, Heidelberg (2014)
15. Dziembowski, S.: Intrusion-resilience via the bounded-storage model. In: Proceedings of the 3rd Theory of Cryptography Conference (2006)
16. Crescenzo, G.D., Lipton, R.J., Walfish, S.: Perfectly secure pass- word protocols in the bounded retrieval model. In: Proceedings of the 3rd Theory of Cryptography Conference (2006)
17. Raz, R.: Extractors with weak random seeds. In: Proceedings of the 37th Annual ACM Symposium on Theory of Computing, pp. 11–20 (2005)
18. Barak, B., Kindler, G., Shaltiel, R., Sudakov, B., Wigderson, A.: Simulating independence: New constructions of condensers, Ramsey graphs, dispersers, and extractors. In: Proceedings of the 37th Annual ACM Symposium on Theory of Computing, pp. 1–10 (2005)
19. Zuckerman, D.: Linear degree extractors and the inapproximability of max clique and chromatic number. In: Theory of Computing, pp. 103–128 (2007)
20. Dodis, Y., Ostrovsky, R., Reyzin, L., Smith, A.: Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. SIAM Journal on Computing 38, 97–139 (2008)
21. Guruswami, V., Umans, C., Vadhan, S.: Unbalanced expanders and randomness extractors from Parvaresh-Vardy codes. Journal of the ACM 56(4) (2009)
22. Nisan, N., Zuckerman, D.: Randomness is linear in space. Journal of Computer and System Sciences 52(1), 43–52 (1996)
23. Schulman, L.J., Zuckerman, D.: Asymptotically good codes correcting insertions, deletions, and transpositions. IEEE Transactions on Information Theory 45(7), 2552–2557 (1999)