

Non-Malleable Extractors with Short Seeds and Applications to Privacy Amplification

Gil Cohen*

Ran Raz*

Gil Segev†

Abstract

Motivated by the classical problem of privacy amplification, Dodis and Wichs (STOC '09) introduced the notion of a *non-malleable extractor*, significantly strengthening the notion of a *strong extractor*. A non-malleable extractor is a function $\text{nmExt} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ that takes two inputs: a weak source W and a uniform (independent) seed S , and outputs a string $\text{nmExt}(W, S)$ that is nearly uniform given the seed S as well as the value $\text{nmExt}(W, S')$ for any seed $S' \neq S$ that may be determined as an arbitrary function of S .

The first explicit construction of a non-malleable extractor was recently provided by Li, Wooley and Zuckerman (arXiv:1102.5415 '11). Their extractor works for any weak source with min-entropy rate $1/2 + \delta$, where $\delta > 0$ is an arbitrary constant, and outputs up to a linear number of bits, but suffers from two drawbacks. First, the length of its seed is linear in the length of the weak source (which leads to privacy amplification protocols with high communication complexity). Second, the construction is conditional: when outputting more than a logarithmic number of bits (as required for privacy amplification protocols) its efficiency relies on a longstanding conjecture on the distribution of prime numbers.

In this paper we present an *unconditional* construction of a non-malleable extractor with *short seeds*. For any integers n and d such that $2.01 \cdot \log n \leq d \leq n$, we present an explicit construction of a non-malleable extractor $\text{nmExt} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$, with $m = \Omega(d)$, and error exponentially small in m . The extractor works for any weak source with min-entropy rate $1/2 + \delta$, where $\delta > 0$ is an arbitrary constant. Moreover, our extractor in fact satisfies an even more general notion of non-malleability: its output $\text{nmExt}(W, S)$ is nearly uniform given the seed S as well as the values $\text{nmExt}(W, S_1), \dots, \text{nmExt}(W, S_t)$ for several seeds S_1, \dots, S_t that may be determined as an arbitrary function of S , as long as $S \notin \{S_1, \dots, S_t\}$.

By instantiating the framework of Dodis and Wichs with our non-malleable extractor, we obtain the first 2-round privacy amplification protocol for min-entropy rate $1/2 + \delta$ with asymptotically optimal entropy loss and poly-logarithmic communication complexity. This improves the previously known 2-round privacy amplification protocols: the protocol of Dodis and Wichs whose entropy loss is not asymptotically optimal, and the protocol of Li, Wooley and Zuckerman whose communication complexity is linear.

*Department of Computer Science and Applied Mathematics, Weizmann Institute of Science, Rehovot 76100, Israel. Email: {gil.cohen, ran.raz}@weizmann.ac.il. Research supported by Israel Science Foundation (ISF) grant.

†Microsoft Research, Mountain View, CA 94043, USA. Email: gil.segev@microsoft.com.

Contents

1	Introduction	1
1.1	Constructions of Non-Malleable Extractors	2
1.2	Privacy Amplification via Non-Malleable Extractors	3
1.3	Our Results	4
1.3.1	Explicit Construction of a Non-Malleable Extractor	4
1.3.2	Applications to Privacy Amplification	4
2	Overview of Our Results	5
2.1	The Non-Malleable Extractor	5
2.2	The Privacy Amplification Protocol	6
3	Preliminaries	7
3.1	Flat Sources	7
3.2	Explicit Constructions of Strong Seeded-Extractors	8
3.3	Fooling Linear Tests of Bounded Size	8
3.4	Privacy Amplification Protocols	8
3.5	Message Authentication Codes	9
3.6	Basic Claims in Probability Theory	10
4	A Central Lemma from [Raz05]	10
5	A Simple Lemma about Graphs	12
6	A Conditional Parity Lemma	13
7	Proof of Main Theorem	15
8	The Privacy Amplification Protocol	20
	References	24

1 Introduction

Randomness extractors are functions that extract nearly uniform bits from biased random sources. Among the wide variety of settings in which randomness extractors play an instrumental role is the classical problem of *privacy amplification* ([BBR88, Mau92, BBCM95]). This problem considers a setting in which two parties, Alice and Bob, begin by sharing a secret $W \in \{0, 1\}^n$ whose distribution may be far from uniform. The parties interact over a public communication channel in the presence of an adversary, Eve, and would like to securely agree on a nearly uniform secret $R \in \{0, 1\}^m$.

In various applications, the secret W is often chosen, for example, as a human-memorizable password or some biometric data, both of which are typically of rather low min-entropy, or even as a truly uniform secret which may have been partially leaked to Eve. In this paper we consider the information-theoretic setting of the problem where no computational assumptions are made (in particular, Eve is assumed to be computationally unbounded).

Formally, a *source of randomness* (or simply, a *source*) of length n is a random variable W of length n bits. We say that a source is *weak* if its distribution is not uniform. The standard measure for the amount of randomness contained in a source is its *min-entropy*: an n -bit random variable W has min-entropy at least k if for every $w \in \{0, 1\}^n$ the probability that $W = w$ is at most 2^{-k} . In this case we say that W is an (n, k) -*source*. We define the *min-entropy rate* of an (n, k) -source as the ratio k/n .

Strong Extractors. In the presence of a *passive* adversary that is assumed to only observe the communication channel between the parties, any *strong extractor* provides an elegant solution to the privacy amplification problem. Informally, a strong (seeded-)extractor is a function $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ that takes two inputs, a weak source W and an independent uniform seed S , and outputs a string $\text{Ext}(W, S)$ that is nearly uniform given the seed S . Using a strong extractor, Alice simply sends Bob a seed S that is chosen uniformly at random, and they both compute $R = \text{Ext}(W, S)$ which is guaranteed to be nearly uniform from Eve's point of view.

Formally, we denote by U_n the uniform distribution over $\{0, 1\}^n$, and for a random variable X over $\{0, 1\}^n$ we denote by (X, U_m) the joint distribution of X and an independent random variable that is uniformly distributed over $\{0, 1\}^m$. In general, many times we will confuse notations between random variables and their distributions. We measure the distance between distributions by the \mathcal{L}_1 norm. Two distributions, X and Y , are ϵ -close if $\|X - Y\|_1 = \sum_s |\Pr[X = s] - \Pr[Y = s]| \leq \epsilon$. This measure of distance is also referred to as *statistical distance*. More precisely, the statistical distance, denoted by $\text{SD}(X, Y)$ is defined as $\frac{1}{2}\|X - Y\|_1$. Given an (n, k) -source W with an unknown distribution, it well-known that if $k \leq n - 1$ one cannot deterministically extract even one non-constant bit from W (unless additional information about the distribution of W is given). Seeded-extractors overcome this barrier by extracting nearly uniform randomness from a weak source, using an additional number of truly random bits, called *seed*. A seeded-extractor is *strong* if its output is almost independent of the seed.

Definition 1.1 (Seeded-Extractor). *A function $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is a (k, ϵ) -seeded-extractor if for every (n, k) -source W and an independent random variable S uniformly distributed over $\{0, 1\}^d$, the distribution of $\text{Ext}(W, S)$ is ϵ -close to U_m . A (k, ϵ) -seeded-extractor is strong if for X and S as above, the distribution of $(\text{Ext}(X, S), S)$ is ϵ -close to (U_m, U_d) .*

Non-Malleable Extractors. In the presence of an *active* adversary that fully controls the communication channel between the parties, however, privacy amplification is significantly more chal-

lenging. One of the main reasons is that in addition to preventing Eve from learning essentially any information on the resulting secret R , a secure privacy amplification protocol should also prevent Eve from causing the parties to output different secrets R and R' .

In this light, extensive research has been devoted for designing privacy amplification protocols that are secure under active attacks (see [Mau97, MW97, Wol98, MW03, RW03, DKRS06, DW09, KR09, CKOR10] and the references therein), with the natural goal of optimizing the efficiency of such protocols. The main measures of efficiency in this setting are the *entropy loss* (defined as the difference between the entropy of the weak secret W and the length of the resulting secret R), the *communication complexity*, and the *round complexity* (i.e. the number of rounds).

A major progress in the design of privacy amplification protocols was recently made by Dodis and Wichs [DW09]. Their approach relies on introducing the new and elegant notion of a *non-malleable extractor*, significantly strengthening the notion of a strong extractor. Informally, a non-malleable extractor is a function $\text{nmExt} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ that takes two inputs: a weak source W and an independent uniform seed S , and outputs a string $\text{nmExt}(W, S)$ that is nearly uniform given the seed S as well as the value $\text{nmExt}(W, S')$ for any seed $S' \neq S$ that may be determined as an arbitrary function of S .

Definition 1.2 (Adversarial Function). *Let $\mathcal{A} : \{0, 1\}^d \rightarrow \{0, 1\}^d$. We say that \mathcal{A} is an adversarial function if it has no fixed points. That is, for every $s \in \{0, 1\}^d$ it holds that $\mathcal{A}(s) \neq s$.*

Definition 1.3 (Non-Malleable Extractor). *A function $\text{nmExt} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is a (k, ϵ) -non-malleable extractor if for every (n, k) -source W and independent random variable S uniformly distributed over $\{0, 1\}^d$, and for every adversarial function $\mathcal{A} : \{0, 1\}^d \rightarrow \{0, 1\}^d$,*

$$\|(\text{Ext}(W, S), \text{Ext}(W, \mathcal{A}(S)), S) - (U_m, \text{Ext}(W, \mathcal{A}(S)), S)\|_1 \leq \epsilon.$$

In this paper we also consider a natural generalization of non-malleable extractors, in which the adversary has the value of the extractor not only on *one* correlated seed $\mathcal{A}(S)$ of her choice, but rather on *many* correlated seeds $\mathcal{A}_1(S), \dots, \mathcal{A}_t(S)$ of her choice.

Definition 1.4 (t -Adversarial Function). *Let $t \in \mathbb{N}$. Let $\mathcal{A} : \{0, 1\}^d \rightarrow \{0, 1\}^{td}$. We think of the output of \mathcal{A} as t concatenated binary strings, each of length d . That is, we think of $\mathcal{A}(s)$ as $\mathcal{A}(s) = (\mathcal{A}_1(s), \dots, \mathcal{A}_t(s))$, where for all $i \in [t]$, \mathcal{A}_i is a function of the form $\mathcal{A}_i : \{0, 1\}^d \rightarrow \{0, 1\}^d$. We say that \mathcal{A} is a t -adversarial function if for every $i \in [t]$ the function \mathcal{A}_i is an adversarial function.*

Definition 1.5 (t -Non-Malleable Extractor). *A function $\text{nmExt} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is a (k, ϵ) - t -non-malleable extractor if for every (n, k) -source W and independent random variable S uniformly distributed over $\{0, 1\}^d$, and for every t -adversarial function $\mathcal{A} : \{0, 1\}^d \rightarrow \{0, 1\}^{td}$,*

$$\|(\text{Ext}(W, S), \text{Ext}(W, \mathcal{A}_1(S)), \dots, \text{Ext}(W, \mathcal{A}_t(S)), S) - (U_m, \text{Ext}(W, \mathcal{A}_1(S)), \dots, \text{Ext}(W, \mathcal{A}_t(S)), S)\|_1 \leq \epsilon.$$

1.1 Constructions of Non-Malleable Extractors

Although the approach of Dodis and Wichs [DW09] seems very promising, they were in fact unable to present an explicit construction of a non-malleable extractor (not even one with poor parameters). They showed, however, using the probabilistic method, that such extractors, with excellence

parameters, exist. More specifically, Dodis and Wichs proved the existence of a (k, ϵ) -non-malleable extractor $\text{nmExt} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$, as long as

$$\begin{aligned} d &> \log(n - k - 1) + 2 \log(1/\epsilon) + 5, \\ k &> 2m + 2 \log(1/\epsilon) + \log d + 6. \end{aligned}$$

Recently, Li, Wooley, and Zuckerman [LWZ11] presented the first explicit construction of a non-malleable extractor. They showed that an extractor introduced by Chor and Goldreich in the context of two-source extractors [CG88] is non-malleable as long as the weak source has min-entropy rate $1/2 + \delta$ for an arbitrarily small constant $\delta > 0$. Their extractor outputs up to a linear number of bits, but suffers from two drawbacks. First, the seed used by their extractor is of length $d = \Omega(n)$ bits, even for the purpose of extracting a single bit. Second, the construction is conditional: when outputting more than a logarithmic number of bits (as required for privacy amplification protocols) its efficiency relies on a longstanding conjecture on the distribution of prime numbers.

1.2 Privacy Amplification via Non-Malleable Extractors

Using a non-malleable extractor, Dodis and Wichs constructed the first 2-round privacy amplification protocol for any min-entropy rate that is secure against active attacks¹. Specifically, Dodis and Wichs demonstrated that the idea underlying the simple privacy amplification protocol discussed above for passive attacks can be implemented also in the setting of active attacks. Moreover, when instantiating their approach with a non-malleable extractor that enjoys sufficiently good parameters (as well as with an essentially optimal strong extractor), the resulting privacy amplification protocol in turn enjoys asymptotically optimal entropy loss $O(\log n + \log(1/\epsilon))$ and communication complexity $O(\log n + \log(1/\epsilon))$, where ϵ is the security parameter of the protocol (i.e., the protocol error).

As discussed above, Dodis and Wichs were in fact unable to present an explicit construction of a non-malleable extractor. Nevertheless, they were still able to construct an explicit privacy amplification protocol by introducing the weaker notion of a *look-ahead extractor*², for which they were able to provide an explicit construction. Using their look-ahead extractor, Dodis and Wichs constructed an explicit 2-round privacy amplification protocol with entropy loss $\beta k + O(\log^2 n + \log^2(1/\epsilon))$, for an arbitrarily small constant $\beta > 0$, and communication complexity $O(\log^2 n + \log^2(1/\epsilon))$, both of which are somewhat far from optimal³.

Li, Wooley, and Zuckerman [LWZ11] showed that when instantiating the protocol of Dodis and Wichs with their explicit non-malleable extractor one obtains a 2-round privacy amplification protocol for weak sources of min-entropy rate $1/2 + \delta$, for an arbitrarily small constant $\delta > 0$, with entropy loss $O(\log n + \log(1/\epsilon))$. However, since the seed of the extractor is of length $\Omega(n)$ bits, the resulting privacy amplification protocol suffers from communication complexity of $\Omega(n)$ bits.

Thus, although the approach of Dodis and Wichs [DW09] for privacy amplification indeed seems very promising, due to the difficulties in constructing explicit non-malleable extractors, the resulting protocols are still rather far from optimal either in their entropy loss or communication complexity.

¹They also showed that 1-round protocols do not exist when the weak secret W has min-entropy $k \leq n/2$, and are inherently inefficient in terms of communication when $n/2 < k \ll n$.

²Informally, a look-ahead extractor is a function $\text{laExt} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ that takes two inputs: a weak source W and a uniform seed S , and outputs a string $\text{laExt}(W, S)$ whose any suffix is nearly uniform given the seed S and the complementing prefix of $\text{laExt}(W, S')$ for some seed $S' \neq S$ that may be determined as an arbitrary function of S . Note that any non-malleable extractor is in particular also a look-ahead extractor.

³In fact, the dependency on the min-entropy k of the weak source can be eliminated from the entropy loss in their protocol. This can be done, for example, by instantiating the strong extractor in their protocol using Theorem 3.4 instead of Theorem 3.3.

1.3 Our Results

1.3.1 Explicit Construction of a Non-Malleable Extractor

In this paper we present an *unconditional* construction of a non-malleable extractor with *short seeds*. We prove the following theorem:

Theorem 1.6 (Main Theorem). *For any integers n and d such that $2.01 \log n \leq d \leq n$, and for any constant $\delta > 0$, there exists an explicit $((1/2 + \delta) \cdot n, 2^{-m})$ -non-malleable extractor $\text{nmExt}: \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$, with $m = \Omega(d)$.*

In particular, setting $d = 2.01 \log n$ yields the first explicit construction of a non-malleable extractor that uses a seed of length $O(\log n)$ bits⁴. This should be compared with the extractor of [LWZ11] that uses a seed of length $\Omega(n)$ bits. This improvement in the seed length is crucial for the communication complexity of the resulting privacy amplification protocols.

In addition, setting $d = n$ yields the first *unconditional* explicit construction of a non-malleable extractor that outputs $\Omega(n)$ bits. This should be compared with the extractor of [LWZ11] whose efficiency relies on an unproven conjecture (when outputting $\omega(\log n)$ bits). In fact, our extractor is the first non-malleable extractor that outputs $\omega(\log n)$ bits unconditionally.

Our main result is in fact more general. We show an explicit construction of a t -non-malleable extractor with essentially optimal parameters⁵, as long as the min-entropy rate of the weak-source is any constant larger than $1/2$ and the output length is shorter than the seed length.

Theorem 1.7 (Main Theorem – Generalized). *For any integers n , d , m and t , and for any $0 < \delta < 1/2$ such that*

$$\begin{aligned} d &\geq \frac{23}{\delta} \cdot tm + 2 \log n, \\ n &\geq \frac{160}{\delta} \cdot tm, \\ \delta &\geq 10 \cdot \frac{\log(nd)}{n}, \end{aligned}$$

there exists an explicit $((1/2 + \delta) \cdot n, 2^{-m})$ - t -non-malleable extractor $\text{nmExt}: \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$.

We note the nice symmetry between m and t in Theorem 1.7.

1.3.2 Applications to Privacy Amplification

By instantiating the framework of Dodis and Wichs [DW09] with the non-malleable extractor from Theorem 1.6 we obtain an explicit 2-round privacy amplification protocol for weak sources of min-entropy rate $1/2 + \delta$ for an arbitrarily small constant $\delta > 0$. The protocol offers a trade-off between its entropy loss and communication complexity, resulting from instantiating it with different explicit constructions of strong extractors. Specifically, it offers asymptotically optimal entropy loss $O(\log n + \log(1/\epsilon))$ with communication complexity $O(\min \{\log^2 n + \log n \cdot \log(1/\epsilon), n\})$, or entropy loss $\beta n + O(\log n + \log(1/\epsilon))$ for an arbitrarily small constant $\beta > 0$ with communication complexity $O(\log n + \log(1/\epsilon))$. In particular, we prove the following theorem:

⁴The constant 2.01 can, in fact, be replaced by any constant strictly greater than 2.

⁵We made no attempt to optimize the constants in the theorem as they depend on each other.

Theorem 1.8. *For any integer n , constant $\delta > 0$, and security parameter $\epsilon = 2^{-O(n)}$, there exists an explicit and efficient 2-round privacy amplification protocol for $(n, (1/2 + \delta)n)$ -sources with entropy loss $O(\log n + \log(1/\epsilon))$, and communication complexity $O(\min\{\log^2 n + \log n \cdot \log(1/\epsilon), n\})$.*

This is the first explicit 2-round privacy amplification protocol for min-entropy rate $1/2 + \delta$ with asymptotically optimal entropy loss and poly-logarithmic communication complexity. This should be compared to the previously known 2-round protocols: the protocol of Dodis and Wichs whose entropy loss is not asymptotically optimal, and the protocol of Li, Wooley and Zuckerman whose communication complexity is linear in the length of the weak secret.

2 Overview of Our Results

In this section we overview the main ideas underlying our constructions⁶. We begin with the construction of the non-malleable extractor in Section 2.1, and then proceed with the resulting privacy amplification protocol in Section 2.2.

2.1 The Non-Malleable Extractor

Raz [Raz05] gave an explicit construction of seeded-extractors⁷, based on small probability spaces of 0-1 random variables that have small bias for linear tests of bounded size (see Section 3.3). We begin by describing the construction of these extractors, starting with extractors that output one bit, and then turn to describe our approach.

The Extractor of Raz [Raz05]. Let $D = 2^d$, and let Z_1, \dots, Z_D be 0-1 random variables that are ϵ -biased for linear tests of size k that can be constructed using n random bits. We define $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}$ by $\text{Ext}(w, s) = Z_s(w)$. That is, $\text{Ext}(w, s)$ is the value of the random variable Z_s when using w as the value of the n bits needed to produce Z_1, \dots, Z_D . In other words, w is used to choose the point in the probability space, and s is used to choose the variable from Z_1, \dots, Z_D that we evaluate.

Extracting many bits is done similarly: Let $D = m \cdot 2^d$, and let Z_1, \dots, Z_D be 0-1 random variables that are ϵ -biased for linear tests of size k that can be constructed using n random bits. We interpret the set of indices $\{1, \dots, D\}$ as the set $\{(i, s) : i \in [m], s \in \{0, 1\}^d\}$. We define $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ by $\text{Ext}_i(w, s) = Z_{(i,s)}(w)$, where $\text{Ext}_i(w, s)$ denotes the i^{th} bit of $\text{Ext}(w, s)$. In other words, w is used to choose the point in the probability space, and the pair (i, s) is used to choose the variable from Z_1, \dots, Z_D that we evaluate.

Raz showed that the above extractor, based on *any* small probability space of 0-1 random variables that have small bias for linear tests of bounded size, is an excellent extractor. Specifically, using any of the probability spaces from [AGHP92], one gets a $((1/2 + \delta) \cdot n, 2^{-m})$ -seeded extractor $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with $m = \min(\Omega(\delta n), \Omega(d))$ as long as $d = \Omega(\log n)$. Moreover, Raz showed that this extractor is a strong seeded-extractor.

Our Approach. In this paper we show that the extractor of Raz is in fact non-malleable with essentially the same parameters. Moreover, we show that it is t -non-malleable with optimal dependency on t . We now present the proof strategy. For simplicity, we focus on the case $m = t = 1$ (though we do not prove this special case in the paper separately, but rather give a proof for general

⁶In this section we use some well known notions. The formal definitions can be found in Section (Section 3), which an unfamiliar reader might prefer to read first.

⁷One can find in [Raz05] a construction of two-sources extractors, as well as other types of pseudo-random objects.

m and t). The proof strategy for the t -non-malleability of the extractor that extracts m bits follows by the same logic, but it is more technical.

Assume for a contradiction that Ext as defined above is not non-malleable. This implies the existence of a weak-source W and an adversarial function $\mathcal{A} : \{0, 1\}^d \rightarrow \{0, 1\}^d$ such that for a typical seed $s \in \{0, 1\}^d$, the value $\text{Ext}(W, s)$ is correlated to $\text{Ext}(W, \mathcal{A}(s))$. We can then find a large set of seeds $S \subseteq \{0, 1\}^d$ such that for every $s \in S$, the random variable $Y_s = \text{Ext}(W, s) \oplus \text{Ext}(W, \mathcal{A}(s))$ is biased.

At this point we consider the directed graph $G = (S \cup \mathcal{A}(S), E)$ where $E = \{(s, \mathcal{A}(s)) : s \in S\}$. We note that G has no self loops, but it might be the case that G contains cycles. We prove the existence of a large subset $S' \subseteq S$ such that the induced graph of G by $S' \cup \mathcal{A}(S')$ is acyclic. To this end we prove a simple lemma about graphs (see Section 5).

For every $s \in S'$, define $Y'_s = Z_s \oplus Z_{\mathcal{A}(s)}$. In the next step of the proof we prove that the set of random variables $(Y'_s)_{s \in S'}$ is ϵ -biased for linear tests of size at most $k/2$. This follows easily by the acyclicity of the above mentioned graph and by the fact that for every $s \in S'$, it holds that Y'_s is a parity of two random variables from a probability space that ϵ -fools linear tests of size k (see Claim 7.2).

Now we consider the extractor that is built upon the random variables $(Y'_s)_{s \in S'}$ as described in the beginning of the section (where the $(Y'_s)_{s \in S'}$ play the role of $(Z_i)_{i=1}^d$). The result of Raz, which holds for *any* probability space that fools linear tests of bounded size, implies that this is a good seeded-extractor. This yields a contradiction (for an appropriate choice of parameters) when feeding the weak-source W to this extractor, because the random variables $(Y_s)_{s \in S'}$ are all biased (see Section 7).

2.2 The Privacy Amplification Protocol

As discussed in Section 1.3.2, by instantiating the framework of Dodis and Wichs [DW09] with our non-malleable extractor we obtain the first explicit 2-round privacy amplification protocol for weak sources of min-entropy rate $1/2 + \delta$, for an arbitrarily small constant $\delta > 0$, with asymptotically optimal entropy loss and poly-logarithmic communication complexity. In what follows we first overview the main idea underlying the Dodis-Wichs protocol, and then discuss the parameters that we obtain by instantiating it with our non-malleable extractor.

The Dodis-Wichs Protocol. In the presence of a *passive* adversary that is assumed to only observe the communication channel between the parties, the privacy amplification problem is well-understood. Specifically, any strong extractor Ext yields the following elegant solution: Alice sends Bob a uniform seed S for Ext , and they both compute $R = \text{Ext}(W, S)$, where W is their shared weak secret. The property of the strong extractor guarantees that the resulting value R is nearly uniform from the adversary's point of view.

The main idea underlying the approach of Dodis and Wichs is that a non-malleable extractor nmExt can be used for implementing the above elegant solution in the presence of an *active* adversary. Specifically, the non-malleable extractor is used for authenticating the seed S , and as long as the communication complexity involved in the authentication is rather small, then W still has sufficient min-entropy which can be extracted as $R = \text{Ext}(W, S)$.

For authenticating the seed S , in the first round of the protocol Alice chooses a uniform seed Y for a non-malleable extractor nmExt , and computes a key $\text{key} = \text{nmExt}(W, Y)$ for a one-time message-authentication code MAC. The adversary may modify Y to any value Y' , and in this case Bob might compute a different authentication key $\text{key}' = \text{nmExt}(W, Y')$. Then, in the second round of the protocol, Bob samples a uniform seed S' for a strong extractor Ext , and sends it to Alice

together with the authentication tag $\sigma' = \text{MAC}_{\text{key}'}(S')$. At this point Bob concludes his part of the protocol by outputting the value $R' = \text{Ext}(W, S')$. The adversary may modify the pair (S', σ') to any pair (S, σ) , and Alice verifies that $\sigma = \text{MAC}_{\text{key}}(S)$. If the verification fails then Alice aborts, and otherwise Alice outputs $R = \text{Ext}(W, S)$.

Note that if the adversary does not modify the seed Y that is chosen by Alice, then Alice and Bob share the same authentication key $\text{key} = \text{key}'$, which is nearly uniform from the adversary's point of view. Thus, the adversary has only a negligible probability of computing a valid authentication tag σ for any seed $S \neq S'$. In addition, if the adversary does modify the seed Y to a different seed Y' , then the property of the non-malleable extractor guarantees that the authentication key key computed by Alice is nearly uniform from the adversary's point of view, even if she receives key' (and, in particular, if she receives σ' which is a deterministic function of S' and key'). Thus, again, the adversary has only a negligible probability of computing a valid authentication tag σ for any seed S with respect to key (and this holds even if $S = S'$). These two observations then easily imply the security of the protocol.

Our Instantiation. In the Dodis-Wichs protocol the output $\text{key} = \text{nmExt}(W, Y)$ of the non-malleable extractor is used as a key for a one-time message authentication code. It is well known (see Theorem 3.7) that explicit and efficient constructions of message-authentication codes exist with keys and authentication tags of length $O(\log n + \log(1/\epsilon))$ bits, where n is the length of the authenticated message and ϵ is the security parameter.

Using our non-malleable extractor we can set its seed Y to be of length $O(\log n + \log(1/\epsilon))$ bits. Then, one can instantiate the strong extractor Ext with any explicit construction, where we choose the one provided by Guruswami, Umans and Vadhan [GUV09] (see Theorem 3.4). It extracts $(1/2 + \delta)n - O(\log n + \log(1/\epsilon))$ bits from the weak source W using a seed S of length $O(\log^2 n + \log n \cdot \log(1/\epsilon))$ bits. When dealing with a very small security parameter ϵ one can instead use the extractor provided by the leftover hash lemma (see Theorem 3.2) for extracting the same number of bits using a seed of length n bits. Thus, our protocol has entropy loss $O(\log n + \log(1/\epsilon))$, and communication complexity $O(\min \{\log^2 n + \log n \cdot \log(1/\epsilon), n\})$.

3 Preliminaries

The logarithm in this paper is always taken base 2. For every natural number $n \geq 1$, define $[n] = \{1, 2, \dots, n\}$. We assume for simplicity that the min-entropy b of an (n, b) -source is always an integer $\leq n$. We sometimes abuse notation and syntactically treat random variables and their distribution as equal, specifically, we denote by U_m a random variable which is uniformly distributed over $\{0, 1\}^m$, and that is independent of all other random variables in context.

3.1 Flat Sources

Let X be an (n, b) -source. We say that the source X is flat if it is uniformly distributed over a set $S_X \subseteq \{0, 1\}^n$ of size 2^b . The following lemma, proved by Chor and Goldreich [CG88], shows that the distribution of any (n, b) -source is a convex combination of distributions of flat (n, b) -sources. Hence, in most cases, it will be enough to consider flat sources rather than general weak sources.

Lemma 3.1. *The distribution of any (n, b) -source is a convex combination of distributions of flat (n, b) -sources.*

3.2 Explicit Constructions of Strong Seeded-Extractors

For instantiating our privacy amplification protocol we rely on the following explicit constructions of strong seeded-extractors. The first is known as the leftover hash lemma due to Impagliazzo, Levin and Luby [ILL89], and the second and third are due to Guruswami, Umans and Vadhan [GUV09].

Theorem 3.2 ([ILL89]). *For all integers $n \geq k > m > 0$ and for any $\epsilon > 0$ such that $m \leq k - 2 \log(1/\epsilon)$, there is an explicit construction of a strong (k, ϵ) -extractor $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$, where $d = O(n)$.*

Theorem 3.3 ([GUV09]). *For every constant $\beta > 0$, for all integers $n \geq k > m > 0$ such that $m \leq (1 - \beta)k$, and for any $\epsilon > 0$, there is an explicit construction of a strong (k, ϵ) -extractor $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$, where $d = O(\log n + \log(1/\epsilon))$.*

Theorem 3.4 ([GUV09]). *For all integers $n \geq k > m > 0$, and for any $\epsilon > 0$ such that $m \leq k - 2 \log(1/\epsilon) - O(1)$, there is an explicit construction of a strong (k, ϵ) -extractor $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$, where $d = \log n + O(\log k \cdot \log(k/\epsilon))$.*

3.3 Fooling Linear Tests of Bounded Size

A random variable Z over $\{0, 1\}$ is ϵ -biased if $\text{bias}(Z) = |\Pr[Z = 0] - \Pr[Z = 1]| \leq \epsilon$, that is, if its distribution is ϵ -close to uniform. A sequence of 0-1 random variables Z_1, \dots, Z_N is ϵ -biased for linear tests of size k if the exclusive-or of any nonempty set of cardinality at most k of these variables is ϵ -biased, that is, for any nonempty $\tau \subseteq [N]$, such that $|\tau| \leq k$, the random variable $Z_\tau = \bigoplus_{i \in \tau} Z_i$ is ϵ -biased. We say in this case, that the sequence Z_1, \dots, Z_N ϵ -fools linear tests of size k .

Explicit constructions of small probability spaces on N random variables that are ϵ -biased for linear tests of size k were given in [NN93], [AGHP92]. In particular, [AGHP92] showed that for every $k, N \geq 2$, variables Z_1, \dots, Z_N as above can be explicitly constructed using $2 \cdot \lceil \log(1/\epsilon) + \log k + \log \log N \rceil$ random bits.

3.4 Privacy Amplification Protocols

Our definition of a privacy amplification protocol (also known as an information-theoretic key-agreement protocol) follows that of Dodis and Wichs [DW09]. In a privacy amplification protocol, two parties, Alice and Bob, begin by sharing a weak secret $W \in \{0, 1\}^n$, that is, a string sampled from a weak-source W . The parties interact over a public communication channel in the presence of an adversary, Eve, and would like to securely agree on a nearly uniform secret $R \in \{0, 1\}^m$. In this paper we consider the information-theoretic setting of the problem where no computational assumptions are made (in particular, Eve is assumed to be computationally unbounded), and the weak secret W may be sampled from any publicly known distribution subject to a pre-specified min-entropy rate. In addition, we assume that Eve is an *active* adversary that fully controls the communication channel between the parties.

At the beginning of the protocol Alice and Bob each have candidate keys R_A and R_B , respectively, which are initially set to the special value \perp . At some point during the execution of the protocol one party can reach a `KeyDerived` state and the other party can reach a `KeyConfirmed` state. Upon reaching either of these states, a party sets its candidate key to some m -bit value and does not modify it afterwards. Informally, the `KeyDerived` and `KeyConfirmed` states should be interpreted as follows:

1. If Alice reaches the `KeyDerived` state, then she possesses a uniformly random candidate key R_A , which remains private no matter how Eve acts during the remainder of the protocol execution. However, she is not sure if her key is shared with Bob, or if Bob is even involved in the protocol execution at all.
2. If Bob reaches the `KeyConfirmed` state and obtains a candidate key R_B , then Alice must have been involved in the protocol execution, she must have reached the `KeyDerived` state, and the two parties share the same key $R_A = R_B$ which is nearly uniform from Eve's point of view.

Definition 3.5 (Privacy amplification protocol). *In an (n, k, m, ϵ) -privacy amplification protocol Alice and Bob share a weak secret $W \in \{0, 1\}^n$ and have candidate keys $R_A, R_B \in \{0, 1\}^m \cup \{\perp\}$, respectively. For any strategy employed by Eve, we denote by V_E the random variable corresponding to the transcript of the protocol execution as seen by Eve. We require that for any weak secret W with min-entropy at least k the protocol satisfies the following three properties:*

1. **Correctness:** *If Eve is passive then one party reaches the `KeyDerived` state, the other party reaches the `KeyConfirmed` state, and $R_A = R_B$.*
2. **Privacy:** *Denote by KeyDerived_A and KeyDerived_B the indicators of the events in which Alice and Bob reach the `KeyDerived` state, respectively. Then, for any adversarial strategy employed by Eve the following two properties hold:*
 - (a) *If $\Pr[\text{KeyDerived}_A] > 0$ then $\text{SD}((R_A, V_E \mid \text{KeyDerived}_A), (U_m, V_E \mid \text{KeyDerived}_A)) \leq \epsilon$.*
 - (b) *If $\Pr[\text{KeyDerived}_B] > 0$ then $\text{SD}((R_B, V_E \mid \text{KeyDerived}_B), (U_m, V_E \mid \text{KeyDerived}_B)) \leq \epsilon$.*
3. **Authenticity:** *Denote by KeyConfirmed_A and KeyConfirmed_B the indicators of the events in which Alice and Bob reach the `KeyConfirmed` state, respectively. Then, for any adversarial strategy employed by Eve it holds that*

$$\Pr[(\text{KeyConfirmed}_A \vee \text{KeyConfirmed}_B) \wedge R_A \neq R_B] \leq \epsilon .$$

Given an (n, k, m, ϵ) -privacy amplification protocol we refer to $k - m$ as its *entropy loss*, and to ϵ as its *security parameter*.

3.5 Message Authentication Codes

One-time message authentication codes (MACs) provide assurance to the receiver of a message that it was sent by a specified legitimate sender, even in the presence of an active and computationally unbounded adversary who controls the communication channel. A message-authentication code for messages of length n , keys of length ℓ , and authentication tags of length τ is defined via a family of deterministic and efficiently computable functions $\{\text{MAC}_{\text{key}} : \{0, 1\}^n \rightarrow \{0, 1\}^\tau\}_{\text{key} \in \{0, 1\}^\ell}$. In terms of security the requirement is that any adversary that obtains an authentication tag on a single message m of her choice with respect to a uniform key, should have only a negligible probability of computing a valid authentication tag on a different message with respect to the same key.

Definition 3.6. *A family $\{\text{MAC}_{\text{key}} : \{0, 1\}^n \rightarrow \{0, 1\}^\tau\}_{\text{key} \in \{0, 1\}^\ell}$ of deterministic and efficiently computable functions is an ϵ -secure one-time message authentication code if for any message $m \in \{0, 1\}^n$ and function $A : \{0, 1\}^\tau \rightarrow \{0, 1\}^n \times \{0, 1\}^\tau$ it holds that*

$$\Pr_{\text{key} \leftarrow \{0, 1\}^\ell} [\text{MAC}_{\text{key}}(m') = \sigma' \wedge m \neq m' \mid (m', \sigma') = A(\text{MAC}_{\text{key}}(m))] \leq \epsilon .$$

For the construction of our privacy amplification protocol we rely on the existence of message-authentication codes with the following well-known parameters (see, for example, [KR09]):

Theorem 3.7. *For any n and $\epsilon > 0$ there exists an explicit ϵ -secure message-authentication code $\{\text{MAC}_{\text{key}} : \{0, 1\}^n \rightarrow \{0, 1\}^\tau\}_{\text{key} \in \{0, 1\}^\ell}$, where $\tau \leq \log n + \log(1/\epsilon)$ and $\ell \leq 2\tau$.*

3.6 Basic Claims in Probability Theory

The following simple claims will be used in our proofs. Claim 3.8 is a simple Markov-like inequality.

Claim 3.8. *Let X be a random variable over the real interval $[0, 1]$. Let $\mu = \mathbb{E}[X]$. Then, $\Pr[X \geq \mu/2] \geq \mu/2$.*

Proof.

$$\mu = \mathbb{E}[X] \leq \frac{\mu}{2} \cdot \Pr[X < \mu/2] + 1 \cdot \Pr[X \geq \mu/2] \leq \frac{\mu}{2} + \Pr[X \geq \mu/2].$$

Therefore, $\Pr[X \geq \mu/2] \geq \mu/2$. ■

Claim 3.9. *Let X be a random variable over $\{0, 1\}^m$. Let Y, Z be two random variables. Then,*

$$\| (X, Y, Z) - (U_m, Y, Z) \|_1 = \mathbb{E}_{z \sim Z} \| (X, Y) |_{Z=z} - (U_m, Y) |_{Z=z} \|_1.$$

Proof.

$$\begin{aligned} & \| (X, Y, Z) - (U_m, Y, Z) \|_1 = \\ & \sum_{x, y, z} \left| \Pr((X, Y, Z) = (x, y, z)) - \Pr((U_m, Y, Z) = (x, y, z)) \right| = \\ & \sum_{x, y, z} \Pr(Z = z) \cdot \left| \Pr((X, Y) = (x, y) | Z = z) - \Pr((U_m, Y) = (x, y) | Z = z) \right| = \\ & \sum_z \Pr(Z = z) \cdot \sum_{x, y} \left| \Pr((X, Y) = (x, y) | Z = z) - \Pr((U_m, Y) = (x, y) | Z = z) \right| = \\ & \mathbb{E}_{z \sim Z} \| (X, Y) |_{Z=z} - (U_m, Y) |_{Z=z} \|_1. \end{aligned}$$
■

The following is a standard lemma regarding conditional min-entropy (see, for example, [NZ96, MW97]):

Lemma 3.10. *Let X and Y be random variables, and let \mathcal{Y} denote the support of Y . Then, for any $\epsilon > 0$ it holds that*

$$\Pr_{y \leftarrow Y} [H_\infty(X|Y = y) \geq H_\infty(X) - \log |\mathcal{Y}| - \log(1/\epsilon)] \geq 1 - \epsilon.$$

4 A Central Lemma from [Raz05]

The following lemma is one of the main components that were used in the construction of two-sources-extractors in [Raz05]. We state the lemma for the special case of seeded-extractors, and prove it for completeness.

Lemma 4.1. *Let $D = 2^d$. Let Z_1, \dots, Z_D be 0-1 random variables that are ϵ -biased for linear tests of size k' that are constructed using n random bits. Define $\text{Ext}^{(1)}: \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}$ by $\text{Ext}^{(1)}(w, s) = Z_s(w)$, that is, $\text{Ext}^{(1)}(w, s)$ is the random variable Z_s , when using w as the value of the n bits needed to produce Z_1, \dots, Z_D . Then, for any $0 < \delta < 1/2$ and even integer $k \leq k'$ such that $k \cdot (1/\epsilon)^{1/k} \leq D^{1/2}$, the function $\text{Ext}^{(1)}$ is a $((1/2 + \delta) \cdot n, \gamma)$ -seeded-extractor, with*

$$\gamma = \left(\epsilon \cdot 2^{(1/2 - \delta)n+1} \right)^{1/k}.$$

Proof. Let W be a $(n, (1/2 + \delta) \cdot n)$ -source. Let S be a random variable that is uniformly distributed over $\{0, 1\}^d$ and is independent of W . We will show that the distribution of $\text{Ext}^{(1)}(W, S)$ is γ -close to uniform. As in [CG88], it is enough to consider the case where W is uniformly distributed over a set $W' \subseteq \{0, 1\}^n$ of size $2^{(1/2 + \delta)n}$. For every $w \in \{0, 1\}^n$ and $s \in \{0, 1\}^d$ denote

$$e(w, s) = (-1)^{Z_s(w)}.$$

Claim 4.2. *For any $r \in [k]$ and any different $s_1, \dots, s_r \in \{0, 1\}^d$,*

$$\sum_{w \in \{0, 1\}^n} \prod_{j=1}^r e(w, s_j) \leq \epsilon \cdot 2^n.$$

Proof.

$$\sum_{w \in \{0, 1\}^n} \prod_{j=1}^r e(w, s_j) = \sum_{w \in \{0, 1\}^n} \prod_{j=1}^r (-1)^{Z_{s_j}(w)} = \sum_{w \in \{0, 1\}^n} (-1)^{Z_{s_1}(w) \oplus \dots \oplus Z_{s_r}(w)},$$

and since $Z_{s_1}(w) \oplus \dots \oplus Z_{s_r}(w)$ is ϵ -biased, the last sum is at most $\epsilon \cdot 2^n$. ■

Denote by $\gamma(W, S)$ the expectation of $e(W, S)$. We will show that $|\gamma(W, S)| \leq \gamma$. Obviously, this means that $\text{Ext}^{(1)}(W, S)$ is γ -close to uniform, as required.

By the definition

$$2^{(1/2 + \delta)n} \cdot 2^d \cdot \gamma(W, S) = \sum_{w \in W'} \sum_{s \in \{0, 1\}^d} e(w, s).$$

Hence, by a convexity argument and since k is even,

$$\begin{aligned} 2^{(1/2 + \delta)n} \cdot \left(2^d \cdot \gamma(W, S) \right)^k &\leq \sum_{w \in W'} \left(\sum_{s \in \{0, 1\}^d} e(w, s) \right)^k \leq \\ &\sum_{w \in \{0, 1\}^n} \left(\sum_{s \in \{0, 1\}^d} e(w, s) \right)^k = \sum_{w \in \{0, 1\}^n} \sum_{s_1, \dots, s_k \in \{0, 1\}^d} \prod_{j=1}^k e(w, s_j) \\ &= \sum_{s_1, \dots, s_k \in \{0, 1\}^d} \sum_{w \in \{0, 1\}^n} \prod_{j=1}^k e(w, s_j). \end{aligned}$$

We will break the sum over $s_1, \dots, s_k \in \{0, 1\}^d$ into two sums. The first sum is over $s_1, \dots, s_k \in \{0, 1\}^d$ such that at least one s_j is different than all other elements in $\{s_1, \dots, s_k\}$, and the second sum is over $s_1, \dots, s_k \in \{0, 1\}^d$ such that every s_j is identical to at least one other element in $\{s_1, \dots, s_k\}$. The number of summands in the first sum is trivially bounded by $2^{d \cdot k}$, and by

Claim 4.2 each summand is bounded by $2^n \cdot \epsilon$. The number of summands in the second sum is bounded by $2^{d \cdot k/2} \cdot (k/2)^k$, and each summand is trivially bounded by 2^n . Hence,

$$\begin{aligned} 2^{(1/2+\delta) \cdot n} \cdot 2^{d \cdot k} \cdot \gamma(W, S)^k &\leq 2^n \cdot \epsilon \cdot 2^{d \cdot k} + 2^n \cdot 2^{d \cdot k/2} \cdot (k/2)^k \\ &\leq 2 \cdot 2^n \cdot \epsilon \cdot 2^{d \cdot k}, \end{aligned}$$

where the last inequality follows by the assumption that $k \cdot (1/\epsilon)^{1/k} \leq D^{1/2}$. That is,

$$|\gamma(W, S)| \leq \left(\epsilon \cdot 2^{(1/2-\delta)n+1} \right)^{1/k}.$$

■

5 A Simple Lemma about Graphs

The following simple lemma about graphs is another ingredient we need for the proof of the main theorem.

Lemma 5.1. *Let $G = (V, E)$ be a directed graph without self-loops. Assume that the out-degree of each vertex is exactly t , where parallel edges are allowed. Let $w: V \rightarrow \mathbb{R}$ be a weight function on the vertices of G . Denote by ω the average vertex weight, that is, $\omega = \frac{1}{|V|} \cdot \sum_{v \in V} w(v)$. Then, there exists a subset of the vertices $V' \subseteq V$, such that the induced graph $H = (V', E')$ of G by the set of vertices V' has the following properties:*

1. H is acyclic,
2. The average vertex weight of H is at least $\omega/(t+1)$, that is, $\frac{1}{|V'|} \cdot \sum_{v \in V'} w(v) \geq \omega/(t+1)$,
3. $|V'| \geq |V|/(t+1)$.

Proof. We construct H by a greedy algorithm. During the running of the algorithm, every vertex in V has one of the following statuses: *available*, *chosen* or *forbidden*. We say that a vertex is available if it has an available status. Similarly, we say that a vertex is chosen / forbidden if it has a chosen / forbidden status. For every vertex $v \in V$ we denote by $\text{status}(v)$ the status of the vertex v . For every vertex v , let $v^+ = \{u \in V : (v, u) \in E\}$. The greedy algorithm is defined as follows:

1. For every vertex $v \in V$ initialize $\text{status}(v) \leftarrow$ available.
2. While there exists an available vertex,
 - (a) Let v be an available vertex such that $w(v) \geq w(v')$ for any available vertex v' .
 - (b) Set $\text{status}(v) \leftarrow$ chosen.
 - (c) For every vertex $v' \in v^+$, if $\text{status}(v') =$ available set $\text{status}(v') \leftarrow$ forbidden.
3. Return $V' = \{v : \text{status}(v) = \text{chosen}\}$.

Assume for contradiction that H contains a cycle C , that is, C is a cycle of chosen vertices. Let v be the first chosen vertex in C . Let $v' \in v^+$ be the vertex that follows v in C . At the time v was chosen, v' was available, and so the algorithm set the status of v' to forbidden. A contradiction is then met as a forbidden vertex is never chosen and so v' cannot be in C .

We now prove property 2. Once the algorithm terminates, the status of every vertex is either chosen or forbidden. For every chosen vertex v , let $v^A \subseteq v^+$ be the set of vertices that were available at the time v was chosen. The vertices of the graph G can be partitioned as follows:

$$V = \bigcup_{v \in V'} (\{v\} \cup v^A). \quad (5.1)$$

By Equation (5.1) and by the fact that all (at most t) vertices in v^A have a weight which is no more than $w(v)$, we get

$$\begin{aligned} \sum_{v \in V} w(v) &= \sum_{v \in V'} \left(w(v) + \sum_{v' \in v^A} w(v') \right) \\ &\leq \sum_{v \in V'} \left(w(v) + w(v) \cdot |v^A| \right) \\ &\leq (t+1) \sum_{v \in V'} w(v). \end{aligned}$$

Hence,

$$\begin{aligned} \frac{1}{|V|} \cdot \sum_{v \in V'} w(v) &\geq \frac{1}{|V|} \cdot \frac{1}{t+1} \cdot \sum_{v \in V} w(v) \\ &\geq \frac{1}{t+1} \cdot \frac{1}{|V|} \cdot \sum_{v \in V} w(v) \\ &= \frac{\omega}{t+1}. \end{aligned}$$

This proves property 2. By Equation (5.1)

$$|V| = \sum_{v \in V'} (1 + |v^A|) \leq \sum_{v \in V'} (1 + t) = (t+1) \cdot |V'|,$$

which proves property 3. ■

6 A Conditional Parity Lemma

The following lemma is a generalization of the Parity Lemma (usually attributed to Vazirani, see for example [NN93]). A similar lemma appears in [LWZ11]. Let Z be a random variable over $\{0, 1\}^{m+n}$. Lemma 6.1 states that given the suffix of length n of Z , one can bound the statistical distance between the remaining length m prefix of Z and the uniform distribution in terms of appropriate biases. Setting $n = 0$ yields the Parity Lemma.

Lemma 6.1. *Let X be a random variable over $\{0, 1\}^m$. Let Y be a random variable over $\{0, 1\}^n$. Then,*

$$\|(X, Y) - (U_m, Y)\|_1 \leq \left(\sum_{\substack{\emptyset \neq \sigma \subseteq [m] \\ \tau \subseteq [n]}} \text{bias}(X_\sigma \oplus Y_\tau)^2 \right)^{1/2}.$$

We derive two corollaries from Lemma 6.1.

Corollary 6.2. Let X be a random variable over $\{0, 1\}^m$. Let Y be a random variable over $\{0, 1\}^n$. Then,

$$\|(X, Y) - (U_m, Y)\|_1 \leq \sum_{\substack{\emptyset \neq \sigma \subseteq [m] \\ \tau \subseteq [n]}} \text{bias}(X_\sigma \oplus Y_\tau).$$

Corollary 6.3. Let X be a random variable over $\{0, 1\}^m$. Let Y be a random variable over $\{0, 1\}^n$. Then,

$$\|(X, Y) - (U_m, Y)\|_1 \leq ((2^m - 1) \cdot 2^n)^{1/2} \cdot \max_{\substack{\emptyset \neq \sigma \subseteq [m] \\ \tau \subseteq [n]}} \text{bias}(X_\sigma \oplus Y_\tau).$$

Deriving both corollaries from Lemma 6.1 can be done by applying basic norms inequalities.

Proof of Lemma 6.1. Let $D \in \mathbb{R}^{2^{m+n}}$. We index the entries of D by strings of length $m+n$ bits. For $x \in \{0, 1\}^m$ and $y \in \{0, 1\}^n$, we define

$$D(xy) = \Pr((X, Y) = (x, y)) - \Pr((U_m, Y) = (x, y)),$$

where xy is the concatenation of x and y . By Parseval and basic norms inequalities,

$$\begin{aligned} \|(X, Y) - (U_m, Y)\|_1^2 &= \left(\sum_{\substack{x \in \{0, 1\}^m \\ y \in \{0, 1\}^n}} |\Pr((X, Y) = (x, y)) - \Pr((U_m, Y) = (x, y))| \right)^2 \\ &= \left(\sum_{\substack{x \in \{0, 1\}^m \\ y \in \{0, 1\}^n}} |D(xy)| \right)^2 \leq 2^{m+n} \cdot \sum_{\substack{x \in \{0, 1\}^m \\ y \in \{0, 1\}^n}} D(xy)^2 \\ &= 2^{m+n} \cdot \|D\|_2^2 = 2^{2(m+n)} \cdot \|\widehat{D}\|_2^2. \end{aligned} \quad (6.1)$$

Claim 6.4. For every $\sigma \subseteq [m]$ and $\tau \subseteq [n]$ ⁸,

$$|\widehat{D}(\sigma\tau)| = \begin{cases} 2^{-(m+n)} \cdot \text{bias}(X_\sigma \oplus Y_\tau), & \sigma \neq \emptyset; \\ 0, & \sigma = \emptyset. \end{cases}$$

Proof. For every $\sigma \subseteq [m]$ and $\tau \subseteq [n]$,

$$\begin{aligned} \widehat{D}(\sigma\tau) &= \frac{1}{2^{m+n}} \cdot \sum_{\substack{x \in \{0, 1\}^m \\ y \in \{0, 1\}^n}} (-1)^{\langle \sigma\tau, xy \rangle} \cdot D(xy) \\ &= \frac{1}{2^{m+n}} \cdot \sum_{\substack{x \in \{0, 1\}^m \\ y \in \{0, 1\}^n}} (-1)^{\langle \sigma\tau, xy \rangle} \cdot \Pr((X, Y) = (x, y)) \\ &\quad - \frac{1}{2^{m+n}} \cdot \sum_{\substack{x \in \{0, 1\}^m \\ y \in \{0, 1\}^n}} (-1)^{\langle \sigma\tau, xy \rangle} \cdot \Pr((U_m, Y) = (x, y)). \end{aligned}$$

We note that

$$\sum_{\substack{x \in \{0, 1\}^m \\ y \in \{0, 1\}^n}} (-1)^{\langle \sigma\tau, xy \rangle} \cdot \Pr((U_m, Y) = (x, y)) = \sum_{y \in \{0, 1\}^n} (-1)^{\langle \tau, y \rangle} \cdot \Pr(Y = y) \cdot \frac{1}{2^m} \sum_{x \in \{0, 1\}^m} (-1)^{\langle \sigma, x \rangle}.$$

⁸We slightly abuse notation and identify sets in $[m]$ with their characteristic vectors over $\{0, 1\}^m$.

For $\sigma \neq \emptyset$ it holds that $\sum_{x \in \{0,1\}^m} (-1)^{\langle \sigma, x \rangle} = 0$. Hence, for $\sigma \neq \emptyset$,

$$\begin{aligned} |\widehat{D}(\sigma\tau)| &= \left| \frac{1}{2^{m+n}} \cdot \sum_{\substack{x \in \{0,1\}^m \\ y \in \{0,1\}^n}} (-1)^{\langle \sigma\tau, xy \rangle} \cdot \Pr((X, Y) = (x, y)) \right| \\ &= \frac{1}{2^{m+n}} \cdot \text{bias}(X_\sigma \oplus Y_\tau). \end{aligned}$$

For $\sigma = \emptyset$ it holds that $\sum_{x \in \{0,1\}^m} (-1)^{\langle \sigma, x \rangle} = 2^m$. Therefore, for $\sigma = \emptyset$,

$$\begin{aligned} \widehat{D}(\sigma\tau) &= \frac{1}{2^{m+n}} \cdot \sum_{\substack{x \in \{0,1\}^m \\ y \in \{0,1\}^n}} (-1)^{\langle \tau, y \rangle} \cdot \Pr((X, Y) = (x, y)) \\ &\quad - \frac{1}{2^{m+n}} \cdot \sum_{y \in \{0,1\}^n} (-1)^{\langle \tau, y \rangle} \cdot \Pr(Y = y) \\ &= \frac{1}{2^{m+n}} \cdot \sum_{y \in \{0,1\}^n} (-1)^{\langle \tau, y \rangle} \left(\sum_{x \in \{0,1\}^m} \Pr((X, Y) = (x, y)) - \Pr(Y = y) \right) \\ &= 0. \end{aligned}$$

■

By Equation (6.1) and Claim 6.4,

$$\begin{aligned} \|(X, Y) - (U_m, Y)\|_1^2 &\leq 2^{2(m+n)} \cdot \|\widehat{D}\|_2^2 = 2^{2(m+n)} \cdot \sum_{\substack{\sigma \subseteq [m] \\ \tau \subseteq [n]}} \widehat{D}(\sigma\tau)^2 \\ &= \sum_{\substack{\emptyset \neq \sigma \subseteq [m] \\ \tau \subseteq [n]}} \text{bias}(X_\sigma \oplus Y_\tau)^2, \end{aligned}$$

which concludes the proof of the lemma. ■

7 Proof of Main Theorem

To ease the reading we restate the main theorem.

Theorem 7.1 (Theorem 1.7 – Restated). *For any integers n, d, m and t , and for any $0 < \delta < 1/2$ such that*

$$\begin{aligned} d &\geq \frac{23}{\delta} \cdot tm + 2 \log n, \\ n &\geq \frac{160}{\delta} \cdot tm, \\ \delta &\geq 10 \cdot \frac{\log(nd)}{n}, \end{aligned}$$

there exists an explicit $((1/2 + \delta) \cdot n, 2^{-m})$ - t -non-malleable extractor $\text{nmExt}: \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$.

Proof of Theorem 7.1. Let $D = m \cdot 2^d$. Let $k' = \lceil \delta n / 8 \rceil$. Let $\epsilon = 2^{-n/2+r}$, where $r = 1 + \log(k') + \log \log(D)$. The explicit construction we present is the extractor constructed in [Raz05]. We now describe it. Let Z_1, \dots, Z_D be 0-1 random variables that are ϵ -biased for linear tests of size k' that are constructed using n random bits. It is easy to verify that

$$n \geq 2 \cdot \lceil \log(1/\epsilon) + \log k' + \log \log D \rceil,$$

and so by [AGHP92] (see Section 3.3) such a construction is indeed possible.

We think of the set of indices $[D]$ as the set $\{(i, s) : i \in [m], s \in \{0, 1\}^d\}$. We define $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ by $\text{Ext}_i(w, s) = Z_{(i,s)}(w)$, where $\text{Ext}_i(w, s)$ denotes the i^{th} bit of $\text{Ext}(w, s)$. In other words, w is used to choose the point in the probability space and i, s are used to choose the variable from Z_1, \dots, Z_D that we evaluate.

Let S be a random variable uniformly distributed over $\{0, 1\}^d$. Assume for contradiction that Ext is not a $((1/2 + \delta) \cdot n, 2^{-m})$ - t -non-malleable-extractor. Then, there exists a source W of length n and min-entropy $(1/2 + \delta) \cdot n$, and a t -adversarial-function $\mathcal{A} : \{0, 1\}^d \rightarrow \{0, 1\}^{td}$ such that,

$$\begin{aligned} & \|(\text{Ext}(W, S), \text{Ext}(W, \mathcal{A}_1(S)), \dots, \text{Ext}(W, \mathcal{A}_t(S)), S) \\ & - (U_m, \text{Ext}(W, \mathcal{A}_1(S)), \dots, \text{Ext}(W, \mathcal{A}_t(S)), S)\|_1 > 2^{-m}. \end{aligned} \quad (7.1)$$

As in [CG88], we may assume that W is uniformly distributed over a set $W' \subseteq \{0, 1\}^n$ of size $2^{(1/2+\delta)n}$.

For every $s \in \{0, 1\}^d$ let X_s be the random variable $\text{Ext}(W, s)$. By Equation (7.1) and Claim 3.9,

$$\mathbb{E}_{s \sim S} \left[\|(X_s, X_{\mathcal{A}_1(s)}, \dots, X_{\mathcal{A}_t(s)}) - (U_m, X_{\mathcal{A}_1(s)}, \dots, X_{\mathcal{A}_t(s)})\|_1 \right] > 2^{-m}.$$

Hence, by Corollary 6.2,

$$\sum_{\substack{\emptyset \neq \sigma \subseteq [m] \\ \tau_1, \dots, \tau_t \subseteq [m]}} \mathbb{E}_{s \sim S} \left[\text{bias} \left((X_s)_\sigma \oplus \left(\bigoplus_{i \in [t]} (X_{\mathcal{A}_i(s)})_{\tau_i} \right) \right) \right] > 2^{-m}.$$

Let $\sigma^*, \tau_1^*, \dots, \tau_t^* \subseteq [m]$ be the indices of (one of) the largest summands in the above sum. For every $s \in \{0, 1\}^d$, let

$$Y_s = (X_s)_{\sigma^*} \oplus \left(\bigoplus_{i \in [t]} (X_{\mathcal{A}_i(s)})_{\tau_i^*} \right).$$

Then,

$$\mathbb{E}_{s \sim S} [\text{bias}(Y_s)] > 2^{-(t+2)m}.$$

Let $G = (V, E)$ be a directed graph with $V = \{0, 1\}^d$ and $E = \{(s, \mathcal{A}_j(s)) : s \in V, j \in [t]\}$. Since \mathcal{A} is a t -adversarial-function, G has no self-loops. Equip G with a weight function on the vertices $w : V \rightarrow \mathbb{R}$ that is defined as follows: For every $s \in V$, $w(s) = \text{bias}(Y_s)$.

By Lemma 5.1, there exists a subset $V' \subseteq V$ such that the induced graph, H , of G by V' has the properties mentioned in that lemma. In particular, by properties 2 and 3,

$$\mu \triangleq \frac{1}{|V'|} \cdot \sum_{s \in V'} \text{bias}(Y_s) > \frac{2^{-(t+2)m}}{t+1},$$

and

$$|V'| \geq \frac{|V|}{t+1} = \frac{2^d}{t+1}.$$

By Claim 3.8, there exists a subset $S' \subseteq V'$ such that

$$|S'| \geq \frac{\mu}{2} \cdot |V'| \geq \frac{2^{d-(t+2)m-1}}{(t+1)^2},$$

and for all $s \in S'$

$$\text{bias}(Y_s) \geq \frac{\mu}{2} > \frac{2^{-(t+2)m-1}}{t+1}.$$

Let $S'_0 = \{s \in S' : \Pr[Y_s = 0] \geq 1/2\}$. Let $S'_1 = S \setminus S'_0$. There exists $b \in \{0, 1\}$ such that $|S'_b| \geq |S'|/2$. Denote S'_b by S'' . Then,

$$|S''| \geq \frac{|S'|}{2} \geq \frac{2^{d-(t+2)m-2}}{(t+1)^2},$$

and for all $s \in S''$

$$\Pr(Y_s = b) - \Pr(Y_s \neq b) > \frac{2^{-(t+2)m-1}}{t+1}. \quad (7.2)$$

Define a random variable $Y_{S''}$ over $\{0, 1\}$ as follows: To sample a bit from $Y_{S''}$, uniformly sample a string s from S'' , and then independently sample a string w uniformly from W' . The sampled value is $Y_s(w)$. We have that

$$\begin{aligned} \text{bias}(Y_{S''}) &= |\Pr(Y_{S''} = 0) - \Pr(Y_{S''} = 1)| \\ &= \frac{1}{|S''|} \cdot \left| \sum_{s \in S''} \Pr(Y_s = 0) - \Pr(Y_s = 1) \right| \\ &= \frac{1}{|S''|} \cdot \sum_{s \in S''} |\Pr(Y_s = 0) - \Pr(Y_s = 1)| \\ &> \frac{2^{-(t+2)m-1}}{t+1}, \end{aligned} \quad (7.3)$$

where the inequality and equality before it follow by Equation (7.2).

For every $s \in S''$, let

$$Y'_s = \left(\bigoplus_{j \in \sigma^*} Z_{(j,s)} \right) \oplus \left(\bigoplus_{i \in [t]} \left(\bigoplus_{j \in \tau_i^*} Z_{(j, \mathcal{A}_i(s))} \right) \right).$$

Claim 7.2. *The set of random variables $\{Y'_s\}_{s \in S''}$ ϵ -fools linear tests of size $k'/((t+1)m)$.*

Proof. Let $A \subseteq S''$ be a nonempty set of size at most $\ell \triangleq k'/((t+1)m)$. Let $Y_A = \bigoplus_{s \in A} Y'_s$. We note that Y_A is a linear combination of random variables from $(Z_{(i,s)})_{i \in [m], s \in \{0,1\}^d}$, composed of at most k' summands. Since $(Z_{(i,s)})_{i \in [m], s \in \{0,1\}^d}$ ϵ -fools linear tests of size k' , it is enough to show that this linear combination is non-trivial, that is, it suffices to prove that Y_A is not the constant 0 random variable.

Assume for contradiction that $Y_A = 0$. Let e be an arbitrary element in σ^* . Such an element exists as $\sigma^* \neq \emptyset$. Let s_1 be an arbitrary element in A . Such an element exists as $A \neq \emptyset$. The random variable $Z_{(e,s_1)}$ is therefore a summand in Y_A .

By the assumption that $Y_A = 0$, $Z_{(e,s_1)}$ must appear an even number of times as a summand in Y_A . In particular, $Z_{(e,s_1)}$ must appear at least one more time as a summand in Y_A . Therefore, there exists some $s_2 \in A$ and $i_2 \in [t]$ such that $\mathcal{A}_{i_2}(s_2) = s_1$.

Since \mathcal{A} is an adversarial-function, $s_2 \neq s_1$ and so the random variable $Z_{(e,s_2)}$ is a summand in Y_A which is different than $Z_{(e,s_1)}$. Following the same logic as above, since $Y_A = 0$, the random variable $Z_{(e,s_2)}$ must appear an even number of times as a summand in Y_A . In particular, it must appear at least one more time. Hence, there exists some $s_3 \in A$ and $i_3 \in [t]$ such that $\mathcal{A}_{i_3}(s_3) = s_2$.

We continue this way to get a sequence s_1, s_2, s_3, \dots of elements of A until two elements in the sequence are equal. Since A is finite, this is bound to happen. However, in such case, a directed cycle in the graph H is implied, contradicting its acyclicity. ■

Let k be the largest even integer that is not larger than $k'/((t+1)m)$.

Claim 7.3.

$$\frac{1}{10} \cdot \frac{\delta n}{(t+1)m} \leq k \leq \frac{1}{5} \cdot \frac{\delta n}{(t+1)m}$$

Proof. As k is the largest even integer that is not larger than $k'/((t+1)m)$,

$$k \in \left\{ \left\lfloor \frac{k'}{(t+1)m} \right\rfloor - 1, \left\lfloor \frac{k'}{(t+1)m} \right\rfloor \right\}.$$

Therefore,

$$k \leq \left\lfloor \frac{k'}{(t+1)m} \right\rfloor \leq \frac{k'}{(t+1)m} = \frac{\lceil \delta n / 8 \rceil}{(t+1)m} \leq \frac{1}{8} \cdot \frac{\delta n}{(t+1)m} + \frac{1}{(t+1)m} \leq \frac{1}{5} \cdot \frac{\delta n}{(t+1)m},$$

where the last inequality follows by the assumption that $\delta \geq 160 \cdot tm/n \geq 40/(3 \cdot n)$. As for the lower bound on k ,

$$k \geq \left\lfloor \frac{k'}{(t+1)m} \right\rfloor - 1 \geq \frac{k'}{(t+1)m} - 2 = \frac{\lceil \delta n / 8 \rceil}{(t+1)m} - 2 \geq \frac{1}{8} \cdot \frac{\delta n}{(t+1)m} - 2 \geq \frac{1}{10} \cdot \frac{\delta n}{(t+1)m},$$

where, again, the last inequality follows by the assumption that $\delta \geq 160 \cdot tm/n \geq 80(t+1)m/n$. ■

We apply Lemma 4.1 on the random variables $\{Y'_s\}_{s \in S''}$ ⁹. The following claim confirms that the assumption of Lemma 4.1 is indeed met with the k defined above.

Claim 7.4.

$$k \cdot \left(\frac{1}{\epsilon}\right)^{1/k} \leq \left(\frac{2^{d-(t+2)m-2}}{(t+1)^2}\right)^{1/2}. \quad (7.4)$$

Proof. Taking the $\log(\cdot)$ of both sides of Equation (7.4) and rearranging the terms, we see it is enough to prove that

$$d \geq (t+2)m + 2 \log k + \frac{2}{k} \log \frac{1}{\epsilon} + 2 \log(t+1) + 2. \quad (7.5)$$

By Claim 7.3

$$\frac{2}{k} \cdot \log \frac{1}{\epsilon} = \frac{2}{k} \cdot \left(\frac{n}{2} - r\right) \leq \frac{n}{k} \leq \frac{10(t+1)m}{\delta}, \quad (7.6)$$

and

$$\log k \leq \log \left(\frac{\delta n}{5(t+1)m}\right) \leq \log \left(\frac{n}{20}\right). \quad (7.7)$$

⁹For simplicity of presentation we assume $|S''|$ is a power of 2. The exact same result can be obtained regardless of this assumption.

By Equations (7.5), (7.6) and (7.7), it is enough to show that

$$d \geq \left(\frac{10(t+1)}{\delta} + t + 2 \right) m + 2 \log n + 2 \log(t+1) + 2 - 2 \log 20.$$

Since for all $t \geq 1$

$$\frac{22}{\delta} \cdot t \geq \frac{10(t+1)}{\delta} + t + 2$$

and

$$2 \log t \geq 2 \log(t+1) + 2 - 2 \log 20,$$

it is enough to prove that

$$d \geq \frac{22}{\delta} \cdot tm + 2 \log n + 2 \log t.$$

The above equation holds as we assume

$$d \geq \frac{23}{\delta} \cdot tm + 2 \log n. \quad \blacksquare$$

Consider the weak-source W . By Lemma 4.1, the distribution of $\text{Ext}^{(1)}(W, S'')$ is γ -biased, for $\gamma = (\epsilon \cdot 2^{1+(1/2-\delta)n})^{1/k} = 2^{(1+r-\delta n)/k}$. However, we note that $\text{Ext}^{(1)}(W, S'')$ has the same distribution as $Y_{S''}$. In particular, both random variables have the same bias. Equation (7.3) yields that

$$2^{(1+r-\delta n)/k} \geq \text{bias}(\text{Ext}^{(1)}(W, S'')) = \text{bias}(Y_{S''}) > \frac{2^{-(t+2)m-1}}{t+1}. \quad (7.8)$$

We conclude the proof of Theorem 7.1 by the following claim, that stands in contradiction to Equation (7.8).

Claim 7.5.

$$2^{(1+r-\delta n)/k} < \frac{2^{-(t+2)m-1}}{t+1}.$$

Proof. It is enough to prove that

$$\frac{\delta n}{k} > (t+2)m + \log(4(t+1)) + \frac{r}{k}.$$

By Claim 7.3, it is enough to show that

$$(4t+3)m > \log(4(t+1)) + \frac{r}{k}.$$

Since for all $n \geq 2$ (indeed $n \geq 160 \cdot tm/\delta \geq 320$),

$$k' = \left\lceil \frac{\delta n}{8} \right\rceil \leq \frac{n}{2},$$

we have that

$$r = 1 + \log k' + \log \log D = \log(2k'(d + \log m)) \leq \log(ndm). \quad (7.9)$$

By Equation (7.9) and Claim 7.3 we have that

$$\frac{r}{k} \leq 10(t+1)m \frac{\log(ndm)}{\delta n}.$$

It is therefore enough to prove that

$$(4t + 3)m > \log(4(t + 1)) + 10(t + 1)m \frac{\log(ndm)}{\delta n}. \quad (7.10)$$

To prove Equation (7.10) for all $t \geq 1$, it is enough to show that

$$m \geq \frac{3}{7} + \frac{20}{7} \cdot m \cdot \frac{\log(ndm)}{\delta n},$$

which holds if

$$n \geq \frac{5}{\delta} \cdot \log(ndm). \quad (7.11)$$

Since $m < n$ (indeed, by the second assumption of Theorem 7.1, $m \leq \delta n / (160t) \leq n / 320$), Equation (7.11) holds by the third assumption of Theorem 7.1. ■

8 The Privacy Amplification Protocol

In the section we present the protocol that is obtained by instantiating the Dodis-Wichs protocol [DW09] with our non-malleable extractor. Given the length n of the weak source and parameters k , ϵ' , ϵ_{nmExt} , ϵ_{Ext} , and ϵ_{MAC} that we will fix later, the protocol relies on the following building blocks:

1. A non-malleable $(k, \epsilon_{\text{nmExt}})$ -extractor $\text{nmExt} : \{0, 1\}^n \times \{0, 1\}^{d_1} \rightarrow \{0, 1\}^\ell$ (see Definition 1.3).
2. A strong $(k - (d_1 + \ell) - \log(1/\epsilon'), \epsilon_{\text{Ext}})$ -extractor $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^{d_2} \rightarrow \{0, 1\}^m$ (see Definition 1.1).
3. An ϵ_{MAC} -secure message authentication code $\{\text{MAC}_{\text{key}} : \{0, 1\}^{d_2} \rightarrow \{0, 1\}^\tau\}_{\text{key} \in \{0, 1\}^\ell}$ (see Definition 3.6).

The protocol is described in Figure 1. The following theorem was proved in [DW09], and we provide here its proof for completeness.

Theorem 8.1. *Let nmExt , Ext and MAC be as specified above. Then for any integers n and $k < n$, the protocol in Figure 1 is a 2-round (n, k, m, ϵ) -privacy amplification protocol, with communication complexity $d_1 + d_2 + \tau$, where $\epsilon = \max\{\epsilon' + \epsilon_{\text{Ext}}, \epsilon_{\text{nmExt}} + \epsilon_{\text{MAC}}\}$.*

For obtaining explicit protocols we instantiate the building blocks nmExt , and MAC with those provided by Theorem 1.6, and Theorem 3.7, respectively. In addition, we instantiate the strong extractor Ext by either one of those provided by Theorem 3.2, Theorem 3.3, or Theorem 3.4. We obtain the following two theorems:

Theorem 8.2. *For any integer n , constant $\delta > 0$, and security parameter $\epsilon = 2^{-O(n)}$, there exists an explicit and efficient 2-round $(n, (1/2 + \delta)n, m, \epsilon)$ -privacy amplification protocol with entropy loss $O(\log n + \log(1/\epsilon))$, and communication complexity $O(\min\{\log^2 n + \log n \cdot \log(1/\epsilon), n\})$.*

Theorem 8.3. *For any integer n , constants δ and β such that $1/2 + \delta > \beta > 0$, and security parameter $\epsilon = 2^{-O(n)}$, there exists an explicit and efficient 2-round $(n, (1/2 + \delta)n, m, \epsilon)$ -privacy amplification protocol with entropy loss $\beta n + O(\log n + \log(1/\epsilon))$, and communication complexity $O(\log n + \log(1/\epsilon))$.*

Shared input: Alice and Bob share a sample from an (n, k) -source W .

The protocol:

1. Alice samples $Y \leftarrow \{0, 1\}^{d_1}$ uniformly at random, sends it to Bob, and computes $\text{key} = \text{nmExt}(W, Y)$.
2. Denote by Y' the value received by Bob, who then computes $\text{key}' = \text{nmExt}(W, Y')$.
3. Bob samples $S' \leftarrow \{0, 1\}^{d_2}$ uniformly at random, computes $\sigma' = \text{MAC}_{\text{key}'}(S')$, and sends the pair (S', σ') to Alice.
4. Bob reaches the **KeyDerived** state and outputs $R_B = \text{Ext}(W, S')$.
5. Denote by (S, σ) the pair received by Alice. If $\sigma = \text{MAC}_{\text{key}}(S)$ then Alice reaches the **KeyConfirmed** state and outputs $R_A = \text{Ext}(W, S)$. Otherwise, Alice outputs $R_A = \perp$.

Figure 1: The Dodis-Wichs privacy amplification protocol.

In the remainder of this section we prove Theorems 8.1, 8.2, and 8.3.

Proof of Theorem 8.1. The correctness of the protocol and the parameters specified in the theorem follow directly from the description of the protocol. Thus, it only remains to argue the privacy and authenticity properties of the protocol. Since no assumptions are made on the computational capabilities of Eve, we assume without loss of generality that Eve is deterministic. Specifically, this implies that the value Y' is a deterministic function of the value Y , and then the pair (S, σ) is a deterministic function of the vector (Y, S', σ') . Therefore, without loss of generality, we refer to the view of Eve in the protocol as the vector $V_E = (Y, S', \sigma')$.

We argue the privacy property of the protocol in Lemma 8.4 and the authenticity property of the protocol in Lemma 8.5. For arguing the privacy of the protocol note that Bob always reaches the **KeyDerived** state (i.e., $\Pr[\text{KeyDerived}_A] = 0$ and $\Pr[\text{KeyDerived}_B] = 1$). Therefore we only need to bound $\text{SD}((R_B, V_E \mid \text{KeyDerived}_B), (U_m, V_E \mid \text{KeyDerived}_B))$.

Lemma 8.4 (Privacy). *It holds that*

$$\text{SD}((R_B, V_E \mid \text{KeyDerived}_B), (U_m, V_E \mid \text{KeyDerived}_B)) \leq \epsilon' + \epsilon_{\text{Ext}} .$$

Proof. The protocol specifies that Bob always reaches the **KeyDerived** state and outputs the value $R_B = \text{Ext}(W, S')$. In addition, recall that Eve's view consists of $V_E = (Y, S', \sigma')$. Therefore,

$$\begin{aligned} & \text{SD}((R_B, V_E \mid \text{KeyDerived}_B), (U_m, V_E \mid \text{KeyDerived}_B)) \\ &= \text{SD}((R_B, V_E), (U_m, V_E)) \\ &= \text{SD}((\text{Ext}(W, S'), Y, S', \sigma'), (U_m, Y, S', \sigma')) \\ &\leq \text{SD}((\text{Ext}(W, S'), Y, S', \text{key}'), (U_m, Y, S', \text{key}')) , \end{aligned}$$

where the last inequality follows from the fact that the value σ' is a deterministic function of the values S' and key' . Thus, since the value S' is uniformly distributed and independent of W , Y , and key' , in order to complete the argument we only need to prove that with high probability W has sufficient min-entropy given the pair (Y, key') . This follows from the fact that the latter pair (Y, key') is of total length $d_1 + \ell$ bits. Formally, Lemma 3.10 implies that with probability $1 - \epsilon'$ over the choice of $(y, \kappa') \leftarrow (Y, \text{key}')$ it holds that

$$\begin{aligned} H_\infty(W \mid Y = y, \text{key}' = \kappa') &\geq H_\infty(W) - (d_1 + \ell) - \log(1/\epsilon') \\ &\geq k - (d_1 + \ell) - \log(1/\epsilon') . \end{aligned}$$

In turn, the fact that Ext is a strong $(k - (d_1 + \ell) - \log(1/\epsilon'), \epsilon_{\text{Ext}})$ -extractor yields

$$\begin{aligned} & \text{SD}((R_B, V_E \mid \text{KeyDerived}_B), (U_m, V_E \mid \text{KeyDerived}_B)) \\ & \leq \text{SD}((\text{Ext}(W, S'), Y, S', \text{key}'), (U_m, Y, S', \text{key}')) \\ & \leq \epsilon' + \epsilon_{\text{Ext}} . \end{aligned}$$

■

Lemma 8.5 (Authenticity). *It holds that*

$$\Pr[(\text{KeyConfirmed}_A \vee \text{KeyConfirmed}_B) \wedge R_A \neq R_B] \leq \epsilon_{\text{nmExt}} + \epsilon_{\text{MAC}} .$$

Proof. The protocol specifies that only Alice may reach the KeyConfirmed state, and therefore

$$\Pr[(\text{KeyConfirmed}_A \vee \text{KeyConfirmed}_B) \wedge R_A \neq R_B] = \Pr[\text{KeyConfirmed}_A \wedge R_A \neq R_B] .$$

We now consider two cases: one in which $S = S'$ (i.e., Eve does not modify S') and the other in which $S \neq S'$ (i.e., Eve does modify S').

Case 1: $S = S'$. In this case Alice either reaches the KeyConfirmed state and outputs $R_A = \text{Ext}(W, S) = \text{Ext}(W, S') = R_B$ or does not reach the KeyConfirmed state and outputs $R_A = \perp$. This implies that

$$\Pr[\text{KeyConfirmed}_A \wedge R_A \neq R_B \mid S = S'] = 0 .$$

Case 2: $S \neq S'$. For Alice to reach the KeyConfirmed state Eve must compute a valid authentication tag σ on S with respect to the authentication key key . This implies that

$$\begin{aligned} \Pr[\text{KeyConfirmed}_A \wedge R_A \neq R_B \mid S \neq S'] & \leq \Pr[\text{KeyConfirmed}_A \mid S \neq S'] \\ & \leq \Pr[\sigma = \text{MAC}_{\text{key}}(S) \mid S \neq S'] \end{aligned}$$

For analyzing this case we consider two subcases: one in which $Y' = Y$ (i.e., Eve does not modify Y) and the other in which $Y' \neq Y$ (i.e., Eve does modify Y). We show that

$$\begin{aligned} \Pr[\sigma = \text{MAC}_{\text{key}}(S) \mid S \neq S' \wedge Y' = Y] & \leq \epsilon_{\text{nmExt}} + \epsilon_{\text{MAC}} \\ \Pr[\sigma = \text{MAC}_{\text{key}}(S) \mid S \neq S' \wedge Y' \neq Y] & \leq \epsilon_{\text{nmExt}} + \epsilon_{\text{MAC}} . \end{aligned}$$

Case 2.1: $Y = Y'$. In this case Alice and Bob share the same authentication key $\text{key} = \text{nmExt}(W, Y) = \text{nmExt}(W, Y') = \text{key}'$, which we will show to be statistically-close to a uniform authentication key due to the fact that the view, V_E , of Eve cannot significantly reduce the min-entropy of W . Therefore, the security of the message authentication code guarantees that even after viewing the authentication tag $\sigma' = \text{MAC}_{\text{key}}(S')$ she has only a negligible probability of computing a valid authentication tag $\sigma = \text{MAC}_{\text{key}}(S)$ for any $S \neq S'$.

Formally, the facts that: (1) nmExt is in particular a strong $(k, \epsilon_{\text{nmExt}})$ -extractor, (2) S' is independent of W and Y , and (3) Y is uniformly distributed, guarantee that

$$\text{SD}((\text{key}, Y, S'), (U_\ell, Y, S')) \leq \epsilon_{\text{nmExt}} .$$

Therefore, the probability that Eve (after viewing $\sigma' = \text{MAC}_{\text{key}}(S')$) computes a valid authentication tag σ on any $S \neq S'$ with respect to authentication key key differs by at most ϵ_{nmExt}

from the probably ϵ_{MAC} that Eve computes a valid authentication tag σ for any $S \neq S'$ with respect to a uniformly and independently chosen authentication key:

$$\Pr [\sigma = \text{MAC}_{\text{key}}(S) \mid S \neq S' \wedge Y' = Y] \leq \epsilon_{\text{nmExt}} + \epsilon_{\text{MAC}} .$$

Case 2.2: $Y \neq Y'$. In this case Eve views an authentication tag $\sigma' = \text{MAC}_{\text{key}'}(S')$ with respect to the authentication key $\text{key}' = \text{nmExt}(W, Y')$, and has to compute an authentication tag $\sigma = \text{MAC}_{\text{key}}(S)$ for some $S \neq S'$ with respect to the authentication key $\text{key} = \text{nmExt}(W, Y)$. The property of the non-malleable extractor nmExt guarantees that even if Eve was in fact given the authentication key key' then from her point of view, the authentication key key is ϵ_{nmExt} -close to an independently and uniformly chosen key. For such a key Eve can compute such an authentication tag σ with probability at most ϵ_{MAC} (and this in fact holds even if $S = S'$).

Formally, the facts that: (1) nmExt is a non-malleable $(k, \epsilon_{\text{nmExt}})$ -extractor, (2) S' is independent of W, Y , and key' , and (3) Y is uniformly distributed, guarantee that

$$\text{SD}((\text{key}, Y, \text{key}', S'), (U_\ell, Y, \text{key}', S')) \leq \epsilon_{\text{nmExt}}$$

which implies

$$\Pr [\sigma = \text{MAC}_{\text{key}}(S) \mid S \neq S' \wedge Y' \neq Y] \leq \epsilon_{\text{nmExt}} + \epsilon_{\text{MAC}} .$$

Combining cases 2.1 and 2.2 we obtain

$$\Pr [\sigma = \text{MAC}_{\text{key}}(S) \mid S \neq S'] \leq \epsilon_{\text{nmExt}} + \epsilon_{\text{MAC}} .$$

■
■

Proof of Theorem 8.2. Given the length n of the weak source, the constant δ , and the security parameter ϵ , we let $\epsilon' = \epsilon_{\text{Ext}} = \epsilon_{\text{nmExt}} = \epsilon_{\text{MAC}} = \epsilon/2$, and instantiate our protocol with the following explicit constructions:

1. Theorem 1.6 guarantees a non-malleable $((1/2 + \delta)n, \epsilon_{\text{nmExt}})$ -extractor $\text{nmExt} : \{0, 1\}^n \times \{0, 1\}^{d_1} \rightarrow \{0, 1\}^\ell$, where $d_1 = \Theta(\log n + \log(1/\epsilon))$ and $\ell = \Theta(\log n + \log(1/\epsilon))$.
2. Theorem 3.2 guarantees a strong $((1/2 + \delta)n - (d_1 + \ell) - \log(1/\epsilon'), \epsilon_{\text{Ext}})$ -extractor $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^{d_2} \rightarrow \{0, 1\}^m$, where $d_2 = \Theta(n)$ and $m = (1/2 + \delta)n - \Theta(\log n + \log(1/\epsilon))$. In addition, Theorem 3.4 guarantees such a strong extractor where $d_2 = \Theta(\log^2 n + \log n \cdot \log(1/\epsilon))$, and therefore we in fact have $d_2 = \Theta(\min \{\log^2 n + \log n \cdot \log(1/\epsilon), n\})$
3. Theorem 3.7 guarantees an ϵ_{MAC} -secure MAC $\{\text{MAC}_{\text{key}} : \{0, 1\}^{d_2} \rightarrow \{0, 1\}^\tau\}_{\text{key} \in \{0, 1\}^\ell}$, where $\tau = \Theta(\log n + \log(1/\epsilon))$ and $\ell = \Theta(\log n + \log(1/\epsilon))$.

By combining the above explicit constructions, the resulting privacy amplification protocol has security parameter $\max\{\epsilon' + \epsilon_{\text{Ext}}, \epsilon_{\text{nmExt}} + \epsilon_{\text{MAC}}\} = \epsilon$, entropy loss $(1/2 + \delta)n - m = \Theta(\log n + \log(1/\epsilon))$, and communication complexity $d_1 + d_2 + \tau = \Theta(\min \{\log^2 n + \log n \cdot \log(1/\epsilon), n\})$. ■

Proof of Theorem 8.3. The proof is identical to the proof of Theorem 8.2, where the only difference is that we instantiate the strong extractor Ext using the one provided by Theorem 3.3. Specifically, for any constants δ and β such that $1/2 + \delta > \beta > 0$, Theorem 3.3 guarantees a strong $((1/2 + \delta)n - (d_1 + \ell) - \log(1/\epsilon'), \epsilon_{\text{Ext}})$ -extractor $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^{d_2} \rightarrow \{0, 1\}^m$, where $d_2 = \Theta(\log n + \log(1/\epsilon))$ and $m = \left(1 - \frac{\beta}{1/2 + \delta}\right) ((1/2 + \delta)n - (d_1 + \ell) - \log(1/\epsilon'))$. In turn, the resulting privacy amplification protocol has entropy loss $(1/2 + \delta)n - m = \beta n + \Theta(\log n + \log(1/\epsilon))$, and communication complexity $d_1 + d_2 + \tau = \Theta(\log n + \log(1/\epsilon))$. ■

References

- [AGHP92] N. Alon, O. Goldreich, J. Håstad, and R. Peralta. Simple construction of almost k-wise independent random variables. *Random Structures and Algorithms*, 3(3):289–304, 1992.
- [BBCM95] C. H. Bennett, G. Brassard, C. Crépeau, and U. M. Maurer. Generalized privacy amplification. *IEEE Transactions on Information Theory*, 41(6):1915–1923, 1995.
- [BBR88] C. H. Bennett, G. Brassard, and J. Robert. Privacy amplification by public discussion. *SIAM Journal on Computing*, 17(2):210–229, 1988.
- [CG88] B. Chor and O. Goldreich. Unbiased bits from sources of weak randomness and probabilistic communication complexity. *SIAM Journal on Computing*, 17(2):230–261, 1988.
- [CKOR10] N. Chandran, B. Kanukurthi, R. Ostrovsky, and L. Reyzin. Privacy amplification with asymptotically optimal entropy loss. In *Proceedings of the 42nd Annual ACM Symposium on Theory of Computing*, pages 785–794, 2010.
- [DKRS06] Y. Dodis, J. Katz, L. Reyzin, and A. Smith. Robust fuzzy extractors and authenticated key agreement from close secrets. In *Advance in Cryptology – CRYPTO ’06*, pages 232–250, 2006.
- [DW09] Y. Dodis and D. Wichs. Non-malleable extractors and symmetric key cryptography from weak secrets. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing*, pages 601–610, 2009.
- [GUV09] V. Guruswami, C. Umans, and S. Vadhan. Unbalanced expanders and randomness extractors from Parvaresh–Vardy codes. *Journal of the ACM*, 56(4):1–34, 2009.
- [ILL89] R. Impagliazzo, L. A. Levin, and M. Luby. Pseudo-random generation from one-way functions. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing*, pages 12–24, 1989.
- [KR09] B. Kanukurthi and L. Reyzin. Key agreement from close secrets over unsecured channels. In *Advance in Cryptology – EUROCRYPT ’09*, pages 206–223, 2009.
- [LWZ11] X. Li, T. D. Wooley, and D. Zuckerman. Non-malleable extractors via character sums. Manuscript available at <http://arxiv.org/abs/1102.5415>, 2011.
- [Mau92] U. M. Maurer. Protocols for secret key agreement by public discussion based on common information. In *Advance in Cryptology – CRYPTO ’92*, pages 461–470, 1992.

- [Mau97] U. M. Maurer. Information-theoretically secure secret-key agreement by NOT authenticated public discussion. In *Advance in Cryptology – EUROCRYPT '97*, pages 209–225, 1997.
- [MW97] U. M. Maurer and S. Wolf. Privacy amplification secure against active adversaries. In *Advance in Cryptology – CRYPTO '97*, pages 307–321, 1997.
- [MW03] U. M. Maurer and S. Wolf. Secret-key agreement over unauthenticated public channels III: Privacy amplification. *IEEE Transactions on Information Theory*, 49(4):839–851, 2003.
- [NN93] J. Naor and M. Naor. Small-bias probability spaces: Efficient constructions and applications. *SIAM Journal on Computing*, 22(4):838–856, 1993.
- [NZ96] N. Nisan and D. Zuckerman. Randomness is linear in space. *Journal of Computer and System Sciences*, 52(1):43–52, 1996.
- [Raz05] R. Raz. Extractors with weak random seeds. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 11–20, 2005.
- [RW03] R. Renner and S. Wolf. Unconditional authenticity and privacy from an arbitrarily weak secret. In *Advance in Cryptology – CRYPTO '03*, pages 78–95, 2003.
- [Wol98] S. Wolf. Strong security against active attacks in information-theoretic secret-key agreement. In *Advance in Cryptology – ASIACRYPT '98*, pages 405–419, 1998.